Permanent and Transitory Components of GNP and Stock Prices

John H. Cochrane

The Quarterly Journal of Economics, Volume 109, Issue 1 (Feb., 1994), 241-265.

Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28199402%29109%3A1%3C241%3APATCOG%3E2.0.CO%3B2-K

PERMANENT AND TRANSITORY COMPONENTS
OF GNP AND STOCK PRICES*

JOHN H. COCHRANE

This paper uses two-variable autoregressions to characterize transitory components in GNP and stock prices. Shocks to GNP holding consumption constant are almost entirely transitory, and account for large fractions of the variance of GNP growth. If consumption does not change, consumers must think that any GNP change is transitory. The facts that the consumption/GNP ratio forecasts GNP growth and that consumption is nearly a random walk drive this result. An implication is that consumption provides a good estimate of the "trend" in GNP. Prices and dividends behave similarly: shocks to prices holding dividends constant are almost entirely transitory.

I. INTRODUCTION

A recent voluminous literature has examined the long-run properties of GNP and stock prices, with surprising results. One would expect that GNP reverts to "potential GNP" or some other trend following a shock. Yet many studies have found no mean-reversion, especially in postwar U. S. GNP.[1] This view obviously challenges a broad spectrum of macroeconomic theories designed to produce and understand transitory fluctuations.

Conventional wisdom once held that stock prices are random walks (martingales) that display *no* mean-reversion. Yet a large number of recent studies have instead *found* mean-reversion or transitory components in stock prices.[2] Depending on the author's tastes, these findings are interpreted as evidence for "fads"— irrational investor behavior—or as evidence for as-yet unmodeled time-variation in real investment opportunities.

To examine long-run properties of GNP and stock prices, I focus on simple two-variable autoregressions: GNP and consump-

1. This view has a long history, starting at least with Fisher [1925] and McCulloch [1975]. More recently, among others, see Nelson and Plosser [1982], Campbell and Mankiw [1987], and Cogley [1990]. Clark [1987] and Cochrane [1988] find some evidence for univariate mean reversion in GNP, but they are the exception.

2. Fama [1991] reviews predictability of stock returns.

tion growth are regressed on their lags and the lagged consumption/ GNP ratio; stock returns and dividend growth are regressed on their lags and the lagged dividend/price ratio.

The results of these VARs confirm and characterize statistically and economically important transitory components in GNP and stock prices. For example, transitory shocks account for an estimated 70–80 percent of the variance of GNP growth and 57 percent of the variance of annual stock returns.

Why do I find such strong transitory variation in postwar U. S. GNP data, where others have not? Transitory variation requires long-horizon forecastability. If GNP growth is unforecastable, GNP follows a random walk. If a negative shock signals many quarters of above-average growth, then GNP contains a transitory component that reverts to its mean. Most of the GNP literature is based on *univariate* forecasts. It documents that lagged GNP growth forecasts future GNP growth poorly. The consumption/ GNP ratio is a much more potent forecaster of long-horizon GNP growth. Thus, it can imply much larger transitory variation.

The forecastability of GNP growth can be documented by regressing it on a variety of variables. The consumption/GNP ratio is special, because it is stable over long time periods (consumption and GNP are *cointegrated*), while consumption is nearly a random walk. As a result, if GNP is more than its customary ratio to consumption, GNP must be forecast to decline until the ratio is reestablished. In this way, consumption defines the "trend" in GNP.

It is natural to interpret these features of the data via the simple permanent income model. The model predicts that consumption is a random walk and that consumption and total income are cointegrated. If consumption does not change, consumers must think any fluctuation in GNP is transitory. Thus, by observing consumption, we separate GNP into permanent and transitory components, as viewed by consumers. Of course, the permanent income model can be statistically rejected, since consumption is not exactly a random walk. Nonetheless, it is a good approximation that helps to describe and interpret the VAR results.

Similar points hold for dividends and stock prices. Dividend/ price ratios forecast returns better than lagged returns, while dividends are nearly a random walk. Therefore, price shocks with constant dividends are almost entirely transitory. As detailed below, the present value model plus the hypothesis that managers smooth dividends predicts and interprets these features of the data. In addition, this model suggests that we interpret the price

shocks as "discount rate" shocks and the dividends shocks as "earnings" shocks.

II. MEASURING THE PERSISTENCE OF GNP SHOCKS IN POSTWAR U. S. DATA

A. VAR and Impulse-Response Functions

Table I presents a vector autoregression of log GNP and log nondurable + services consumption growth on their lags and the lagged log consumption/GNP ratio. It also presents a univariate

TABLE I

CONSUMPTION AND GNP REGRESSIONS

y_t denotes log real GNP. c_t denotes log nondurable + services consumption. Δ denotes first differences, $\Delta y_t = y_t - y_{t-1}$. Data sample 1947:1–1989:3.

1. Vector autoregression							
Left-hand variable		Right-hand variable					R^2
	const.	$c_{t-1} - y_{t-1}$	Δc_{t-1}	Δc_{t-2}	Δy_{t-1}	Δy_{t-2}	p-value[b]
Δc_t coeff.	−0.43	−0.02	0.07	−0.02	0.09	−0.02	0.06
t-stat.	−0.49	−1.23	0.90	−0.19	1.91	−0.40	0.7
p-value[a] (percent)		0.29					
Δy_t coeff.	5.19	0.08	0.52	0.16	0.22	0.14	0.27
t-stat.	3.49	3.45	3.81	1.12	2.74	1.89	0.0
p-value[a] (percent)		1.8					

2. Variance decomposition				
Due to	Variance of			
	Δc_t	Δy_t	$\Delta c_t - E_{t-1}\Delta c_t$	$\Delta y_t - E_{t-1}\Delta y_t$
"Permanent" c shock	97	30	100	15
"Temporary" y shock	3	70	0	85

3. Univariate autoregression						
Left-hand variable		Right-hand variable				R^2
Δy_t	const.	Δy_{t-1}	Δy_{t-2}	Δy_{t-3}	Δy_{t-4}	p-value[b]
Coefficient	0.56	0.33	0.19	−0.11	−0.11	0.18
t-statistic	4.82	4.17	2.39	−1.37	−1.36	0.0

a. Percent of replications with a coefficient farther from 0, under the null that the coefficient = 0 and c/y has a unit root (bootstrap).
b. Percent probability value of an F-test for the joint significance of the right-hand variables.

autoregression of GNP growth on lagged GNP growth. Both regressions are generous; for example, the last lag is statistically insignificant.

The regressions verify that GNP growth is predictable. In particular, the t-statistic of GNP growth on the lagged consumption/ GNP ratio is 3.45, verifying the hunch that the ratio helps to forecast GNP.[3] Consumption growth is slightly predictable in the VAR, but with a much lower R^2 than GNP growth. Most of this slight predictability comes from the first-order serial correlation of quarterly consumption growth and the first lag of GNP, as noted by the permanent income literature [Flavin 1981; Campbell and Mankiw 1989].

