

NBER WORKING PAPER SERIES

WHAT IS AN OIL SHOCK?

James D. Hamilton

Working Paper 7755
<http://www.nber.org/papers/w7755>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2000

This paper is based on research supported by the NSF under Grant No. SBR-9707771 and the Department of Energy. It is adapted from (and replaces) an earlier paper circulated under the title, "Analysis of the Transmission of Oil Price Shocks through the Macroeconomy." I thank Julian Betts and Michael Dotsey for helpful comments. All data and software used in this paper can be downloaded free of charge from <http://weber.ucsd.edu/~jhamilto>. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2000 by James D. Hamilton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What is an Oil Shock?
James D. Hamilton
NBER Working Paper No. 7755
June 2000
JEL No. E3

ABSTRACT

This paper uses a flexible approach to characterize the nonlinear relation between oil price changes and GDP growth. The paper reports clear evidence of nonlinearity, consistent with earlier claims in the literature-- oil price increases are much more important than oil price decreases, and increases have significantly less predictive content if they simply correct earlier decreases. An alternative interpretation is suggested based on estimation of a linear functional form using exogenous disruptions in petroleum supplies as instruments. The evidence suggests that oil shocks matter because they disrupt spending by consumers and firms on certain key sectors.

James D. Hamilton
Department of Economics, 0508
University of California, San Diego
La Jolla, CA 92093-0508
and NBER
jhamilton@ucsd.edu

A large body of research suggests that there is a significant effect of energy supply disruptions on economic activity. A clear negative correlation between energy prices and aggregate measures of output or employment has been reported by Rasche and Tatom (1977, 1981), Hamilton (1983), Burbidge and Harrison (1984), Santini (1985, 1994), Gisser and Goodwin (1986), Rotemberg and Woodford (1996), Daniel (1997), Raymond and Rich (1997) and Caruth, Hooker, and Oswald (1998), among others. Analyses of microeconomic data sets at the level of individual industries, firms, or workers also demonstrate significant correlations between oil price shocks and output, employment, or real wages (Keane and Prasad, 1996; Davis, Loungani, and Mahidhara, 1996; Davis and Haltiwanger, 1997; Lee and Ni, 1999). Nevertheless, the suggestion that oil price shocks contribute directly to economic downturns remains controversial, in part because the correlation between oil prices and economic activity appears to be much weaker in data obtained since 1985; (see Hooker, 1996).

A number of authors have attributed this instability of the empirical relation between oil prices and output to misspecification of the functional form. Loungani (1986), Davis (1987a,b), Mork (1989), Lee, Ni and Ratti (1995), Hamilton (1996), Davis, Loungani, and Mahidhara (1996), Davis and Haltiwanger (1997), and Balke, Brown, and Yücel (1999), among others, have suggested that the relation between oil prices and economic activity is nonlinear. Insofar as there has been a shift in the process generating oil prices, a linear approximation to the relation between oil prices and economic activity will appear unstable over time, even if the underlying nonlinear relation is stable.

One problem with suggesting that this is indeed what happened is that there is an un-

bounded universe of alternative nonlinear specifications. How does one decide which nonlinear specification is the right one to use, and how can we distinguish between a statistically significant nonlinear relation and the outcome of determined data-mining?

This paper applies a methodology recently developed by Hamilton (2000) to address these questions. This approach provides a valid test of the null hypothesis of linearity against a broad range of alternative nonlinear models, consistent estimation of what the nonlinear relation looks like, and formal comparison of alternative nonlinear models. The results generate strong support for the claim of a nonlinear relation along the lines suggested in the literature: oil price increases affect the economy whereas decreases do not, and increases that come after a long period of stable prices have a bigger effect than those that simply correct previous decreases. Although the data do not allow us to distinguish clearly between the particular nonlinear specifications examined in detail here, they do provide a clear basis for ruling out certain economic hypotheses of how oil shocks might affect the macroeconomy.

The paper also shows that very similar results can be obtained with a linear instrumental variables regression, in which oil prices are presumed to have a linear effect on the economy and identifiable exogenous disruptions in world petroleum supplies are used as instruments. It appears that the nonlinear transformations favored by the data amount to something very similar to focusing on only those oil price movements that were caused by exogenous military conflicts. The paper suggests that the time series for the magnitude of war-induced petroleum supply disruptions can be used to estimate linear reduced-form regressions as an alternative to using a nonlinear transformation of the oil price series itself.

The plan of the paper is as follows. The first section discusses why an investigation of the linearity of the relationship is important both for econometric inference and economic interpretation. Section 2 reviews the methodology applied in this paper. Empirical results are presented in Section 3. Section 4 suggests an alternative interpretation of the results in terms of instrumental variable estimation. Conclusions are offered in Section 5.

1 Why functional form matters.

Many economic analyses of the effects of oil shocks¹ begin with a production function relating output to inputs of capital, labor, and energy. An exogenous decrease in the supply of energy reduces output directly by lowering productivity and indirectly to the extent that lower wages induce movement along a labor supply schedule (Rasche and Tatom, 1977, 1981; Kim and Loungani, 1992), changes in business markups (Rotemberg and Woodford, 1996), or capacity utilization rates (Finn, 1997). These models imply that the log of real GDP should be linearly related to the log of the real price of oil. One implication of this linearity is that if the price of oil goes down, then output should go up; if an oil price increase brings about a recession, then an oil price decline should induce an economic boom by the same mechanism operating in the reverse direction.²

¹ Useful reviews of the different mechanisms by which oil shocks could affect economic performance are provided by Bohi (1989) and Mork (1994).

² Atkeson and Kehoe (1999) offered an interesting extension of this class of models by assuming putty-clay investment technology. In their formulation, an oil price decrease still produces an increase in output, though the output boom from a large oil price decline is smaller in magnitude than the output decline that follows an oil price increase of the same logarithmic magnitude.

These models view recessions as supply driven rather than demand driven. According to these models, an oil price increase produces a recession because it makes cars more costly to manufacture. This seems contrary to reports in the trade and business press, in which the problem is invariably perceived as a reduction in the number of cars consumers are willing to buy.

A number of early analyses focused instead on demand-side effects of an oil price increase. In these models, an increase in oil prices would increase the overall price level, which, given the Keynesian assumption of rigid wages, reduces employment. Examples of such models include Pierce and Enzler (1974), Solow (1980), and Pindyck (1980). Mork and Hall (1980) demonstrated the potential for interactive effects between wage rigidities and supply-side effects. These models again all maintain the existence of a linear relation between the log of the price of oil and the log of GDP, so that again an oil price decline is expected to produce an economic boom.

These models also have the characteristic that there is nothing all that special about oil. The basic economic inefficiency is the familiar Keynesian mismatch between the aggregate wage and the aggregate price level, and oil price disruptions are just one of many developments that might contribute to such a mismatch.

Surely the price and availability of gasoline matter for car sales not simply because they affect the overall price level but further because they are key inputs in how cars get used. Is your next car going to be a small foreign car or a large SUV? Your decision depends in part on what you think about gasoline availability. If you are very unsure about where gas

prices are headed, you might be inclined to postpone a new purchase until you have a better idea of where the market stands.

Energy prices and availability may be quite relevant for a host of other durable goods purchases, including housing. How long a commute to work are you willing to put up with? How energy-efficient should your appliances, windows, and insulation be? What equipment and industrial techniques should a firm build a new factory around? When energy prices and availability are as uncertain as they were in early 1974, it is rational to postpone such commitments until better information is available.

Oil shocks may matter for short-run economic performance precisely because of their ability temporarily to disrupt purchases of large-ticket consumption and investment goods. A major disruption in oil supplies makes people uncertain about the future, with the result that spending on cars, housing, appliances, and investment goods temporarily falls. A variety of microeconomic evidence suggests that oil shocks have substantial potential to exert such effects. Bresnahan and Ramey (1993) documented that the oil shocks of 1974 and 1980 caused a significant shift in the mix of demand for different size classes of automobiles with an attendant reduction in capacity utilization at U.S. automobile plants. Sakellaris (1997) found that changes in the stock market valuation of different companies in response to the 1974 oil shock were significantly related to the vintage of their existing capital. Davis and Haltiwanger (1997) discovered a dramatic effect of oil price shocks on the rate of job loss in individual economic sectors, with the job destruction rising with capital intensity, energy intensity, product durability, and plant age and size. See also Bernanke (1983), Loungani

(1986), Davis (1987a,b), Hamilton (1988a,b), Santini (1992), and Davis, Loungani, and Mahidhara (1996) for related evidence and discussion.

These studies have further noted that, if allocative disturbances are indeed the mechanism whereby oil shocks affect economic activity, then there is no reason to expect a linear relation between oil prices and GDP. An oil price increase will decrease demand for some goods but possibly increase demand for others. If it is costly to reallocate labor or capital between sectors, the oil shock will be contractionary in the short run. Note moreover that an oil price decrease could also be contractionary in the short run. A price decrease also depresses demand for some sectors, and unemployed labor is not immediately shifted elsewhere. Furthermore, if it is primarily the postponement of purchases of energy-sensitive big-ticket items that produces the downturn, then an oil price decrease could in principle be just as contractionary as an oil price increase.

Of course, an oil price decrease is not all bad news, by virtue of the production function and inflation effects noted earlier. But surely it is unreasonable to assume that an oil price decrease would produce an economic boom that mirrors the recession induced by an oil price increase.

As a simple statistical illustration of how the specification of functional form can matter in practice, consider the following example. Let y_t denote the growth rate of real GDP and let o_t denote the percentage change in the price of oil. Let us assume that the effect of oil prices on output is given by

$$y_t = f(o_t) + \varepsilon_t \tag{1.1}$$

where ε_t is a regression error term. Suppose that every one percent increase in oil prices produces a β percent decrease in real GDP, but that a decrease in oil prices has no effect on GDP. Then the function $f(o_t)$ takes the form

$$f(o_t) = \begin{cases} \alpha & \text{if } o_t \leq 0 \\ \alpha - \beta o_t & \text{if } o_t > 0 \end{cases}. \quad (1.2)$$

I simulated 50 observations from equations (1.1) and (1.2) with a critical precondition— all of the o 's in the sample were positive.³ Thus for this particular sample the data satisfy the classic linear regression assumptions. Ordinary least squares (OLS) regression produces excellent inference about the values of α and β (standard errors in parentheses):

$$y_t = \underset{(0.47)}{2.29} - \underset{(0.044)}{0.117} o_t. \quad (1.3)$$

Oil prices are inferred to have a strong and statistically significant effect on the economy, with a t -statistic of -2.66. The simulated data, estimated relation, and true relation are displayed in Figure 1.

