OIL AND THE MACROECONOMY REVISITED*

Mark A. Hooker
Federal Reserve Board, Stop 71
20th & C St., NW
Washington, DC 20551
mhooker@frb.gov

August 1999

JEL Classification Codes E32, E37
Key words: Oil Shocks, Macroeconomic Fluctuations, Granger Causality, Asymmetry

Abstract:

The relationship between oil price shocks and U.S. macroeconomic fluctuations advocated by Hamilton (1983) broke down in the 1980s amidst a new regime of highly volatile oil price movements. Several authors have argued that asymmetric and nonlinear transformations of oil prices restore that relationship, and thus that the economy responds asymmetrically and nonlinearly to oil price shocks. In this paper, I show that this is only part of the story: the two leading such transformations do not in fact Granger cause output or unemployment in the post-1980 period without further refinements, and they derive much of their apparent success from data in the 1950s. If output is expressed in year-over-year changes, which are smoother than the usual quarterly changes, and the equations exclude variables like interest rates and inflation, then asymmetric and nonlinear oil prices predict output, but not unemployment, while the real level of oil prices predicts unemployment, but not output. I interpret this evidence as supportive of significant oil price effects on the macroeconomy which a) are at relatively low frequencies, b) are indirect, through variables like interest rates and inflation, c) can induce departures from Okun's law, and d) changed qualitatively around 1980.

* An earlier version of this paper, titled "Exploring the Robustness of the Oil Price-Macroeconomy Relationship," was prepared for the U.S. Department of Energy Conference on the Macroeconomic Impacts of Oil Price Shocks, Washington DC, 1996. Financial support from Oak Ridge National Laboratory is gratefully acknowledged. I thank Andrew Filardo, Jeff Fuhrer, Prakash Loungani, Brian Sack, Eric Swanson, and participants at that conference and the IAEE conference in Quebec City, May 1998, for helpful comments. The views expressed are those of the author, and do not necessarily represent those of the Federal Reserve Board. Tel: (202) 452-3424; fax: (202) 452-2301.

1. Introduction

Oil price shocks receive considerable attention for their presumed macroeconomic consequences. However, despite substantial research, we are still far from a consensus about the channels through which oil prices influence the economy and the magnitudes of their effects. The oil price swings of the past few years have been substantial, making an understanding of those effects especially important from a policy perspective. For example, oil prices fell by more than 50% from the fall of 1996 to the end of 1998, while Phillips curve and other traditional models—even those with energy shock terms—consistently overpredicted inflation. It would be useful for monetary policymakers to know whether the reversal of those price declines in mid-1999 will lead to higher inflation pressures. Moreover, that reversal was comparable in magnitude to some of the large increases of the 1970s, which most economists believe contributed substantially to the deep recessions that followed. Finally, the Department of Energy maintains several months of oil imports (roughly $10 billion worth, at mid-1999 prices) in its Strategic Petroleum Reserve to "protect the domestic U.S. economy from the impact of energy supply disruptions," at an annual cost of nearly $200 million.[1]

James Hamilton (1983,1985) played a major role in convincing economists that oil price increases generally, and not just the OPEC supply disturbances of the 1970s, are important contributors to recessions.[2] Ironically, at the same time that Hamilton's arguments were gaining acceptance, the evidence for them was breaking down: Lee, Ni and Ratti (1995), Hooker (1996), and others have shown that oil prices typically fail to Granger cause macro variables when data samples are extended past the mid-1980s.

The breakdown of Hamilton's Granger-causal relationship roughly coincided with major changes in the behavior of oil prices. As Figure 1 shows, prices were very stable until 1973 (note the narrow *y*-axis range in the upper panel). OPEC raised oil prices dramatically in the 1970s, but this may be viewed as a continuation of the pattern of occasional, discrete nominal increases. The 1980s brought both nominal price decreases, beginning in 1981, and wide swings, following the market collapse in late 1985.

1 From the DOE fiscal year 1999 Congressional budget request, available at www.doe.gov.
2 Hamilton's arguably most persuasive piece of evidence was that oil prices strongly Granger cause output and unemployment in samples that end *before* 1973.

Several researchers have argued that falling and volatile oil prices *caused* the breakdown of Hamilton's original specification by revealing asymmetric and nonlinear characteristics of the true oil price-macroeconomy relationship, and supported this position with evidence that correspondingly asymmetric and nonlinear transformations of oil prices continue to Granger cause output in samples up through the present. The two leading such proposals have been made by Lee, Ni and Ratti (1995) and Hamilton (1996). Bernanke, Gertler and Watson (1997) recently endorsed this position, using Hamilton's transformation as an instrument to identify the effects of systematic monetary policy on output, and Davis and Haltiwanger (1999), who analyze job creation and destruction in response to (their own) measure of oil price changes, "view the evidence for asymmetric responses to oil price ups and downs as well established (for the United States)".

In this paper, I reevaluate the empirical evidence for these transformations, and the corresponding case that oil prices affect the macroeconomy in asymmetric and nonlinear ways. I find that Lee, Ni and Ratti's and Hamilton's oil price series do not in fact Granger cause output or unemployment in the post-1980 data that they were designed to fit, without further refinements to the equations. Furthermore, these series derive much of their success from a single observation in the 1950s.

Two sufficient refinements are that output be expressed in year-over-year changes, which are smoother than the usual quarterly series, and that the equations exclude variables like interest rates and inflation. With these conditions, asymmetric and nonlinear oil prices do Granger cause output, but not unemployment, while the real level of oil prices Granger causes unemployment, but not output. I interpret this evidence, and some additional out-of-sample forecasting results, as supportive of significant oil price-macroeconomy effects which a) are at annual rather than quarterly frequencies, b) are indirect, through variables like interest rates and inflation, c) can induce departures from Okun's law, and d) changed qualitatively around 1980. The remainder of the paper establishes these empirical results and fleshes out their interpretation.

