

NBER WORKING PAPER SERIES

NONLINEARITIES AND THE MACROECONOMIC EFFECTS OF OIL PRICES

James D. Hamilton

Working Paper 16186

<http://www.nber.org/papers/w16186>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

July 2010

I thank Rob Vigfusson for graciously sharing his data and helping me follow his code. An earlier version of this paper was circulated under the title, "Yes, the Response of the U.S. Economy to Energy Prices is Nonlinear." The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 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.

Nonlinearities and the Macroeconomic Effects of Oil Prices
James D. Hamilton
NBER Working Paper No. 16186
July 2010
JEL No. E32,Q43

ABSTRACT

This paper reviews some of the literature on the macroeconomic effects of oil price shocks with a particular focus on possible nonlinearities in the relation and recent new results obtained by Kilian and Vigfusson (2009).

James D. Hamilton
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
jhamilton@ucsd.edu

1 Overview.

I noted in a paper published in the *Journal of Political Economy* in 1983 that at that time, 7 out of the 8 postwar U.S. recessions had been preceded by a sharp increase in the price of crude petroleum. Iraq's invasion of Kuwait in August 1990 led to a doubling in the price of oil in the fall of 1990 and was followed by the ninth postwar recession in 1990-91. The price of oil more than doubled again in 1999-2000, with the tenth postwar recession coming in 2001. Yet another doubling in the price of oil in 2007-2008 accompanied the beginning of recession number 11, the most recent and frightening of the postwar economic downturns. So the count today stands at 10 out of 11, the sole exception being the mild recession of 1960-61 for which there was no preceding rise in oil prices.

Oil shocks could affect the economy through their consequences for both supply and demand. On the supply side, consider a firm whose output Y depends on inputs of capital K , labor N , and energy E :

$$Y = F(K, N, E).$$

Suppose that the capital stock is fixed in the short run and that wages adjust instantly to ensure that labor demand equals a fixed supply \bar{N} . Then if X denotes the price of energy relative to the price of output,

$$\frac{\partial Y}{\partial X} = \frac{\partial F}{\partial E} \frac{\partial E}{\partial X}. \tag{1}$$

Multiplying (1) by X/Y results in

$$\frac{\partial \ln Y}{\partial \ln X} = \frac{\partial F}{\partial E} \frac{E}{Y} \frac{\partial \ln E}{\partial \ln X}. \tag{2}$$

If the marginal product of energy equals its relative price ($\partial F/\partial E = X$), then the first terms on the right side of (2) will be recognized as the energy expenditure share

$$\frac{\partial F}{\partial E} \frac{E}{Y} = \frac{EX}{Y} = \gamma$$

where γ denotes the firm's spending on energy relative to the value of its total output.

Letting lower-case letters denote natural logarithms, (2) can be written

$$\frac{\partial y}{\partial x} = \gamma \frac{\partial e}{\partial x}. \quad (3)$$

In other words, the elasticity of output with respect to the relative price of energy would be the energy expenditure share γ times the price-elasticity of energy demand.

The energy expenditure share is a small number. In 2009, the U.S. consumed about 7.1 billion barrels of petroleum products, which at the current \$80/barrel price of crude corresponds to a value around \$570 billion. This would represent only 4% of U.S. GDP. Moreover, the short-run price-elasticity of petroleum demand is extremely small (Dahl, 1993), so that expression (3) implies an output response substantially below 4%. For this reason, models built around this kind of mechanism, such as Kim and Loungani (1992), imply that oil shocks could only have made a small contribution to historical downturns. Note also that (3) implies a linear relation between y and x ; an oil price decrease should increase output by exactly the same amount that an oil price increase of the same magnitude would decrease output.

To account for larger effects, it would have to be the case that either K or N also adjust in response to the oil price shock. Finn (2000) analyzed the multiplier effects that result if

firms adjust capital utilization rates in order to minimize depreciation expenses. Leduc and Sill (2004) incorporated this utilization effect along with labor adjustments resulting from sticky wages. Again these models imply a linear relation between y and x , though Atkeson and Kehoe's (1999) treatment of putty-clay investment technology produces some nonlinear effects.