Table I also presents probability values for t-statistics on the lagged consumption/GNP ratio obtained from a bootstrap in which that coefficient is zero. The other coeffiicients are reestimated imposing this null. Since the consumption/GNP ratio has a unit root under this null, the distribution of the t-statistic is nonstandard. Nonetheless, the coefficient is again significant at normal levels.[4]

The top panel of Figure I presents impulse-response functions of the consumption-GNP VAR. The VAR errors are orthogonalized so that consumption does not respond contemporaneously to a GNP shock. Equivalently, current consumption growth is included in the GNP growth regression. I discuss this assumption below.

Note three features of the impulse-response functions. First, the eventual response of the two series to each shock is the same. If

3. Private GNP—GNP less government purchases—paints an even stronger picture. The consumption/private GNP ratio ignores long-term fluctuations in government purchases, and responds more strongly to business-cycle related fluctuations in investment. Using private GNP, the conventional t-statistic rises from 3.39 to 4.49. Cochrane [1990] contains tables and plots with private GNP. King, Plosser, Stock, and Watson [1991] use private GNP. Assigning a fraction of government purchases to consumption, as in Christiano and Eichenbaum [1992], produces similar results.

4. Under this null, the estimated coefficients are biased to positive values. Thus, the negative consumption coefficient is "more significant" than the larger, positive GNP coefficient. Again, private GNP gives stronger results: the bootstrap p-value is less than 1/2000 (the number of replications).

Augmented Dickey-Fuller tests also reject the unit root null for the consumption/ GNP ratio at better than 1 percent levels. For private GNP, the statistic is three times larger than the 1 percent critical value.

I include these tests to verify that the consumption/GNP ratio forecasts GNP growth. Stationarity of the consumption/GNP ratio is best regarded as an a priori assumption based on common sense and basic economic theory. Blough [1992] and Cochrane [1991] warn of the dangers of pretesting for unit roots or cointegration and then imposing that form in subsequent analysis. Their point applies with special force for series like the consumption/GNP ratio. This ratio *should* be stationary but slow-moving Hence it should be hard to reject unit roots, so it can forecast slow mean reversion in GNP.
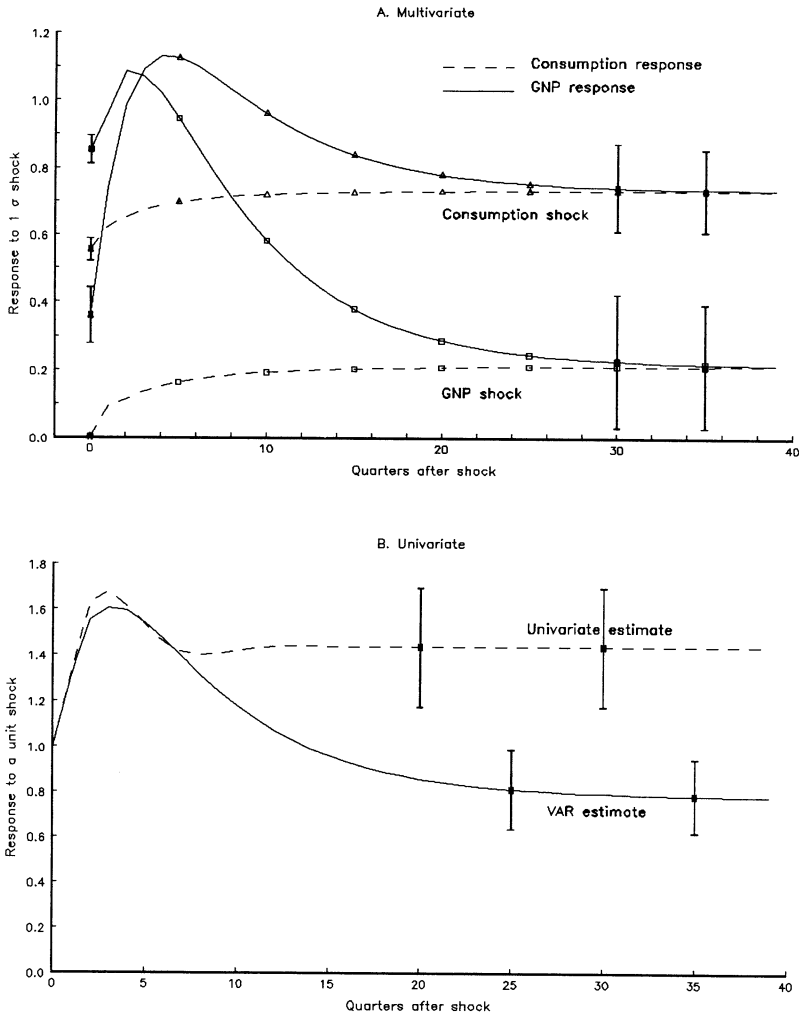
FIGURE I

Impulse-Response Functions for Consumption and GNP

Top panel: Response of consumption (c) and GNP (y) to one-standard-deviation shocks in the consumption-GNP VAR (Table I).

Bottom panel: Response of GNP to a unit univariate GNP shock. The "univariate estimate" is based on the regression of GNP growth on past GNP growth, panel 3 of Table I; the "VAR estimate" is based on the consumption-GNP VAR, panel 1 of Table I.

Bars show bootstrap one standard error bands.

the two responses did not end up at the same value, the consumption/GNP ratio would not be restored following a shock. This feature results from the inclusion of the consumption/GNP ratio on the right-hand side, and the fact that it enters with nonzero coefficients.

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. The slight rise in the long-run consumption impulse-response is statistically insignificant. Also, it disappears with slight changes in specification, such as changes in the sample period, annual data, or with private GNP. GNP has a hump-shaped response to the consumption shock. Fama [1993] interprets this pattern as the lagged response of investment to a wealth shock.

Third, and most important, look at the responses to a GNP shock. This is a shock to GNP that does not contemporaneously shock consumption (by the orthogonalization assumption). This shock has a small impact on consumption at any horizon, but the response of GNP to this shock is almost completely transitory.

How important are the transitory components of GNP? Figure I gives some indication. The figure plots the responses to one-standard-deviation shocks. Thus, the relative size of the two GNP impulse-response functions gives some measure of the relative importance of the two shocks. More formally, Table I includes a variance decomposition of GNP and consumption growth. Almost all of the variance of consumption growth is due to the permanent, or consumption shock. But 70 percent of the variance of GNP growth and 85 percent of the variance of one-step-ahead GNP prediction errors are due to the transitory, or GNP shock.[5] And this calculation does not include the transitory variation in GNP seen in the hump-shaped response to the consumption shock. Thus, not only does there *exist* a transitory component to GNP, but also it is economically important.

The pattern of the impulse-response functions has a natural permanent income interpretation. (The Appendix derives the relevant predictions of the PIH model more formally.) The model predicts that consumption should be a random walk, which is pretty much what we find. It also predicts that if consumption does not change, then consumers must regard a movement in GNP as transitory. Thus, it predicts that a GNP shock with no contempora-

---

5. With private GNP, the fractions rise to 80 percent and 89 percent.

neous consumption change (i.e., orthogonalized as above) should be completely transitory. Again, this is pretty much what we find.