Now let us double the sample size but allow both positive and negative values for o_t , as displayed in Figure 2. The OLS regression estimates now turn out to be

$$y_t = \underset{(0.16)}{1.65} - \underset{(0.015)}{0.029} o_t. \quad (1.4)$$

The result of using the larger sample is that oil prices are only imputed to have 1/4 as big an effect as they seemed to have in the smaller sample, and this effect is no longer statistically

³ The simulation used $\alpha = 2$, $\beta = 0.1$, $u_t \sim N(0, 2.25)$, and $o_t \sim N(0, 100)$, where (o_t, y_t) pairs were thrown out if $o_t \leq 0$ until 50 pairs of observations were generated. These values were chosen to correspond to rough magnitudes that might characterize the actual U.S. relation.

significant. The reason is that if equation (1.2) represents the true model and if the sample includes negative values for o , then regression (1.4) is misspecified and is not providing a consistent estimate of the parameter β .

Mork (1989) argued that this is essentially what is going on with the historical U.S. experience.⁴ In the postwar data up until 1980, there was very little experience with falling oil prices, so that the sample was essentially like that in Figure 1. After 1980, however, there are a lot of observations of big oil price decreases, so that the sample becomes more like that shown in Figure 2. The result is that when more observations are added to the sample, the estimated effect becomes smaller in magnitude and loses statistical significance.

Specifically, to get away from simulated data and turn to the actual numbers, consider the results of regressing each quarter's GDP growth (y_t) on four lags of GDP growth and four lags of the percent change in the nominal price of crude petroleum (o_{t-j}). When this regression is estimated for data from 1949:II to 1980:IV, the result is

$$\begin{aligned}
 y_t = & 1.19 + 0.20 y_{t-1} + 0.06 y_{t-2} - 0.09 y_{t-3} - 0.20 y_{t-4} \\
 & \quad \quad \quad (0.19) \quad (0.09) \quad (0.09) \quad (0.09) \quad (0.09) \\
 & - 0.003 o_{t-1} - 0.030 o_{t-2} - 0.036 o_{t-3} - 0.064 o_{t-4}. \quad (1.5) \\
 & \quad \quad \quad (0.027) \quad (0.027) \quad (0.027) \quad (0.028)
 \end{aligned}$$

The coefficient on o_{t-4} has a t -statistic of -2.3. The estimated values imply that a 10% increase in oil prices will result four quarters later in a level of GDP that is 1.4% lower than it otherwise would be. Because of the imposed linearity, the regression also requires that a 10% decrease in oil prices will result in a 1.4% higher level of GDP.

⁴ Actually, Mork estimated separate coefficients for oil price increases and decreases, and found that the coefficients on decreases were insignificantly different from zero. The argument as presented in the text here jumps to the conclusion in a single step. Mork, Olsen and Mysen (1994) reported more qualified support for this idea in the experience of other countries.

When the same regression is reestimated using data from 1949:II to 1999:IV, the result is

$$\begin{aligned}
y_t = & \underset{(0.12)}{0.74} + \underset{(0.07)}{0.28} y_{t-1} + \underset{(0.07)}{0.12} y_{t-2} - \underset{(0.07)}{0.06} y_{t-3} - \underset{(0.07)}{0.12} y_{t-4} \\
& - \underset{(0.007)}{0.004} o_{t-1} - \underset{(0.007)}{0.005} o_{t-2} - \underset{(0.007)}{0.007} o_{t-3} - \underset{(0.007)}{0.018} o_{t-4}.
\end{aligned} \tag{1.6}$$

The coefficient on o_{t-4} is about 1/4 of its value in the smaller sample, though it remains statistically significant at the 5% level. The reason the coefficient on o_{t-4} is much smaller in the larger sample is that an oil price decrease of 10% does not add 1.4% to the level of GDP. In order for a linear relation to be consistent with what happened after the oil price declines during 1980-1994, a smaller coefficient is needed.

As evidence in support of Mork's claim that the historical regressions (1.5) and (1.6) are reflecting the same factors as the simulated regression (1.3) and (1.4), consider imposing the functional form (1.2) directly. Define

$$o_t^+ = \begin{cases} 0 & \text{if } o_t \leq 0 \\ o_t & \text{if } o_t > 0 \end{cases}. \tag{1.7}$$

When o_{t-j} in (1.6) is replaced by o_{t-j}^+ , the estimated relation over 1949:II to 1999:IV is

$$\begin{aligned}
y_t = & \underset{(0.14)}{0.96} + \underset{(0.07)}{0.25} y_{t-1} + \underset{(0.07)}{0.10} y_{t-2} - \underset{(0.07)}{0.08} y_{t-3} - \underset{(0.07)}{0.14} y_{t-4} \\
& - \underset{(0.009)}{0.012} o_{t-1}^+ - \underset{(0.010)}{0.007} o_{t-2}^+ - \underset{(0.010)}{0.014} o_{t-3}^+ - \underset{(0.010)}{0.030} o_{t-4}^+.
\end{aligned} \tag{1.8}$$

The estimated effects of oil price increases are considerably larger and more statistically significant than those implied by the linear relation (1.6).

Although the functional form in equation (1.2) seems consistent with the empirical evidence in regressions (1.5), (1.6), and (1.8), Mork's interpretation has recently been challenged on two grounds. First, Hooker (1996) argued that even the asymmetric relation (1.8) offers a relatively poor fit to data since 1986. Second, if the mechanism is indeed that an oil price increase causes postponement of certain major purchases, then equation (1.2) is surely too crude—consumers' behavior should be based not just on whether oil prices went up, but further on what they believe the increase means for the future. Most of the quarters in which oil prices went up since 1986 followed a quarter in which oil prices had gone down even more, so that the increases were simply partial corrections to a chronic downward trend. Several authors have suggested alternative functional forms that might better represent the true relation. Ferderer (1996) argued that oil price volatility itself depresses spending. Lee, Ni, and Ratti (1995) suggested that what matters is how surprising an oil price increase is based on the observed recent changes.⁵ Hamilton (1996) claimed that the key question is whether the oil price increase is big enough to reverse any decreases observed in the immediately preceding quarters. Davis, Loungani, and Mahidhara (1996) and Davis and Haltiwanger (1997) focused on whether the oil price increase was sufficient to raise the price above its previous 5-year average.

All of these specifications have a certain plausibility. One logical way to sort out the

⁵ Lee, Ni and Ratti constructed a variable $e_t^* = e_t/\sqrt{h_t}$ where e_t is the error in forecasting the real price of oil based on past observations and $\sqrt{h_t}$ is the standard error of this forecast as estimated by a GARCH model. The GARCH specification assumes that e_t is distributed $N(0, h_t)$, so the probability of observing an increase in oil prices during quarter t as large or larger than what was actually observed is $1 - \Phi(e_t^*)$, where $\Phi(\cdot)$ is the cumulative distribution function for a standard Normal variate. Thus the statement that the effect of an oil price increase depends on e_t^* is equivalent to the statement that the effect of an oil price increase depends on how surprising that increase is, given the recent behavior of oil prices.

various alternatives would be to leave the function $f(o_t)$ in equation (1.1) totally unrestricted, and let the data tell us which of the various nonlinear alternatives is best supported by the data. This paper pursues that idea using a flexible approach to nonlinear modeling recently suggested by Hamilton (2000). The basic technique is described in the following section.

2 A flexible approach to nonlinear inference.

Consider a nonlinear regression model of the form

$$y_t = \mu(\mathbf{x}_t) + \boldsymbol{\delta}'\mathbf{z}_t + \varepsilon_t \quad (2.1)$$

where y_t is a scalar dependent variable, \mathbf{x}_t and \mathbf{z}_t are k - and p -dimensional vectors of explanatory variables, and ε_t is an error term. The form of the function $\mu(\cdot)$ is unknown, and we seek to represent it using a flexible class. On the other hand, there may be some subset of variables \mathbf{z}_t for which the researcher is willing to assume linearity, and, if so, significant efficiency gains can be obtained by imposing this restriction. In the application below, $\mathbf{z}_t = (y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4})'$ contains lagged values of GDP growth while $\mathbf{x}_t = (o_{t-1}, o_{t-2}, o_{t-3}, o_{t-4})'$ contains lagged changes in oil prices. The approach suggested by Hamilton (2000) is to view the function $\mu(\cdot)$ itself as the outcome of a random field.⁶ That is, if $\boldsymbol{\tau}_1$ denotes an arbitrary, nonstochastic k -dimensional vector, then the value of the function $\mu(\cdot)$ evaluated at $\boldsymbol{\tau}_1$, denoted $\mu(\boldsymbol{\tau}_1)$, is regarded as a random variable. Hamilton

⁶ The general approach of viewing $\mu(\cdot)$ as stochastic has a very old and extensive tradition in the statistics literature. The particular form for the random field used here is, to my knowledge, original— see Hamilton (2000) for references.

(2000) treats this random variable as being Normally distributed with mean $\alpha_0 + \boldsymbol{\alpha}'\boldsymbol{\tau}_1$ and variance λ^2 , where α_0 , $\boldsymbol{\alpha}$, and λ are population parameters to be estimated. Note that if $\lambda = 0$, then model (2.1) becomes a simple linear regression model $y_t = \alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \varepsilon_t$. The larger λ , the more the model (2.1) is allowed to deviate from a linear regression model.