Davis (1987a, 1987b) stressed the role of specialized labor and capital in the transmission mechanism. If the marginal product of labor falls in a particular sector, it can take time before workers relocate to something more productive, during which transition the economy will have some unemployed resources. In my 1988 paper, unemployment could result not just from workers who are in transition between sectors but also from workers who are simply waiting until conditions in their sector once again improve. In such models, idle labor and capital rather than decreased energy use as in (1) account for the lost output. Moreover, these effects are clearly nonlinear. For example when energy prices fell in 1985, some workers in the oil-producing sector were forced to find other jobs. As a result, it is possible in principle for aggregate output to fall temporarily in response to an oil price decrease just as it does for an oil price increase.

An alternative mechanism operates through the demand side. An increase in energy prices leaves consumers with less money to spend on non-energy items and leaves an oil-importing country with less income overall. If a consumer tries to purchase the same quantity of energy E in response to an increase in the relative price given by ΔX , then saving or expenditures on other items must fall by $E \cdot \Delta X$, with a proportionate effect on

demand given by

$$\frac{\Delta Y^d}{Y} = \frac{E}{Y} \Delta X.$$

If as in a fixed-price Keynesian model demand Y^d is the limiting determinant of total output, we would have

$$\frac{\partial y}{\partial x} = \gamma,$$

so that by this mechanism the effect once again is linear and bounded¹ by the expenditure share γ .

Specialization of labor and capital could also be important for the transmission of demand effects as well. Demand for less fuel-efficient cars would be influenced not just by the consequences of an oil price increase for current disposable income but also by consideration of future gasoline prices over the lifetime of the car. Bernanke (1983) noted that uncertainty per se could lead to a postponement of purchases for capital and durable goods. A shift in demand away from larger cars seems to have been a key feature of the macroeconomic response to historical oil shocks (Bresnahan and Ramey, 1993; Edelstein and Kilian, 2009; Hamilton, 2009; Ramey and Vine, 2010), and Bresnahan and Ramey (1993) and Ramey and Vine (2010) map out in detail exactly how specialization of labor and capital in the U.S. automobile industry amplified the effects of historical oil price shocks and introduced nonlinearities of the sort anticipated by the sectoral-shifts hypothesis. In the model of Hamilton (1988), shifts in the advantages between sectors resulting from supply effects (greater production costs for sector 1 as a result of higher energy prices) or demand effects (less demand

¹ Price adjustment would make this effect smaller whereas the traditional Keynesian multiplier could make it bigger.

for the output of sector 1 as a result of higher energy prices) have identical macroeconomic consequences, operating in either case through idled labor in the disadvantaged sector.

In terms of empirical evidence on nonlinearity, Loungani (1986) demonstrated that oil-induced sectoral imbalances contributed to fluctuations in U.S. unemployment rates. Mork (1989) found that oil price increases have different predictive implications for subsequent U.S. GDP growth than oil price decreases. Other studies also reporting evidence that nonlinear forecasting equations do better include Lee, Ni and Ratti (1995), Balke, Brown, and Yücel (2002), and Hamilton (1996, 2003). Both Carlton (2010) and Ravazzolo and Rothman (2010) confirmed these predictive improvements using real-time data. Ferderer (1996) and Elder and Serletis (2010) demonstrated that oil-price volatility predicts slower GDP growth, implying that oil price decreases include some contractionary implications. Davis and Haltiwanger (2001) found nonlinearities in the effects of oil prices on employment at the individual plant level for U.S. data. Herrera, Lagalo, and Wada (2010) found a strong nonlinear response of U.S. industrial production to oil prices, with the biggest effects in industries the use of whose products by consumers is energy intensive. A nonlinear relation between oil prices and subsequent real GDP growth has also been reported for a number of OECD countries by Cuñado and Pérez de Gracia (2003), Jiménez-Rodríguez and Sánchez (2005), Kim (2009), and Engemann, Kliesen and Owyang (2010).

By contrast, a prominent recent study by Kilian and Vigfusson (2009) found little evidence of nonlinearity in the relation between oil prices and U.S. GDP growth. In the next section I explore why they seem to have reached a different conclusion from many of the

previous researchers mentioned above. Sections 3 and 4 note some of the further implications of their results for inference about nonlinear dynamic relations.