B. Comparison with a Univariate Estimate

The bottom panel of Figure I presents univariate impulse-response functions for GNP—the response of GNP to a change in GNP growth not forecast by past GNP growth. The figure includes the response implied by the univariate autoregression presented in Table I, and the univariate impulse-response function implied by the VAR. (See the Appendix for construction.) The latter are two consistent estimates of the same object. However, they may differ in finite samples, and because the univariate process implied by a given order VAR is a different order ARMA than the directly estimated univariate process.

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 to only about 1.4. This univariate behavior of postwar GNP is found by many authors using a wide variety of univariate techniques (for example, Campbell and Mankiw [1987]). It results from the positive serial correlation of GNP growth at short horizons, captured in the autoregression in Table I.

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]

Why does the question "how persistent are shocks to GNP?" lead to such widely different answers? There are three reasons. First, the consumption/GNP ratio forecasts long-term GNP growth better than lagged GNP growth. Second, one can isolate two shocks in a bivariate system, and they are *different objects* than univariate shocks. A univariate GNP shock is a movement in GNP not forecast by lagged GNP. A bivariate GNP shock is a movement in GNP not forecast by lagged GNP *and consumption*. One of the bivariate shocks has permanent effects, and one has transitory effects. The single univariate shock mixes the permanent and transitory effects. Third, cointegration means that the long-run properties of consumption and GNP are the same. Since consumption is nearly a random walk, its long-run properties are almost the

---

6. Using private GNP, the long-run response estimated from the VAR is otherwise similar, but declines to 0.4 rather than about 0.8, and is much more precisely estimated. The basic pattern of the graph is robust to this and other variations in specification.

same as its short-run properties, and thus easy to estimate. The cointegrated VAR measures long-run properties of GNP by looking at consumption, which is why it can reveal more *univariate* persistence than the GNP autoregression.[7]

C. Orthogonalization and Interpretation of the Shocks

To calculate the impulse-response functions in Figure I, I identified a shock to GNP with no contemporaneous change in consumption. The purpose of this assumption is to create a permanent and a transitory shock. As explained above, the permanent income interpretation suggests this orthogonalization, since it predicts that changes in GNP with consumption (contemporaneously) held constant will be transitory.

An obvious alternative is to *force* the long-run response to one shock to zero, following Blanchard and Quah [1989]. Since the long-run response to the GNP shock is close to zero (Figure I), this orthogonalization produces very similar results.

In fact, if consumption were a pure random walk, then the conventional c first and Blanchard-Quah orthogonalizations would produce *exactly* the same result (see the Appendix). The intuition is simple: if consumption is a random walk, and if a shock does not change consumption contemporaneously, then that shock does not affect long-run consumption. Since consumption and GNP are cointegrated, the shock does not affect long-run GNP, which is the Blanchard-Quah assumption. The conventional orthogonalization has the slight advantage in that it does not require an estimate of the spectral density at frequency zero.

The contemporaneous correlation of the nonorthogonalized shocks is about 0.3, so other orthogonalizations (y first, for example) do influence the results. Responses to such shocks are linear combinations of the responses in Figure I. They typically do not change the shape of impulse-response functions much, but they have mixtures of permanent and transitory responses. Again, the aim is to find an interpretable orthogonalization that *does* produce permanent and transitory shocks, so there is not much reason to pursue these alternative orthogonalizations.

---

7. In the frequency domain, cointegration means that the spectral density of consumption growth and GNP growth must be the same at frequency zero. The spectral density of GNP growth implied by the univariate autoregression and the VAR are similar at high frequencies. But near frequency zero, the VAR GNP spectral density drops down to the substantially lower spectral density of consumption growth, implying lower univariate persistence measures. See Watson [1990] and Cochrane [1990] for plots.

I do not attach labels to the shocks beyond "permanent" or "consumption" shock and "transitory" or "GNP" shock. VAR shocks are linear forecast errors. For these VAR shocks to recover the underlying shocks to the economy, we must impose an economic model in which there are exactly two true shocks: one shock has no permanent effects; and the model's responses to the two shocks is linear and invertible.[8] Within the context of simple models that satisfy the above conditions, Blanchard and Quah [1989] and Shapiro and Watson [1988] can label their shocks "supply" and "demand."[9] But without such a model, the VAR says nothing about the number, persistence, relative importance, or identity ("supply" versus "demand," preferences versus technology, money versus fiscal policy, etc.) of the underlying shocks to the economy. For example, suppose that there is a single underlying technology shock, as in real business cycle models. The VAR may decompose this single shock into two components: a "consumption" or "permanent" component, and a "GNP" or "transitory" component. Conversely, if there are many underlying shocks with various persistences, the VAR will still find two shocks, corresponding to permanent and transitory movements in GNP.

D. VAR Specification

A variety of related VARs have been run. Blanchard and Quah [1989] use a detrended unemployment rate to help forecast GNP. Like the consumption/GNP ratio, the unemployment rate is a "level" variable that can indicate higher or lower than average GNP growth for many subsequent quarters. The consumption/GNP VAR improves on this idea:[10] the consumption/GNP ratio does not need to be detrended; it is directly linked to GNP growth; it provides almost the same result with conventional versus

8. Lippi and Reichlin [1993] examine non-invertible representations of the Blanchard-Quah VAR and note the sensitivity of the impulse response function and variance decomposition to the choice of representation.

9. Cassou and Mittnik [1991] criticize Blanchard and Quah's economic model. They argue that in alternative models, "demand" variables such as money growth can have permanent effects, while "supply" variables such as technology shocks can have very little permanent effect.

10. Evans [1989] also uses unemployment to forecast GNP. Fama [1993] uses ratios to uncover transitory movement in GNP—he essentially studies the consumption shock of the VAR in detail. Cochrane and Sbordone [1988] use the consumption/GNP and dividend/price ratios, but they characterize persistence with variance ratios rather than impulse-response functions. Harvey and Stock [1988] and King, Plosser, Stock, and Watson [1991] present the most similar cointegrated GNP VARs. Blanchard [1993] uses consumption to help forecast GNP through the 1990–1991 recession.

long-run orthogonalization; and its results are robust to extra variables, as explained next.

Variables beyond the consumption/GNP ratio are useful for forecasting short-term GNP growth, and the forecasting literature therefore typically uses larger systems.[11] However, further variables or longer lags just add wiggles to the short end of the impulse-response functions, with little impact on the *long-run* behavior of GNP.

There is a good reason. If consumption were a pure random walk—i.e., if consumption growth were unpredictable by any variable—then including any other variable would leave the long-run properties of consumption and GNP completely unaffected. In particular, the response of GNP to a shock with constant consumption would still be completely transitory. In the permanent income interpretation, consumption summarizes all consumers' information about long-run GNP, so other variables are superfluous.[12] Since consumption growth is not very predictable, even when many other variables are included, this result holds as an approximation in the data.

I use the bivariate systems to focus on long-run behavior in this paper. A model designed to produce good short- and long-term forecasts should include short-term forecasting variables as well as lagged consumption/GNP.

