

NBER WORKING PAPER SERIES

ARE OUTPUT FLUCTUATIONS TRANSITORY?

John Y. Campbell

N. Gregory Mankiw

Working Paper No. 1916

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1986

We are grateful to Ben Bernanke, Christina Romer, David Romer, Jim Stock, and Mark Watson for helpful discussions. The research reported here is part of the NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Working Paper # 1916
May 1986

Are Output Fluctuations Transitory?

ABSTRACT

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, ten or twenty years. Our examination of quarterly post-war United States data leads us to be skeptical about this implication. We find that a unexpected change in real GNP of one percent should change one's forecast by over one percent over a long horizon. While it is obviously imprudent to make definitive judgments regarding theories on the basis of one stylized fact alone, we believe that the great persistence of output shocks documented in this paper is an important and often neglected feature of the data that should more widely be used for evaluating theories of economic fluctuations.

Professor John Y. Campbell
Department of Economics
Dickinson Hall
Princeton University
Princeton, NJ 08544

N. Gregory Mankiw
NBER
1050 Massachusetts Avenue
Cambridge, MA 02138

I. Introduction

Robert Lucas begins his classic article "Understanding Business Cycles" 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 a 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 from this natural rate.

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, ten or twenty 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 post-war United States data leads us to be skeptical about this implication. In particular, we find that a one percent innovation in real GNP should change one's forecast of GNP over a long horizon by over one 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 great persistence of output shocks is an important and often neglected feature of the data that should more widely be used 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 three 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. Our results

¹Schwert [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.

indeed suggest that mixed ARIMA models are preferred to pure autoregressive models with the same number of parameters.

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 non-stationarity. Our test thus provides a natural complement to standard tests of non-stationarity.

Third, we try to direct attention away from the question of the existence of a unit root in real GNP, and towards the question of 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.

²It 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.

II. Methodology

Suppose real GNP falls one percent lower than one would have expected from its past history. How much should one change one's forecast of GNP for ten or twenty 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 10 or 20 years, the detrended series displays a considerable effect of today's innovation. However, in samples of typical 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 seriously biased toward zero. With 100 observations, as might be the case with post-war 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, by 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 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 :

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

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 $1-L$:

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

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. 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 of the differenced representation 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 $B(L) = \phi(L)^{-1} \tilde{\theta}(L)$, 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

³The autoregressive representation for the model includes an infinite number of parameters which 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.

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 to any more general specification.

Another approach is to choose the optimal parameterization using either the Schwarz [1978] 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 minimize

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

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

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

where T is the number of observations. Since our sample includes 152 observations and $\ln(152)$ 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 not general

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

agreement about which criterion is best. Second, it is not clear whether these criteria will perform well for our purposes, since 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 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]; here we summarize the approach.

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

$$(5) \alpha_t = T\alpha_{t-1} + R\eta_t$$

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

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 $\alpha^2 P_0$. In our application we subtract the sample mean from the data and set $a_0 = 0$. P_0 is given by

$$(6) \text{Vec}(P_0) = [I - T \otimes T]^{-1} \text{Vec}(RR')$$

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, $v_t = \Delta Y_t - \Delta Y_{t|t-1}$. Finally one updates for the next round, setting $a_t = a_{t|t-1} + P_{t|t-1}z'v_t/f_t$ and $P_t = P_{t|t-1} - P_{t|t-1}z'z'P_{t|t-1}/f_t$.

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

$$(7) \quad -(T/2)\log(2\pi) - (T/2)\log(\sigma^2) - \left(\frac{1}{2}\right) \sum_{t=1}^T \log f_t - \left(\frac{1}{2\sigma^2}\right) \sum_{t=1}^T \nu_t^2 / f_t$$

We maximize this likelihood function using a method of scoring with modified step size (Berndt, Hall, Hall and Hausman [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) given the parameters. The impulse response at horizon k , A_k , is just

$$(8) \quad A_k = z' T^k R$$

To compute the impulse response at horizon k in levels, B_k , one simply uses equation (4) 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 $\sqrt{(d'\Gamma d)}$ where d is the vector of derivatives of the function with respect to the parameters. We compute d numerically.