2 Testing for nonlinearity.

Let y_t denote the rate of growth of real GDP, x_t the change in the price of oil, and \tilde{x}_t a proposed known nonlinear function of oil prices. The null hypothesis that the optimal one-period-ahead forecast of y_t is linear in past values of x_{t-i} is quite straightforward to state and test: we just use OLS to estimate the forecasting regression

$$y_t = \alpha + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^p \beta_i x_{t-i} + \sum_{i=1}^p \gamma_i \tilde{x}_{t-i} + \varepsilon_t \quad (4)$$

and test whether $\gamma_1 = \dots = \gamma_p = 0$. As noted above, a large number of papers have tested such a hypothesis and rejected it. Kilian and Vigfusson's paper might leave the impression that these earlier tests were somehow misspecified or insufficiently powerful, and that the reason Kilian and Vigfusson reach a different conclusion from previous researchers is that they are proposing superior tests. Such a result would be surprising if true. For Gaussian ε_t in (4), OLS produces maximum likelihood estimates which are asymptotically efficient, and the OLS F test is the likelihood ratio test with well-known desirable properties. That some new test could be more powerful than the standard OLS test seems unlikely, and certainly if the OLS test rejects and the new test does not, the reconciliation cannot be based on the assertion that the new test is more powerful. Kilian and Vigfusson also include in their analysis some standard OLS tests, which offer further support for their conclusion that the relation appears to be linear. But insofar as these are the same OLS tests that have already

produced rejections of the null hypothesis in previous studies, the difference in conclusions must come from a different data set or differences in the specification of the basic forecasting regression (4), and not from any superior properties of the new tests proposed in their paper.

Most of their paper explores the case in which \tilde{x}_t is given by $x_t^+ = \max\{0, x_t\}$, the alternative hypothesis of interest taken to be that oil price increases have different economic effects from oil price decreases. This particular specification is one that previous researchers have found to be unstable over earlier data sets (e.g., Hooker, 1996; Hamilton, 2003), so it is unsurprising that Kilian and Vigfusson find that such a relation does not perform well on their sample either. My earlier investigation (Hamilton, 2003) concluded that the nonlinearities can be captured with a specification in which what matters is whether oil prices make a new 3-year high:

$$x_t^\# = \max\{0, X_t - \max\{X_{t-1}, \dots, X_{t-12}\}\}$$

for X_t the log level of the oil price. Below I reproduce the coefficients as reported in equation (3.8) of Hamilton (2003):

$$\begin{aligned} y_t = & \frac{0.98}{(0.13)} + \frac{0.22y_{t-1}}{(0.07)} + \frac{0.10y_{t-2}}{(0.07)} - \frac{0.08y_{t-3}}{(0.07)} - \frac{0.15y_{t-4}}{(0.07)} \\ & - \frac{0.024x_{t-1}^\#}{(0.014)} - \frac{0.021x_{t-2}^\#}{(0.014)} - \frac{0.018x_{t-3}^\#}{(0.014)} - \frac{0.042x_{t-4}^\#}{(0.014)}. \end{aligned} \quad (5)$$

If one adds the linear terms $\{x_{t-1}, x_{t-2}, x_{t-3}, x_{t-4}\}$ to this regression and calculates the OLS χ^2 test of the hypothesis that the coefficients on $\{x_{t-1}^\#, x_{t-2}^\#, x_{t-3}^\#, x_{t-4}^\#\}$ are zero using the original data set, the result is a $\chi^2(4)$ statistic of 16.93, with a p -value of 0.002. The last entry of Kilian and Vigfusson's Table 4 reports the OLS χ^2 test on a similar specification

for their data set which results in a p -value of 0.046. Clearly it must be differences in the specification and data set between the two papers, rather than differences in the testing methodology, that accounts for the different findings. There are a number of differences that could explain the higher p -value obtained by Kilian and Vigfusson.

Different data sets. In my earlier analysis, t in (5) ran from 1949:Q2 to 2001:Q3 (or 210 total observations), whereas in Kilian and Vigfusson's analysis, t runs from 1974:Q4 to 2007:Q4 (or 133 total observations). One would expect to see a higher p -value in a shorter sample, since fewer observations make it harder to reject any hypothesis. In addition, it is possible that there has been a structural change since 2001, so that the earlier proposed nonlinear relation (5) does a poorer job with more recent data.