The other item one needs to know about the random field $\mu(\cdot)$ is how the random variable $\mu(\boldsymbol{\tau}_1)$ is correlated with $\mu(\boldsymbol{\tau}_2)$, for $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ again arbitrary k -dimensional vectors. We assume that $\mu(\boldsymbol{\tau}_1)$ is uncorrelated with $\mu(\boldsymbol{\tau}_2)$ if $\boldsymbol{\tau}_1$ is sufficiently far away from $\boldsymbol{\tau}_2$, specifically, that

$$E\{[\mu(\boldsymbol{\tau}_1) - \alpha_0 - \boldsymbol{\alpha}'\boldsymbol{\tau}_1][\mu(\boldsymbol{\tau}_2) - \alpha_0 - \boldsymbol{\alpha}'\boldsymbol{\tau}_2]\} = 0$$

if $(1/2) \left[\sum_{i=1}^k g_i^2 (\tau_{i1} - \tau_{i2})^2 \right]^{1/2} > 1$, where τ_{i1} denotes the i th element of the vector $\boldsymbol{\tau}_1$ and g_1, g_2, \dots, g_k are k additional population parameters to be estimated. The closer that $\boldsymbol{\tau}_1$ gets to $\boldsymbol{\tau}_2$, specifically, the smaller the value of the scalar $h_{12} = (1/2) \left[\sum_{i=1}^k g_i^2 (\tau_{i1} - \tau_{i2})^2 \right]^{1/2}$, the higher the correlation between $\mu(\boldsymbol{\tau}_1)$ and $\mu(\boldsymbol{\tau}_2)$, with the correlation going to unity as h_{12} goes to zero. If the nonlinear part of the model includes $k = 4$ explanatory variables, then the correlation is assumed to be given by

$$\text{Corr}(\mu(\boldsymbol{\tau}_1), \mu(\boldsymbol{\tau}_2)) = \begin{cases} H_4(h_{12}) & \text{if } 0 \leq h_{12} \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

where

$$H_4(h_{12}) = 1 - (2/\pi) \left[(2/3)h_{12}(1 - h_{12}^2)^{3/2} + h_{12}(1 - h_{12}^2)^{1/2} + \sin^{-1}(h_{12}) \right]. \quad (2.2)$$

See Hamilton (2000) for the motivation behind this specification. Note that $H_k(\cdot)$ is not a parameterization of the functional relation $\mu(\cdot)$ itself, but rather a parameterization of the

correlation between the random variables $\mu(\boldsymbol{\tau}_1)$ and $\mu(\boldsymbol{\tau}_2)$; a given realization of $\mu(\cdot)$ from this random process can take on any of a variety of different forms, and this is what gives the approach its flexibility. The parameter g_i governs the likely variability of the nonlinear function $\mu(\boldsymbol{\tau})$ as the value of τ_i varies; as g_i becomes small, the value of $\mu(\boldsymbol{\tau})$ changes little when τ_i changes. If $g_i = 0$, then the function $\mu(\boldsymbol{\tau})$ is linear with respect to τ_i .

The above specification can be written in the form

$$\begin{aligned} y_t &= \alpha_0 + \boldsymbol{\alpha}' \mathbf{x}_t + \boldsymbol{\delta}' \mathbf{z}_t + \lambda m(\mathbf{x}_t) + \varepsilon_t \\ &= \alpha_0 + \boldsymbol{\alpha}' \mathbf{x}_t + \boldsymbol{\delta}' \mathbf{z}_t + u_t \end{aligned} \quad (2.3)$$

where $m(\cdot)$ denotes the realization of a scalar-valued Gaussian random field with mean zero, unit variance, and covariance function given by (2.2) and where $u_t = \lambda m(\mathbf{x}_t) + \varepsilon_t$. If the regression error ε_t is assumed to be i.i.d. $N(0, \sigma^2)$ and if the regressors $(\mathbf{x}'_t, \mathbf{z}'_t)$ are strictly exogenous, then this specification implies a GLS regression model of the form

$$\mathbf{y} | \mathbf{X} \sim N(\mathbf{X}\boldsymbol{\beta}, \mathbf{P}_0 + \sigma^2 \mathbf{I}_T)$$

where $\mathbf{y} = (y_1, y_2, \dots, y_T)'$, T is the sample size, \mathbf{X} is a $(T \times (1 + k + p))$ matrix whose t th row is given by $(1, \mathbf{x}'_t, \mathbf{z}'_t)$, $\boldsymbol{\beta}$ is the $(1 + k + p)$ -dimensional vector $(\alpha_0, \boldsymbol{\alpha}', \boldsymbol{\delta}')$, \mathbf{I}_T is the $(T \times T)$ identity matrix, and \mathbf{P}_0 is a $(T \times T)$ matrix whose row s , column t element is given by $\lambda^2 H_k(h_{st}) \delta_{[h_{st} < 1]}$ where $h_{st} = (1/2) \left[\sum_{i=1}^k g_i^2 (x_{is} - x_{it})^2 \right]^{1/2}$, x_{is} denotes the value of the i th explanatory variable for observation s , and the function $H_k(\cdot)$ is as specified in (2.2) when $k = 4$. Nonlinearity of the functional form $\mu(\cdot)$ implies a correlation between u_t and u_s , the residuals of the linear specification, whenever \mathbf{x}_t and \mathbf{x}_s are close together.

The population parameters of the model thus consist of the linear part of the regression function $(\alpha_0, \boldsymbol{\alpha}, \boldsymbol{\delta})$, the variance of the nonlinear regression error σ^2 , the parameter governing the overall importance of the nonlinear component λ^2 , and the parameters governing the variability of the nonlinear component with respect to each explanatory variable (g_1, g_2, \dots, g_k) . The nonlinear regression function $m(\cdot)$ itself does not involve any parameters but instead is regarded as a random outcome whose probability law is to be modeled along with the observed data as described above. Conditional on the parameters, the likelihood function is simply that of a GLS Gaussian regression, and numerical Bayesian methods described in Hamilton (2000) can be used to evaluate the posterior distribution of any statistics of interest; the specific priors used in this study are detailed in Appendix A. Conditional on the parameters, the optimal inference of the value of the unobserved function $\mu(\mathbf{x}^*)$ at some arbitrary point \mathbf{x}^* is given by

$$\hat{\mu}(\mathbf{x}^*) = \boldsymbol{\alpha}'\mathbf{x}^* + \mathbf{q}'(\mathbf{P}_0 + \sigma^2\mathbf{I}_T)^{-1}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \quad (2.4)$$

where \mathbf{q} denotes a $(T \times 1)$ vector whose t th element is given by $\lambda^2 H_k(h_t^*) \delta_{[h_t^* < 1]}$ for $h_t^* = (1/2) \left[\sum_{i=1}^k g_i^2 (x_{it} - x_i^*)^2 \right]^{1/2}$ where x_{it} denotes the i th element of \mathbf{x}_t and x_i^* denotes the i th element of \mathbf{x}^* . The inference thus modifies the linear estimate $\boldsymbol{\alpha}'\mathbf{x}^*$ by taking a linear combination of residuals u_t for those observations with \mathbf{x}_t close to \mathbf{x}^* . Hamilton shows that the inference $\hat{\mu}(\mathbf{x}^*)$ converges to the true value $\mu(\mathbf{x}^*)$ for $\mu(\cdot)$ any function from a broad class of continuous functions satisfying a certain smoothness condition. Monte Carlo investigation by Christian Dahl (1998) confirms that the procedure is useful in small samples for a variety of nonlinear time series models. One can calculate a 95% confidence interval for this inference

by generating values of α_0 , $\boldsymbol{\alpha}$, $\boldsymbol{\delta}$, σ , λ , and \mathbf{g} from their posterior distributions and calculating the inference (2.4) along with its known standard error for each given parameter vector, and examining the resulting distribution of inferences.

The framework also suggests a simple test of the null hypothesis that the true relation is linear ($H_0 : \lambda = 0$). Hamilton suggests fixing the smoothing parameters g_i for purposes of this test on the basis of the sample standard deviation of the i th explanatory variable,

$$g_i = 2 \left[k \left(T^{-1} \sum_{t=1}^T (x_{it} - \bar{x}_i)^2 \right) \right]^{-1/2}.$$

Using these values for g_i , construct the $(T \times T)$ matrix \mathbf{H} whose row t , column s element is given by

$$H_k \left\{ (1/2) \left[g_1^2 (x_{1t} - x_{1s})^2 + g_2^2 (x_{2t} - x_{2s})^2 + \dots + g_k^2 (x_{kt} - x_{ks})^2 \right]^{1/2} \right\} \quad (2.5)$$

where $H_k(\cdot)$ is given in expression (2.2) when $k = 4$, or by zero when the argument of $H_k(\cdot)$ exceeds unity. Next perform an OLS linear regression of y_t on \mathbf{x}_t , \mathbf{z}_t and a constant, $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$, calculating the OLS residuals $\hat{\boldsymbol{\varepsilon}}$, regression squared standard error, $\tilde{\sigma}^2 = (T - k - p - 1)^{-1} \hat{\boldsymbol{\varepsilon}}' \hat{\boldsymbol{\varepsilon}}$, and $(T \times T)$ projection matrix $\mathbf{M} = \mathbf{I}_T - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$. Then calculate the following function of the OLS residuals:

$$\nu^2 = \frac{[\hat{\boldsymbol{\varepsilon}}' \mathbf{H} \hat{\boldsymbol{\varepsilon}} - \tilde{\sigma}^2 \text{tr}(\mathbf{M}\mathbf{H}\mathbf{M})]^2}{\tilde{\sigma}^4 (2 \text{tr}\{[\mathbf{M}\mathbf{H}\mathbf{M} - (T - k - p - 1)^{-1} \mathbf{M} \text{tr}(\mathbf{M}\mathbf{H}\mathbf{M})]^2\})}. \quad (2.6)$$

If the OLS residuals for observation t are similar to those for other observations with similar \mathbf{x} 's, then ν^2 will be large and evidence against linearity is obtained. Hamilton shows that ν^2 has an asymptotic $\chi^2(1)$ distribution under the null hypothesis that the true relation is

linear. Dahl's (1998) Monte Carlo evidence establishes that the test has good small-sample size and strong power against a variety of nonlinear alternatives.

3 Empirical results.

The series used for real output y_t is the quarterly growth rate of chain-weighted real GDP.⁷ The oil price series o_t is 100 times the quarterly logarithmic growth rate of the nominal crude oil producer price index, seasonally unadjusted.⁸ The sample used for estimation (not including the lagged initial values for conditioning) runs from $t = 1949:II$ to $1999:IV$, for a total of $T = 203$ usable observations.

The test statistic ν^2 of the null hypothesis of linearity has a value of 41.51, which for a $\chi^2(1)$ variable implies overwhelming rejection of the null hypothesis that the relation between oil prices and GDP growth is linear. Bayesian posterior estimates and their standard errors⁹ for the flexible nonlinear alternative are as follows:

$$\begin{aligned}
 y_t = & \frac{0.58}{(0.16)} - \frac{0.0030}{(0.0070)} o_{t-1} - \frac{0.0042}{(0.0070)} o_{t-2} - \frac{0.0062}{(0.0070)} o_{t-3} - \frac{0.0174}{(0.0076)} o_{t-4} \\
 & + \frac{0.24}{(0.07)} y_{t-1} + \frac{0.11}{(0.07)} y_{t-2} - \frac{0.07}{(0.07)} y_{t-3} - \frac{0.14}{(0.07)} y_{t-4} \\
 & + 0.94 \left[\frac{0.35}{(0.13)} m\left(\frac{0.09}{(0.11)} o_{t-1}, \frac{0.09}{(0.12)} o_{t-2}, \frac{0.08}{(0.10)} o_{t-3}, \frac{0.08}{(0.09)} o_{t-4}\right) + v_t \right]
 \end{aligned} \tag{3.1}$$

where $v_t \sim N(0, 1)$ and $m(\cdot)$ denotes an unobserved realization from a Gaussian random

⁷ Data were downloaded from the Bureau of Economic Analysis web page (<http://www.bea.doc.gov/bea/dn1.htm>) on April 24, 2000.