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 post-war, seasonally adjusted, quarterly data from 1947:1 to 1985:1. We consider all ARMA models for the difference of log GNP with up to three AR parameters and three MA parameters. There are thus sixteen models under consideration for GNP growth, the simplest being white noise, the most complex the ARMA(3,3).

Table 1 presents the model selection criteria for the sixteen models. The value of the likelihood function points toward the ARMA(2,2) model. Both the ARMA(1,2) and the ARMA(2,1) are rejected in favor of this more general alternative. Moreover, one cannot reject the ARMA(2,2) specification in favor of an ARMA(2,3), an ARMA(3,2), or an ARMA(3,3). It appears that classical statistical inference leads one to an ARMA(2,2) specification.

The Akaike and Schwarz criteria lead to different conclusions, however. The Akaike criterion suggests an ARMA(2,3), and possibly more moving average parameters than we have estimated. The Schwarz criterion, which penalizes extra parameters more heavily, suggests the much more parsimonious ARMA(1,0) specification.⁵ 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 ARMA(1,0) for $k=1$, the ARMA(0,2) for $k=2$, the ARMA(0,3) for $k=3$, and the

⁵Interestingly, this is the process that Deaton [1985] suggests for labor income growth and Watson [1985] for GNP growth. Watson goes on to argue for an unobservable components model that implies a restricted ARMA(2,2) representation.

ARMA(2,2) for $k=4$.

Table 1 also presents 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 can be very strongly rejected when it is imposed on the most parsimonious models we consider, with up to one autoregressive parameter and two moving average parameters. The rejections are weaker when we impose the constraint on higher order models; and in two cases, ARMA(2,1) and ARMA(1,3), our unconstrained estimates obey the constraint almost exactly. However, neither of these models is preferred, given the number of their parameters, and it remains true that for our preferred ARMA(2,2) model one can reject the unit root constraint at about the 6 percent level if one compares the likelihood ratio to the Chi-squared distribution with one degree of freedom.

Table 2 presents the parameter estimates for the different unconstrained models. Again, the ARMA(2,2) appears to describe the data well. Both the second AR and the second MA parameters are significant, while additional parameters that might be added are insignificant.

Table 3 presents the implied impulse response function for the level of log GNP for each specification. While the particular parameter estimates in Table 2 sometimes appear to differ substantially, the implied impulse response functions in Table 2 appear almost unanimous. With two exceptions, the impulse response increases above one and settles between 1.5 and 2.0 at about the eighth quarter and remains there even at ten or twenty years. That is, a one percent innovation in real GNP increases the univariate forecast of GNP by

over one percent over any foreseeable horizon.

The two exceptions to this statement are of course the models which are estimated to have unit moving average roots. The estimated ARMA(1,3) is fairly persistent, with an impulse response of 1.6 even at a five-year horizon, and 0.39 at ten years. However, the likelihood function for this model is extremely flat, both locally as reflected in the enormous computed standard errors, and globally as shown by the fact that coefficient values similar to those estimated for the ARMA(1,2) process have a likelihood within 0.05 of the reported values. The estimated ARMA(2,1) dies away more rapidly, but here again the likelihood function is relatively flat.

The evidence of persistence which we find in the quarterly post-war GNP data is robust to change in sample period and frequency of 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. (Indeed, the unit root in the ARMA(2,1) process does not obtain in the pre-1972 period.) When we examine post-war 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.

IV. Econometric Issues

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 in samples of typical size.⁶ 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 (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 roots, however, these results break down. Sargan and Bhargava [1983] have shown that in the first-order moving average case with a unit parameter there exists a local maximum of the likelihood function whose distance from unity is of order $(1/T)$ where T is the sample size; but this local maximum is not distributed asymptotically normal. It has a probability mass at exactly unity, and Sargan and Bhargava show how to compute this mass for any sample size.