Different measure of oil prices. In my original analysis, x_t was based on the producer price index for crude petroleum, whereas Kilian and Vigfusson use the refiner acquisition cost for imported oil. The values of these two measures are compared in the top two panels of Figure 1. The RAC is not available prior to 1974, and Kilian and Vigfusson imputed values back to 1971. The two oil price measures are very similar after 1983, but are somewhat different in the 1970s. Most notably, according to RAC, the first oil shock of 1974:Q1 was three times the size of that seen in any other quarter of the 1970s, and there was very little change in oil prices in 1981:Q1. By contrast, the PPI registers the shocks of 1974:Q1, 1979:Q2-Q3, and 1981:Q1 as similar events. If one thought that a key factor in the transmission mechanism to the U.S. economy involved the price consumers paid for gasoline, the PPI may provide a better measure, since the CPI also represents these three

shocks as having similar magnitude (see the bottom panel of Figure 1). In any case, it is certainly possible that for such different measures of oil prices, the functional form of the optimal forecast could differ.

Different price adjustment. Another difference is that (5) used for x_t the nominal change in the price of oil, whereas Kilian and Vigfusson subtract the percentage change in the consumer price index in their definition of x_t . They argue correctly that most economic theories would involve the real rather than the nominal price of oil. On the other hand, if the nonlinearity represents threshold responses based on consumer sentiment, it is possible that these thresholds are defined in nominal terms. I would also note that empirical measurement of “the” aggregate price level is problematic, and deflating by a particular number such as the CPI introduces a new source of measurement error, which could lead to a deterioration in the forecasting performance. In any case, it is again quite possible that there are differences in the functional form of forecasts based on nominal instead of real prices.

Inclusion of contemporaneous regressors. The χ^2 statistics in Kilian and Vigfusson’s Table 4 are in fact not based on the forecasting regression (4), but instead come from testing $\gamma_0 = \gamma_1 = \dots = \gamma_p$ in

$$y_t = \alpha + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=0}^p \beta_i x_{t-i} + \sum_{i=0}^p \gamma_i \tilde{x}_{t-i} + \varepsilon_t. \quad (6)$$

Kilian and Vigfusson suggest that this second test would have more power than the first, though again I find that an odd claim, since the two regressions are asking different questions. Equation (4) is asking a forecasting question: can GDP growth in quarter t be predicted on the basis of variables known at the end of quarter $t - 1$? Equation (6) is estimating

something else, which perhaps has a structural interpretation for some proposed model of the economy. If the conditional expectation of y_t given current and past oil prices in fact does not involve the current oil price (so that $\beta_0 = \gamma_0 = 0$), then the two regressions would be asking the same question. In that case, a test based on (4) would have to be more powerful since it does not require the estimation and testing of auxiliary parameters whose true value is in fact zero.

Number of lags. My original regression (5) used $p = 4$ lags, whereas Kilian and Vigfusson have used $p = 6$ lags throughout. If the truth is $p = 4$, estimating and testing the additional lags will result in a reduction in power. On the other hand, it might be argued that an optimal linear forecast of y_t requires more than 4 lags of y_{t-i} and x_{t-i} , and that omitting the extra lags accounts for the apparent success of a nonlinear specification (since $x_{t-4}^\#$ incorporates some additional information about x_{t-i} for $i > 4$).

The contribution of each factor. Table 1 identifies the role of each of these differences in turn, by changing one element of the specification at a time and seeing what effect it has on the results. The first row gives the p -value for the last entry reported in Kilian and Vigfusson's Table 4, while the second row gives the p -value for my specification on the original data set. The third row isolates the effect of the choice of sample period alone, by estimating my original specification using the sample period adopted by Kilian and Vigfusson. Instead of a p -value of 0.002 obtained for the original sample, the p -value is only 0.013 on the new data set. Is this because the sample is shorter, or because the relation has changed? One can test for the latter possibility by using data for 1949:Q2-2007:Q4 to re-

estimate equation (5) and test the hypothesis that the coefficients on $\{x_{t-1}^{\#}, x_{t-2}^{\#}, x_{t-3}^{\#}, x_{t-4}^{\#}\}$ were different subsequent to 2001:Q4 compared with those prior to 2001:Q4. One fails to reject the null hypothesis of no change in these coefficients ($\chi^2(4) = 5.03, p = 0.284$). Thus a key explanation for why Kilian and Vigfusson find weaker evidence of nonlinearity is that they have used a shorter sample.