⁸ The monthly WPI0561 series was converted to quarterly by using end-of-period values. Data from 1947:II to 1974:I are from Hamilton (1982). Data from 1974:II to 1999:IV are from Citibase, downloaded from <http://ssdc.ucsd.edu/citibase> on April 24, 2000.

⁹ Based on 20,000 draws from the importance sampling density described in Hamilton (2000).

field with mean zero, unit variance, and correlations given by (2.2). The innovation ε_t in (2.3) is written here as $\sigma = 0.94$ times v_t , and the parameter λ in (2.2) is written as σ times the parameter ζ , whose estimate is 0.35. Each of the four lags of oil price changes exerts an overall negative effect on output growth as indicated by the linear coefficients, though only the coefficient on o_{t-4} is statistically significant. And although the nonlinear component overall makes a highly significant contribution (as evidenced by the t -statistic for $\zeta = 0$ or the LM test), it is not clear from (3.1) which lags on oil prices are most important for this nonlinearity.

Given any particular values for the vector \mathbf{g} — for example, given the posterior means $\hat{\mathbf{g}} = (0.09, 0.09, 0.08, 0.08)'$ — one can use (2.2) and (2.5) to calculate the value of $H_4(\cdot)$ associated with any pair of observations on \mathbf{x}_t and \mathbf{x}_s . For a given value of λ — for example, $\hat{\lambda} = 0.33$ — one can then calculate the row t , column s element of the matrix \mathbf{P}_0 as $\lambda^2 H_4(\cdot)$. Given values for $\boldsymbol{\alpha}$ as well, one can calculate a value for (2.4) for any \mathbf{x}^* of interest, which represents the econometrician's inference as to the value of the conditional mean $\mu(\mathbf{x}^*)$ when the explanatory variables take on the value represented by \mathbf{x}^* and when the parameters are known to take on these specified values. By generating values of \mathbf{g} and other parameters from the posterior distribution whose mean and standard deviation are reported in (3.1), we generate a range of estimates of $\mu(\mathbf{x}^*)$, and the mean of this range then represents the econometrician's posterior inference as to the value of $\mu(\mathbf{x}^*)$.

As a first step for seeing what the nonlinear function $\mu(\cdot)$ looks like, I fixed the values of o_{t-2} , o_{t-3} , and o_{t-4} equal to their sample means and examined the consequences of changing

o_{t-1} alone, that is, I set $\mathbf{x}^* = (x_1, \bar{o}, \bar{o}, \bar{o})$ and evaluated the Bayesian posterior expectation of (2.4) for various values of x_1 . Figure 3 plots the result as a function of x_1 along with 95% confidence intervals. The intervals are narrowest for values of x_1 closest to the sample mean, for these represent the values of \mathbf{x}^* for which we can have the greatest confidence about the inference. The implied function is nonlinear, suggesting that if oil prices either increase or decrease after three quarters of stability, the forecast calls for slightly slower GDP growth than if oil prices had remained stable, though increases are worse news than decreases.

Figure 4 answers the analogous question, fixing o_{t-1} , o_{t-3} , and o_{t-4} equal to their sample means and varying the value of o_{t-2} . Both it and Figure 5 (the effect of o_{t-3} in isolation) give a similar impression to that of Figure 3. Figure 6 (the effect of o_{t-4} alone) is somewhat more dramatic, suggesting that decreases in oil prices four quarters earlier have essentially no consequences for current GDP growth, whereas oil price increases significantly reduce expected GDP growth. The figure indicates that Mork's (1989) asymmetric specification (1.2) is exactly the form suggested by the data, at least as far as describing the consequences of changing o_{t-4} alone are concerned.

To get an impression about interactive effects, I calculated how the apparent consequences of o_{t-3} are affected by different values of o_{t-4} . Figure 7 compares the three functions $\hat{\mu}(\bar{o}, \bar{o}, x_3, 0)$, $\hat{\mu}(\bar{o}, \bar{o}, x_3, 5)$, and $\hat{\mu}(\bar{o}, \bar{o}, x_3, -5)$, plotted as a function of x_3 . The first relation is represented by the solid line, which is essentially the same as the mean value plotted in Figure 5. The second relation shows how the effect of an x_3 percent oil price increase 3 quarters ago would be different if oil prices had also increased 5% the quarter before that.

This is represented by the lower, short-dashed line in Figure 7. The line is uniformly lower than the solid line— an oil price increase 4 quarters earlier definitely causes one to lower the forecast for GDP growth, regardless of the value of o_{t-3} . Even so, if one compares the solid and short-dashed line at any given x_3 , the slope of the short-dashed line is less steep than the solid line— an oil price increase 4 quarters earlier reduces the additional information content of any change, up or down, in quarter $t - 3$. The top, long-dashed line plots the predicted GDP growth for quarter t when $o_{t-3} = x_3$ and $o_{t-4} = -5$. If oil prices went down 4 quarters earlier, this has little consequences for forecasting GDP if it was followed by $o_{t-3} = 0$. If a 5% decrease was followed by another change, either up or down, one should downweight the otherwise contractionary signal implied by the oil price change in period $t - 3$. Overall, Figure 7 supports the view of Lee, Ni, and Ratti (1995) and Hamilton (1996) that previous turbulence in oil prices causes the marginal effect of any given oil price change to be reduced.

Another way to make this point is to look at how the three lines in Figure 7 would be predicted to appear under alternative parametric nonlinear models. The upper left panel of Figure 8 shows the values for $\tilde{\mu}(\bar{o}, \bar{o}, x_3, 0)$, $\tilde{\mu}(\bar{o}, \bar{o}, x_3, 5)$, and $\tilde{\mu}(\bar{o}, \bar{o}, x_3, -5)$ for $\tilde{\mu}(\mathbf{x}^*) = \tilde{\alpha}_0 + \tilde{\alpha}'\mathbf{x}^*$ the fitted values from the simple linear regression (1.6). Each plot is a straight line with slope -0.007, and changing o_{t-4} by ± 5 induces a parallel shift of the line by ∓ 0.090 . The two key ways in which this figure differs from what appears in the flexible nonlinear summary of the data (Figure 7) are that, according to the flexible inference, negative values of o_{t-3} imply lower GDP growth rather than higher values as required by the linear specification, and the flexible inference suggests that negative values of o_{t-4} have a much smaller effect

than do positive values of o_{t-4} .

The upper right panel of Figure 8 plots the analogous three relations for Mork's specification (1.8). The solid line plots $\mu^+(\bar{o}, \bar{o}, x_3, 0)$, which is a horizontal line for negative values of x_3 and a line with slope -0.014 for positive values of x_3 . The plot of $\mu^+(\bar{o}, \bar{o}, x_3, -5)$ is exactly the same relation, since under Mork's specification, negative values for o_{t-4} are completely irrelevant. The plot of $\mu^+(\bar{o}, \bar{o}, x_3, 5)$ is a vertical downward shift of these curves by 0.15. The key difference between this panel of Figure 8 and what appears to be in the data in Figure 7 is that, contrary to Mork's specification, an oil price decrease in $t - 4$ appears to mitigate somewhat the effects of an oil price increase in $t - 3$.

Hamilton (1996) argued that an oil price increase of 10% that comes immediately after an oil price decrease of 20% would do little to alarm consumers or deter them from purchasing gas-guzzling vehicles. He suggested looking at the net amount by which oil prices have gone up over the past year as a better measure than the amount by which oil prices go up in any given quarter. His measure of the net oil price increase, o_t^\dagger , is defined as the amount by which oil prices in quarter t exceed their peak value over the previous 12 months; if they do not exceed the previous peak, then o_t^\dagger is taken to be zero. OLS estimates of a relation based on this measure are as follows:

$$\begin{aligned}
y_t = & \frac{0.99}{(0.13)} + \frac{0.23}{(0.07)} y_{t-1} + \frac{0.09}{(0.07)} y_{t-2} - \frac{0.08}{(0.07)} y_{t-3} - \frac{0.14}{(0.07)} y_{t-4} \\
& - \frac{0.013}{(0.012)} o_{t-1}^\dagger - \frac{0.013}{(0.013)} o_{t-2}^\dagger - \frac{0.022}{(0.014)} o_{t-3}^\dagger - \frac{0.045}{(0.014)} o_{t-4}^\dagger.
\end{aligned} \tag{3.2}$$

Estimates of $\mu^\dagger(\bar{o}, \bar{o}, x_3, 0)$, $\mu^\dagger(\bar{o}, \bar{o}, x_3, -5)$, and $\mu^\dagger(\bar{o}, \bar{o}, x_3, 5)$ based on this OLS regression

are plotted in the lower left panel of Figure 8 under the assumption that oil prices had been steady prior to period $t - 4$. The first relation (solid line) is a horizontal line for negative values of o_{t-3} and a line with slope -0.022 for positive values of o_{t-3} . If instead $o_{t-4} = -5$ (dashed line), oil price increases only matter in $t - 3$ to the extent they exceed 5%, causing the horizontal line to be extended 5% before turning down. By contrast, if $o_{t-4} = 5$ (alternate-dashed line), the original relation is everywhere shifted down by 0.225. These features are roughly consistent with what appears in the flexibly estimated nonlinear relation (Figure 7).