III. "DETRENDING" OR "CYCLICALLY ADJUSTING" GNP

Since the above results indicate a large transitory component in GNP, they suggest that we can meaningfully decompose GNP into "trend" and "stationary" or "cyclical" components. Beveridge and Nelson [1981] suggest an attractive way to do this. They define the trend in GNP as the level GNP will reach after all transitory dynamics work themselves out. Equivalently, the trend in GNP is GNP plus all expected future GNP growth: if GNP growth is expected to be higher than average in the future, GNP is below trend; if it is expected to be below average, GNP is above trend. Formally, the trend z_t is defined by

$$z_t = \lim_{k \to \infty} E_t(y_{t+k} - k\mu) = y_t + \sum_{j=1}^{\infty} E_t(\Delta y_{t+j} - \mu),$$

11. For example Sims [1980] and Litterman and Weiss [1985]. Shapiro and Watson [1988] add short-run forecasting variables to the Blanchard-Quah VAR.
12. Hansen, Roberds, and Sargent [1992] show how consumption summarizes information about permanent income.

where $\mu \equiv E(\Delta y)$. Figure II presents log GNP and the Beveridge-Nelson stochastic trend, constructed from the estimated consumption-GNP VAR of Table I (see the Appendix for construction). This stochastic trend responds to long-run movements in GNP growth

A. Multivariate
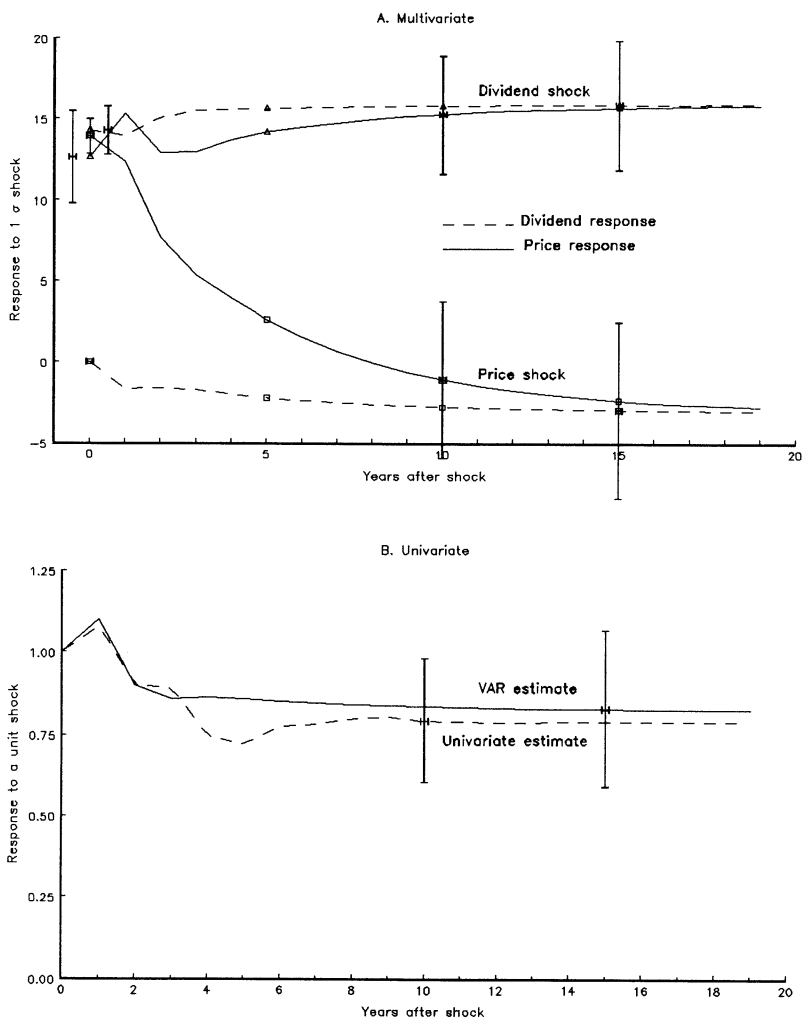

B. Univariate

FIGURE II

Beveridge-Nelson GNP Trend and Consumption

The "VAR trend" is the Beveridge-Nelson [1981] trend calculated from the consumption-GNP VAR (Table I) as GNP plus all expected future above-average growth in GNP. "CN&S + Mean Ratio" gives log nondurable + services consumption plus the mean log GNP/consumption ratio. The graph is limited to 1963–1990 for clarity.

during the seventies and eighties, 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 less the mean log GNP/consumption ratio (see the Appendix). Figure II also plots this quantity. Since consumption growth is poorly forecastable, this quantity is almost the same as the Beveridge-Nelson trend.

Thus, *consumption* provides a good measure of the trend in GNP, since it measures *consumers'* expectations of long-run GNP. In place of the traditional potential GNP calculation, "what would GNP be if the unemployment rate were x percent," this measure is "what would GNP be if the consumption/GNP ratio was at its historical mean." This measure could easily be used in place of Hodrick-Prescott filtering, or other decompositions based on univariate representations. Since univariate representations of GNP typically show little transitory variation, this procedure will isolate much larger cyclical components (see Cochrane [1990] for plots of one example).

IV. STOCK PRICES AND DIVIDENDS

A. Results

Table II presents a VAR of log dividends and log prices,[13] and a univariate regression of log returns on lagged log returns. The data are from the CRSP value-weighted NYSE portfolio. They are annual to avoid the seasonal in dividends. Again, more lags and other right-hand variables (term premium, default premium, interest rate) add more wiggles to the short-run impulse-response functions and raise the one-step-ahead forecast R^2, but do not alter the pattern of the long-run impulse-response functions.

The results in Table II are similar to the previous results for GNP and consumption, with dividends in place of consumption and prices in place of GNP. The dividend/price ratio forecasts returns much more strongly than it forecasts dividend growth, so *prices* rather than dividends adjust to bring the ratio back to its mean.[14]

---

13. I define the "price" variable to be cumulated returns, so that price growth regressions are return forecasting regressions, and so that movements in the "price" response can be interpreted as variation in expected returns. Since dividends are poorly forecastable, results in which the price variable is formed by cumulating the CRSP price-only return (P_t/P_{t-1}) are very similar.

14. The coefficient of returns on the d/p ratio is significant using conventional distribution theory. However, the bootstrap only rejects a zero coefficient at a 9.2 percent level. Similarly, an augmented Dickey-Fuller test rejects at the 2.5 percent level with one lag but at 10 percent with two lags. The dividend/price ratio is more slow-moving than the consumption/GNP ratio, so the greater difficulty of rejecting a unit root is expected.

TABLE II
DIVIDEND AND PRICE REGRESSIONS

d_t denotes log dividends and p_t denotes log price (cumulated return) on the value-weighted NYSE portfolio. Δ denotes first difference; Δp is the log return. Data sample 1927–1988.

