



Are Output Fluctuations Transitory?

John Y. Campbell; N. Gregory Mankiw

The Quarterly Journal of Economics, Vol. 102, No. 4 (Nov., 1987), 857-880.

Stable URL:

<http://links.jstor.org/sici?sici=0033-5533%28198711%29102%3A4%3C857%3AAOFT%3E2.0.CO%3B2-F>

The Quarterly Journal of Economics is currently published by The MIT Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/mitpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

ARE OUTPUT FLUCTUATIONS TRANSITORY?*

JOHN Y. CAMPBELL AND N. GREGORY MANKIW

According to the conventional view of the business cycle, fluctuations in output represent temporary deviations from trend. The purpose of this paper is to question this conventional view. If fluctuations in output are dominated by temporary deviations from the natural rate of output, then an unexpected change in output today should not substantially change one's forecast of output in, say, five or ten years. Our examination of quarterly postwar United States data leads us to be skeptical about this implication. The data suggest that an unexpected change in real GNP of 1 percent should change one's forecast by over 1 percent over a long horizon.

I. INTRODUCTION

Robert Lucas begins his classic article, "Understanding Business Cycles" [1977], with the question, "Why is it that, in capitalist economies, aggregate variables undergo repeated fluctuations about trend, all of essentially the same character?" Many textbooks introduce macroeconomics with a graph of real GNP together with a trend line, implying that the purpose of macroeconomic theory is to explain the deviations of production from the trend. Implicit both in Lucas's question and in such a picture is the notion that output fluctuations are transitory. Certainly this view is implicit in the standard explanation of the business cycle: the natural rate of output grows at a more or less constant rate, while output fluctuations represent temporary deviations.

The purpose of this paper is to question this conventional view. In particular, we examine one simple implication for the univariate properties of economic time series. If fluctuations in output are dominated by temporary deviations from the natural rate, then an innovation in output should not substantially change one's forecast of output in, say, five or ten years. Over a long horizon, the economy should return to its natural rate; the time series for output should be trend-reverting.

Our examination of quarterly postwar United States data leads us to be skeptical about this implication. In particular, we estimate a number of models in which a 1 percent innovation to real GNP should change one's forecast of GNP over a long horizon by over 1

*We are grateful to Ben Bernanke, Olivier Blanchard, Angus Deaton, Robert King, Christina Romer, David Romer, G. William Schwert, Jim Stock, and Mark Watson for helpful discussions and comments on earlier drafts. We acknowledge financial support from the National Science Foundation.

percent. While we find some evidence of short-run dynamics that makes GNP different from a random walk with drift, the long-run implications of our estimates suggest that shocks to GNP are largely permanent.

Our goal here is to establish a stylized fact against which macroeconomic theories can be measured. It is obviously imprudent to make definitive judgments regarding theories on the basis of one stylized fact alone. Nonetheless, we believe that the substantial persistence of output shocks is an important and often neglected feature of the postwar data that should be used more widely for evaluating theories of economic fluctuations. Most of this paper is aimed at establishing the high degree of persistence. In the last section we briefly discuss the extent to which prominent theories of the business cycle are consistent with our finding.

The research presented here builds on the work of Nelson and Plosser [1982]. These authors show that for a number of macroeconomic time series, measured annually over periods of 60 to 120 years, one cannot reject the existence of a unit root in the series' autoregressive representation. That is, one cannot reject that some fraction of an innovation in the series is permanent. Nelson and Plosser also argue for a simple MA(1) representation of real output growth. Our work extends theirs in four ways.

First, we estimate general ARIMA models for real GNP growth. Pure autoregressive and pure moving average models are highly restrictive.¹ More general ARIMA models with relatively few parameters may be better able to capture the dynamics that characterize economic time series.

Second, we show how to test the null hypothesis that a time series is stationary around a deterministic trend. In contrast, Nelson and Plosser test and fail to reject the null hypothesis of nonstationarity. Our test thus provides a natural complement to standard tests of nonstationarity.

Third, we present a nonparametric estimate of persistence, proposed recently by Cochrane [1986]. When applied to postwar quarterly data, this nonparametric procedure confirms the conclusion from the ARIMA models that innovations to real GNP are highly persistent.

Fourth, we try to direct attention away from the question of the existence of a unit root in real GNP, and toward the question of

1. Schwert [1985] demonstrates that omitted moving average components can have serious effects on tests for the presence of unit roots in time series, and often are not well proxied by extra autoregressive terms.

its quantitative importance for GNP behavior.² As we show below, a time series can contain a unit root, while an innovation today has only little effect on one's long-run forecast. Our results suggest not only that a unit root is present, but also that it is essential to understanding economic dynamics.

The organization of the paper is as follows. In Section II we discuss our method for studying persistence by computing the long-run impact of innovations to univariate ARMA models. In Section III we apply the method to quarterly postwar data and, for comparison, to annual postwar data and annual data 1869–1984. Section IV discusses econometric issues that arise in estimating ARMA models. Section V presents nonparametric estimates of persistence, and Section VI concludes. An Appendix gives technical details on maximum likelihood estimation of ARMA models.

II. MEASURING THE LONG-RUN IMPACT OF INNOVATIONS

Suppose that real GNP falls 1 percent lower than one would have expected from its past history. How much should one change one's forecast of GNP for five or ten years ahead?

In this section we address some methodological issues that arise in formulating a convincing answer to this question.

Detrending and Differencing

The first feature of GNP or similar economic data that becomes apparent to any user is that it has historically drifted upward. GNP was higher in 1960 than in 1950, still higher in 1970, and higher again in 1980. The macroeconometrician must deal with this upward drift in some way. Perhaps the most standard approach (e.g., Blanchard [1981]) is to detrend the data before analysis.

It may be obvious that detrending the data is not well suited for our purposes. Detrending forces the resulting series to be trend-reverting, so that today's innovation has no ultimate effect on output. Thus, detrending presupposes the answer to our question at an infinite horizon.

Of course, it could still be the case that at a large but finite horizon of five or ten years, the detrended series displays a considerable effect of today's innovation. However, in samples of typical

2. It has been pointed out to us that most economists would probably be more uncertain about their forecast of GNP at a 100 year horizon than their forecast at a 50 year horizon. It follows that most economists implicitly believe that log GNP is not stationary around a trend, and perhaps that it has a unit root. But the presence of a unit root does not determine the answer to our question.

size, detrending gives a seriously biased answer to our question, even at a finite horizon, when the time series actually has a unit root. A simple example illustrates this pitfall. Suppose that Y_t , such as the log of GNP, followed a random walk with drift:

$$Y_t = \alpha + Y_{t-1} + \epsilon_t,$$

where α is the drift term, representing long-run growth. If one detrends the Y_t series and then estimates an AR(1) process, the coefficient is severely biased toward zero [Nelson and Kang, 1981]. With 100 observations, as might be the case with postwar quarterly data, Mankiw and Shapiro [1985] show in a Monte Carlo study that the median value of the autoregressive term is 0.91. If one used this biased estimate to answer our question, one would note that $(0.91)^{40} = 0.02$ is a small number and erroneously conclude that innovations in Y_t have little information on Y_{t+40} .