2. The breakdown and some fix-ups

The breakdown of Hamilton's original (1983) oil price-macroeconomy relationship is illustrated in Figure 2, with p-values from tests that oil prices Granger cause output growth in quarterly data samples which begin in 1950:2 and expand from 1970:3 to 1998:4.[3] The Figure shows that oil prices, in real levels or nominal log-differences, are significant at the 10% level in *all* of the samples that end before 1983, and at better than the 5% level in most cases.[4] However, as the samples expand into the 1980s, significance dissipates rapidly, with the typical *p*-value for a sample ending in the late 1980s or 1990s well above 0.5. This pattern is quite robust, as similar results obtain using unemployment, industrial production, and other measures of real macroeconomic activity, other lag lengths, and other sets of conditioning variables.[5] It also obtains with Mork's (1989) asymmetric oil price, the first entry in the "fix-ups" literature.

Several authors have argued that this breakdown reflects the greater power to reject misspecified equations brought by the increased variation in oil prices in the 1980s and 1990s. For example, Lee, Ni and Ratti (1995) (hereafter, LNR) and Hamilton (1996, 1999) argue that oil price effects are asymmetric and nonlinear, and that the pre-1980 period, with almost no (nominal) price decreases and not much volatility, provided little information for identifying functional forms. LNR and Hamilton support their position by constructing asymmetric and nonlinear transformations of oil prices which continue to Granger cause macroeconomic variables in samples up through the present.

LNR focus on volatility, maintaining that "an oil shock is likely to have greater impact in an environment where oil prices have been stable than in an environment where oil price movement has been frequent and erratic" (p. 42), because price changes in a volatile environment are likely to be soon reversed. The associated economic mechanisms involve fixed costs and uncertainty: small increases and increases within volatile periods will not have much of an effect if they do not push agents across *S,s* bands or generate enough new uncertainty to delay irreversible investments. LNR capture these features with a GARCH-based oil price transformation that scales estimated oil price shocks by their conditional variance.

---

3 The beginning date is determined by data availability (the full set of variables is available from 1948:2 and 8 lags are used in the VAR), and the length of the shortest sample by degrees of freedom (it contains 82 quarters, or twice the number of parameters being estimated). The VARs include output growth and the oil price, and GDP deflator inflation, the 3-month Treasury bill rate, and import price inflation as conditioning variables (the same specification as in Hamilton (1983), with wage changes dropped to conserve degrees of freedom. The correlation between import price inflation and nominal oil price changes in the sample is 0.54). See the data appendix for details.

4 Hamilton used nominal log-differences in his original specification; since GDP deflator inflation is included in the VAR, the results are essentially identical using real log-differences.

5 The particular cases in which I find no breakdown are discussed in Section 4.

Hamilton also advocates a variation on this theme, emphasizing uncertainty effects. He argues that "[i]f one wants a measure of how unsettling an increase in the price of oil is likely to be for the spending decisions of consumers and firms, it seems more appropriate to compare the current price of oil with where it has been over the previous year rather than during the previous quarter alone" (1996, p. 216). Hamilton thus bases his transformation on the percentage increase from the previous year's high rather than the previous quarter's value.

Both LNR and Hamilton set oil price decreases to zero in their transformations. The first to argue that oil price decreases have no macroeconomic consequences, and thus should be ignored, was Mork (1989). Both LNR and Hamilton (1999) sketch theories of why oil price decreases might have small macroeconomic effects on net. The transmission mechanism must be at least two-channeled, with one channel involving contractionary macroeconomic responses to price movements in either direction. Such effects might arise from substantial costs of reallocating resources across sectors, or from investment-inhibiting uncertainty generated by oil price fluctuations. The second channel would involve conventional symmetric responses. When oil prices rise, the two effects reinforce one another, but could largely offset when oil prices fall.

In order to evaluate these nonlinearity and asymmetry hypotheses, I construct versions of LNR's and Hamilton's oil price series, denoted SOPI for 'scaled oil price increases' and NOPI for 'net oil price increases'. Both series are based on Mork's (1989) oil price, which makes an adjustment for the price controls of the early 1970s; it is described in the data appendix and Mork's paper. For SOPI, a GARCH(1,1) model is estimated using LNR's specification:

(1a) $$z_t = \alpha + \sum_{j=1}^{4} \beta_j z_{t-j} + \varepsilon_t , \quad \varepsilon_t \mid I_t \sim N(0, h_t) ,$$

where

(1b) $$h_t = \gamma_0 + \gamma_1 \varepsilon_{t-1}^2 + \gamma_2 h_{t-1} ,$$

and z_t is Mork's series (in quarterly growth rates) updated through the present. SOPI is then computed using the estimated ε's and h's:

(1c) $$SOPI_t = \max(0, \hat{\varepsilon}_t / \sqrt{\hat{h}_t}) .$$

NOPI is constructed according to Hamilton (1996), as the percentage increase from the previous year's (quarterly) high price if that is positive and zero otherwise:

(2) $$NOPI_t = \max(0,100*\{\ln(o_t) - \ln[\max(o_{t-1}, o_{t-2}, o_{t-3}, o_{t-4})]\}),$$

where o_t is Mork's oil price in levels.

Setting price decrease observations to zero gives both of these series an extreme degree of asymmetry. SOPI is also nonlinear and time-dependent, as the effects of a given oil price change will generally not be proportional to the size of the shock or independent of shocks in other periods. For example, a small shock will be scaled up if it occurs in a quiet period, and scaled down if it occurs in a volatile period. NOPI displays a threshold type of nonlinearity and time-dependence: if a shock is not large enough to bring prices up above the previous year's high, then it is scaled down to zero; otherwise it is counted as is.

The constructed SOPI and NOPI series are plotted in Figure 3. One feature which stands out is the dramatic degree to which SOPI rescales the data. According to LNR, the two most important oil shocks of the postwar period occurred in 1953 and 1957, while the OPEC shocks of 1973 and 1979 rank only slightly larger than ones in 1952, 1969, and 1995. By contrast, the effects of the NOPI filter are mostly to "thin" the data, replacing some small increases with zeros, while leaving the relative magnitudes of the larger increases relatively intact.