Lee, Ni, and Ratti (1995) pursued a related idea based on a GARCH representation of oil prices. Let o_t^R denote the real change in oil prices ($o_t^R = o_t - \Delta \ln(p_t)$ for p_t the GDP deflator for quarter t). Lee, Ni and Ratti reported the following GARCH parameter estimates¹⁰

$$o_t^R = -0.4965 + 0.436 o_{t-1}^R - 0.401 o_{t-2}^R + 0.244 o_{t-3}^R - 0.238 o_{t-4}^R + e_t \quad (3.3)$$

$$e_t = \sqrt{h_t} \nu_t \quad \nu_t \sim N(0, 1) \quad (3.4)$$

$$h_t = 1.49 + 2.208 e_{t-1}^2 + 0.197 h_{t-1}. \quad (3.5)$$

This highly explosive GARCH process has root 2.405 much greater than one, with oil price volatility leading to rapid increases in the variance and periods of calm bringing it quickly back down within sample. The average value for e_t^2 when calculated from (3.3) in my sample is 122, and the average value for $\sqrt{h_t}$ when generated from equation (3.5) starting

¹⁰ Figures from Table 3 in Lee, Ni, and Ratti have been converted from the annual rates used by these authors for comparability with the quarterly rates used throughout this paper.

from $h_0 = 100$ is 10.54. Following Lee, Ni and Ratti,¹¹ I defined the volatility-adjusted real oil price increase o_t^\ddagger to be $o_t^R/\sqrt{h_t}$ when the latter is positive and zero otherwise. I used h_t generated from (3.3)-(3.5) to construct o_t^\ddagger and obtained the following OLS estimates:

$$\begin{aligned}
y_t = & \underset{(0.13)}{1.05} + \underset{(0.07)}{0.20} y_{t-1} + \underset{(0.07)}{0.11} y_{t-2} - \underset{(0.07)}{0.09} y_{t-3} - \underset{(0.07)}{0.12} y_{t-4} \\
& - \underset{(0.11)}{0.19} o_{t-1}^\ddagger - \underset{(0.11)}{0.06} o_{t-2}^\ddagger - \underset{(0.11)}{0.35} o_{t-3}^\ddagger - \underset{(0.11)}{0.44} o_{t-4}^\ddagger.
\end{aligned} \tag{3.6}$$

The lower right panel of Figure 8 was then constructed as follows. Suppose that h_{t-4} was equal to 122 (the average value for e_t^2 in sample) and that, if o_{t-4} had equalled its average value in sample (0.4), then e_{t-4} would have been zero. Then to calculate the predicted consequences of changing o_{t-3} when o_{t-4} was not the sample average of 0.4 but instead was equal to -5, the implied value for h_{t-3} would be

$$h_{t-3} = 1.49 + (2.208)(5.4)^2 + (0.197)(122) = 89.90$$

whose square root is 9.48. Positive values for x_3 thus imply a value for o_{t-3}^\ddagger of $x_3/9.48$ when $o_{t-4} = -5$. The plot of $\mu^\ddagger(\bar{o}, \bar{o}, x_3, -5)$ is thus a horizontal line for negative values of x_3 and a line with slope $-0.35/9.48$ for positive values of x_3 (the short-dashed line in the lower right panel of Figure 8).

If instead the oil price change in period $t - 3$ followed a quarter of constant rather than falling oil prices, we have

$$h_{t-3} = 1.49 + (2.208)(0.4)^2 + (0.197)(122) = 25.88$$

¹¹ These authors used $e_t/\sqrt{h_t}$ rather than $o_t^R/\sqrt{h_t}$ as their measure. Results should be similar, and o_t^R is used here for compatibility with the other results in this paper.

for which $\mu^\ddagger(\bar{o}, \bar{o}, x_3, 0)$ is again a horizontal line for negative values of x_3 but now the solid line with the steeper slope of $-0.35/5.09$ for positive values. Finally, if $o_{t-4} = 5$, the relation is everywhere shifted down by $-0.44 \times 5 \div \sqrt{122} = -0.20$, but the marginal effect of o_{t-3} is now smaller owing to a higher value for h_{t-3} caused by the oil price increase in $t - 4$. The flatter slope implies that the graphs of $\mu^\ddagger(\bar{o}, \bar{o}, x_3, 0)$ and $\mu^\ddagger(\bar{o}, \bar{o}, x_3, 5)$ eventually cross for sufficiently high values of x_3 . One does not see such crossing in Figure 7, though the ability of the flexible inference to detect such a feature is likely to be quite weak, and the broad pattern of the bottom right panel of Figure 8 is otherwise quite consistent with that in Figure 7.

There is a more formal statistical basis for comparing the nonlinear dynamics implied by alternative specifications with what appears in the data from the flexible inference procedure used here. Note that Mork's specification (1.8), Hamilton's specification (3.2), and Lee, Ni, and Ratti's specification (3.6), although nonlinear functions of oil prices, are linear functions of the parameters, that is, they all can be described as a linear regression model of the form¹²

$$y_t = \alpha_0 + \boldsymbol{\delta}' \mathbf{z}_t + \varepsilon_t \quad (3.7)$$

for a suitable specification of \mathbf{z}_t . For example, equation (3.2) is a special case of (3.7) with $\mathbf{z}_t = (y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4}, o_{t-1}^\dagger, o_{t-2}^\dagger, o_{t-3}^\dagger, o_{t-4}^\dagger)'$. As such, one can test directly whether such a specification for \mathbf{z}_t adequately captures any nonlinearity that appears in the data by

¹² This discussion ignores the fact that o_t^\ddagger is a generated regressor. Mitigating any concerns about this is the fact that the parameters used to construct o_t^\ddagger , although estimated, were estimated using a different sample from that used here. In any case, no such concerns apply to interpreting the statistical results for o_t^+ and o_t^\dagger .

comparing (3.7) with the more general model

$$y_t = \alpha_0 + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t) + \varepsilon_t$$

for $\mathbf{x}_t = (o_{t-1}, o_{t-2}, o_{t-3}, o_{t-4})$ and $m(\cdot)$ a realization of the random field whose correlations are characterized by (2.2). A test of the null hypothesis $\lambda = 0$ is now a test of whether the definition of \mathbf{z}_t adequately captures the nonlinear dependence of y_t on o_{t-j} . This is simply a special case of the test already described in (2.6). By changing the definition of \mathbf{z}_t , the test has been adapted from testing the null hypothesis of linearity to testing the null hypothesis that the nonlinearity takes on a particular known form.

The results of this specification test are $\nu^2 = 0.31$ for Mork's formulation (1.8), $\nu^2 = 0.63$ for Hamilton's formulation (3.2), and $\nu^2 = 0.66$ for Lee, Ni and Ratti's formulation (3.6). Given a 5% critical value for a $\chi^2(1)$ test of 3.84, all three formulations are accepted by this test. Thus, although the general approach employed here has the advantage that it is flexible enough to estimate a broad class of nonlinear functions consistently, it does not have enough power to distinguish statistically between the three specific nonlinear alternatives examined here. We can say categorically that GDP does not depend linearly on lagged oil prices, on the basis of the overwhelming rejection of (1.6) by the ν^2 test reported earlier. We can further say qualitatively that the nature of the nonlinearity is an asymmetry in which oil price increases matter substantially more than oil price decreases and in which the sensitivity of the economy to oil price changes seems to be dampened by previous volatility of oil prices. However, choosing between alternative nonlinear formulations is not something one can do with any great confidence based on the statistics reported here.

As a final check, consider the relation between the oil price transformations and the apparent instability of the forecasting equation. I regressed o_t on a constant and four of its own lags and tested the null hypothesis that all five coefficients changed at some date t_0 . The top panel of Figure 9 plots the p -value for this F test for every possible t_0 between 1959:II and 1989:IV as a function of t_0 . This p -value stays near 0.05 throughout the 1980's, though this is not statistically significant evidence based on Andrews's (1993) test that there has been a structural change in the process generating oil prices. If the relation between oil prices and GDP is nonlinear, this would show up as a structural change in a linear regression of GDP on lagged oil prices. The second panel plots the p -value for the null hypothesis that the four coefficients on lagged oil prices in equation (1.6) changed at date t_0 , and this relation indeed appears to break down in the 1980's, though the minimum value for this p -value is still not statistically significant according to Andrews's test.¹³ Stronger evidence for instability appears when additional lags and explanatory variables are included in regression (1.6) (Hooker, 1996, 1997).

The bottom three panels of Figure 9 plot the corresponding tests for stability of the coefficients on o_t^+ , o_t^\dagger , and o_t^\ddagger in equations (1.8), (3.2), and (3.6), respectively. Any of the three specifications eliminate all statistically significant evidence of instability, although Hooker (1997) reports some evidence of instability in larger vector systems. We leave the broader question of the stability of the connection between oil prices, monetary policy, and economic activity as a topic for future research, but conclude that the question of the

¹³ The largest χ^2 statistic is 13.61, compared with Andrews's 5% critical value of 15.84.

functional form appropriate for (2.1) has been adequately addressed here.

4 A Linear Instrumental Variable Interpretation.

We framed the investigation in Section 3 in the form of a pure question about forecasting, to wit, What is the correct functional form for describing the conditional expectation of GDP growth y_t conditional on lagged GDP growth and lagged oil price changes? The answer to this forecasting question could be given a causal interpretation if one were persuaded that oil price changes are exogenous.

Hamilton (1983, 1985) argued that, over the period 1948-1972, oil prices were indeed exogenous with respect to the U.S. economy, on the basis of institutional, historical, and statistical evidence. Institutionally, there was a particular reason why endogenous factors had no effect on oil prices during this period. U.S. oil producing states had commissions that actively regulated the quantity of oil that could be produced by each field, the most important of which was the Texas Railroad Commission (TRC). Whenever demand went up, the TRC would increase the amount of production it allowed, and when demand went down, the TRC would decrease the amount of production, thus preventing demand changes from causing any change in price. The only events that did change the price were exogenous disruptions to supply. Historically, one can unambiguously identify the events behind major oil price movements in this period, and they are clearly political and military developments that have little to do with the U.S. macroeconomy. Finally, in terms of statistical evidence, Granger-causality tests uncover no U.S. macroeconomic variables that could have predicted

oil price changes over this period, bolstering the case for the claim of exogeneity.

None of these arguments applies to post-1973 data, however. The Texas Railroad Commission ceased to be relevant once the Middle East oil producers became the dominant factor in the world petroleum market. Today oil prices respond quite dramatically to demand conditions with constant adjustment and readjustment. Statistically, oil prices certainly are predictable from U.S. macroeconomic developments in post-1973 data (Barsky and Kilian, 1998).

If one is interested in a causal interpretation rather than a simple forecasting equation, it would be useful to isolate the component of the post-1973 oil price movements that could be attributed to strictly exogenous events. To do so, I have developed a quantitative version of the dummy-variable approach used by Dotsey and Reid (1992) and Hoover and Perez (1994). There are a number of historical episodes in which military conflicts produced dramatic and unambiguous effects on the petroleum production from particular fields. For example, in July 1990, Iraq and Kuwait had been producing 5.3 million barrels of oil daily. After Iraq invaded Kuwait, production from these two countries stopped altogether, representing an 8.8% drop in world petroleum production that is straightforward to measure and can unambiguously be attributed entirely to this exogenous military event. Appendix B discusses four other Middle East military conflicts, with the consequences for world petroleum supply of each episode summarized in Table 1.