	1. Vector autoregression						
Left-hand variable			Right-hand variable				R^2
	const.	$d_{t-1}-p_{t-1}$	Δd_{t-1}	Δd_{t-2}	Δp_{t-1}	Δp_{t-2}	p-value[b]
Δd_t coeff.	20.01	0.038	0.046	0.062	−0.082	0.040	0.038
t-stat.	0.78	0.47	0.25	0.34	−0.65	0.32	32.0
p-value[a] (percent)		20.7					
Δp_t coeff.	78.65	0.225	0.060	−0.086	0.114	−0.090	0.140
t-stat.	2.34	2.11	0.25	−0.36	0.68	−0.55	1.2
p-value[a] (percent)		9.9					

	2. Variance decomposition			
Due to			Variance of	
	Δd_t	Δp_t	$\Delta d_t - E_{t-1}\Delta d_t$	$\Delta p_t - E_{t-1}\Delta p_t$
"Permanent" d shock	99	43	100	45
"Temporary" p shock	1	57	0	55

	3. Univariate autoregression				
Left-hand variable					
Δp_t	const.	Right-hand variable			R^2
		Δp_{t-1}	Δp_{t-2}	Δp_{t-3}	Δp_{t-4}	p-value[b]
Coefficient	11.91	0.075	−0.179	0.015	−0.18	.061
t-statistic	3.46	0.57	−1.37	0.12	−1.37	49.7

a. Percent of replications with a coefficient farther from 0, under the null that the coefficient = 0 and d/p has a unit root (bootstrap).
b. Percent probability value of an F-test for the joint significance of the right-hand variables.

The R^2 of the dividend growth forecasting regression is lower than the R^2 of the return forecasting regression, and the bivariate return regression has a higher R^2 than the univariate return

In postwar data the dividend/price ratio forecasts both returns and dividend growth more strongly. The t-statistics rise from 2.11 to 4.00, and 0.78 to 2.71 respectively. Fama and French [1988b] and Hodrick [1992] also find more statistically significant forecasts of returns from dividend/price ratios in postwar data. Given this evidence, the strong a priori reasons to believe the dividend/price ratio is stationary, and the dangers noted above of pretesting for unit roots, I assume the dividend/price ratio is stationary and proceed to study the implications of this assumption.
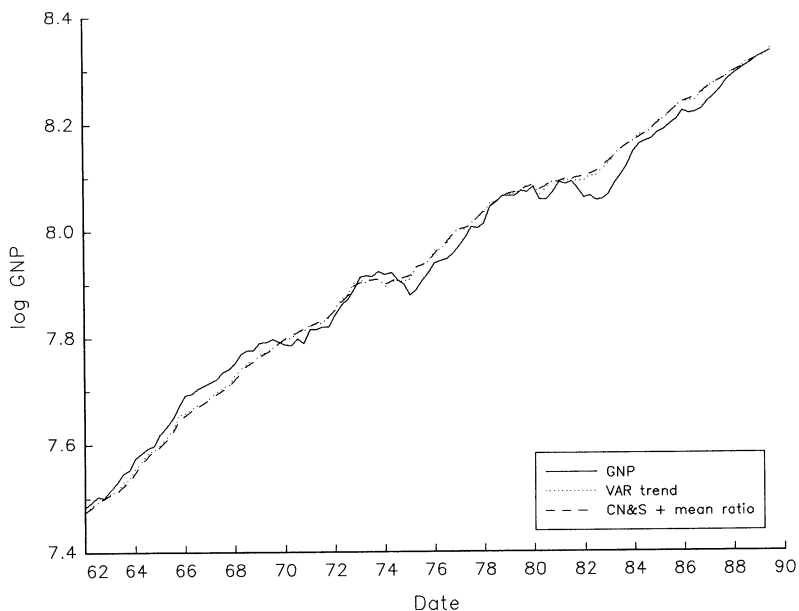
FIGURE III

Impulse-Response Functions for Dividends and Prices

Top panel: Response of dividends (d) and prices (p) to one-standard-deviation shocks in the dividend-price VAR (Table II).

Bottom panel: Response of price to a unit univariate price shock. The "univariate estimate" is based on the regression of returns on past returns, panel 3 of Table II; the "VAR estimate" is based on the dividend-price VAR, panel 1 of Table II.

Bars show bootstrap one-standard-error bands.

regression. Dividends look a lot like a random walk, as do returns when regressed only on lagged returns.[15]

The top panel of Figure III presents VAR impulse-response functions. The pattern is similar to that of the consumption-GNP impulse-response in Figure I. In response to a dividend shock, prices and dividends move immediately to their long-run values. On the other hand, a price shock with no movement in dividends

15. The return forecasting literature has focused on long-horizon returns. One can infer long-horizon properties from a VAR. (See Hodrick [1992].) The long-horizon return R^2 implied by the VAR in Table II rises to a peak of 0.24 at a seven-year horizon and then gradually declines back to 0.15 at a twenty-year horizon. Again, greater predictability occurs in the postwar sample. The variance decomposition also varies with horizon. Since the dividend shock is permanent, it gradually accounts for more of the variance of longer-horizon dividend growth and returns.

has a completely transitory effect on prices and no effect on dividends. As the swings in the dividend/price ratio are longer, the transitory movement in prices occurs over a longer horizon (half-life of about five years) than was the case for GNP (half-life of about one–two years).[16]

The variance decomposition finds that 57 percent of the variance of returns and 55 percent of the variance of one-step-ahead return forecast errors are attributed to the price shock. Thus, the transitory component of price is economically as well as statistically significant.

The bottom panel of Figure III presents univariate impulse-response functions for stock prices. Here, the VAR and univariate estimates look the same. There is little evidence for univariate mean reversion no matter how estimated; evidence for mean reversion in prices (or predictability in returns) comes when one isolates a transitory multivariate shock. (The point estimates that suggest univariate predictability come from long-horizon regressions as in Fama and French [1988a] or variance ratios as in Poterba and Summers [1988]. Both estimates are equivalent to much longer AR's with restrictions on the coefficients.)

B. Interpretation: Present Value-Dividend Smoothing Model

The present value model together with the hypothesis that managers smooth dividends can be used to interpret the price/dividend VAR. This section presents the ideas verbally, while the Appendix contains a more formal treatment.

The present value model states that the price/dividend ratio depends on expectations of future discount rates and dividend growth rates.[17] (Prices rise on news of high future long-term

---

16. In postwar data the long-run impulse-response functions are quite similar. Dividend growth is more forecastable in postwar data, so its impulse-response function is not flat. However, it does not end up far from where it starts, so the price shock is still almost completely permanent.

17. With time-varying discount rates, the present value relation is

$$P_t = E_t\left[\sum_{j=1}^{\infty} \left(\prod_{k=1}^{j} m_{t+k}\right) D_{t+j}\right],$$

where P = price, D = dividend, and m = stochastic discount factor ($m = 1/(1 + \text{discount rate})$). Dividing by D_t and letting $\Delta D_t = D_t/D_{t-1}$,

$$\frac{P_t}{D_t} = E_t\left[\sum_{j=1}^{\infty} \prod_{k=1}^{j} m_{t+k}\,\Delta D_{t+k}\right].$$

Thus, price/dividend ratios are related to future discount factors m and dividend growth rates ΔD. "Time-varying discount rates" or "time-varying expected returns" do not require time-varying interest rates, since risk premiums may vary over time.

dividends, or on news of lower long-term discount rates or expected returns.) Therefore, if dividend growth and discount rates (expected returns) are stationary, the present value model implies that the price/dividend ratio is stationary.