Figure 4 repeats the expanding-sample Granger causality tests from Figure 2 using SOPI and NOPI. The contrast is striking: SOPI achieves such high confidence levels throughout the sample that the *p*-values are often indistinguishable from the *x*-axis, even with the very narrow scaling. NOPI is less successful; again the *p*-values rise sharply when the samples extend into the 1980s, often exceeding 10%, but they remain well below those recorded with the standard oil series in Figure 2.[6]

3. How well do the fix-ups explain the data?

The results in Figure 4 confirm LNR's and Hamilton's findings that SOPI and NOPI Granger cause output in samples up through the present, with SOPI particularly impressive. In this section, I examine the robustness of these results in three dimensions. First, I identify a particularly influential point early in the sample, and examine the sensitivity of the Granger causality results to excluding that point. I next compute Granger causality tests on post-1980

---

6 These significance levels should not be interpreted too literally, because the highly nonlinear nature of SOPI and NOPI invalidates the assumption of normal errors which underlies the distribution of the test statistics.

samples, which the series were ostensibly designed to fit. Finally, I compute forecasts of unemployment over the 1990s which use actual realized values of SOPI and NOPI. If the macroeconomy responds to oil price shocks in the asymmetric and nonlinear ways that SOPI and NOPI embody, then knowledge of SOPI and NOPI realizations ought to significantly improve those forecasts.

A. Sensitivity to particular observations

Thoma and Gray (1998) showed that money-income Granger causality test results which include the commercial paper-Treasury bill rate spread are very sensitive to the presence or absence of the single datapoint 1974:12, when industrial production growth exhibited its largest negative value (the CP-Tbill spread reached its own high five months earlier). In the present dataset, 1957:1 is a similarly influential datapoint. While oil prices increased by only about 10% in that quarter—which appears as a modest jump in the top panel of Figure 1—it was from a very stable level, and thus a low conditional variance. As shown in Figure 3, SOPI scales that 10% increase into a spike nearly twice as large as any other postwar observation. A deep recession began in August, with the 11 percent decline in GDP in the first quarter of 1958 the worst single quarter of the postwar period.

Following Thoma and Gray (1998), I repeat the expanding-sample Granger causality tests above, with 1957:1 removed, and plot the *p*-values in Figure 5.[7] Though the change in dataset is minor, the results are dramatically different from those in Figure 4, particularly for SOPI: now, neither series Granger causes output at the 5% level in more than a handful of the samples, and in most cases the *p*-values are far above 10% (note the changed axis scale)!

B. Granger causality tests on post-1980 data

As emphasized above, the SOPI and NOPI filters were motivated by the volatility of post-1980 oil prices, and designed to extract the oil price movements from that period which matter for macroeconomic fluctuations. However, the expanding-sample Granger causality tests are dominated by data from the earlier periods of stable and rising oil prices; even in the full dataset,

---

7 I remove the 1957:1 observation on all variables before taking lags, so it does not appear anywhere in the equations. Thus the tests differ from those in Figure 4 by one of between 82 and 203 observations. Note that all of the datapoints are used in construction of SOPI and NOPI; 1957:1 is only omitted from the VAR.

nearly two-thirds of the observations are from before 1981. With 72 post-1980 quarters available through the end of 1998, reasonable estimates of the predictive power of SOPI and NOPI in the period of falling and volatile oil prices should now be possible.[8]

I exercise some caution with regard to the beginning of the samples, because there were several large oil price increases right around 1980, and two recessions in 1980 and 1981-82. For example, the Energy Department's composite refiner's acquisition cost of oil, used by Mork (1989) and many others, was about $13.70 per barrel in March of 1979, $26.85 in March of 1980, and $37.28 in March of 1981, generating (annualized) quarter-over-quarter increases of more than 50% in 1979:2, 1979:3, 1980:1, and 1981:1. Starting the sample in 1981:1 would thus exclude most of these shocks, and perhaps bias the results against finding effects of oil price increases, so I use three different sample beginnings—1979:1, 1980:1, and 1981:1.[9] Furthermore, given rather short samples and a concern with evaluating the robustness of these variables' predictive power, I vary the lag lengths and use both output growth and the unemployment rate as dependent variables as well. The additional included variables are as before, described in footnote 3. With 18 Granger causality tests for each oil price specification, I focus on overall patterns rather than the individual test results.

Table 1 reports *p*-values for the hypothesis that all of the oil price coefficients equal zero in the output or unemployment equation of the VAR from this battery of tests. As the Table shows, these oil price specifications do quite poorly. SOPI is not significant at the 10% level in *any* of the 18 tests, and NOPI is only significant at that level in three of the 18. In many of the tests, they are not even close to significant: the median *p*-values are 0.54 for SOPI, and 0.22 for NOPI.[10] One may object that there are few oil price shocks in these samples, and so SOPI and NOPI are not given a fair chance. However, as Figure 3 shows, the 1990s contain several price spikes, which in Hamilton's measure are comparable to all but the 1973 price hike, and many of the samples also include the large shocks, described above, in 1979-81. It appears that the significance in Granger causality tests of SOPI, to a large degree, and NOPI, to a lesser degree,

---

8 If the VARs are stable across the full sample, analyses of subsamples are still valid, but may have low power. The large number of rejections of non-Granger causality found in particular cases should mitigate this concern. In fact, Chow-type tests for structural breaks around 1980 do tend to indicate breaks.

9 These are the ranges before taking lags. For example, with the 1979:1-1998:4 sample and 4 lags, the dependent variable runs from 1980:1-1998:4. Thus the two large 1979 shocks are included in all of the 1979-98 samples.

comes from a tight fit in the *pre*-1980 data—and particularly the 1957-58 recession—rather than an ability to match oil price movements to macroeconomic fluctuations in the 1980s and 1990s.

C. Oil price-augmented out-of-sample forecasts

Most empirical work on the oil price-macroeconomy relationship has relied on Granger causality tests like those discussed above. However, additional types of evidence may be useful in evaluating different oil price transformations and their associated transmission mechanisms. In this subsection, I examine the out-of-sample forecasting performance of SOPI and NOPI over the 1990s. That is a particularly interesting period, containing a sharp oil spike and a recession in the early 1990s, and large movements in unemployment and oil prices in late 1990s as well. If the economy responds to oil price shocks in the asymmetric and nonlinear ways that SOPI and NOPI embody, then out-of-sample forecasts augmented with their realized values should better track macroeconomic fluctuations. I focus on unemployment, which is a much smoother variable than output, so that forecast accuracy may be gauged by "eyeball" as well as by more formal criteria.

The forecasts are generated from the VAR used in the Granger causality tests, which contains the five variables listed in footnote 3 and eight lags, estimated up through 1990:4. Stopping the estimation there leaves most of the recessionary increase in unemployment for the model to predict. Forecasts are generated from the VAR one quarter at a time, according to the chain rule of forecasting, substituting actual values of SOPI and NOPI for the model's predictions.[11] The forecast horizon is the end of the available data, 1998:4.