The same problem arises when using time as an explanatory variable in a regression. As first noted by Frisch and Waugh [1933], including a time trend in a regression is numerically identical to detrending all the variables. Hence, because of the above argument, we avoid the use of time trends throughout this paper.

A second response to the upward drift in log GNP is to difference the series. The differenced series, the growth rate of real GNP, appears stationary, allowing one to invoke asymptotic distribution theory. We therefore begin with the differenced series as the primary data.

Two issues arise, however, in using differenced data. First, does differencing the data presuppose the answer to our question? The answer is no, as the following example illustrates. Suppose that Y_t follows an IMA(1,1) process:

$$Y_t - Y_{t-1} = \alpha + \epsilon_t - \theta\epsilon_{t-1}.$$

Then a unit impulse in Y_t changes one's forecast of Y_{t+n} by $(1 - \theta)$ regardless of n . Hence, depending on the value of θ , news about current GNP could have a large or small effect on one's forecast of GNP in ten years. Assuming a unit root is therefore consistent with both great and little long-run persistence.

Second, if Y in fact does not have a unit root but is stationary around a trend, does differencing the data bias our conclusions toward finding excessive persistence? The answer is again, no. This result is discussed below.

Impulse Response Functions

We model the change in log GNP as a stationary ARMA process. That is,

$$(1) \quad \phi(L) \Delta Y_t = \theta(L) \epsilon_t,$$

where

$$\phi(L) = 1 - \phi_1 L - \phi_2 L^2 - \dots - \phi_p L^p,$$

and

$$\theta(L) = 1 + \theta_1 L + \theta_2 L^2 + \dots + \theta_q L^q,$$

This equation can be rearranged to arrive at the moving average representation (or impulse response function) for ΔY_t :

$$(2) \quad \begin{aligned} \Delta Y_t &= \phi(L)^{-1} \theta(L) \epsilon_t \\ &= A(L) \epsilon_t. \end{aligned}$$

If the change in log GNP is stationary, then $\sum_{i=0}^{\infty} A_i^2$ is finite, implying that the limit of A_i as i approaches infinity is zero. In other words, stationarity of the differenced series implies that an innovation does not change one's forecast of growth over a long horizon.

We can derive the moving average representation for the level of Y_t by inverting the difference operator $1 - L$:

$$(3) \quad \begin{aligned} Y_t &= (1 - L)^{-1} A(L) \epsilon_t \\ &= B(L) \epsilon_t, \end{aligned}$$

where

$$(4) \quad B_i = \sum_{j=0}^i A_j.$$

Of course, Y_t need not be stationary, and thus B_i need not approach zero as i approaches infinity. Instead, the limit of B_i is the infinite sum of A_j coefficients, which can also be written as $A(1)$. The value of B_i for large i is exactly what we wish to estimate, since it measures the response of Y_{t+i} to an innovation at time t .

The above representation keeps open the possibility that the level of log GNP is stationary around a deterministic linear trend. In this case, the moving average representation of the difference has a unit root, that is, $\theta(L) = (1 - L)\tilde{\theta}(L)$, where $\tilde{\theta}(L)$ is the moving average component of the process in levels. Thus, if the level process is ARMA(p, q), then the differenced process will be ARMA($p, q+1$). (This implies that allowing for stationarity

requires at least one moving average parameter.³) Direct computation shows that $B(L) = \phi(L)^{-1} \hat{\theta}(L)$, as expected. Hence, modeling ΔY_t as a stationary ARMA process leaves open the question of whether Y_t is stationary.

Parameterization

To estimate the ARMA process, we must choose the parameterization, that is, the number of AR and MA parameters. One approach, suggested by classical statistical methods, is to include as many parameters as are significant at standard levels of inference. We report below the likelihood values for a variety of parameterizations; simple likelihood ratio tests can be used to compare any specification with any more general specification.

Another approach is to choose the optimal parameterization using either the Schwarz [1987] criterion or the Akaike [1974, 1976] criterion.⁴ Both rules involve choosing the parameterization with the maximum likelihood after imposing a penalty for the number of parameters. The two rules differ in the size of the penalty. In particular, the Akaike criterion tells us to maximize

$$-2 \ln L - 2k,$$

where L is the likelihood, and $k = p + q$ is the number of parameters. The Schwarz criterion tells us to maximize

$$-2 \ln L - k \ln T,$$

where T is the number of observations. Since our sample includes 155 observations and $\ln(155)$ is about five, the Schwarz criterion penalizes extra parameters much more heavily.

Note that both criteria are based on the principle that for any given number of parameters ($p + q$), a higher likelihood indicates a better model. A robust strategy, therefore, is to prefer, given the total number of parameters, the ARMA model with the greatest likelihood.

While we report the values of both the Schwarz and the Akaike criteria, we do not rely exclusively on this strategy. First, there is no general agreement about which criterion is best. Second, it is not clear whether these criteria will perform well for our purposes, since

3. The autoregressive representation for the model includes an infinite number of parameters that do *not* die out to zero. Schwert [1985] shows that even if the moving average component does not contain a unit root, long autoregressive representations need not provide good approximations.

4. See Nefci [1982] for a discussion of these criteria.

they were not designed with our question in mind. We therefore report results for a variety of parameterizations to gauge to what extent our conclusions are robust.

Estimation

A key problem in estimating a time series model with moving average parameters is that innovations in the series are not identifiable, even if the parameters of the model are known. Because the autoregressive representation of the model is infinite, in any finite sample the innovation sequence depends on pre-sample information. When the moving average roots are strictly less than unity, the process is called invertible. In this case the dependence on pre-sample information decreases through the sample and can be ignored altogether in large samples. Simple approximate estimators for ARMA models are available which exploit this fact, for example, by assuming that all pre-sample innovations are zero.

Unfortunately, these simple methods do not work well for ARMA processes with moving average roots equal or close to unity. It is known that they tend to produce estimates of the MA parameters whose roots are seriously biased away from unity (see, for example, Plosser and Schwert [1977], Davidson [1981], and Harvey [1981]).

Accordingly, we use an exact maximum likelihood estimation method which explicitly recognizes that the innovation sequence is unobservable. We use a Kalman filter to build up the log likelihood function of the model as a sum of conditional log likelihoods. Full details are given in Harvey [1981]; we summarize the approach in the Appendix.