I then constructed a series Q_t which is defined as the magnitude of the production shortfall identified in Table 1 if an indicated episode began in quarter t and is zero otherwise, and

claim that Q_t is a valid exogenous instrument for disturbances to world petroleum supplies. Note that in each case the measure Q_t is based on how much petroleum production fell in the affected countries only. In each episode, petroleum production increased somewhere else in the world to make up in part for the lost production in the affected regions. However, it seems clear that the latter represents an endogenous response to the crisis rather than a separate exogenous event. If one wants a measure of the magnitude of the exogenous shock itself (which set off both an endogenous supply response elsewhere as well as any possible macroeconomic consequences), it is clear that Q_t is the correct measure to use in any reduced-form evaluation. Note, however, that the statistical consequences of a given shock of $Q_t = \Delta$ measure the effect of the usual historical case in which a magnitude Δ shock to petroleum production in one particular region is partly cushioned by increased production in some other part of the world, as opposed to finding the answer to the hypothetical (and counterfactual) question, What would be the effects if production were exogenously decreased by the amount Δ in some region and there was no possibility of making up this lost production elsewhere?

Although we have a clear measure of the timing and magnitude of the initial supply disruption, the endogenous response of supply increases following these shocks makes it difficult to measure how long the crisis persisted. For this reason, we simply use the size and date of the initial shock as the exogenous explanatory variable and allow unrestricted lags of this variable to summarize the cumulative economic consequences of this initial shock over time.

The suggestion is then that the fitted values of a regression of o_t on a constant and

$Q_t, Q_{t-1}, \dots, Q_{t-4}$ represent a component of oil price changes that can unambiguously be attributed to exogenous events. These fitted values \hat{o}_t are plotted in the bottom panel of Figure 10, along with the actual oil price changes o_t (top panel). Oil prices have been extremely volatile since 1986, with most increases coming shortly after even bigger decreases. The net oil price series o_t^\dagger (second panel) nets out most of these changes, leaving a series that is in fact similar in appearance to the fitted values \hat{o}_t . Thus the nonlinear transformation that produced o_t^\dagger from current and lagged values of o_t seems to be doing something rather similar to attempting to isolate the exogenous component of oil price changes by linearly regressing o_t on current and lagged values of Q_t .

If one thought that the true causal relation between y_t and o_t were linear, but that much of the historical movement in oil prices was caused by endogenous factors, the correct approach would be to estimate a relation of the form of (1.6) using as instruments $Q_{t-1}, Q_{t-2}, \dots, Q_{t-8}, y_{t-1}, \dots, y_{t-4}$, and a constant. The result of this instrumental variable estimation is

$$\begin{aligned}
y_t = & 0.93 + 0.20 y_{t-1} + 0.09 y_{t-2} - 0.05 y_{t-3} - 0.16 y_{t-4} \\
& \quad (0.16) \quad (0.09) \quad (0.09) \quad (0.09) \quad (0.09) \\
& - 0.025 o_{t-1} - 0.051 o_{t-2} - 0.011 o_{t-3} - 0.065 o_{t-4}. \tag{4.1} \\
& \quad (0.028) \quad (0.025) \quad (0.025) \quad (0.029)
\end{aligned}$$

The estimates in (4.1) are remarkably similar to those estimated by OLS for the pre-1980 data used in (1.5), over which period I have argued that essentially all oil price changes were exogenous, and likewise quite similar to estimates based on the nonlinear net oil price series o_t^\dagger in (3.2). Figure 11 plots the implied consequences for y_{t+j} if oil prices were to increase by 10% at date t and remain at this new level, as predicted by (3.2), the nonlinear relation

based on o_t^\dagger (bottom panel), and (4.1), the linear instrumental variables estimates. The predicted effects are quite similar in terms of both magnitude and dynamics.

One can test formally whether the nonlinear transformations of oil prices add anything beyond that contained in the exogenous military shock component isolated in the bottom panel of Figure 10 by looking at the reduced form underlying (4.1). Specifically, I regressed y_t on a constant, four lags of y_{t-j} , and eight lags of Q_{t-j} . I then added four lags of oil price increases (o_t^+) and accepted the null hypothesis that these last coefficients were all zero (p -value = 0.42). Likewise, four lags of o_t^\dagger add nothing to the reduced-form regression (p -value = 0.43). On the other hand, one rejects that the coefficients on the Lee, Ni and Ratti transformation (o_t^\ddagger) are all zero (p -value = 0.013). The results suggest that the predictive power of the nonlinear transformations o_t^+ and o_t^\dagger could be attributed to their ability to filter out influences on oil prices that do not come from these five particular military conflicts. The same can not be said of o_t^\ddagger .

Finally, the instrumental variable estimate seems to have solved at least some of the structural instability of the linear estimates. Figure 12 plots the p -value for a test of the null hypothesis that the coefficients on o_{t-1}, \dots, o_{t-4} in (4.1) changed at some date t_0 as a function of t_0 . The hypothesis that the Suez coefficients are different from the others would be rejected at the 5% level, though not using Andrews's critical values given the cumulative nature of the test. Moreover, there appears to be no problem with stability over the other four episodes.

Comparing the second and fourth panel of Figure 10, the only major shock to the net oil

price increase series that was not associated with a large exogenous supply disruption is the most recent episode. The Asian financial crisis was associated with a drop in world oil prices of over 50% during 1997 and 1998. The Asian rebound and a concerted OPEC effort to roll back production brought prices back up during 1999. The price increases from 1999:II on set new annual highs, with the result that both the net oil price increase specification (3.2) and the Lee, Ni, and Ratti specification (3.6) would predict U.S. GDP growth under 1% for the first quarter of 2000, slipping into recession by the second quarter.

On the other hand, if one extends the historical filter slightly farther, by the end of 1999, oil prices had only reached back up to the values seen at the start of 1997 before the Asian crisis. Hence, if one asks whether a shock does anything other than reverse the trend of the the previous 36 months, as in Panel 3 of Figure 10 (instead of the previous 12 months, as used in Panel 2), the interpretation would be that there was no oil shock during 1999. This alternative nonlinear transformation, like o_t^\dagger and o_t^\ddagger , exhibits a highly statistically significant correlation with GDP growth that passes the Andrews test for stability over time and could be judged to have captured statistically the nonlinearity between oil prices and GDP as determined by the ν^2 test statistic in (2.5) of 1.19.

If the U.S. economy avoids a recession during 2000, the indicated conclusion would seem to be that the essential correlation is between GDP growth and the bottom panel of Figure 10, that is, between GDP growth and the events summarized in Table 1. If one views the bottom panel as a nonlinear transformation of the top panel, there are many ways to approximately do this, of which o_t^\dagger , o_t^\ddagger , or the third panel of Figure 10 represent three

possibilities. It may prove very difficult to come up with compelling statistical evidence for distinguishing among these.

On the other hand, the fact that there is a plethora of candidates is not an excuse to dismiss them all. The null hypothesis of no nonlinear relation was rejected categorically on the basis of the ν^2 test, and the hypothesis that the coefficients on oil in the linear IV regression (4.1) are all zero is rejected at the 1% level. Clearly there is something here. Just as clearly, however, it is not a simple dependence of the economy on the level of oil prices, as would be suggested by production-function-based accounts such as Rasche and Tatom (1977, 1981), Kim and Loungani (1992), Rotemberg and Woodford (1996), or Mork and Hall (1980). Instead, the key events appear to be triggered by a disruption in supplies or surprise in pricing, as suggested by dislocation-based interpretations such as Loungani (1986), Davis (1987a,b), or Hamilton (1988a).

What guide does this analysis suggest for applied research? Use of oil prices themselves as an exogenous instrument or disturbance is certainly called into question. In its place, the simplest and most robust alternative might be the Dotsey and Reid (1992) or Hoover and Perez (1994) dummy variables for exogenous oil supply shocks, or the refinement suggested here of using current and lagged values of Q_t in place of oil prices themselves, that is, estimating the reduced-form regression on the basis of which instrumental-variables estimates such as (4.1) would be constructed.

5 Conclusion

There is no question that a linear regression of output on lagged oil prices exhibits instability over time. Some researchers have attributed this to the fact that the true relation is nonlinear. Others have questioned whether a specification hunt over the class of all possible nonlinear relations has simply produced a particular nonlinear relation that spuriously appears to be significant and stable.

This paper approached this question using a framework that explicitly parameterizes the set of nonlinear relations investigated and takes into account the uncertainty about functional form in conducting hypothesis tests. The evidence appears quite strong that one should use a nonlinear function of oil price changes if the goal is to forecast GDP growth. When one looks at this nonlinear relation from a flexible, unrestricted framework, the functional form looks very much like what has been suggested in earlier parametric studies. In particular, it is quite clear from the data that oil price increases are much more important for predicting GDP than are decreases, and that oil price changes are less useful for forecasting if they follow a period of earlier volatile price changes, as suggested by earlier researchers. However, there is not enough information in the historical experience to choose one particular functional form unambiguously over another, though any of the nonlinear representations that have been suggested in the literature is clearly superior to a linear specification.

The paper further suggested a linear instrumental variable interpretation of at least part of this phenomenon. One can clearly identify five military conflicts in the Middle East that have significantly disrupted world petroleum supplies. If the magnitudes of these disruptions

are used as an instrument for oil price changes, the predictions of a linear IV regression are very similar to those of the nonlinear specifications. It thus appears that part of the success of the nonlinear specifications is that they filter out many of the endogenous factors that have historically contributed to changes in oil prices.

Given the abundant evidence that exogenous disruptions in petroleum supplies lead one to predict lower GDP, the natural conclusion is that the disruptions are an important factor in causing the economic downturns. The results presented here make it difficult to attribute this effect to movement along an aggregate production function. The evidence instead suggests that oil shocks matter because they disrupt certain categories of spending by consumers and firms.