If managers smooth dividends by setting dividends equal to the discounted value of earnings (discounted at the risk-free rate), then dividends follow a random walk just as consumption does in the PIH model. In turn, a stationary price/dividend ratio and random walk in dividends imply that the price shock is completely transitory. Thus, the model predicts the impulse-response pattern of Figure III.

In addition, the present value–dividend smoothing model suggests that we label the shocks as an "earnings" shock and a "discount rate" or "expected return" shock. In the model an increase in dividends signals an increase in long-term earnings, as perceived by managers. Since discount rates or expected returns have not changed, prices rise simultaneously. Dividends, following a random walk, are expected to remain at the new level, so prices are too. Thus, the earnings shock gives rise to simultaneous, permanent shocks to prices and dividends, just like the "dividend" shock in the VAR.

On the other hand, if there is a decrease in expected returns (risk premiums, not risk-free rates), prices rise with no concurrent change in dividends. As expected, returns revert to their mean, and prices revert too. Thus, a discount rate or expected return shock gives rise to a transitory price shock with no change in dividends, just like the "price" shock in the VAR.

Of course, the dividend smoothing model can be statistically rejected, just like the PIH model. Dividend growth is statistically predictable. However, the result that price shocks with constant dividends are transitory requires only that the instantaneous response of dividends to any shock equals their limiting response; the impulse-response functions can wiggle on their way, and dividend growth may be predictable (see the Appendix). Much of dividend predictability is of this type, and so has little effect on the long-run results. Also, as found extensively in the variance bound literature,[18] almost all variation in price/dividend ratios is in fact due to changing forecasts of expected returns rather than changing forecasts of dividend growth. In this sense, the *economic* impor-

---

18. See Campbell and Shiller [1988] and Cochrane [1992] for evidence beyond that presented here.

tance of dividend forecastability is small, even though it is *statistically* significant. Furthermore, much evidence against the simple dividend smoothing model comes from individual firm data. These price and dividend movements largely define the idiosyncratic or diversifiable component of firm returns, so index dividends, as used here, are much less forecastable. In any case, the dividend smoothing model, like the PIH model, is a useful approximation that helps us to interpret the dividend/price VAR.

Comparing the price/dividend and consumption/GNP VARs, one might think that prices should take the place of consumption, since prices are a forecast of future dividends as consumption is a forecast of future income. If expected returns were constant, this is in fact what we would see: price/dividend ratios would forecast long-term *dividend growth* and not returns, and *prices* would be a random walk. However, expected returns are not constant, and aggregate long-run dividends happen to be nearly unpredictable. Hence, price, like GNP, is the series with the interesting temporary component; dividends, like consumption, is the near-random walk that defines the "trend;" and the PIH analogy is that dividends equal "permanent earnings," not that prices equal "permanent dividends."

APPENDIX

Subsection A formalizes the relations between time-series representations discussed in the text. It shows that the following three statements are equivalent: (1) c and y are cointegrated and c is a random walk, (2) c measures the trend in y, and (3) conventional (Sims) orthogonalization with c first is equivalent to Blanchard-Quah orthogonalization. Subsections B and C discuss the predictions of the permanent income and present value–dividend smoothing models. Subsection D details the VAR estimation and construction of impulse-response functions.

A. Relation Between Statistical Assumptions

Let y denote log GNP or log stock prices, and c denote log consumption or log dividends. Let $I(t)$ denote an information set available to consumers at time t, including at least all lags of c_t and y_t. $E_t(\cdot)$ denotes $E(\cdot \,|\, I(t))$.

I assume that $(1 - L)y_t$ is stationary with zero mean. (It is easy to generalize to nonzero means.) Stationarity implies that there exists a moving average representation $(1 - L)y_t = A(L)w_t$, where

w_t is a vector of unforecastable random variables that generates the information set $I(t)$. Thus, I maintain

ASSUMPTION A1. (Stationarity) $(1 - L)y_t$ is stationary, or
$(1 - L)y_t = A(L)w_t$.

Next, we formalize three possible characterizations of the (c, y) process. First, y and c may be cointegrated. Cointegration requires that $(1 - L)c_t$ is stationary, and thus has a moving average representation; and it implies that c and y have the same long-run response to any shock. Thus, cointegration means

ASSUMPTION A2. (Cointegration) A1 and $y_t - c_t$ stationary, or
$(1 - L)c_t = B(L)w_t$, with $B(1) = A(1)$.

Second, the long-run response of c to shocks might equal its immediate response.

ASSUMPTION A3.
$$c_t - E_{t-1}(c_t) = \lim_{j \to \infty} E_t c_{t+j} - \lim_{j \to \infty} E_{t-1} c_{t+j} \qquad \text{or} \qquad B_0 = B(1).$$

A3 is weaker than a random walk. The impulse-response of a process that obeys A3 can wiggle on the way to infinity, and Δc can be predictable.

Third, innovations in c might measure innovations in long-run y:

ASSUMPTION A4.
$$c_t - E_{t-1}(c_t) = \lim_{j \to \infty} E_t y_{t+j} - \lim_{j \to \infty} E_{t-1} y_{t+j}, \qquad \text{or} \qquad B_0 = A(1).$$

These assumptions are related by

PROPOSITION 1. Given stationarity (A1) and cointegration (A2), innovations in c measure innovations in long-run y (A4) if and only if the immediate and long-run responses of c to a shock are the same (A3).

Proof of Proposition 1. Given $B(1) = A(1)$, $B_0 = B(1)$ iff $B_0 = A(1)$.

A pure random walk in c is sufficient for A3 and hence Proposition 1. It also implies that c *equals* long-run y. Formally, consider

ASSUMPTION A5. (Random walk) $c_t = E_t(c_{t+1})$, or $B(L) = B_0$.

ASSUMPTION A6. (c equals long-run y) $c_t = \lim\limits_{j\to\infty} E_t(y_{t+j})$, or $B(L) = A(1)$.

PROPOSITION 2. Given stationarity (A1) and cointegration (A2), c is a random walk (A5) if and only if c equals the long-term forecast of y (A6).

Proof of Proposition 2. First, A2, A5 $\Rightarrow$ A6. $B(1) = A(1)$, $B(L) = B_0$ imply that $B(L) = A(1)$. Second, A6 $\Rightarrow$ A2, A5. $B(L) = A(1)$ implies that $B(L) = B_0$ and $B(1) = A(1)$.

The text claimed that Blanchard-Quah and conventional orthogonalization might be the same. The conditions are specified in

PROPOSITION 3. Cointegration (A2) and property (A3) imply that Blanchard-Quah [1989] and Sims [1980] orthogonalization with c first are equivalent.