III. RESULTS

We estimate the ARMA process (1) for the differenced series and calculate the implied impulse response function for the level of the series (B_i 's) using real GNP data for the United States. We use 1982 base year, seasonally adjusted, quarterly data from 1947:1 to 1985:4.⁵ We consider all ARMA models for the difference of log real GNP with up to three AR parameters and three MA parameters. There are thus 16 models under consideration for GNP growth, the simplest being white noise, the most complex the ARMA(3,3).

5. These data are identical to those used by Clark [1987]. An earlier version of the present paper [Campbell and Mankiw, 1986] used 1972 base year data, with fairly similar results.

TABLE I
MODEL SELECTION CRITERIA, $\Delta \ln$ REAL GNP

Number of AR parameters (p)	Number of MA parameters (q)			
	0	1	2	3
0	952.558	968.034	979.983	981.305
	(952.558)	(966.034)	(975.983)	(975.305)
	(952.558)	(962.991)	(969.896)	(966.175)
		[523.322]	[691.069]	[805.201]
1	974.678	976.112	980.952	982.297
	(972.678)	(972.112)	(974.952)	(974.297)
	(969.635)	(966.025)	(965.822)	(962.123)
		[953.152]	[969.224]	[982.297]
2	977.522	979.294	983.797	984.372
	(973.522)	(973.294)	(975.747)	(974.254)
	(967.435)	(964.163)	(963.573)	(959.155)
		[977.234]	[979.772]	[984.372]
3	981.392	982.725	984.428	985.842
	(975.322)	(974.725)	(974.428)	(973.842)
	(966.262)	(962.551)	(959.211)	(955.581)
		[981.773]	[983.122]	[985.842]

For each model, we report $2 \ln L$ (Akaike Criterion $= 2 \ln L - 2k$), (Schwarz Criterion $= 2 \ln L - k \ln T$), [$2 \ln L$ of model restricted to have a unit moving average root].

Table I reports the selection criteria for the 16 models. Three models stand out as particularly worthy of attention. The MA(2) is not rejected at the 5 percent level by any more general model and is selected by both the Akaike and Schwarz criteria. The AR(1) has almost as high a value of the Schwarz criterion.⁶ Finally, the ARMA(2,2) has almost as high a value of the Akaike criterion and nests the other two models; the MA(2) can be rejected against it at the 15.2 percent level, and the AR(1) at the 2.8 percent level. If we adopt the robust strategy of choosing the model with the highest likelihood given the number of parameters k , we are led to adopt the AR(1) for $k = 1$, the MA(2) for $k = 2$, the AR(3) for $k = 3$, and the ARMA(2,2) for $k = 4$.

Our main interest is not in selecting one particular model for real GNP growth but in reaching a conclusion about the persistence of GNP fluctuations. We use several devices for doing this. First, we

6. Interestingly, this is the process that Deaton [1985] suggests for labor income growth and Watson [1986] for GNP growth. Watson goes on to argue for an unobservable components model that implies a restricted ARMA(2,2) representation. Clark [1987] estimates a similar model.

present in Table I the maximum likelihood obtainable for each model, under the constraint that the moving average parameters sum to minus one, or equivalently that the limit of the impulse response function in levels is zero. The table shows that this constraint causes a very large drop in likelihood when it is imposed on the most parsimonious models we consider, with up to one autoregressive parameter and two moving average parameters. The drop in the likelihood is smaller when we impose the constraint on higher order models (for example, a standard likelihood ratio test would reject the constrained ARMA(2,2) at the 4.6 percent level

TABLE II
MODEL PARAMETER ESTIMATES, $\Delta \ln$ REAL GNP

Model p, q	ϕ_1	ϕ_2	ϕ_3	θ_1	θ_2	θ_3
0,1	—	—	—	0.261* (0.072)	—	—
0,2	—	—	—	0.305* (0.073)	0.268* (0.075)	—
0,3	—	—	—	0.323* (0.077)	0.324* (0.074)	0.107 (0.080)
1,0	0.363* (0.070)	—	—	—	—	—
1,1	0.522* (0.178)	—	—	-0.179 (0.210)	—	—
1,2	0.239 (0.258)	—	—	0.083 (0.264)	0.237* (0.100)	—
1,3	0.950* (0.037)	—	—	-0.678* (0.122)	-0.041 (0.098)	-0.281* (0.081)
2,0	0.314* (0.073)	0.134 (0.074)	—	—	—	—
2,1	-0.134 (0.340)	0.313* (0.126)	—	0.455 (0.353)	—	—
2,2	0.626* (0.216)	-0.462* (0.171)	—	-0.325 (0.201)	0.592* (0.122)	—
2,3	1.280* (0.239)	-0.329 (0.240)	—	-0.992* (0.273)	0.232 (0.200)	-0.241* (0.098)
3,0	0.336* (0.076)	0.184* (0.075)	-0.156 (0.080)	—	—	—
3,1	0.801* (0.393)	0.039 (0.156)	-0.221* (0.078)	-0.481 (0.402)	—	—
3,2	0.628* (0.191)	-0.631* (0.225)	0.109 (0.134)	-0.310 (0.168)	0.737* (0.195)	—
3,3	1.578* (0.258)	-0.986* (0.417)	0.370 (0.197)	-1.299* (0.279)	0.830* (0.320)	-0.531* (0.140)

Standard errors are in parentheses. An asterisk indicates significance at the 5 percent level.

against the general ARMA(2,2)); and in the ARMA(1,3), ARMA(2,3), and ARMA(3,3) models, our unconstrained estimates obey the constraint exactly.

Tables II, III, and IV present the estimated parameters, roots, and impulse response functions for the unconstrained models. The parameters in Table II appear to be somewhat unstable across models, particularly when one compares the models that have a unit moving average root with those that do not. In Table III, however, the reason for this becomes clear. The models with a unit moving average root also have an autoregressive root that is very close to unity, and that almost cancels the moving average root. The remaining roots are similar to those of the lower-order models.

This phenomenon of near cancellation of roots is also apparent in the shape of the likelihood function for the ARMA(1,3), ARMA(2,3), and ARMA(3,3) models. These models have a second peak in the likelihood function with no unit moving average root and almost the same likelihood (the difference in twice the log likelihood is only 0.118 for the ARMA(2,3)). The impulse response function for the second peak is similar to that for the lower-order models.