Appendix A

The Bayesian priors used in this study consist of (i) a gamma prior for σ^{-2} ,

$$p(\sigma^{-2}) = \frac{\xi^\nu}{\Gamma(\nu)} \sigma^{-2(\nu-1)} \exp[-\xi \sigma^{-2}],$$

with $\nu = 0.25$ and $\xi = (\nu s_y^2/2)$ for s_y^2 the sample variance of y ; (ii) a Gaussian distribution for $\boldsymbol{\beta}$ conditional on σ^{-2} ,

$$p(\boldsymbol{\beta}|\sigma^{-2}) = \frac{1}{(2\pi\sigma^2)^{(p+k+1)/2}} |\mathbf{M}|^{-1/2} \exp\left[\left(\frac{-1}{2\sigma^2}\right) (\boldsymbol{\beta} - \mathbf{m})' \mathbf{M}^{-1} (\boldsymbol{\beta} - \mathbf{m})\right],$$

where the first element of \mathbf{m} is the sample mean of y_t and all other elements of \mathbf{m} are zero, and where $\mathbf{M} = T(\mathbf{X}'\mathbf{X})^{-1}$, so that the prior has the weight of a single observation on (y_t, \mathbf{x}'_t) ; and a lognormal prior for each element of (\mathbf{g}', ζ) , where $\zeta = \lambda/\sigma$:

$$p(\mathbf{g}, \zeta) = \frac{1}{\sqrt{2\pi}} \exp\left\{-[\ln(\zeta)]^2/2\right\} \prod_{i=1}^k \frac{1}{\sqrt{2\pi\tau_i g_i}} \exp\left[\frac{-[\ln(g_i) - \vartheta_i]^2}{2\tau_i^2}\right],$$

where for $i = 1, \dots, k$ we set $\tau_i = 1$ and $\vartheta_i = -\ln\left(\sqrt{k s_i^2}\right)$ for $s_i^2 = T^{-1} \sum_{t=1}^T (x_{it} - \bar{x}_i)^2$ and \bar{x}_i the sample mean of the i th explanatory variable.

Appendix B

This appendix details the calculations behind Table 1.

1956. On October 29, 1956, Israeli troops invaded Egypt, followed by French and British. In the ensuing crisis, oil tankers were prevented from using the Suez Canal. The major pipeline that carried oil from Iraq through Syria was sabotaged, and exports of Middle East oil to Britain and France were blockaded.¹⁴ Overall, Middle East production fell by 1.7 million barrels per day (1.7 mbd) between October and November of 1956, or 10.1% of total world crude production of 16.8 mbd.¹⁵

1973. War between Israel and its neighbors again broke out on October 17, 1973, causing the Arab members of OPEC to announce an embargo of oil shipments to countries perceived to be sympathetic to Israel. Production of oil in these nations fell from 19.865 mbd in September 1973 to 15.528 mbd in November, or a loss of 7.8% of 1973 world production of 55.679 mbd.¹⁶

1978. The Iranian revolution in October 1978 led to a drop of Iranian production from 6.093 mbd in September 1978 to 0.729 mbd by January 1979, or 8.9% of 1978 world production of 60.158 mbd.

¹⁴ *Oil and Gas Journal*, November 12, 1956, pp. 122-125.

¹⁵ *Oil and Gas Journal*, April 1, 1957, p. 96.

¹⁶ Data for 1973-1990 are from *Historical Monthly Energy Review, 1973-1992*, U.S. energy Information Administration, U.S. Government Printing Office, 1994.

1980. The onset of war between Iran and Iraq led to a fall in Iraqi production from 3.240 mbd in July 1980 to 0.143 mbd by October 1980, while Iranian production fell from 1.699 mbd in July to 0.510 in October. The combined drop represents 7.2% of 1980 world oil production of 59.599 mbd.

1990. On August 1, 1990, Iraq invaded Kuwait. Kuwaiti production had been 1.858 mbd in July 1990 while Iraq had been producing 3.454 mbd. Production effectively ceased in both countries, representing an 8.8% drop from 1990 world production of 60.471 mbd.

Table 1

Exogenous disruptions in world petroleum supply.

Date	Event	Drop in world production
Nov. 1956	Suez Crisis	10.1%
Nov. 1973	Arab-Israel War	7.8%
Dec. 1978	Iranian Revolution	8.9%
Oct. 1980	Iran-Iraq War	7.2%
Aug. 1990	Persian Gulf War	8.8%

Source: See Appendix B.

References

- Andrews, Donald W. K. (1993), "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 61, pp. 821-856.
- Atkeson, Andrew, and Patrick J. Kehoe (1999), "Models of Energy Use: Putty-Putty Versus Putty-Clay," *American Economic Review*, 89, 1028-1043.
- Balke, Nathan S., Stephen P. A. Brown, and Mine Yücel (1999). "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" Working Paper, Federal Reserve Bank of Dallas.
- Barsky, Robert B., and Lutz Kilian (1998), "Money, Stagflation, and Oil Prices: A Re-Interpretation," Mimeo, University of Michigan.
- Bernanke, Ben S. (1983), "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, 98, pp. 85-106.
- Bohi, Douglas R. (1989), *Energy Price Shocks and Macroeconomic Performance*, Washington D.C.: Resources for the Future.
- Bresnahan, Timothy F., and Valerie A. Ramey (1993), "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry," *American Economic Review Papers and Proceedings*, 83, no. 2, pp. 213-218.
- Burbidge, John, and Alan Harrison (1984), "Testing for the Effects of Oil-Price Rises Using Vector Autoregressions," *International Economic Review*, 25, pp. 459-484.
- Carruth, Alan A., Mark A. Hooker, and Andrew J. Oswald (1998), "Unemployment Equilibria and Input Prices: Theory and Evidence from the United States," *Review of Economics and Statistics*, 80, pp. 621-628.
- Dahl, Christian (1998), "An Investigation of Tests for Linearity and the Accuracy of Flexible Nonlinear Inference," working paper, University of Aarhus.
- Daniel, Betty C. (1997), "International Interdependence of National Growth Rates: A Structural Trends Analysis," *Journal of Monetary Economics*, 40, pp. 73-96.
- Davis, Steven J. (1987a), "Fluctuations in the Pace of Labor Reallocation," in K. Brunner and A. H. Meltzer, eds., *Empirical Studies of Velocity, Real Exchange Rates, Unemployment and Productivity*, Carnegie-Rochester Conference Series on Public Policy, 24, Amsterdam: North Holland.
- (1987b), "Allocative Disturbances and Specific Capital in Real Business Cycle theories," *American Economic Review Papers and Proceedings*, 77, no. 2, pp. 738-751.
- , and John Haltiwanger (1997), "Sectoral Job Creation and Destruction Responses to Oil Price Changes and Other Shocks," working paper, University of Chicago.
- , Prakash Loungani, and Ramamohan Mahidhara (1996), "Regional Labor Fluctuations: Oil Shocks, Military Spending, and Other Driving Forces," working paper, University of Chicago.

- Dotsey, Michael, and Max Reid (1992), "Oil Shocks, Monetary Policy, and Economic Activity," *Economic Review of the Federal Reserve Bank of Richmond*, 78/4, pp. 14-27.
- Engle, Robert F. (1982), "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, pp. 987-1007.
- Ferderer, J. Peter (1996), "Oil Price Volatility and the Macroeconomy: A Solution to the Asymmetry Puzzle," *Journal of Macroeconomics*, 18 (1996), pp. 1-16.
- Finn, Mary G. (1997), "Perfect Competition and the Effects of Energy Price Increases on Economic Activity," working paper, Wharton School, University of Pennsylvania.
- Gisser, Micha, and Thomas H. Goodwin (1986), "Crude Oil and the Macroeconomy: Tests of Some Popular Notions," *Journal of Money, Credit, and Banking*, 18, pp. 95-103.
- Hamilton, James D. *The Macroeconomic Effects of Petroleum Supply Disruptions*, Ph.D. Thesis, U.C. Berkeley, 1982.
- ____ (1983), "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91, pp. 228-248.
- ____ (1985), "Historical Causes of Postwar Oil Shocks and Recessions," *Energy Journal*, 6, pp. 97-116.
- ____ (1988a), "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96, pp. 593-617.
- ____ (1988b), "Are the Macroeconomic Effects of Oil-Price Changes Symmetric? A Comment," *Carnegie-Rochester Conference Series on Public Policy*, 28, pp. 369-378.
- ____ (1996), "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics*, 38, pp. 215-220.
- ____ (2000), "A Parametric Approach to Flexible Nonlinear Inference," working paper, UCSD, <http://weber.ucsd.edu/~jhamilto>. Forthcoming, *Econometrica*.
- Hooker, Mark A. (1996), "What Happened to the Oil Price-Macroeconomy Relationship?," *Journal of Monetary Economics*, 38, pp. 195-213.
- ____ (1997), "Exploring the Robustness of the Oil Price-Macroeconomy Relationship," Working paper, Board of Governors of the Federal Reserve System.
- Hoover, Kevin D., and Stephen J. Perez (1994), "Post Hoc Ergo Propter Once More: An Evaluation of 'Does Monetary Policy Matter?' in the Spirit of James Tobin," *Journal of Monetary Economics*, 34, pp. 89-99.
- Keane, Michael P., and Eswar Prasad (1996), "The Employment and Wage Effects of Oil Price Changes: A Sectoral Analysis," *Review of Economics and Statistics*, 78, pp. 389-400.
- Kim, In-Moo, and Prakash Loungani (1992), "The Role of Energy in Real Business Cycle Models," *Journal of Monetary Economics*, 29, no. 2, pp. 173-189.