Figure 6 plots the forecasts using SOPI and NOPI against the actual unemployment rate over this period. Both forecasts predict an increase in unemployment in 1991, although NOPI expected a much deeper, and SOPI much milder, contraction. It should be noted that this recession is generally viewed as having been hard to predict. Both forecasts fare considerably worse after 1994. They both predict a mild but extended recession beginning in 1995, with unemployment returning to the 7% range by late 1997. Actual unemployment, of course,

---

10 Hamilton (1996) found that NOPI Granger caused output in 1948-1973 and 1948-1994, but not 1973-1994.

11 Burbidge and Harrison (1984) refer to such forecasts as "base plus oil" historical decompositions. Using actual values of the oil series gets away from the awkward question of how to implement the SOPI and NOPI filters with the model's forecasts for oil prices. The forecasts are not very sensitive to variations in the VARs.

continued to fall steadily from its peak in 1992, and dropped another 1-1/2 percentage points between 1994 and 1998.

4. Elements of the oil price-macroeconomy relationship

The results in the previous section provide little support for LNR's and Hamilton's hypothesis that an asymmetric and nonlinear transmission mechanism explains "what happened to the oil price-macroeconomy relationship". In this section, I identify some conditions which *do* seem to produce reliable oil price-macroeconomy correlations in post-1980 data.

My starting point is the support that Carruth, Hooker, and Oswald (1998) (hereafter, CHO) found for a model relating unemployment to the real level of oil prices. The economic model in that paper is based on the simple idea that when firms' costs rise, they must pay lower wages (the output market is competitive), which is only possible in an efficiency wage framework with a higher equilibrium unemployment rate. Changes in oil prices, an important and exogenous input to production, thus drive the "natural rate" of unemployment. CHO find support for this model in Granger causality tests, and in out-of-sample unemployment forecasts using a cointegration/error correction model.

A. Granger causality tests

CHO found strong evidence of Granger causality from the real level of oil prices to unemployment in their full sample (1954:2-1995:2) as well as the subsamples 1954-1973 and 1973-1995, in contrast to the standard breakdown result. One difference between their tests and those in Figure 2 and Table 1, other than sample coverage, is that the CHO tests are in bivariate and trivariate (including real long-term interest rates) equations, while the tests above are in multivariate equations that include import price inflation, GDP deflator inflation, and 3-month Treasury bill rates.

It turns out that the inclusion or exclusion of these conditioning variables is very important. The final pair of columns in Table 1 report *p*-values from the same 'post-1980' Granger causality tests as in the earlier columns, using the real oil price level (the PPI for crude oil divided by the PPI for all commodities). As the Table shows, this oil price does not contain much information useful for predicting movements in output and unemployment beyond that contained in lags of interest rates, import price inflation, and GDP deflator inflation. However, parallel results from *bivariate* tests, where only output or unemployment and the oil price are included, shown in Table 2, are very different. In those tests, CHO's results come through clearly: real oil prices Granger cause unemployment in nearly all of the possible cases at the 10% level, and in many cases at the 1% level.

Exclusion of conditioning variables from the VARs is not sufficient to achieve Granger causality of output, however: Table 2 shows that the improvement for SOPI and NOPI, and for real oil levels with output, is much more modest. Again, the work in CHO provides a useful reference point. As suggested by that paper's *equilibrium* unemployment rate characterization, oil prices affect the macroeconomy at fairly low frequencies, while output growth is measured as quarterly changes in real GDP, which are quite noisy. A simple way to extract a lower-frequency component of output is to use its year-over-year changes. Surprisingly, this small modification makes a big difference: Table 3 shows that SOPI and NOPI strongly Granger cause annual output growth in almost all of the bivariate cases.

B. Out-of-sample forecasts

The Granger causality test results in Table 2 provide some support for the position that real oil prices drive unemployment fluctuations. To bolster that argument, and for comparison with the forecasts in Section 3, I use a version of the CHO model to forecast unemployment over 1991-1998. The ECM equation is estimated in first differences, with lag lengths chosen by simple t-statistic and R^2 considerations:

$$\text{(3a)}\quad \Delta u_t = \alpha + \beta_1 \Delta u_{t-1} + \beta_2 \Delta u_{t-2} + \gamma_1 \Delta o_{t-1} + \gamma_2 \Delta o_{t-2} + \gamma_3 \Delta o_{t-3} + \delta_1 \Delta r_{t-1} + \delta_2 \Delta r_{t-2} + \theta_1 ECM_{t-1} + \varepsilon_t;$$

where u_t is the unemployment rate, o_t is the real oil price level, r_t is the real interest rate (5-year Treasury yield less contemporaneous GDP deflator inflation, representing non-wage input costs), and ECM_t is the residual from the cointegrating regression[12]

12 Cointegration analysis requires that each variable be integrated and is usually preceded by pretesting for stationarity, while the unemployment rate is often taken to be stationary because it is bounded by 0 and 1. However, given that unemployment is highly serially correlated and a unit root is often not rejected, I believe that it may sensibly be regarded as "pseudo I(1)", and the resulting forecasts judged on their own merit.

(3b) $$u_t = a + b_1 o_t + b_2 r_t + e_t.$$

The estimated parameters of the cointegrating regression and ECM equation are reported in Table 4. The cointegrating relationship indicates that real oil price and interest rate levels are positively correlated with movements in the unemployment rate; the oil price is highly significant while the interest rate is not. The ECM equation indicates that when unemployment is away from its "natural rate," it tends to move towards that rate (the ECM term has a negative and significant coefficient).

Forecasts from this model are constructed in a way parallel to those in the previous section: Both the cointegrating regression and the ECM equation are estimated using data up through 1990, and then forecasts are recursively generated from 1991:1 out to the end of 1998. Actual values of real oil prices again are used, to see how well they help keep the model on track, and the predicted unemployment rates are recursively substituted into equation (3a) in both the lag and ECM terms.[13]

The forecasts are plotted in Figure 7. Like the equations using SOPI and NOPI, this model has only modest success in capturing the increase in unemployment from the 1990-91 recession. However, the ECM equation *is* able to translate the large decline in oil prices after 1991 into predictions of a sustained decline in unemployment which tracks the actuals over much of the period. It is also fooled far less than the SOPI and NOPI equations by the upward movements in oil prices in 1994-96, predicting a less than 1 percentage point rise in unemployment and a return to the downtrend in 1997. At the end of the forecast period, the ECM model is off by less than a percentage point, compared with more than two percentage points for the SOPI and NOPI forecasts. For the horizon as a whole, the ECM model has a root mean squared error of 0.86, vs. 1.29 for SOPI and 1.17 for NOPI, more than a 25% improvement, consistent with the argument that oil price decreases have had favorable and significant effects on unemployment.