The impulse response functions in Table IV are of two types. For the models with no unit moving average root, the impulse response increases above one and settles between 1.3 and 1.9 at about the eighth quarter, remaining there even at ten or twenty

TABLE III
ROOTS OF ESTIMATED MODELS, $\Delta \ln$ REAL GNP

Model p,q	AR roots	MA roots
0,1	—	-3.831
0,2	—	$-0.569 \pm 1.846i$
0,3	—	$-3.045, 0.008 \pm 1.752i$
1,0	2.755	—
1,1	1.916	5.587
1,2	4.185	$-0.175 \pm 2.047i$
1,3	1.053	$1.000, -0.573 \pm 1.797i$
2,0	1.801, -4.144	—
2,1	2.014, -1.586	-2.198
2,2	$0.677 \pm 1.306i$	$0.274 \pm 1.270i$
2,3	2.808, 1.082	$1.000, -0.018 \pm 2.038i$
3,0	$-1.852, 1.516 \pm 1.079i$	—
3,1	$-2.281, 1.229 \pm 0.688i$	2.079
3,2	$5.004, 0.392 \pm 1.296i$	$0.210 \pm 1.146i$
3,3	$1.053, 0.806 \pm 1.385i$	$1.000, 0.282 \pm 1.343i$

TABLE IV
MODEL IMPULSE RESPONSES, In REAL GNP

Model p,q	1	2	4	8	16	20	40	80
0,1	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)	1.261 (0.072)
0,2	1.305 (0.073)	1.573 (0.123)	1.573 (0.123)	1.573 (0.123)	1.573 (0.123)	1.573 (0.123)	1.573 (0.123)	1.573 (0.123)
0,3	1.323 (0.077)	1.647 (0.128)	1.754 (0.170)	1.754 (0.170)	1.754 (0.170)	1.754 (0.170)	1.754 (0.170)	1.754 (0.170)
1,0	1.363 (0.070)	1.496 (0.120)	1.561 (0.161)	1.571 (0.171)	1.571 (0.172)	1.571 (0.172)	1.571 (0.172)	1.571 (0.172)
1,1	1.344 (0.077)	1.523 (0.119)	1.666 (0.202)	1.715 (0.268)	1.719 (0.278)	1.719 (0.279)	1.719 (0.279)	1.719 (0.279)
1,2	1.322 (0.075)	1.635 (0.130)	1.728 (0.206)	1.734 (0.222)	1.734 (0.222)	1.734 (0.222)	1.734 (0.222)	1.734 (0.222)
1,3	1.271 (0.119)	1.488 (0.269)	1.341 (0.572)	1.090 (1.110)	0.721 (1.895)	0.586 (2.177)	0.208 (2.958)	0.026 (3.338)
2,0	1.314 (0.073)	1.547 (0.116)	1.730 (0.201)	1.804 (0.264)	1.812 (0.276)	1.812 (0.276)	1.812 (0.276)	1.812 (0.276)
2,1	1.321 (0.071)	1.591 (0.122)	1.731 (0.198)	1.770 (0.242)	1.772 (0.248)	1.772 (0.248)	1.772 (0.248)	1.772 (0.248)
2,2	1.302 (0.078)	1.621 (0.128)	1.572 (0.193)	1.532 (0.142)	1.517 (0.162)	1.517 (0.160)	1.517 (0.161)	1.517 (0.161)
2,3	1.289 (0.119)	1.561 (0.268)	1.502 (0.596)	1.115 (1.178)	0.592 (1.921)	0.431 (2.140)	0.088 (2.599)	0.004 (2.720)
3,0	1.336 (0.076)	1.632 (0.132)	1.641 (0.207)	1.568 (0.230)	1.571 (0.223)	1.571 (0.222)	1.571 (0.222)	1.571 (0.222)
3,1	1.320 (0.077)	1.614 (0.131)	1.604 (0.206)	1.334 (0.327)	1.364 (0.288)	1.360 (0.297)	1.360 (0.297)	1.360 (0.297)
3,2	1.318 (0.078)	1.624 (0.127)	1.630 (0.210)	1.626 (0.196)	1.595 (0.206)	1.596 (0.203)	1.597 (0.203)	1.597 (0.203)
3,3	1.279 (0.122)	1.563 (0.267)	1.416 (0.602)	1.095 (1.141)	0.720 (1.929)	0.584 (2.213)	0.207 (3.001)	0.026 (3.389)

Standard errors are in parentheses.

years. That is, a 1 percent innovation in real GNP increases the univariate forecast of GNP by over 1 percent over any foreseeable horizon.

The impulse response functions for the ARMA(1,3), ARMA(2,3), and ARMA(3,3) of course behave differently. They die out to zero, but very slowly; after five years 40 percent to 60 percent of a shock is still present in GNP. The standard error on the impulse response for these models is very large. This reflects the near cancellation of roots discussed above.

The evidence of persistence in the quarterly postwar GNP data

is robust to change in the sample, within the postwar period, and to change in the frequency of the data. If we end our sample in 1972, prior to OPEC and to the productivity slowdown, we continue to find impulse response functions above one. When we examine postwar annual data, we cannot reject the hypothesis that the log of real GNP is a random walk with drift. In this case, the impulse response is unity at all horizons.

Cochrane [1986] has recently challenged the view that output fluctuations are highly persistent. He studies a long time series of annual per capita real GNP numbers for the period 1869–1984, constructed by splicing data from Friedman and Schwartz [1982] to the standard postwar series.

When we apply our methods to annual real GNP 1869–1984,⁷ we find much less persistence than in postwar data. Low-order pure AR and MA models have impulse response functions whose limits are above one, but only just: the AR(1) limit is 1.158 (standard error 0.102), the AR(2) limit is 1.077 (standard error 0.101), and so forth. Higher-order pure models have impulse response limits somewhat below one: the AR(3) limit is 0.962 (0.125), and the MA(3) limit is 0.852 (0.126). Higher-order mixed models have limits well below one: the ARMA(2,1) limit is 0.144 (0.258), and the ARMA(2,2) limit is 0.161 (0.247). The mixed models with low limits have considerably higher likelihoods than the models with limits close to one; the white noise, AR(1), AR(2), and MA(1) models are all rejected in favor of the ARMA(2,1) model at about the 1 percent level, and the Akaike criterion also selects the ARMA(2,1) over these models.

The low persistence of the annual series, 1869–1984, is due largely to the pre-1929 data. In the pre-1929 data, even the AR(1) model has a negative coefficient of -0.166 , implying a long-run impulse response of 0.858. Recently the conventional measures of GNP before 1929 have been challenged [Romer, 1986], and we believe that results based largely on these data should be treated with caution.

IV. ECONOMETRIC ISSUES IN ESTIMATING ARMA MODELS

In recent years economists have become more aware of various pitfalls in applied time-series econometrics. Even apparently straightforward procedures can suffer from severe problems of bias

7. We do not deflate by population, in order to maintain comparability with our postwar data. In fact, real GNP and per capita real GNP seem to have very similar time series properties.

in samples of typical size.⁸ In this section, therefore, we review the literature on the asymptotic and finite-sample properties of our estimator, and present a very small Monte Carlo study with 20 simulations of a process with a unit root in the moving average component.

There is a small recent literature analyzing the properties of maximum likelihood estimates of ARMA model parameters [Plosser and Schwert, 1977; Ansley and Newbold, 1980; Davidson, 1981; Harvey, 1981; Pesaran, 1981; Sargan and Bhargava, 1983]. When the moving average roots are strictly less than unity, the maximum likelihood estimator is consistent and asymptotically normal. When there are unit moving average roots, however, these results break down.