Subsequent rows of Table 1 use Kilian and Vigfusson's 1974:Q4-2007:Q4 sample, but change other elements of their specification one at a time. Row 4 uses the real change in the producer price index of crude petroleum in place of the real change in the refiner acquisition cost of imported oil, but otherwise follows Kilian and Vigfusson in all the other details. Using the PPI instead of the RAC would reduce Kilian and Vigfusson's reported p -value from 0.046 to 0.024. Row 5 keeps the RAC, but uses the nominal price rather than the real. This change alone would again have reduced the p -value from 0.046 to 0.028. Row 6 simply omits the contemporaneous term, basing the test on (4) rather than (6), and would be another way to reduce the p -value to 0.027. Finally, row 7 shows that using $p = 4$ instead of $p = 6$ would also increase the evidence of nonlinearity. Furthermore, a test of the null hypothesis that $\beta_5 = \beta_6 = \gamma_5 = \gamma_6 = 0$ in Kilian-Vigfusson's original (6) fails to reject ($\chi^2(4) = 3.90, p = 0.420$), suggesting that this again is another factor in reducing the power of their tests.

To summarize, Kilian and Vigfusson make a number of changes from previous research, including a shorter sample, different oil price measure, different price adjustment, inclusion of contemporaneous terms, and longer lags. Each of these changes, taken by itself, would

lead them to find weaker evidence of nonlinearity than previous research. Taken together, they explain why the overall conclusions of their study differ from most earlier investigations.

Post-sample performance. The nonlinear terms in (4) seem to improve the in-sample fit, but how helpful are they out of sample? If one estimates a purely linear forecasting equation from the Kilian-Vigfusson 1974:Q3 to 2007:Q4 data set and specification,

$$y_t = \alpha + \sum_{i=1}^6 \phi_i y_{t-i} + \sum_{i=1}^6 \beta_i x_{t-i} + \varepsilon_t,$$

and uses these estimated coefficients to predict the one-quarter-ahead GDP growth over the data we've subsequently received for 2008:Q1 to 2009:Q3, the out-of-sample root-mean-squared-error is 1.09. On the other hand, when the estimated relation is allowed to include the Kilian-Vigfusson real RAC net measure $x_t^\#$ as in (4), the out-of-sample RMSE is 0.80, a 27% improvement. Interestingly, if one uses the coefficients from (5) as estimated 1949:Q2 to 2001:Q3 to form a post-sample forecast over the period 2008:Q1 to 2009:Q3, the RMSE is 0.62. These results confirm the impression that nonlinear terms, particularly of the form I proposed in my 2003 paper, are helpful for forecasting U.S. real GDP growth.

3 Censoring bias.

In Section 2 of their paper, Kilian and Vigfusson (2009) demonstrate that if the true relation is linear and one mistakenly estimates a nonlinear specification, the resulting estimates are asymptotically biased. These results parallel the demonstration in Hamilton (2003) that if the true relation is nonlinear and one mistakenly estimates a linear specification, the resulting estimates are asymptotically biased. Both statements are of course true, and are

illustrations of the broader theme that one runs into problems whenever one tries to estimate a misspecified model.

Kilian and Vigfusson suggest one should take the high road of including both linear and nonlinear terms as a general strategy to avoid either problem. While that would indeed work if one had an infinite sample, in practice it is not always better to add more parameters, particularly in a sample as small as that used by Kilian and Vigfusson. After all, the same principle would suggest we include *both* the RAC and PPI as the oil price measure on the right-hand side, since there is disagreement as to which is the better measure, and nonlinear transformations of both the real and nominal magnitudes. Nobody would do that, and nobody should. All empirical research necessarily faces a trade-off between parsimony and generality, and one is forced to choose some point on that trade-off in literally every empirical study that has ever been done. My personal belief is that there are very strong arguments for trying to keep the estimated relations parsimonious. I note for example that equation (5), with 9 estimated parameters, provides a 22% improvement in terms of the out-of-sample RMSE compared with the 19 parameters required for Kilian and Vigfusson's (4).