Proof of Proposition 3. Sims orthogonalization with c first states that the first element of B_0 is zero; i.e., c does not respond to the y shock. Blanchard-Quah orthogonalization states that the first shock has no long-run impact on y (or c); i.e., the first element of $A(1)$ (or $B(1)$) is zero. Cointegration (A2) $A(1) = B(1)$ and A3 $B_0 = B(1)$ imply that the first element of B_0 is zero if and only if the first elements of $A(1)$ or $B(1)$ are zero.

A pure random walk or c = expected long-run y (A5 and A6) are sufficient for Proposition 3, but the weaker conditions A3 or A4 will do. Proposition 3 has an important corollary.

COROLLARY. Cointegration (A2) and A3 imply that shocks to y with c held constant are completely transitory.

This is just an interpretation of the Blanchard-Quah identification condition. Keep in mind that the information set is not limited to past c and y in any of these representations. A3 with respect to a large information set implies that y shocks with constant c are transitory even if many other variables are added to the VAR.

B. Permanent Income Model

PROPOSITION 4. If labor income is stationary in levels or first differences, the permanent income model predicts that consumption is a random walk and that consumption and total income are cointegrated.

Proof of Proposition 4. The random walk implication is well-known; see Hall [1978] or Sargent [1987]. To show cointegration, start with the permanent income decision rule:

$$c_t = rk_t + r\beta \sum_{j=0}^{\infty} \beta^j E_t e_{t+j},$$

where r = interest and discount rate, $\beta = 1/(1 + r)$, e_t = endowment or labor income, k_t = wealth. Total income is $y_t = e_t + rk_t$. Hence,

$$c_t - y_t = rk_t + r\beta \sum_{j=0}^{\infty} \beta^j E_t e_{t+j} - e_t - rk_t = r\beta \sum_{j=0}^{\infty} \beta^j (E_t e_{t+j} - e_t).$$

If e_t is either level- or difference-stationary, then $E_t e_{t+j} - e_t$, and hence $c_t - y_t$, are stationary.

Propositions 1–3 then show that the PIH model implies that

$$c_t = \lim_{j \to \infty} E_t(y_{t+j})$$

or consumption is the Beveridge-Nelson trend in GNP, that Sims orthogonalization with c first should be the same as Blanchard-Quah orthogonalization; i.e., that GNP shocks with constant consumption should be transitory.

Proposition 4 is not true for consumption and *labor* income. The PIH does *not* predict cointegration of consumption and labor income, and it predicts that

$$c_t = rk_t + r\beta \sum_{j=0}^{\infty} \beta^j E_t e_{t+j} \quad \text{not that} \quad c_t = \lim_{j \to \infty} E_t e_{t+j}.$$

Proposition 4 is true for labor income only in the limit $\beta \to 1$. Instead, Proposition 4 describes a permanent income-like (limit as $\beta \to 1$) relation between consumption and *total* income. This relation could be used to test the PIH avoiding the difficulties inherent in separating labor from total income.

The permanent income model predicts that the (per capita) *level* of consumption is a random walk, and that the consumption-GNP *difference* is stationary. I follow Campbell and Mankiw [1989] in applying this model as an approximation to *log* consumption and GNP.

Stochastic growth models (for example, King, Plosser, Stock, and Watson [1991]) also typically imply stationarity of the consumption/GNP ratio, as well as other "great ratios." They do not predict

a random walk in consumption, since they predict a small amount of interest rate variation. However, they do predict less transitory variation in consumption than in investment or GNP. Therefore, consumption still helps VAR estimation of permanent versus transitory components in GNP, but the above representation results only hold approximately.

C. Present Value and Dividend Smoothing Models

If dividend growth and discount rates are strongly stationary, then the dividend/price ratio is strongly stationary. This proposition is discussed in detail in Cochrane [1992] and Craine [1993]. The basic idea comes from staring at the present value relation:

$$\ln\left(\frac{D_t}{P_t}\right) = -\ln\left(E_t \sum_{j=0}^{\infty} \prod_{k=1}^{j} m_{t+k} \Delta D_{t+k}\right),$$

where P_t = price, D_t = dividend, $\Delta D_t = D_t/D_{t-1}$, and m_t = discount rate. Assume that m_t, ΔD_t, and the variables generating E_t are strongly stationary. The right-hand side is then a time-invariant function of stationary variables, so the log dividend-price ratio is strongly stationary.

To motivate a random walk in dividends, we can map the standard dividend-smoothing model into the permanent income framework given above. Suppose that managers can borrow and lend at constant interest rates. (Expected risky returns may still vary over time.) Let k_t = accumulated earnings. Then, if managers set dividends to "permanent" earnings e_t,

$$D_t = rk_t + r\beta\, E_t \sum_{j=0}^{\infty} \beta^j e_{t+j},$$

dividends will follow a random walk, just like consumption. (The right-hand side of the last equation is *not* the present value of the firm, since risk-adjusted or stochastic discount factors must be used to find the latter. For this reason, a separate proof of dividend/price stationarity, given above, is required.) Propositions 2 and 3 then imply that dividends define permanent and transitory components of prices with respect to managers' presumably large information set.

D. VAR Transformations

The VARs in Tables I and II can be represented as

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

where

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

and $A(L)$ is a matrix of lag polynomials, $A(0) = I$.

One can also run $[\Delta c \ c - y]'$, $[\Delta y \ c - y]'$, or $[c \ y]$ on their lags. If $c - y$ is really stationary, these procedures are equivalent up to lag length restrictions. I use the error-correction representation (A.1) to focus on growth rate forecasts, and because it naturally gives rise to an interesting orthogonalization.

I orthogonalize the error terms by Choleski decomposing the variance-covariance matrix of the innovations in the order c, then y. Precisely, I find a triangular matrix R such that $RR' = E(\epsilon_t \epsilon_t') = \Sigma$. Then I define new errors

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

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 impulse-response function, I simulate the response of the VAR (without constants) to the ν_t shocks starting with all variables set to zero. This calculation is most easily performed by writing the VAR in AR(1) notation. For example, a two-lag VAR (without means) is written as

$$
(\text{A.2}) \quad
\begin{bmatrix} \Delta c_t \\ \Delta y_t \\ \Delta c_{t-1} \\ \Delta y_{t-1} \\ y_t - c_t \end{bmatrix}
=
\begin{bmatrix} -\beta_c- \\ -\beta_y- \\ 0 \ 1 \ 0 \ 0 \\ 0 \ 0 \ 1 \ 0 \\ 1 + \beta_y \ -\beta_c \end{bmatrix}
\begin{bmatrix} \Delta c_{t-1} \\ \Delta y_{t-1} \\ \Delta c_{t-2} \\ \Delta y_{t-2} \\ \nu_{t-1} - c_{t-1} \end{bmatrix}
+ R
\begin{bmatrix} \nu_t^c \\ \nu_t^y \\ 0 \\ 0 \\ \nu_t^y - \nu_t^c \end{bmatrix}
$$

$$x_t \qquad = \qquad B \qquad\qquad x_{t-1} \qquad + R \quad \nu_t,$$

where β_c and β_y denote vectors of OLS regression coefficients. This simulation leads to the Wold moving average representation,

$$(\text{A.3}) \qquad\qquad \Delta x_t = \mu + D(L)\nu_t.$$

Cumulative sums of the elements of $D(L)$ are plotted in Figures I and III. For subsequent calculations, I retained 50 terms of the moving average (A.3).