The reason for this can be understood most easily by considering the first-order moving average (MA(1)) case. An MA(1) model with a parameter value of θ^* is observationally equivalent to a model with a parameter value of $1/\theta^*$, because these two models imply the same autocovariances for the series they describe. This "fundamental identification problem" [Plosser and Schwert, 1977] is handled by restricting attention to models with parameter values less than or equal to unity. A parameter equal to unity is on the boundary of the restricted parameter space, and for this reason maximum likelihood estimates of a model with unit parameter do not have the usual asymptotic distribution [Chernoff, 1954]. A similar problem occurs in higher-order models when a moving average root is unity.

Other authors have conducted Monte Carlo simulations to characterize the properties of maximum likelihood estimates of a first-order moving average parameter in finite samples. Ansley and Newbold [1980], Harvey [1981], and Davidson [1981] all report finding a probability mass at exactly unity, particularly large when the true root is unity but also present when it is considerably less than unity. This occurs because the observational equivalence of θ^* and $1/\theta^*$ discussed above causes the sample likelihood function to be flat at $\theta = 1$. In finite samples this local maximum or minimum will often be the global maximum of the sample likelihood function, even if the true root is less than unity. It follows that, in Davidson's words, "the occurrence of boundary estimates in empirical work with the exact maximum likelihood estimator is very weak evidence of over-differencing." The rest of the distribution of the estimator is

8. See, for example, Flavin [1983] and Mankiw and Shapiro [1986].

roughly bell-shaped and centered on the true value when this is less than unity.

These results lead one to expect that a standard likelihood ratio test of the hypothesis that a moving average root is unity will tend to reject the hypothesis at the 5 percent level less than 5 percent of the time when the hypothesis is true. Davidson reports some Monte Carlo results that confirm this expectation. He finds that the test rejects only four times out of 200 when the 5 percent critical value of the Chi-squared distribution with one degree of freedom is used. This suggests that a rejection using this test, of the kind we report for the ARMA(2,2) model of output growth, is legitimate evidence against the hypothesis of a unit moving average root.

Finally, Ansley and Newbold report Monte Carlo results for the computed standard errors of the maximum likelihood parameter estimates. They find that, in samples of size 50 and 100, standard errors are often too small, particularly when the true parameter values display near parameter redundancy (that is, when an autoregressive root almost cancels with a moving average root).

Our interest in this paper is in precisely the difficult case where there may be a unit moving average root, and the time series may display near parameter redundancy. To get a sense of the behavior of our estimator under the hypothesis that log GNP is in fact stationary, we ran a small Monte Carlo experiment and applied our estimator to the first differences of 20 randomly generated series, each with 155 observations, which follow an AR(2) in levels. The first AR parameter was 1.34, and the second was -0.42 —the values estimated by Blanchard [1981]. We estimated an ARMA(2,2) in first differences, an overparameterized model. For each series we conducted a likelihood ratio test of the hypothesis that the moving average terms have a unit root (sum to -1), and we estimated the impulse response at horizon 80 with standard errors. The results are reported in Table V.

The number of runs is of course too small to draw any strong conclusions from the table. However, the results are in line with those reported in the literature. The likelihood ratio test of the unit root restriction does not reject more often than it should under the null hypothesis. Furthermore, for 14 out of the 20 runs, the unrestricted estimate of the root is exactly unity (to two decimal places). The unrestricted estimator has a probability mass at this value for the root. When the moving average root is estimated equal to unity, the impulse responses have extremely large standard

TABLE V
MONTE CARLO STUDY OF MAXIMUM LIKELIHOOD ESTIMATOR

Run	Unrestricted -2*(log L)	Restricted -2*(log L)	Likelihood ratio test statistic	Impulse response	Standard error
1	460.24	460.24	0.00	0.00	(1.42)
2	424.76	424.76	0.00	0.00	(0.93)
3	403.92	403.92	0.00	0.00	(1.33)
4	440.20	440.20	0.00	0.00	(0.91)
5	432.18	432.79	0.61	0.28	(0.28)
6	461.26	461.26	0.00	0.00	(1.30)
7	453.53	454.16	0.63	0.27	(0.23)
8	446.67	446.67	0.00	0.00	(1.69)
9	446.22	446.22	0.00	0.00	(0.71)
10	438.06	438.08	0.02	0.11	(0.49)
11	447.24	447.46	0.22	1.36	(0.15)
12	441.50	441.50	0.00	0.00	(1.01)
13	443.17	443.17	0.00	0.00	(2.94)
14	448.29	448.29	0.00	0.00	(1.88)
15	453.56	453.56	0.00	0.00	(1.26)
16	472.68	474.45	1.77	0.39	(0.28)
17	414.30	414.89	0.59	0.22	(0.21)
18	450.96	451.27	0.31	0.68	(0.26)
19	421.72	421.72	0.00	0.00	(1.48)
20	431.75	431.75	0.00	0.00	(1.22)

Notes. This table reports the results of estimating an ARMA(2,2) in first differences for 20 randomly generated series, each with 150 observations, which are AR(2) in levels with parameters 1.34 and -0.42. The impulse responses are at a horizon of 80 periods.

errors. When it is estimated away from unity, they typically have rather small standard errors; in two cases the implied 95 percent confidence intervals do not include the true value. This result is consistent with Ansley and Newbold's findings.

We conclude from our literature review and our small Monte Carlo study that while there are some statistical difficulties with our estimator, there is no reason to think that these bias us toward rejecting stationarity.⁹ In fact, they offer an explanation for the exact unit root found in the ARMA(1,3), ARMA(2,3), and ARMA(3,3) models for GNP growth. The major caveat from the statistical literature, and our own small Monte Carlo study, is that

9. There would be a bias against finding stationarity if we had used an approximate estimator that sets pre-sample innovations to zero; see, for example, Plosser and Schwert [1977].

standard errors on parameters and impulse response functions may be too small when there is near parameter redundancy.

Clearly it would be desirable to have some results on the distribution of the likelihood ratio test statistic for a unit moving average root; this is a topic that we hope to pursue in future research.

V. AN ALTERNATIVE MEASURE OF PERSISTENCE

Cochrane [1986] has recently proposed an alternative measure of persistence that can be estimated nonparametrically. In this section we show how Cochrane's measure is related to ours; we state a simple formula for its asymptotic standard error; and we present estimates for our quarterly postwar data and for annual data since 1869. The results are consistent with our findings from univariate ARMA models.