5. Discussion

Evidence on the macroeconomic effects of oil price shocks may be summarized as follows.

---

13 Actual values of the real interest rate are used as well. However, allowing the model to know the real interest rate is of little consequence. As CHO found, the model places little weight on this variable, and replacing actual values with various counterfactuals like the average or trend real interest rate leads to very minor changes in the forecasts.

The significant Granger causality from nominal crude oil price changes to output and unemployment found by Hamilton (1983) broke down in the 1980s amidst a new regime of highly volatile oil price movements. Asymmetric and nonlinear transformations of oil prices are not sufficient to restore the relationship, contrary to the claims of their proponents: the two leading candidates, Lee, Ni and Ratti's (1995) and Hamilton's (1996) SOPI and NOPI trans-formations, do not generally Granger cause output or unemployment in post-1980 data, and they owe much of their apparent success to an improved fit in the 1950s data. Finding consistent correlations between macroeconomic activity and oil prices in post-1980 data seems to require some further conditions; a sufficient set is to use bivariate equations and express output in annual changes. Even then, SOPI and NOPI predict output, but not unemployment, while the real oil price level (a simple linear and symmetric specification) predicts unemployment but not output.

Should this evidence be interpreted as supportive of oil prices as important determinants of economic activity after 1980? A skeptical view is that the necessity of such particular conditions for significant results is a symptom of data mining. However, sensible explanations exist for those conditions.

Bivariate equations might yield different results from multivariate ones if the oil price-macroeconomy transmission mechanism is indirect. For example, the Federal Reserve may tighten when oil prices rise to head off incipient inflation. The proximate cause of the ensuing downturn would then be higher interest rates, and oil prices might contribute little marginal explanatory power for output. The breakdown of Granger causality in multivariate systems around 1980 would then be consistent with a change in the joint oil price, monetary policy, real economy relationship around that time. In fact, the Fed began a disinflation campaign with the appointment of Paul Volcker in 1979, which quite likely involved a different policy response to oil price changes than had prevailed in the 1970s. In fact, Bernanke, Gertler and Watson (1997) estimated federal funds rate responses to oil price shocks that were very different across three subsamples spanning 1966-75, 1976-85 and 1986-95. Less directly, Hooker (1999) finds strong evidence of a structural break in the oil price coefficients of some Phillips curve inflation equations around 1980. An investigation of models with monetary policy regime shifts, to see whether they can generate patterns of oil price-macroeconomy behavior like those found in this paper, is a topic for future research.

That oil price effects are at low frequencies, which is consistent with output responding more in annual than quarterly terms, does not seem controversial. Almost all of the (many) theoretical oil price-macroeconomy transmission mechanisms—including changes in productivity owing to different effective capital/labor ratios, induced movements of workers and capital across regions and industries, and effects on the nature and timing of capital investment decisions—involve medium- to long-term processes. The unemployment rate apparently needs no low-frequency filter, as its standard deviation is roughly a third that of quarterly output growth, and a much larger percentage of its variability comes from low-frequency components.

Perhaps the most puzzling result in the paper is that different transformations of oil prices seem to affect output and unemployment. If in fact oil price levels drive the unemployment rate, while only scaled or net oil price increases affect output, then some components of oil price changes induce departures from Okun's law rather than movements consistent with it. For example, sustained oil price decreases, like those observed over 1991-1998, could lead to substantial reductions in unemployment that are reflected only moderately in output growth. A deeper investigation of both the theoretical and empirical aspects of this result is another topic for further research.

Data Appendix

The data used in this study, the source, and any transformations are as follows. Data are available on request.

Treasury bill rate: aggregated from monthly averages to quarterly using the middle month of the quarter.

Inflation: GDP chain-type deflator, seasonally adjusted, entered as 400*log first difference.

Import price inflation: ratio of nominal to chained 1992 dollars' imports, NIPA accounts, seasonally adjusted, entered as 400*log first difference.

Unemployment rate: civilian age 16 and over rate, seasonally adjusted, aggregated from monthly to quarterly using the middle month of the quarter.

Output: in chained 1992 dollars, seasonally adjusted. Quarterly changes are $400*\ln(y_t/y_{t-1})$ and year-over-year changes are $100*(y_t/y_{t-4} - 1)$, where y_t is real GDP in quarter t.

Mork's oil price: through 1972, 400*log first difference of the PPI for crude oil (BLS series WP0561), aggregated using middle month of quarter. From 1974, the same transformation applied to the DOE composite domestic first purchase price (ftp.eia.doe.gov). Intervening quarters are the growth rate of WP0561 multiplied by 1.095 (see Mork (1989)).

SOPI oil price: The parameters of the GARCH model (1a-b) were obtained via maximum likelihood on Mork's series. The recursions to create $SOPI_t$ use the estimated unconditional variance and its square root, respectively, for the initial values of $\hat{h}_t$ *and* $\hat{\varepsilon}_t$ **.**

NOPI oil price: the percentage change in Mork's levels series from the past four quarters' high if that is positive, and zero otherwise (see equation (2)).

Real oil price level: the PPI for crude oil divided by the PPI for all commodities, WP0000, aggregated from monthly to quarterly using the middle month of the quarter.

REFERENCES

Bernanke, Ben S., Mark Gertler, and Mark Watson, "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity* 1:1997, 91-142.

Burbidge, John and Alan Harrison, "Testing for the Effects of Oil-Price Rises using Vector Autoregressions," *International Economic Review* 25 (1984), 459-84.

Carruth, Alan A., Mark A. Hooker and Andrew J. Oswald, "Unemployment Equilibria and Input Prices: Theory and Evidence for the United States," *Review of Economics and Statistics* 80 (1998), 621-28.