I calculated variance decompositions as follows. Express (A.3) as

$$\begin{bmatrix} \Delta c_t \\ \Delta y_t \end{bmatrix} = \mu + \begin{bmatrix} D_{cc}(L) \ D_{cy}(L) \\ D_{yc}(L) \ D_{yy}(L) \end{bmatrix} \begin{bmatrix} v_t^c \\ v_t^y \end{bmatrix}.$$

Since the shocks are all orthogonal with unit variance,

$$\text{var}(\Delta y_t - E_{t-1}\Delta y_t) = D_{yc,0}^2 + D_{yy,0}^2$$

$$\text{var}(\Delta y_t) = \sum_{j=1}^{\infty} D_{yc,j}^2 + \sum_{j=1}^{\infty} D_{yy,j}^2.$$

The first terms on the right-hand side give the variance attributed to the c shock; the second terms give the variance attributed to the y shock.

I constructed spectral densities from (A.3) from the definition,

(A.4) $$S_{\Delta x}(e^{-i\omega}) = D(e^{-i\omega})D(e^{i\omega})'.$$

To find the univariate impulse response function for GNP implied by the VAR, I factored the spectral density of Δy_t. The second row of (A.4) 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. Since the Wold moving average is unique, $a(L)$ gives the univariate impulse-response function.

To obtain the Beveridge-Nelson decomposition, I use the AR(1) form of the VAR, as in (A.2). Then the trend is given by $z_t = y_t + \sum_{j=1}^{\infty} E_t(\Delta y_{t+j})$ or

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

DEPARTMENT OF ECONOMICS, UNIVERSITY OF CHICAGO, N.B.E.R., AND N.O.R.C.

REFERENCES

Beveridge, Stephen, and Charles R. Nelson, "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,* VII (1981), 151–74.
Blanchard, Olivier Jean, "Consumption and the Recession of 1990–1991," *American Economic Review,* LXXXIII (1993) 270–74.
Blanchard, Olivier Jean, and Danny Quah, "The Dynamic Effects of Aggregate Supply and Demand Disturbances," *American Economic Review,* LXXIX (1989), 655–73.
Blough, Stephen R., "The Relationship Between Power and Level for Generic Unit Root Tests in Finite Samples," *Journal of Applied Econometrics,* VII (1992), 295–308.
Campbell, John Y., and N. Gregory Mankiw, "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics,* CII (1987), 857–80.
Campbell, John Y., and N. Gregory Mankiw, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," *NBER Macroeconomics Annual,* Olivier J. Blanchard and Stanley Fischer, eds. (Cambridge MA: MIT Press, 1989).
Campbell, John Y., and Robert J. Shiller, "The Dividend Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies,* I (1988), 195–228.
Cassou, Steven P., and Stefan Mittnik, "Causal Relationships in Real Business Cycles," manuscript, Department of Economics, State University of New York at Stony Brook, 1991.
Christiano, Lawrence J., and Martin S. Eichenbaum, "Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations," *American Economic Review,* LXXXII (1992), 530–50.
Clark, Peter K., "The Cyclical Component in U. S. Economic Activity," *Quarterly Journal of Economics,* CII (1987), 797–814.
Cochrane, John H., "How Big is the Random Walk in GNP?" *Journal of Political Economy,* XCVI (1988), 893–920.
——, "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," NBER Working Paper No. 3427, 1990.
——, "A Critique of the Application of Unit Roots Tests," *Journal of Economic Dynamics and Control,* XV (1991) 275–84.
——, "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies,* V (1992), 243–80.
Cochrane, John H., and Argia M. Sbordone, "Multivariate Estimates of the Permanent Components of GNP and Stock Prices," *Journal of Economic Dynamics and Control,* XII (1988), 255–96.
Cogley, Timothy, "International Evidence on the Size of the Random Walk in Output," *Journal of Political Economy,* XCVII (1990), 501–18.
Craine, Roger, "Asset Prices and Fundamentals: A *q* Test," *Journal of Economic Dynamics and Control,* (1993), forthcoming.
Evans, George, "Output and Unemployment Dynamics in the United States: 1950–1985," *Journal of Applied Econometrics,* IV (1989), 213–37.
Fama, Eugene F., "Efficient Capital Markets II," *Journal of Finance,* XLVI (1991) 1575–1617.
——, "Transitory Variation in Investment and Output," *Journal of Monetary Economics* (1993), forthcoming.
Fama, Eugene F., and Kenneth R. French, "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy,* XCVI (1988a), 246–73.
Fama, Eugene F., and Kenneth R. French, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics,* XXII (1988b), 3–25.
Fisher, Irving, "Our Unstable Dollar and the So-Called Business Cycle," *Journal of the American Statistical Association,* XX (1925), 179–200.
Flavin, Marjorie A., "The Adjustment of Consumption to Changing Expectations of Future Income," *Journal of Political Economy,* LXXXIX (1981), 974–1009.
Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income

Hypothesis: Theory and Evidence," *Journal of Political Economy,* LXXXVI (1978), 971–88.

Hansen, Lars Peter, William Roberds, and Thomas J. Sargent, "Implications of Present Value Budget Balance and of Martingale Models of Consumption and Taxes," Chapter 5 of *Rational Expectations Econometrics,* by Lars Peter Hansen and Thomas J. Sargent, eds. (London and Boulder, CO: Westview Press, 1992).

Harvey, Andrew C., and James H. Stock, "Continuous Time Autoregressive Models with Common Stochastic Trends," *Journal of Economic Dynamics and Control,* XII (1988), 365–84.

Hodrick, Robert J., "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies,* V (1992), 357–86.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson, "Stochastic Trends and Economic Fluctuations," *American Economic Review,* LXXXI (1991), 819–40.

Lippi, Marco, and Lucrezia Reichlin, "The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment," *American Economic Review,* LXXXIII (1993), 644–52.

Litterman, Robert B., and Lawrence Weiss, "Money, Real Interest Rates and Output: A Reinterpretation of Postwar U. S. Data," *Econometrica,* LIII (1985), 129–56.

McCulloch, J. Huston, "The Monte Carlo Cycle in Business Activity," *Economic Inquiry,* XII (1975), 303–21.

Nelson, Charles R., and Charles I. Plosser, "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications," *Journal of Monetary Economics,* X (1982), 139–62.

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

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

Shapiro, Matthew, and Mark Watson, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual,* Olivier Blanchard and Stanley Fischer, eds. (Cambridge MA: MIT Press, 1988), pp. 111–41.

Sims, Christopher A., "Macroeconomics and Reality," *Econometrica,* XLVII (1980), 1–48.

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