Recall that our measure of persistence at an infinite horizon is $A(1)$, the infinite sum of moving average coefficients for the differenced process. Cochrane starts, not from moving average coefficients, but from autocovariances of the differenced process. Writing the j th autocovariance as C_j , and defining $C(z) \equiv \sum_{j=-\infty}^{\infty} C_j z^j$, a well-known result relates autocovariances to moving average coefficients as $C(z) = A(z)A(z^{-1})\sigma_\epsilon^2$. Here, as before, σ_ϵ^2 is the variance of univariate innovations to the differenced process. The variance of the differenced process itself, σ^2 , can also be written as C_0 . Finally, it will be convenient to write the j th autocorrelation of the differenced process as $\rho_j \equiv C_j/C_0$.

Cochrane's measure of persistence, which we shall write as V , is just $V = \rho(1) = C(1)/C_0$, the two-sided infinite sum of autocorrelations. It is straightforward to show that this is related to our measure by

$$V = (\sigma_\epsilon^2/\sigma^2)[A(1)]^2.$$

Defining $R^2 = 1 - \sigma_\epsilon^2/\sigma^2$, the fraction of the variance that is predictable from knowledge of the past history of the process, we have

$$(5) \quad A(1) = \sqrt{\frac{V}{1 - R^2}}.$$

Equation (5) shows that the square root of Cochrane's persistence measure is a lower bound on our measure. The more highly

predictable is the differenced process, the greater is the difference between our measure and the square root of Cochrane's measure.

Cochrane proposes an estimator of V which we shall write as $\hat{V}^k$. The estimator uses sample autocorrelations:

$$\hat{\rho}_j = \left(\frac{T}{T-j} \right) \sum_{t=j+1}^T (\Delta y_t \Delta y_{t-j}) / \sum_{t=1}^T (\Delta y_t)^2.$$

The first k of these are used, weighted according to the following scheme:

$$(6) \quad \hat{V}^k = 1 + 2 \sum_{j=0}^k \left(1 - \frac{j}{(k+1)} \right) \hat{\rho}_j.$$

This "triangular" pattern gives higher-order autocorrelations linearly declining weights, out to the k th autocorrelation. $\hat{V}^k$ is a consistent estimate of V , provided that k increases with the sample size.¹⁰

$\hat{V}^k$ can be interpreted in various ways. Cochrane points out that it is the ratio of the variance of $(k+1)$ -period differences of the process, to the variance of one-period differences. It is also closely related to the autocorrelation-consistent covariance matrix estimator of Newey and West [1985]. Finally, it can be interpreted in terms of the frequency domain, as an estimate of the normalized spectral density at frequency zero which uses a "Bartlett window."¹¹ A result in spectral analysis (see, e.g., Priestley [1982, p. 463]) gives the asymptotic standard error of $\hat{V}^k$ as

$$\text{S.e.}[\hat{V}^k] = \frac{\hat{V}^k}{\sqrt{\frac{3}{4} \frac{T}{(k+1)}}}.$$

where T is the sample size. The usefulness of this standard error in samples of typical size, however, is unclear.

As an estimator of V , $\hat{V}^k$ has a number of advantages. Notably, it has a smaller asymptotic standard error than some other simple estimators, for example, the one which weights all autocorrelations equally to lag k . However, one must be careful not to misinterpret the behavior of $\hat{V}^k$ as k increases to the point where it approaches T :

10. k must not increase too rapidly, however; see Newey and West [1986].

11. As written, (6) can have a negative sample value; this possibility would be ruled out if $\hat{\rho}_j$ were multiplied by $(T-j)/T$. For small values of k relative to T , the correction is negligible.

for $k = T - 1$, $\hat{V}^k$ is identically equal to zero when the sample mean has been subtracted from the data. Intuitively this is because demeaning forces the observations to sum to zero, creating a large number of small negative sample autocorrelations.

In Campbell and Mankiw [1987] we report the results of a Monte Carlo study of the behavior of $\hat{V}^k$ in a sample of size 130. We compare the estimates of $\hat{V}^k$ with $k = 10$ to 100, for a random walk and a stationary AR(2) with coefficients 1.34 and -0.42 . Our main results are three. First, the window size k must be at least 30, and preferably 40 or 50, if one is to be able to discriminate between these two processes. Second, there is severe downward bias in $\hat{V}^k$; for the random walk, the mean of $\hat{V}^k$ is approximately $(T - k)/T$ rather than unity. And finally, there is a great deal of sample variation in $\hat{V}^k$ so one must be cautious in making inferences based on this estimator.

In Table VI we report values of $\hat{V}^k$ for the postwar quarterly real GNP data, along with asymptotic standard errors, using window sizes from 10 to 75. We also report $\hat{A}^k(1)$, the values of $A(1)$ that correspond to these estimates of V , using the square of the first sample autocorrelation, $\hat{\rho}_1^2$, as a conservative estimate of R^2 .

The $\hat{V}^k$ values in Table VI start out well above unity but fall gradually and are below unity for window sizes of 40 and above. Yet $\hat{V}^k$ for the real GNP data are consistently above $(T - k)/T$, so the

TABLE VI
NONPARAMETRIC ESTIMATES OF PERSISTENCE

Window size (k)	$\hat{V}^k$	$\hat{A}^k$
10	1.676 (0.515)	1.392
20	1.302 (0.553)	1.227
30	1.131 (0.584)	1.144
40	0.882 (0.524)	1.010
50	0.689 (0.456)	0.892
60	0.732 (0.530)	0.920
75	0.784 (0.634)	0.952

Standard errors are in parentheses.

TABLE VII
NONPARAMETRIC ESTIMATES OF PERSISTENCE: ANNUAL DATA

Window Size (k)	Sample period					
	1869-1984		1869-1929		1930-1984	
	$\hat{V}^k$	$\hat{A}^k$	$\hat{V}^k$	$\hat{A}^k$	$\hat{V}^k$	$\hat{A}^k$
10	0.588 (0.210)	0.774	0.293 (0.145)	0.549	0.705 (0.367)	0.941
20	0.304 (0.150)	0.557	0.096 (0.065)	0.314	0.248 (0.178)	0.558
30	0.270 (0.162)	0.525	0.102 (0.085)	0.324	0.175 (0.154)	0.470
40	0.182 (0.126)	0.431	—	—	—	—
50	0.168 (0.129)	0.414	—	—	—	—
60	0.081 (0.068)	0.287	—	—	—	—

values are larger than one would expect to find for a random walk in a sample of this size. The implied values of $\hat{A}^k(1)$ are also close to unity at all window sizes. The nonparametric estimates thus confirm our finding that postwar quarterly real GNP appears to be more persistent than a random walk.