Davis, Steven J. and John Haltiwanger, "Sectoral Job Creation and Destruction Responses to Oil Price Changes," NBER Working Paper No. 7095 (1999).

Hamilton, James D., "Oil and the Macroeconomy since World War II," *Journal of Political Economy* 91 (1983), 228-48.

_______, "Historical Causes of Postwar Oil Shocks and Recessions," *Energy Journal* 6 (1985), 97-116.

_______, "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics* 38 (1996), 215-20.

_______, "What is an oil price shock," manuscript, UCSD, 1999.

Hooker, Mark A., "What Happened to the Oil Price-Macroeconomy Relationship?," *Journal of Monetary Economics* 38 (1996), 195-213.

_______, "Inflation, Oil Prices, and Monetary Policy," manuscript (1999).

Lee, Kiseok, Shawn Ni, and Ronald Ratti, "Oil Shocks and the Macroeconomy: The Role of Price Variability," *Energy Journal* 16 (1995), 39-56.

Mork, Knut Anton, "Oil Shocks and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy* 97 (1989), 740-44.

Thoma, Mark and Jo Anna Gray, "Financial Market Variables do not Predict Real Activity," *Economic Inquiry* 36 (1998), 522-39.

Table 1: Multivariate Oil Price Granger Causality Tests

	SOPI		NOPI		RoilL	
	(1)	(2)	(3)	(4)	(5)	(6)
	output	unemp.	output	unemp	output	unemp.
A. Sample 1979:1-1998:4						
4 lags	.15	.26	.24	.46	.22	.30
6 lags	.68	.25	.56	.40	.21	.25
8 lags	.93	.63	.08*	.06*	.53	.33
B. Sample 1980:1-1998:4						
4 lags	.43	.13	.11	.46	.02**	.36
6 lags	.44	.18	.21	.22	.24	.17
8 lags	.82	.73	.02**	.12	.31	.55
C. Sample 1981:1-1998:4						
4 lags	.88	.37	.36	.56	.53	.59
6 lags	.25	.66	.01**	.51	.15	.70
8 lags	.79	.96	.02**	.18	.15	.53

Notes: SOPI and NOPI are as described in the text; RoiL is the PPI for crude oil divided by the PPI for all commodities. Entries in the table are *p*-values for the *F*-test that all oil price coefficients equal zero in the output or unemployment equation of a VAR that also includes the 3-month Treasury bill rate, import price inflation, and GDP deflator inflation. See the data appendix for details. *, **, *** denote significance at the 10, 5, and 1% levels.

Table 2: Bivariate Oil Price Granger Causality Tests

	SOPI		NOPI		RoilL	
	(1)	(2)	(3)	(4)	(5)	(6)
	output	unemp.	output	unemp	output	unemp.
A. Sample 1979:1-1998:4						
4 lags	.00***	.02**	.00***	.19	.06*	.00***
6 lags	.11	.46	.13	.02**	.30	.01**
8 lags	.13	.09*	.04**	.01**	.02**	.00***
B. Sample 1980:1-1998:4						
4 lags	.02**	.13	.01**	.71	.02**	.00***
6 lags	.24	.37	.11	.18	.26	.01***
8 lags	.50	.17	.19	.02**	.40	.03**
C. Sample 1981:1-1998:4						
4 lags	.45	.95	.21	.94	.65	.02**
6 lags	.60	.92	.10	.33	.86	.23
8 lags	.63	.76	.23	.07*	.56	.27

Notes: SOPI and NOPI are as described in the text; RoilL is the PPI for crude oil divided by the PPI for all commodities. Entries in the table are *p*-values for the *F*-test that all oil price coefficients equal zero in the output or unemployment equation of a bivariate VAR. *, **, *** denote significance at the 10, 5, and 1% levels.

Table 3: Oil Price Granger Causality Tests with Year-over-Year Output Growth

	multivariate			bivariate		
	SOPI	NOPI	RoilL	SOPI	NOPI	RoilL
	(1)	(2)	(3)	(4)	(5)	(6)
A. Sample 1979:1-1998:4						
4 lags	.22	.31	.03**	.04**	.03**	.11
6 lags	.89	.80	.78	.02**	.05*	.14
8 lags	.82	.58	.43	.06*	.12	.02**
B. Sample 1980:1-1998:4						
4 lags	.22	.35	.03**	.01**	.09*	.12
6 lags	.21	.26	.74	.01***	.01***	.07*
8 lags	.16	.08*	.24	.03**	.01***	.02**
C. Sample 1981:1-1998:4						
4 lags	.33	.20	.24	.03**	.05**	.88
6 lags	.12	.00***	.51	.30	.01**	.72
8 lags	.31	.00***	.24	.53	.04**	.18

Notes: SOPI and NOPI are as described in the text; RoiL is the PPI for crude oil divided by the PPI for all commodities. Entries in the table are *p*-values for the *F*-test that all oil price coefficients equal zero in the year-over-year output growth equation of a VAR. *, **, *** denote significance at the 10, 5, and 1% levels.

Table 4: A Cointegration-Error Correction Model of Oil Prices and Unemployment

A. Cointegrating Regression: Dependent variable = unemployment rate

Constant	2.876
	(10.50)
Real Oil Price	.054
	(9.40)
Real Interest Rate	.047
	(1.08)

B. Unit Root Tests

	DF	ADF1	ADF2	ADF3	ADF4
Unemployment	-1.78	-2.87*	-3.09**	-2.89*	-2.37
ECM Without Trend	-3.79	-4.62***	-4.67***	-4.17**	-3.70*
ECM With Trend	-4.19	-4.65**	-4.71**	-4.21**	-3.73

C. Error Correction Model: Dependent variable = change in unemployment rate

Constant	.009
	(0.29)
Δu_{t-1}	.405
	(5.13)
Δu_{t-2}	.131
	(1.59)
Δo_{t-1}	.001
	(0.17)
Δo_{t-2}	-.012
	(-1.58)
Δo_{t-3}	.017
	(2.28)
Δrr_{t-1}	-.026
	(-.98)
Δrr_{t-2}	-.014
	(-.53)
ECM_{t-1}	-.140
	(-4.27)

Adjusted R^2 = .29; ser = .37, DW = 2.02.

Notes: *t*-statistics in parentheses. DF and ADF refer to augmented Dickey-Fuller unit root test statistics with the specified number of lags. *, ** , *** denote significance at the 10, 5, and 1% levels, from MacKinnon's (1990) tables.