4 On calculating impulse-response functions.

A separate question raised by Kilian and Vigfusson (2009), and one for which I am in full agreement with their analysis, is how one should calculate impulse-response functions. An expression like (5) is perfectly appropriate for the purpose of forming a one-quarter-ahead

forecast. The conditional expectation function,

$$E(y_t | y_{t-1}, x_{t-1}, y_{t-2}, x_{t-2}, \dots),$$

is what it is, and if the functional form of this expectation is indeed as proposed in (5), then OLS should give optimal parameter estimates. However, if one's goal is to calculate multi-period-ahead forecasts, such as required by an impulse-response function, Kilian and Vigfusson are quite correct that there is a problem with the standard approach of simply adding to (5) a second equation of the form

$$x_t^\# = c + \sum_{i=1}^p b_i x_{t-i}^\# + \sum_{i=1}^p d_i y_{t-i} + u_t, \quad (7)$$

and iterating the one-period-ahead forecasts of the two equations forward assuming iterated linear projections. The problem comes from the fact that while (5) in such a system would be correctly specified, the fitted values of (7) cannot possibly represent the conditional expectation

$$E(x_t^\# | x_{t-1}, y_{t-1}, x_{t-2}, y_{t-2}, \dots), \quad (8)$$

since (7) could generate a negative predicted value for $x_t^\#$, which an optimal forecast (8) would never allow. This point was first noted by Balke, Brown and Yücel (2002), though most researchers have ignored the concern.

Not only is there a problem with applying mechanically the standard linear impulse-response tools in such a setting as a result of the difference between (7) and (8), but at a more fundamental level, researchers need to reflect on the underlying question which they are intending such calculations to answer. A variable such as $x_t^\#$ is nonnegative by definition,

and therefore the conditional expectation (8) must always be a positive number. Thus if one defines an “oil shock” as a deviation from the conditional expectation,

$$u_t = x_t^\# - E(x_t^\# | x_{t-1}, y_{t-1}, x_{t-2}, y_{t-2}, \dots), \quad (9)$$

then there is a range of positive realizations of $x_t^\#$ that are defined to be a “negative oil shock”. More generally, insofar as an impulse-response function is intended to summarize the revision in expectations of future variables associated with a particular realization of (9), as Gallant, Rossi, and Tauchen (1993), Koop, Pesaran and Potter (1996), and Potter (2000) emphasized, such an object is potentially different for every different information set $\{x_{t-1}, y_{t-1}, x_{t-2}, y_{t-2}, \dots\}$ and size of the shock u_t . For small shocks, one would expect from Taylor’s Theorem that a linear representation of the function would be a good approximation around the point of linearization. Kilian and Vigfusson assume that the researcher is interested in calculating the consequences of a one-standard-deviation shock averaged over the history of different dates t . For purposes of answering this question, and particularly given the underlying weak evidence of nonlinearity for their data set and specification, they find limited evidence of nonlinearity in the impulse-response function. Notwithstanding, one could imagine wanting to know the answer to dynamic questions other than this, such as measuring the implications of a big change in the price of oil occurring at some particular date of interest. By construction, certain questions of this form would be much more poorly approximated with a linear function than might be suggested by the summary statistics reported in their paper.

5 Conclusion.

To me, the evidence is convincing that the relation between GDP growth and oil prices is nonlinear. The recent paper by Kilian and Vigfusson (2009) does not challenge that conclusion, but does offer a useful reminder that we need to think carefully about what question we want to ask with an impulse-response function in such a system and cannot rely on off-the-shelf linear methods for an answer.

References

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 (2002) “Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?” *Energy Journal*, 23, pp. 27-52.

Bernanke, Ben S. (1983) “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 98, pp. 85-106.

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.

Carlton, Amelie Benear (2010) “Oil Prices and Real-Time Output Growth.” Working paper, University of Houston.