Finally, we note that one obtains much less persistence if one calculates $\hat{V}^k$ for annual real GNP 1869-1984 (Table VII). In this series, $\hat{V}^k$ is well below unity at 0.588, even for $k = 10$, and falls to 0.081 for $k = 60$.¹² Fluctuations appear particularly transitory before 1929. For the first half of the sample (1869-1929), $\hat{V}^k$ is 0.293 for $k = 10$; while for the second half (1930-1984), it is 0.705 for $k = 10$. At larger window sizes there is less difference across the two halves of the sample, but the persistence estimates remain larger in the second half. Hence, the low level of persistence reported by Cochrane is partly attributable to his use of pre-1929 data.

VI. CONCLUSION

We have estimated standard ARIMA processes for the log of United States real GNP using the standard post-war quarterly time

12. In Table VII we report persistence estimates only for window sizes up to about half the sample size, since small-sample bias dominates the results for larger window sizes.

series. Yet the estimates have a surprising implication: a 1 percent innovation in real GNP should change one's forecast of real GNP by over 1 percent over a long horizon.¹³

This finding should be interpreted with caution, since work on the small sample properties of ARIMA estimation, especially in the presence of unit roots, is only in its infancy. Moreover, the likelihood function is sufficiently flat that it appears impossible to reject the view that output reverts to trend after twenty years. This may explain why imposing prior restrictions, such as those of Clark [1987], can substantially change one's estimate of persistence. Yet we are encouraged by the fact that a nonparametric approach also suggests high persistence, a result we also find for most other major countries [Campbell and Mankiw, 1987b].

Many traditional theories of the business cycle maintain two fundamental premises. First, fluctuations in output are assumed to be driven primarily by shocks to aggregate demand, such as monetary policy, fiscal policy, or animal spirits. Second, shocks to aggregate demand are assumed to have only a temporary effect on output; in the long run the economy returns to the natural rate. These two premises underlie many monetarist and neo-Keynesian theories.

If output fluctuations are highly persistent, both of these premises cannot be maintained. It is not clear, however, which of these two premises should be called into question.

Nelson and Plosser argue that the first premise, that fluctuations are driven by aggregate demand (in particular, monetary disturbances) should be abandoned. They advocate models in which fluctuations are attributable to changes in aggregate supply, such as shifts in the available production technology. Certainly the evidence of persistence presented by Nelson and Plosser and in this paper is consistent with such real business cycle models.

Even if one concludes with Nelson and Plosser that real shocks dominate as a source of output fluctuations, these shocks need not work through the mechanisms highlighted in real business cycle models. The real business cycle model of Kydland and Prescott [1982], for example, is an intertemporal Walrasian model, so that the allocation of resources is always Pareto efficient. It is possible that economic fluctuations are driven by real shocks but that these real shocks affect the economy through some Keynesian channel.

A conclusion as extreme as that of Nelson and Plosser is of

13. We are told that commercial forecasters have long known this result: when forecasts are updated on the basis of new information, real GNP is increased (or decreased) approximately proportionately at all horizons.

course not necessary. One can attribute a major role to supply shocks without completely abandoning a role for demand shocks. For example, suppose that output Y is the sum of two components, supply-driven "trend" Y^T and demand-driven "cycle" Y^C , that are uncorrelated at all leads and lags. Suppose further that ΔY^T is a first-order autoregressive process with parameter ρ and that Y^C is some stationary process. Then one can show that the persistence measure V equals $[(1 + \rho)/(1 - \rho)]\text{var}(\Delta Y^T)/[\text{var}(\Delta Y^T) + \text{var}(\Delta Y^C)]$. The finding that $V \geq 1$ therefore implies that $\text{var}(\Delta Y^C)/\text{var}(\Delta Y^T) \leq 2\rho/(1 - \rho)$. If trend output is approximately a random walk, so that ρ is small, then the finding of great persistence implies that fluctuations in the cycle are small relative to fluctuations in the trend. If the change in trend is highly serially correlated, however, the finding of persistence is consistent with a substantial cyclical component.

The second way to interpret the finding of persistence is to abandon the second premise, the natural rate hypothesis. There are a variety of possible mechanisms through which aggregate demand shocks might have permanent or near permanent effects of the level of output. Models of multiple equilibria (e.g., Diamond [1984]) might explain a long-lasting effect of aggregate demand if shocks to aggregate demand can move the economy between equilibria. Shocks to aggregate demand could have permanent effects if technological innovation is affected by the business cycle. Perhaps models of temporary nominal rigidities (e.g., Fischer [1977]) or misperceptions (e.g., Lucas [1973]) could be reconciled with findings of persistence by abandoning the natural rate hypothesis in favor of some highly potent propagation mechanism.

APPENDIX: MAXIMUM LIKELIHOOD ESTIMATION OF ARMA MODELS

If the change in log GNP, ΔY_t , follows an ARMA(p, q) process, it can be written as one element of a vector Markov process α_t , where α_t obeys

$$\alpha_t = T\alpha_{t-1} + R\eta_t$$

$$(A1) \quad T = \left[\begin{array}{c|c} \phi_1 & \\ \vdots & \\ \vdots & \\ \phi_{m-1} & \\ \hline \phi_m & 0 \dots 0 \end{array} \right] \quad R = \left[\begin{array}{c} 1 \\ \theta_1 \\ \vdots \\ \theta_{m-1} \end{array} \right].$$

Here the ϕ are the AR parameters and the θ are the MA parameters. $m = \max(p, q + 1)$ and $\theta_i = 0$ for $i > q$, $\phi_i = 0$ for $i > p$. The innovation process η_t is assumed to be normal white noise with variance σ^2 . ΔY_t is the first element of α_t , so we have $\Delta Y_t = z'\alpha_t$, where $z' = [1 \ 0 \dots 0]$.

The steady-state distribution of α is normal with mean a_0 and variance $\sigma^2 P_0$. In our application we subtract the sample mean from the data and set $a_0 = 0$. P_0 is given by

$$(A2) \quad \text{vec}(P_0) = [I - T \otimes T]^{-1} \text{vec}(RR').$$

Given a_0 and P_0 , one can compute, for $t = 1, \dots, T$, the following quantities. First, the one-step-ahead *prediction* of α_t , conditional on time $t - 1$ information, is $a_{t|t-1} = Ta_{t-1}$. Then $\Delta Y_{t|t-1} = z'a_{t|t-1}$. The conditional variance-covariance matrix of the errors in the one-step-ahead prediction of α_t is $P_{t|t-1} = TP_{t-1}T' + RR'$, and the conditional variance of the error in the one-step-ahead prediction of ΔY_t is $f_t = z'P_{t|t-1}z$.

Using the observation of ΔY_t , one can compute the prediction error itself, $\nu_t = \Delta Y_t - \Delta Y_{t|t-1}$. Finally, one *updates* for the next round, setting $a_t = a_{t|t-1} + P_{t|t-1}z\nu_t/f_t$ and $P_t = P_{t|t-1} - P_{t|t-1}zz'P_{t|t-1}/f_t$.