These results can be understood intuitively by noting that in the first-order moving average case, values for the parameter of r and $(1/r)$ are observationally equivalent. It follows that the likelihood function has

⁶See, for example, Flavin [1983] and Mankiw and Shapiro [1986].

turning points (local maxima or minima) at $r = 1$ and -1 . Davidson [1981] develops some intuition and discusses higher-order moving averages.

Other authors have conducted Monte Carlo simulations to characterize the properties of the global maximum of the likelihood function in finite samples. Ansley and Newbold [1980], Harvey [1981] and Davidson [1981] all report finding a probability mass at exactly unity for this maximum, particularly large when the true root is unity but also present when it is considerably 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 roughly bell-shaped and centred on the true value when this is less than unity.

Davidson reports some Monte Carlo results for a likelihood ratio test of the hypothesis that the moving average root equals one. He finds that the test rejects too infrequently (4 times out of 200) when the 5 percent critical value of the Chi-Squared distribution with one degree of freedom is used.

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 150 observations, which follow an AR(2) in levels. The first AR parameter was 1.366, and the second was -0.415, the values estimated in Table 2 for the ARMA(2,1) process with a unit root.⁷ 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 4.

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 errors. When it is estimated away from unity, they typically have rather small standard errors; in two cases the implied confidence intervals are well away from the true value. This result is consistent with Ansley and Newbold's findings.

⁷Hence, the assumed process is very close to the one Blanchard [1981] and Kydland and Prescott [1980] suggest as an accurate representation of the data.

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

V. Conclusion

We have estimated standard ARIMA processes for the log of real GNP using the standard post-war quarterly time series. Yet the estimates have a surprising implication: A one percent innovation in real GNP should change one's forecast of real GNP by over one percent over a long horizon.⁸

This finding should be interpreted with caution. The existing econometric literature and our small Monte Carlo experiment give us some confidence in our estimation procedure. Yet work on the small sample properties of ARIMA estimation, especially in the presence of unit roots, is only in its infancy. Applying other statistical procedures, perhaps less parametric, would be useful in determining more fully the sensitivity of our finding to the estimation method.

The finding that a one percent innovation in GNP should change one's long-run forecast of GNP by over one percent has important implications for business cycle theory. In particular, this result is inconsistent with many prominent theories in which output fluctuations are primarily caused by shocks to aggregate demand. Both models based on misperceptions, such as that of Lucas [1973], and models based on long-term nominal contracts, such as that of Fischer [1977], imply that the deviations from trend caused by demand shocks are transitory in nature. This implication does not appear consistent with the time series properties of measured production.

⁸We 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.

One defense of such models is that they explain how demand shocks move the economy away from the natural rate and that other mechanisms make the deviation from the natural rate long-lived. Yet the utility of the natural rate concept, which is central to these theories, is called into question if the forces restoring equilibrium are very slow moving. There is little content to the claim that economic forces return the economy to the natural rate if that return takes substantially more than two decades.

One might also defend such models by claiming that they correctly explain deviations around the natural rate while the output series we examine are dominated by fluctuations in the natural rate itself. Certainly this is a logical possibility. Yet if deviations around the natural rate account for a relatively small amount of output fluctuations, such models are not useful for understanding these fluctuations.⁹

Two lines of research appear more consistent with our results. The first is real business cycles models. If shocks to the production function are the driving force of output fluctuations, as in the model of Kydland and Prescott [1982], then news about output today may well convey information about output over a long horizon. Perhaps future work could shed light on the validity of these models by examining separately employment changes and productivity changes.