Cuñado, Juncal, and Fernando Pérez de Gracia (2003) “Do Oil Price Shocks Matter? Evidence for Some European Countries,” *Energy Economics*, 25, pp. 137–154.

Dahl, Carol A. (1993) “A Survey of Oil Demand Elasticities for Developing Countries,” *OPEC Review* 17(Winter), pp. 399-419.

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. 326-332.

_____, and John Haltiwanger (2001) “Sectoral Job Creation and Destruction Responses to Oil Price Changes,” *Journal of Monetary Economics*, 48, pp. 465-512.

Edelstein, Paul, and Lutz Kilian (2009) “How Sensitive are Consumer Expenditures to Retail Energy Prices?” *Journal of Monetary Economics*, 56(6), pp. 766-779.

Elder, John, and Apostolos Serletis (2010) “Oil Price Uncertainty,” *Journal of Money, Credit and Banking*, forthcoming.

Engemann, Kristie M., Kevin L. Kliesen, and Michael T. Owyang (2010) “Do Oil Shocks Drive Business Cycles? Some U.S. and International Evidence,” working paper, Federal Reserve Bank of St. Louis.

Ferderer, J. Peter (1996) “Oil Price Volatility and the Macroeconomy: A Solution to the Asymmetry Puzzle,” *Journal of Macroeconomics*, 18, pp. 1-16.

Finn, Mary G. (2000) “Perfect Competition and the Effects of Energy Price Increases on Economic Activity,” *Journal of Money, Credit and Banking* 32, pp. 400-416.

Gallant, A. Ronald, Peter E. Rossi, and George Tauchen (1993) “Nonlinear Dynamic Structures,” *Econometrica* 61, pp. 871-908.

Hamilton, James D. (1983) “Oil and the Macroeconomy Since World War II,” *Journal of Political Economy*, 91, pp. 228-248.

_____, (1988) “A Neoclassical Model of Unemployment and the Business Cycle,” *Journal of Political Economy*, 96, pp. 593-617.

_____, (1996) “This is What Happened to the Oil Price Macroeconomy Relation,”

Journal of Monetary Economics, 38, pp. 215-220.

_____ (2003) “What Is an Oil Shock?”, *Journal of Econometrics*, 113, pp. 363-398.

_____ (2005) “What’s Real About the Business Cycle?” *Federal Reserve Bank of St. Louis Review*, July/August (87, no. 4), pp. 435-452.

_____ (2009) “Causes and Consequences of the Oil Shock of 2007-08,” *Brookings Papers on Economic Activity*, Spring 2009, pp. 215-261.

Herrera, Ana María, Latika Gupta Lagalo, and Tatsuma Wada (2010), “Oil Price Shocks and Industrial Production: Is the Relationship Linear?,” working paper, Wayne State University.

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

Jiménez-Rodríguez, Rebeca, and Marcelo Sánchez (2005) “Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries,” *Applied Economics*, 37, pp. 201-228.

Kilian, Lutz, and Robert J. Vigfusson (2009) “Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?” Working Paper, University of Michigan (<http://www-personal.umich.edu/~lkilian/kvsubmission.pdf>).

Kim, Dong Heon (2009) “What is an Oil Shock? Panel Data Evidence,” working paper, Korea University.

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.

Koop, Gary, M. Hashem Pesaran, and Simon M. Potter (1996) “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74, pp. 19-147.

Leduc, Sylvain, and Keith Sill (2004) “A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns,” *Journal of Monetary Economics*, 51, pp. 781-808.

Lee, Kiseok, Shawn Ni 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*, 68, 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.

Potter, Simon M. (2000) “Nonlinear impulse response functions,” *Journal of Economic Dynamics and Control*, 24, pp. 1425-1446.

Ramey, Valerie A., and Daniel J. Vine (2010) “Oil, Automobiles, and the U.S. Economy: How Much Have Things Really Changed?”, working paper, UCSD.

Ravazzolo, Francesco, and Philip Rothman (2010) “Oil and US GDP: A Real-Time Out-of-Sample Examination,” working paper, East Carolina University.