Once one has computed ν_t and f_t for the whole sample $t = 1, \dots, T$, one can form the log likelihood function for the sample as

$$(A3) \quad -\left(\frac{T}{2}\right) \log(2\pi) - \left(\frac{T}{2}\right) \log(\sigma^2) - \left(\frac{1}{2}\right) \sum_{t=1}^T \log f_t - \left(\frac{1}{2}\sigma^2\right) \sum_{t=1}^T \frac{\nu_t^2}{f_t}.$$

We maximize this likelihood function using a method of scoring with modified step size [Berndt et al., 1974]. We compute an asymptotic variance-covariance matrix for the parameters, Γ , as the inverse of the moment matrix of the numerical derivatives of the conditional log likelihoods with respect to the parameters. A model with parameter restrictions can be estimated in a similar manner, and the likelihood ratio computed.

Since the process α_t is Markov, it is straightforward to obtain the impulse response function of the ΔY_t process from equation (1) in the text given the parameters. The impulse response at horizon k , A_k , is just

$$(A4) \quad A_k = z'T^k R.$$

To compute the impulse response at horizon k in *levels*, B_k , one simply uses equation (4) in the text and sums A_i for $i = 0, \dots, k$. This estimate of the impulse response is a nonlinear function of the

parameters, whose limit as k increases is

$$z'[I - T]^{-1}R.$$

Its asymptotic standard error can be estimated as $\sqrt{(d'Td)}$, where d is the vector of derivatives of the function with respect to the parameters. We computed d numerically.

PRINCETON UNIVERSITY
HARVARD UNIVERSITY

REFERENCES

- Akaike, H., "A New Look at Statistical Model Identification," *IEEE Transactions on Automatic Control*, AC-19 (1974), 716-23.
- , "Canonical Correlation Analysis of Time Series and the Use of an Information Criterion," in *System Identification: Advances and Case Studies*, R. K. Mehra and D. G. Lainiotis, eds. (New York: Academic Press, 1976), pp. 52-107.
- Ansley, Craig F., and Paul Newbold, "Finite Sample Properties of Estimators for Autoregressive Moving Average Models," *Journal of Econometrics*, XIII (1980), 1959-83.
- Berndt, Ernst, Bronwyn Hall, Robert Hall, and Jerry Hausman, "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement* (1974), 653-65.
- Blanchard, Olivier, "What Is Left of the Multiplier-Accelerator?" *American Economic Review Papers and Proceedings*, LXXI (1981), 150-54.
- Campbell, John Y., and N. Gregory Mankiw, "Are Output Fluctuations Transitory?" NBER Working Paper No. 1916, May 1986.
- , and —, "Permanent and Transitory Components in Macroeconomic Fluctuations," *American Economic Review Papers and Proceedings*, LXXVII (1987a), 111-17.
- , and —, "International Evidence on the Persistence of Economic Fluctuations," unpublished paper, 1987b.
- Chernoff, H., "On the Distribution of the Likelihood Ratio Test Statistic," *Annals of Mathematical Statistics*, XXV (1954), 573-78.
- Cochrane, John H., "How Big Is the Random Walk in GNP?" unpublished paper, University of Chicago, April 1986.
- Davidson, James E. H., "Problems with the Estimation of Moving Average Processes," *Journal of Econometrics*, XVI (1981), 295-310.
- Deaton, Angus S., "Life-Cycle Models of Consumption: Is the Evidence Consistent With the Theory?" Princeton University, 1985.
- Diamond, Peter A., *A Search-Equilibrium Approach to the Micro Foundations of Macroeconomics* (Cambridge, MA: MIT Press, 1984).
- Fischer, Stanley, "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy*, LXXXV (1977), 191-205.
- Flavin, Marjorie A., "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence," *Journal of Political Economy*, XCI (1983), 929-56.
- Friedman, Milton, and Anna J. Schwartz, *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices and Interest Rates, 1967-1975* (Chicago, IL: University of Chicago Press, 1982).
- Frisch, Ragnar, and Frederick V. Waugh, "Partial Time Regressions as Compared with Individual Trends," *Econometrica*, I (1933), 387-401.
- Harvey, A. C., *Time Series Models* (Oxford: Philip Allan, 1981).
- Kydland, Finn E., and Edward C. Prescott, "A Competitive Theory of Fluctuations and the Feasibility and Desirability of Stabilization Policy," in S. Fischer, ed. *Rational Expectations and Economic Policy* (Chicago, IL: University of Chicago Press, 1980).

- , and —, "Time to Build and Aggregate Fluctuations," *Econometrica*, L (1982), 1345–70.
- Lucas, Robert E., "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review*, LXIII (1973), 326–34.
- , "Understanding Business Cycles," *Carnegie-Rochester Conference on Public Policy*, V (1977), 7–29.
- Mankiw, N. Gregory, and Matthew D. Shapiro, "Trends, Random Walks, and Tests of the Permanent Income Hypothesis," *Journal of Monetary Economics*, XVI (1985), 165–74.
- , and —, "Do We Reject Too Often? Small Sample Properties of Tests of Rational Expectation Models," *Economics Letters*, XX (1986), 139–45.
- Neftci, Salih N., "Specification of Economic Time Series Models Using Akaike's Criterion," *Journal of the American Statistical Association*, LXXXVII (1982), 537–40.
- Nelson, Charles, and Heejaon Kang, "Spurious Periodicity in Inappropriately Detrended Time Series," *Econometrica*, XLIX (1981), 741–51.
- 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.
- Newey, Whitney K., and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," NBER Technical Working Paper No. 55 (April 1986), forthcoming in *Econometrica*.
- Pesaran, M. H., "A Note on the Maximum Likelihood Estimation of Regression Models With First Order Moving Average Errors When Its Root Lies on the Unit Circle," unpublished paper, May 1981.
- Plosser, Charles I., and G. William Schwert, "Estimation of a Non-Invertible Moving Average Process: The Case of Overdifferencing," *Journal of Econometrics*, VI (1977), 199–224.
- Priestley, Maurice B., *Spectral Analysis and Time Series* (London: Academic Press, 1982).
- Romer, Christina, "The Prewar Business Cycle Reconsidered: New Estimates of Gross National Product, 1869–1918," NBER Working Paper No. 1969, July 1986.
- Sargan, J. D., and Alok Bhargava, "Maximum Likelihood Estimation of Regression Models with First Order Moving Average Errors When the Root Lies on the Unit Circle," *Econometrica*, LI (May 1983), 799–820.
- Schwarz, G., "Estimating the Dimension of a Model," *Annals of Mathematical Statistics*, VI (1978), 461–64.
- Schwert, G. William, "Effects of Model Specification on Tests for Unit Roots," University of Rochester, 1985.
- Watson, Mark W., "Univariate Detrending Methods with Stochastic Trends," *Journal of Monetary Economics*, XVIII (1986), 49–75.