The second type of theory consistent with our results is Keynesian theory

⁹It may be possible to reconcile our results with these theories by abandoning the natural rate hypothesis. In particular, it may be possible to add some propagation mechanism to make temporary misperceptions or nominal rigidities have permanent effects.

allowing for the possibility of secular stagnation. For example, the models of Diamond [1984] and Weitzman [1982] exhibit multiple equilibria; if the economy gets stuck in a "bad" equilibrium, there is no force driving the economy back to a Pareto-dominating equilibrium. Since these models do not determine which equilibrium is chosen, it is difficult to discuss dynamics. Yet one suspects that if a shock of some sort moves the economy from one equilibrium to another, it will tend to stay at the new equilibrium rather than returning to the old one. If so, this new line of Keynesian research may be consistent with our empirical findings.

References

- Akaike, H., 1974, A New Look at Statistical Model Identification, IEEE Transactions on Automatic Control, AC-19, 716-723.
- _____, 1976, "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, 52-107).
- Ansley, Craig F. and Paul Newbold, 1980, "Finite Sample Properties of Estimators for Autoregressive Moving Average Models," Journal of Econometrics 13, 159-183.
- Berndt, Hall, Hall and Hausman, 1974, "Estimation and Inference in Nonlinear Structural Models," Annals of Economic and Social Measurement, 653-665.
- Blanchard, Olivier, 1981, "What is Left of the Multiplier-Accelerator?" American Economic Review Papers and Proceedings. 150-154.
- Davidson, James E.H., 1981, "Problems With the Estimation of Moving Average Processes," Journal of Econometrics 16, 295-310.
- Deaton, Angus S., 1985, "Life-Cycle Models of Consumption: Is the Evidence Consistent With the Theory?," Princeton University.
- Diamond, Peter A., A Search-Equilibrium Approach to the Micro Foundations of Macroeconomics (Cambridge, MA: MIT Press, 1984).
- Fischer, Stanley, 1977, "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," Journal of Political Economy 85, 191-205.
- Flavin, Marjorie A., 1983, "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence," Journal of Political Economy 91, 929-956.
- Frisch, Ragnar and Frederick V. Waugh, 1933, "Partial Time Regressions as Compared with Individual Trends," Econometrica 1, 387-401.
- Harvey, A.C., Time Series Models (Oxford: Philip Allan, 1981).
- Kydland, Finn E. and Edward C. Prescott, 1980, "A Competitive Theory of Fluctuations and the Feasibility and Desirability of Stabilization Policy," in ed., S. Fischer, Rational Expectations and Economic Policy (Chicago: University of Chicago Press).
- _____, 1982, "Time to Build and Aggregate Fluctuations," Econometrica 50, 1345-1370.
- Lucas, Robert, E., 1973, "Some International Evidence on Output-Inflation Tradeoffs," American Economic Review 63, 326-334.

- _____, 1977, "Understanding Business Cycles, Carnegie-Rochester Conference on Public Policy 5, 7-29.
- Mankiw, N. Gregory and Matthew D. Shapiro, 1985, "Trends, Random Walks, and Tests of the Permanent Income Hypothesis," Journal of Monetary Economics 16, 165-174.
- Mankiw, N. Gregory and Matthew D. Shapiro, 1986, "Do We Reject Too Often? Small Sample Properties of Tests of Rational Expectation Models," Economic Letters 20, 139-145.
- Neftci, Salih N., 1982, "Specification of Economic Time Series Models Using Akaike's Criterion," Journal of the American Statistical Association 77, 537-540.
- Nelson, Charles R. and Charles I. Plosser, 1982, "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications," Journal of Monetary Economics 10, 139-162.
- Pesaran, M.H., 1981, "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).
- Sargan, J.D. and Alok Bhargava, 1983, "Maximum Likelihood Estimation of Regression Models With First Order Moving Average Errors When the Root Lies on the Unit Circle," Econometrica 51 (May), 799-820.
- Schwarz, G., 1978, "Estimating the Dimension of a Model," Annals of Mathematical Statistics 6, 461-464.
- Schwert, G. William, 1985, "Effects of Model Specification on Tests for Unit Roots," University of Rochester.
- Watson, Mark W., 1985, "Univariate Detrending Methods with Stochastic Trends," Harvard Institute of Economic Research Discussion Paper No. 1158.
- Weitzman, Martin L., 1982, "Increasing Returns and the Foundations of Unemployment Theory," Economic Journal 92, 787-804.