Figure 1: Real Oil Price (PPI for Crude Oil over PPI for all Commodities)

Figure 2: Granger Causality from Standard Measures of Oil Prices to Output

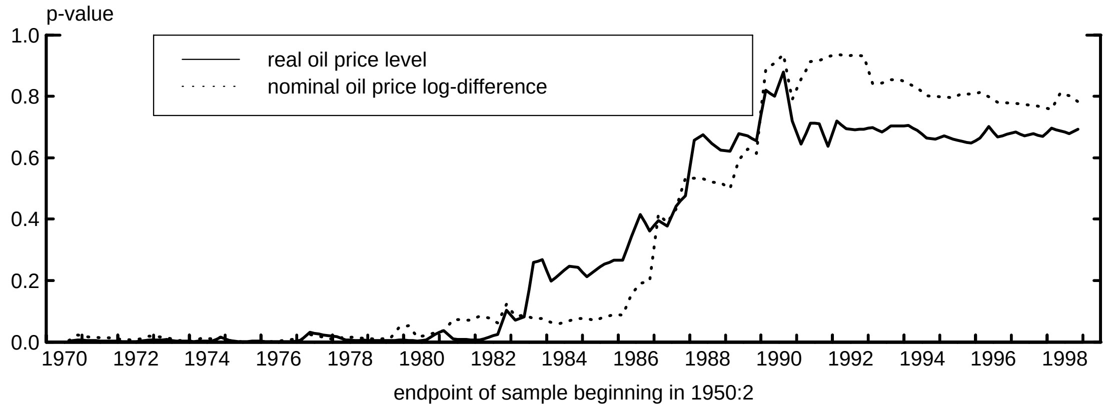

Figure 3a: SOPI Oil Price

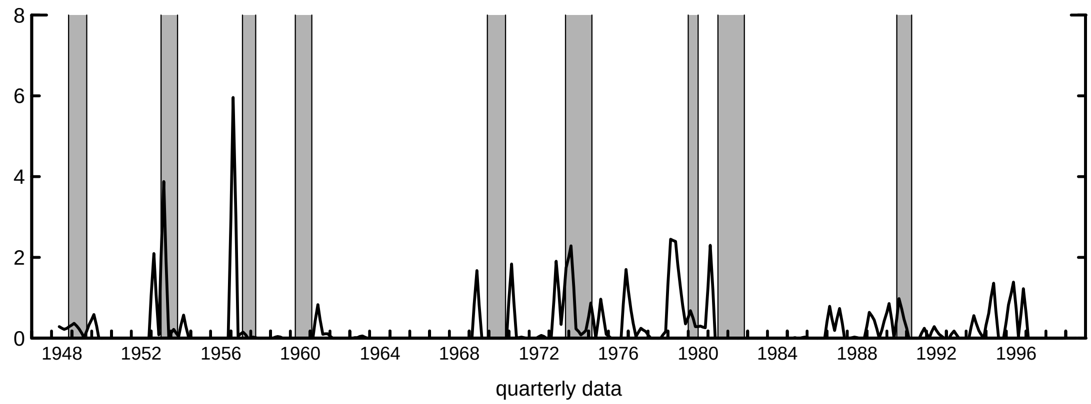

Figure 3b: NOPI Oil Price

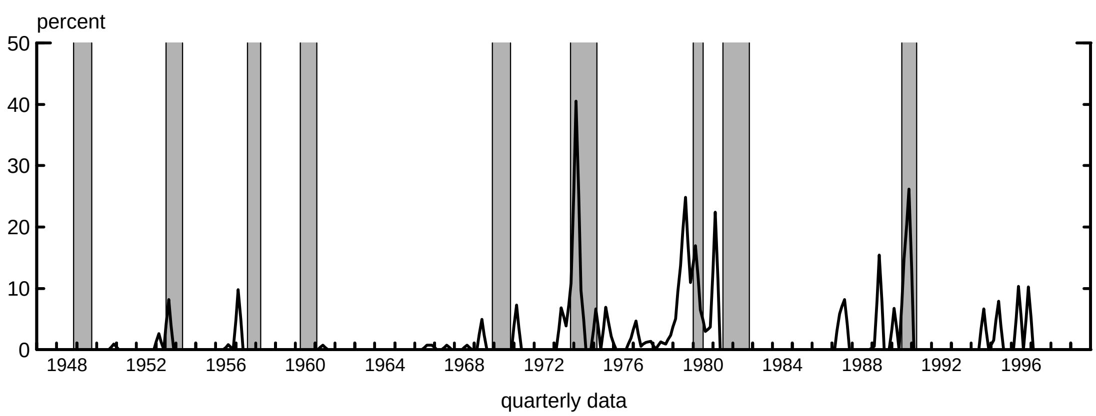