Table 1
P-values for test of null hypothesis of linearity for alternative specifications

	sample	oil measure	price adjustment	contemporaneous	# lags	<i>p</i> -value
(1)	1974:Q4-2007:Q4	RAC	real	include	6	0.046
(2)	1949:Q2-2001:Q3	PPI	nominal	exclude	4	0.002
(3)	1974:Q4-2007:Q4	PPI	nominal	exclude	4	0.013
(4)	1974:Q4-2007:Q4	PPI	real	include	6	0.024
(5)	1974:Q4-2007:Q4	RAC	nominal	include	6	0.028
(6)	1974:Q4-2007:Q4	RAC	real	exclude	6	0.027
(7)	1974:Q4-2007:Q4	RAC	real	include	4	0.036

Notes to Table 1: *P*-values for test that $\gamma_0 = \gamma_1 = \dots = \gamma_p = 0$ in equation (6) (for rows with “include” in contemporaneous column) or test that $\gamma_1 = \dots = \gamma_p = 0$ in equation (4) (for rows with “exclude” in contemporaneous column). Boldface entries in each row indicate those details of the specification that differ from the first row.

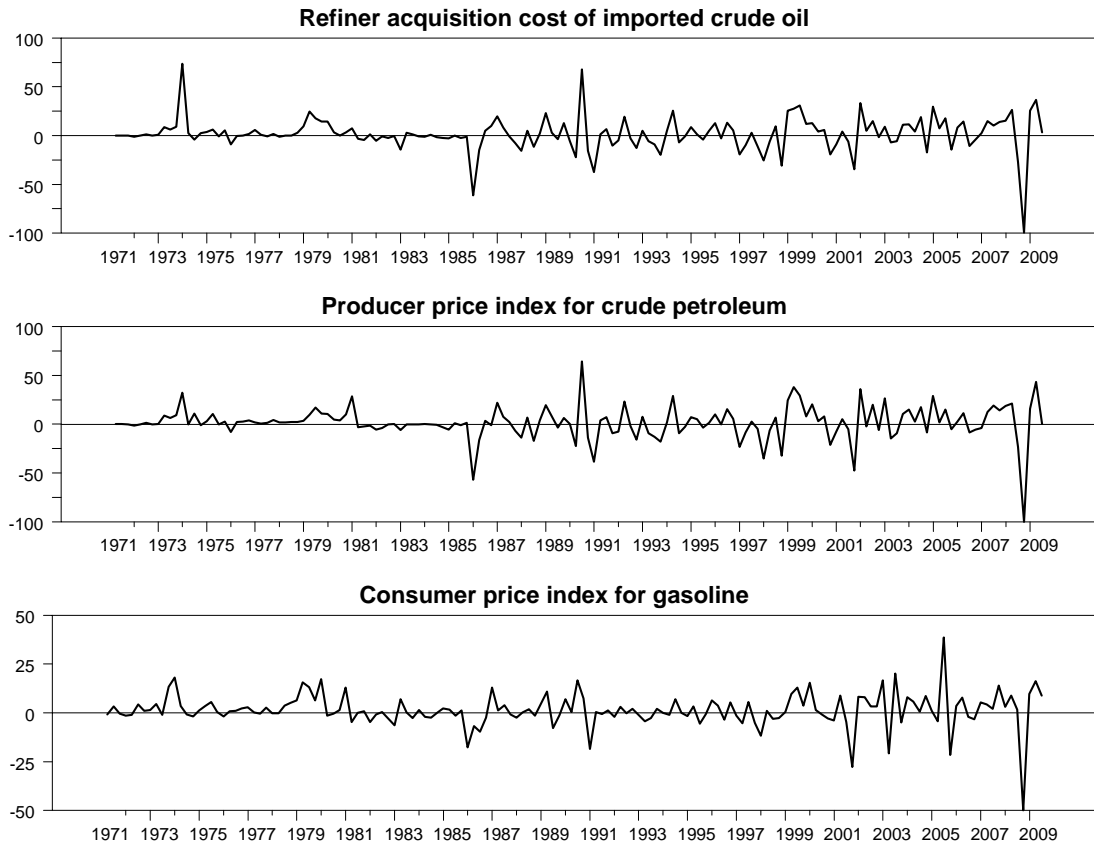


Figure 1. Quarterly percent changes in PPI, RAC, and gasoline CPI. Third series is seasonally adjusted, first two are not.