Table 1
Model Selection Criteria, $\Delta \ln$ Real GNP

Number of AR Parameters (p)	Number of MA Parameters (q)			
	0	1	2	3
0	-1193.477 (-1193.477) (-1193.477)	-1212.637 (-1210.637) (-1207.613) [807.128]	-1221.169 (-1217.169) (-1211.121) [-974.566]	-1224.237 (-1218.237) (-1209.165) [-1072.266]
1	-1218.853 (-1216.637) (-1213.829)	-1219.319 (-1215.319) (-1209.271) [-1194.539]	-1222.498 (-1216.498) (-1207.426) [-1214.709]	-1224.595 (-1216.595) (-1204.499) [-1224.595]
2	-1219.694 (-1215.694) (-1209.646)	-1223.056 (-1217.056) (-1207.984) [-1223.056]	-1228.046 (-1220.046) (-1207.950) [-1224.473]	-1230.468 (-1220.468) (-1205.348) [-1227.642]
3	-1222.731 (-1216.731) (-1207.659)	-1225.568 (-1217.568) (-1205.472) [-1225.321]	-1230.059 (-1220.059) (-1204.939) [-1225.784]	-1231.760 (-1219.760) (-1201.616) [-1230.947]

For each model, we report $-2\ln L$

(Akaike Criterion = $-2\ln L + 2k$)

(Schwarz Criterion = $-2\ln L + k\ln T$)

[-2ln L of model restricted to have a unit moving average root]

Table 2
Model Parameter Estimates, $\Delta \ln$ Real GNP

Model p, q	ϕ_1	ϕ_2	ϕ_3	θ_1	θ_2	θ_3
0,1	--	--	--	0.306* (0.073)	--	--
0,2	--	--	--	0.332* (0.070)	0.243* (0.082)	--
0,3	--	--	--	0.379* (0.075)	0.346* (0.075)	0.180* (0.081)
1,0	0.391* (0.068)	--	--	--	--	--
1,1	0.479* (0.186)	--	--	-0.102 (0.216)	--	--
1,2	0.263 (0.262)	--	--	0.090 (0.263)	0.207 (0.112)	--
1,3	0.931 (0.444)	--	--	0.339 (1.025)	2.464 (1.848)	-3.803 (1.250)
2,0	0.362* (0.075)	0.074 (0.082)	--	--	--	--
2,1	1.366* (0.070)	-0.415* (0.073)	--	-1.000* (0.094)	--	--
2,2	0.585* (0.1757)	-0.529* (0.148)	--	-0.279 (0.159)	0.688* (0.111)	--
2,3	0.183 (0.249)	-0.489* (0.166)	--	0.170 (0.257)	0.743* (0.116)	0.244 (0.142)
3,0	0.372* (0.076)	0.126 (0.089)	-0.140 (0.078)	--	--	--
3,1	1.051* (0.228)	-0.120 (0.143)	-0.190* (0.083)	-0.705* (0.235)	--	--
3,2	0.572* (0.153)	-0.667* (0.183)	0.152 (0.118)	-0.224 (0.131)	0.804* (0.150)	--
3,3	0.155 (0.250)	-0.706* (0.080)	0.301 (0.186)	0.206 (0.274)	0.959* (0.055)	0.075 (0.249)

Standard errors are in parentheses. Asterisk indicates significance at 5 percent level.