Figure 4: Granger Causality from Transformations of Oil Prices to Output

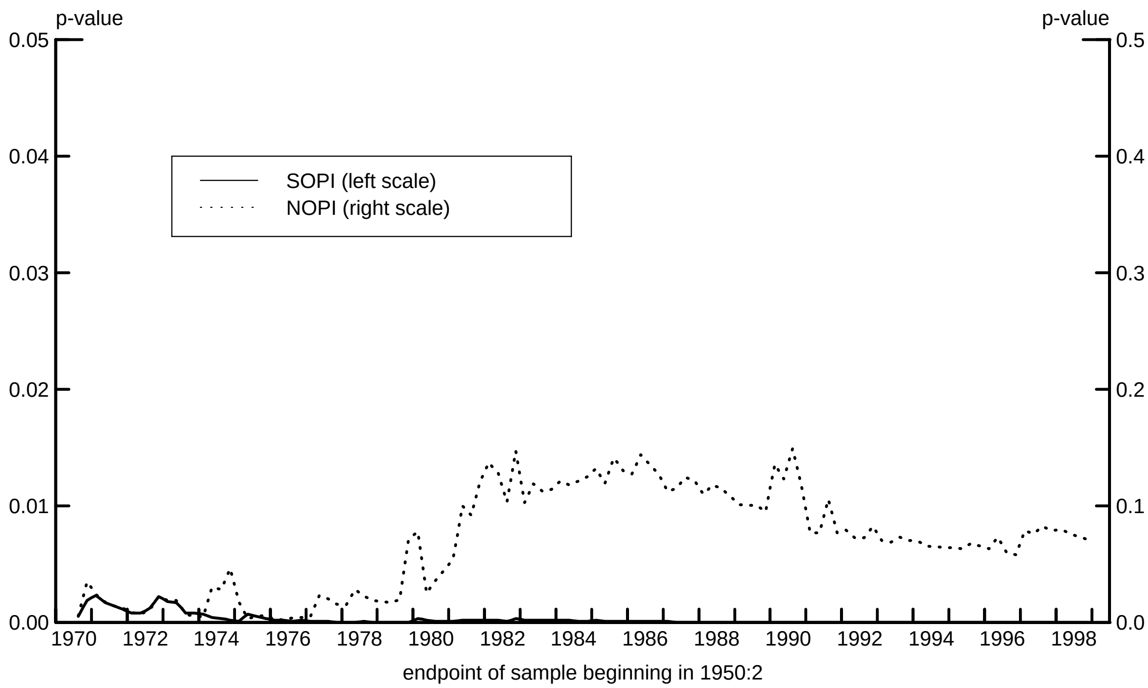

Figure 5: Granger Causality from SOPI and NOPI to Output, 1957:1 deleted

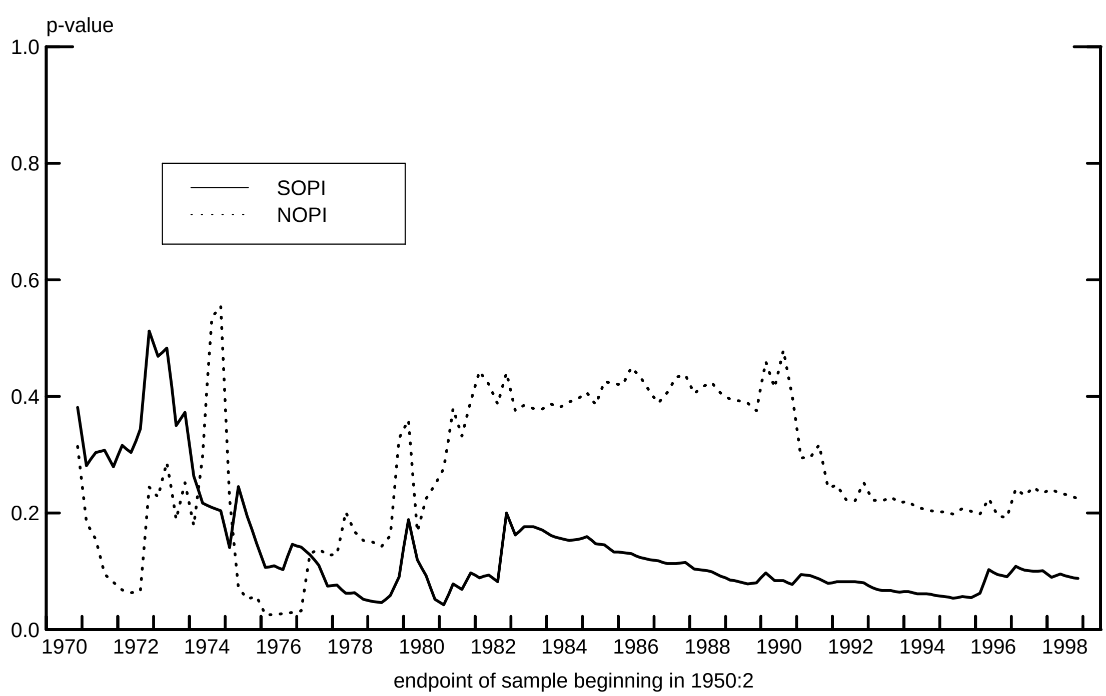

Figure 6: Unemployment rate and forecasts using SOPI and NOPI

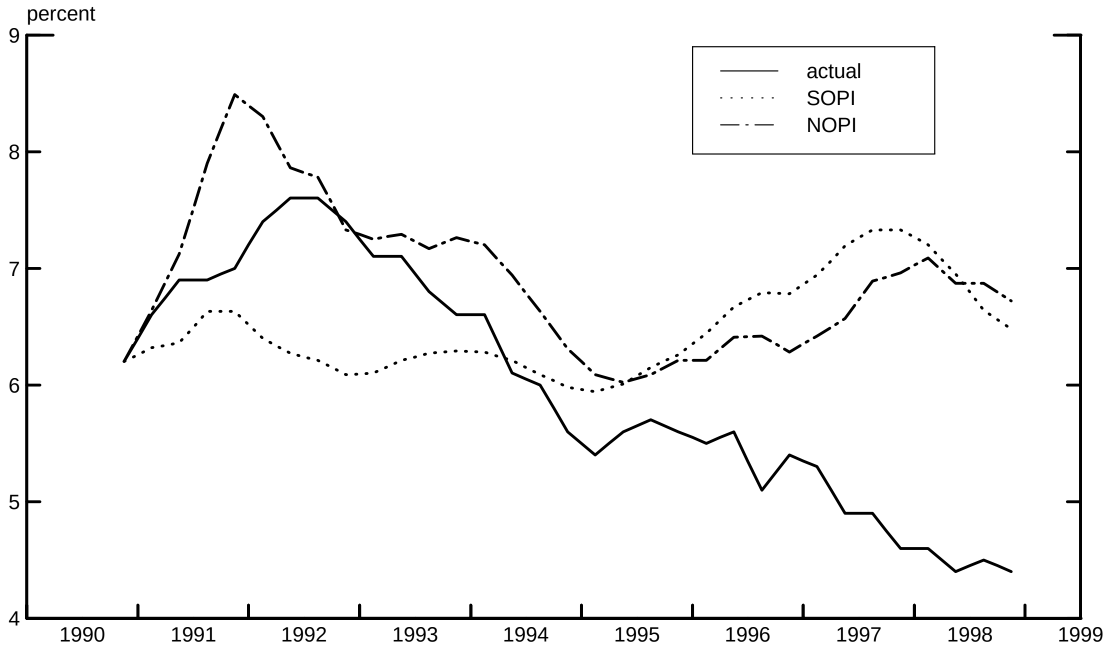

Figure 7: Unemployment rate and forecasts using oil price level