- Lee, Kiseok, and Shawn Ni (1999), "On the Dynamic Effects of Oil Price Shocks" A Study Using Industry Level Data," working paper, University of Missouri, Columbia.
- , ———, and Ronald A. Ratti, (1995) "Oil Shocks and the Macroeconomy: The Role of Price Variability," *Energy Journal*, 16, pp. 39-56.
- Loungani, Prakash (1986), "Oil Price Shocks and the Dispersion Hypothesis," *Review of Economics and Statistics*, 58, pp. 536-539.
- Mork, Knut A. (1989), "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 91, pp. 740-744.
- (1994), "Business Cycles and the Oil Market," *Energy Journal*, 15, special issue, pp. 15-37.
- , and Robert E. Hall (1980), "Energy Prices, Inflation, and Recession, 1974-1975," *Energy Journal*, 1, no. 3, pp. 31-63.
- , Øystein Olsen, and Hans Terje Mysen (1994), "Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries," *Energy Journal*, 15, no. 4, pp. 19-35.
- Pierce, James L., and Jared J. Enzler (1974), "The Effects of External Inflationary Shocks," *Brookings Papers on Economic Activity*, I:1974, pp. 13-61.
- Pindyck, Robert S. (1980), "Energy Price Increases and Macroeconomic Policy," *Energy Journal*, 1, no. 3, pp. 1-20.
- Rasche, R. H., and J. A. Tatom (1977), "Energy Resources and Potential GNP," *Federal Reserve Bank of St. Louis Review*, 59 (June), pp. 10-24.
- , and ——— (1981), "Energy Price Shocks, Aggregate Supply, and Monetary Policy: The Theory and International Evidence." In K. Brunner and A. H. Meltzer, eds., *Supply Shocks, Incentives, and National Wealth, Carnegie-Rochester Conference Series on Public Policy*, vol. 14, Amsterdam: North-Holland.
- Raymond, Jennie E., and Robert W. Rich (1997), "Oil and the Macroeconomy: A Markov State-Switching Approach," *Journal of Money, Credit and Banking*, 29 (May), pp. 193-213. Erratum 29 (November, Part 1), p. 555.
- Rotemberg, Julio J., and Michael Woodford (1996), "Imperfect Competition and the Effects of Energy Price Increases," *Journal of Money, Credit, and Banking*, 28 (part 1), pp. 549-577.
- Sakellaris, Plutarchos (1997), "Irreversible Capital and the Stock Market Response to Shocks in Profitability," *International Economic Review*, 38, pp. 351-379.
- Santini, Danilo J. (1985), "The Energy-Squeeze Model: Energy Price Dynamics in U.S. Business Cycles," *International Journal of Energy Systems*, 5, pp. 18-25.
- (1992), "Energy and the Macroeconomy: Capital Spending After an Energy Cost Shock," in J. Moroney, ed., *Advances in the Economics of Energy and Resources*, vol. 7, Greenwich, CN: J.A.I. Press.

____ (1994), "Verification of Energy's Role as a Determinant of U.S. Economic Activity," in J. Moroney, ed., *Advances in the Economics of Energy and Resources*, vol. 8, Greenwich, CN: J.A.I. Press.

Solow, Robert M. (1980), "What to Do (Macroeconomically) when OPEC Comes," in Stanley Fischer, ed., *Rational Expectations and Economic Policy*, Chicago: University of Chicago Press

Figure Captions

Figure 1. Simulated data when oil price changes are all increases. True relation (dashed line): $y_t = 2 - 0.1o_t + u_t$ where y_t is GDP growth, $o_t \sim N(0, 10^2)$ is the oil price change, and $u_t \sim N(0, 1.5^2)$. Estimated relation (solid line): equation (1.3).

Figure 2. Simulated data when oil prices can go up or down. True relation (dashed line): $y_t = 2 + u_t$ if $o_t \leq 0$ and $y_t = 2 - 0.1o_t + u_t$ if $o_t > 0$. Estimated relation (solid line): equation (1.4).

Figure 3. Effect of oil prices on GDP growth one quarter later. Solid line plots the posterior expectation of the function $\alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t)$ evaluated at $\mathbf{x}_t = (x_1, \bar{o}_{-2}, \bar{o}_{-3}, \bar{o}_{-4})'$ and $\mathbf{z}_t = (\bar{y}_{-1}, \bar{y}_{-2}, \bar{y}_{-3}, \bar{y}_{-4})'$ as a function of x_1 where $\bar{z}_{-j} = T^{-1} \sum_{t=1}^T z_{t-j}$ and where the expectation is with respect to the posterior distribution of $\alpha_0, \boldsymbol{\alpha}, \boldsymbol{\delta}, \lambda$, and $m(\mathbf{x}_t)$ conditional on observation of $\{y_t, \mathbf{x}_t, \mathbf{z}_t\}_{t=1}^T$, with this posterior distribution estimated by Monte Carlo importance sampling with 20,000 simulations. Dashed lines give 95% confidence intervals.

Figure 4. Effect of oil prices on GDP growth two quarters later. Solid line plots the posterior expectation of the function $\alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t)$ evaluated at $\mathbf{x}_t = (\bar{o}_{-1}, x_2, \bar{o}_{-3}, \bar{o}_{-4})'$ and $\mathbf{z}_t = (\bar{y}_{-1}, \bar{y}_{-2}, \bar{y}_{-3}, \bar{y}_{-4})'$ as a function of x_2 .

Figure 5. Effect of oil prices on GDP growth three quarters later. Solid line plots the posterior expectation of the function $\alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t)$ evaluated at $\mathbf{x}_t = (\bar{o}_{-1}, \bar{o}_{-2}, x_3, \bar{o}_{-4})'$ and $\mathbf{z}_t = (\bar{y}_{-1}, \bar{y}_{-2}, \bar{y}_{-3}, \bar{y}_{-4})'$ as a function of x_3 .

Figure 6. Effect of oil prices on GDP growth four quarters later. Solid line plots the posterior expectation of the function $\alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t)$ evaluated at $\mathbf{x}_t =$

$(\bar{o}_{-1}, \bar{o}_{-2}, \bar{o}_{-3}, x_4)'$ and $\mathbf{z}_t = (\bar{y}_{-1}, \bar{y}_{-2}, \bar{y}_{-3}, \bar{y}_{-4})'$ as a function of x_4 .

Figure 7. Effect of oil prices on GDP growth three quarters later for different possible values of o_{t-4} . Each line plots the posterior expectation of the function $\alpha_0 + \boldsymbol{\alpha}'\mathbf{x}_t + \boldsymbol{\delta}'\mathbf{z}_t + \lambda m(\mathbf{x}_t)$ evaluated at $\mathbf{x}_t = (\bar{o}_{-1}, \bar{o}_{-2}, x_3, x_4)'$ and $\mathbf{z}_t = (\bar{y}_{-1}, \bar{y}_{-2}, \bar{y}_{-3}, \bar{y}_{-4})'$ as a function of x_3 . For the solid line, $x_4 = 0$, for the long-dashed line, $x_4 = -5$, and for the short-dashed line, $x_4 = 5$.

Figure 8. Effect of oil prices on GDP growth three quarters later for different possible values of o_{t-4} as predicted by four different models. In each panel, the solid line plots (as a function of x_3) the function $E(y_t | y_{t-1} = \bar{y}_{-1}, y_{t-2} = \bar{y}_{-2}, y_{t-3} = \bar{y}_{-3}, y_{t-4} = \bar{y}_{-4}, o_{t-1} = \bar{o}_{-1}, o_{t-2} = \bar{o}_{-2}, o_{t-3} = x_3, o_{t-4} = x_4)$. For the solid line, $x_4 = 0$, for the dashed line, $x_4 = -5$, and for the alternate-dashed line, $x_4 = 5$. In the upper left panel, the conditional expectation is based on the linear regression (1.6). The other panels (reading clockwise) are based on (1.8), (3.6), and (3.2), respectively.

Figure 9. Evidence of structural change in various relations. Each figure plots the p -value for a test of the null hypothesis of no structural change at date t_0 in the coefficients on lagged oil prices as a function of t_0 ; dashed line denotes $p = 0.05$. Top panel: univariate autoregression to predict oil prices (test for this panel also allows for shift in constant term); second panel: regression (1.6); third panel: regression (1.8); fourth panel: regression (3.2); bottom panel: regression (3.6).

Figure 10. Top panel: oil price changes (o_t). Second panel: net oil price increases (o_t^\dagger). Third panel: value by which oil price rises above its 3-year maximum; zero if below the

3-year maximum. Bottom panel: fitted values from a regression of o_t on a constant and $Q_t, Q_{t-1}, \dots, Q_{t-4}$.

Figure 11. Dynamic effects of a 10% oil price increase. Both panels plot $10 \times \partial E(y_{t+j} | y_t, y_{t-1}, y_{t-2}, y_{t-3}, o_{t+j} = 0, o_{t+j-1} = 0, \dots, o_{t+1} = 0, o_t, o_{t-1}, o_{t-2}, o_{t-3}) / \partial o_t$ as a function of j . For the top model, the conditional expectation is based on (4.1). For the bottom panel, it is based on (3.2).

Figure 12. Evidence of structural change in the instrumental variables estimation. The figure plots the p -value for a test of the null hypothesis of no structural change at t_0 in the coefficients on o_{t-1}, \dots, o_{t-4} in the IV regression (4.1). The short-dashed line (at 5%) represents the critical value that would be appropriate if one had used only a single value for t_0 and the long-dashed represents the critical value that would be used for a 5% test based on the lowest value found for any t_0 .

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8

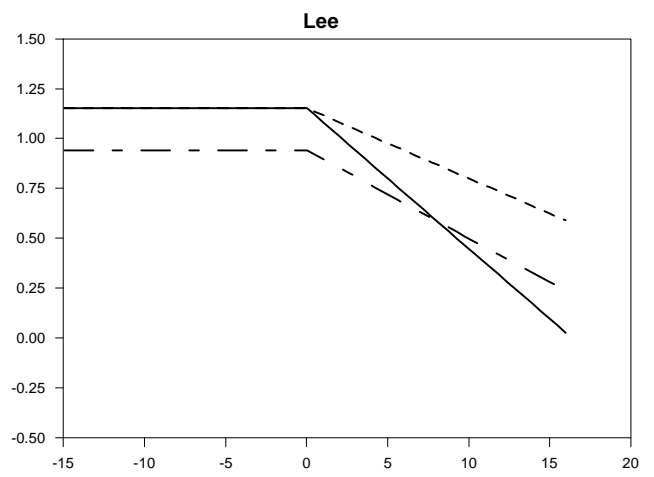
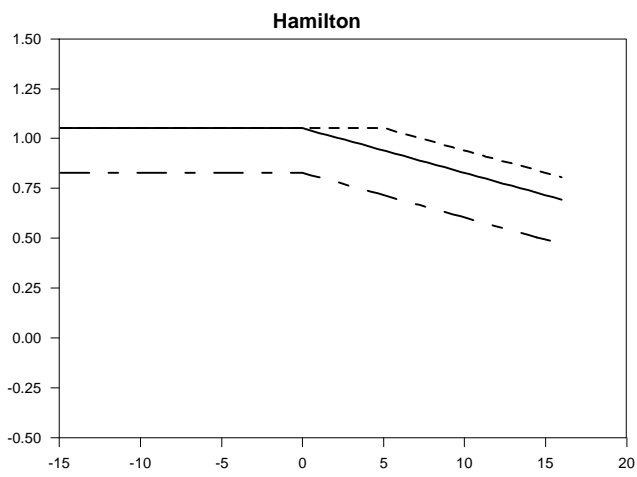




Figure 9






Figure 10


Oil Price Change



Net Oil Price Increase (one year)



Net Oil Price Increase (three years)



First Stage Least Squares





Figure 11

Instrumental Variables



Net Oil Price Increase




Figure 12