Table 3
Model Impulse Responses, In Real GNP

Model p,q	1	2	4	8	16	20	40	80
0,1	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)	1.306 (0.073)
0,2	1.332 (0.070)	1.575 (0.124)	1.575 (0.124)	1.575 (0.124)	1.575 (0.124)	1.575 (0.124)	1.575 (0.124)	1.575 (0.124)
0,3	1.379 (0.075)	1.725 (0.125)	1.905 (0.172)	1.905 (0.172)	1.905 (0.172)	1.905 (0.172)	1.905 (0.172)	1.905 (0.172)
1,0	1.391 (0.068)	1.545 (0.122)	1.628 (0.169)	1.643 (0.184)	1.643 (0.184)	1.643 (0.184)	1.643 (0.184)	1.643 (0.184)
1,1	1.377 (0.077)	1.557 (0.121)	1.685 (0.207)	1.721 (0.262)	1.723 (0.268)	1.723 (0.268)	1.723 (0.268)	1.723 (0.268)
1,2	1.354 (0.072)	1.653 (0.133)	1.753 (0.216)	1.761 (0.235)	1.761 (0.235)	1.761 (0.235)	1.761 (0.235)	1.761 (0.235)
1,3	2.270 (1.032)	5.916 (3.778)	5.126 (11.201)	3.849 (23.177)	2.171 (38.788)	1.630 (43.768)	0.390 (55.041)	0.024 (58.266)
2,0	1.362 (0.075)	1.567 (0.120)	1.720 (0.206)	1.770 (0.260)	1.773 (0.267)	1.773 (0.267)	1.773 (0.267)	1.773 (0.267)
2,1	1.366 (0.102)	1.451 (0.222)	1.332 (0.473)	0.941 (0.920)	0.444 (1.444)	0.305 (1.586)	0.046 (1.860)	0.001 (1.919)
2,2	1.306 (0.073)	1.645 (0.124)	1.523 (0.183)	1.530 (0.123)	1.491 (0.148)	1.494 (0.143)	1.493 (0.144)	1.493 (0.144)
2,3	1.352 (0.074)	1.671 (0.124)	1.669 (0.181)	1.667 (0.154)	1.653 (0.157)	1.652 (0.159)	1.651 (0.159)	1.651 (0.159)
3,0	1.372 (0.076)	1.637 (0.135)	1.625 (0.215)	1.554 (0.228)	1.558 (0.222)	1.558 (0.222)	1.558 (0.222)	1.558 (0.222)
3,1	1.346 (0.078)	1.589 (0.136)	1.544 (0.212)	1.124 (0.324)	1.144 (0.340)	1.142 (0.322)	1.138 (0.332)	1.138 (0.332)
3,2	1.348 (0.071)	1.683 (0.124)	1.688 (0.210)	1.703 (0.193)	1.677 (0.204)	1.674 (0.203)	1.675 (0.201)	1.675 (0.201)
3,3	1.361 (0.080)	1.669 (0.131)	1.755 (0.223)	1.756 (0.252)	1.771 (0.249)	1.781 (0.249)	1.792 (0.250)	1.791 (0.250)

Standard errors are in parentheses.

Table 4
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	429.02	429.02	0.00	0.01	(2.19)
2	413.82	417.06	3.24	1.06	(0.05)
3	381.09	381.09	0.00	0.01	(2.94)
4	428.82	428.82	0.00	0.00	(1.06)
5	419.64	420.80	1.16	0.70	(0.29)
6	426.86	426.86	0.00	0.00	(1.74)
7	447.33	447.33	0.00	0.01	(2.64)
8	414.11	414.72	0.61	0.21	(0.22)
9	425.59	525.61	0.02	0.12	(0.43)
10	442.31	442.35	0.04	1.10	(0.28)
11	400.64	400.64	0.00	0.00	(3.47)
12	443.02	443.83	0.81	0.99	(0.45)
13	419.22	419.22	0.00	0.01	(1.67)
14	446.70	446.70	0.00	0.01	(2.53)
15	403.59	403.59	0.00	0.00	(3.15)
16	395.30	395.30	0.00	0.15	(6.40)
17	382.09	382.09	0.00	0.04	(4.25)
18	416.64	416.64	0.00	0.01	(2.83)
19	442.89	442.89	0.00	0.12	(0.66)
20	405.04	405.04	0.00	0.40	(0.23)

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 1.366 and -0.415. The impulse responses are at a horizon of 80 periods.