Oil Price Shocks and the U.S. Stock Market: Do Sign and Size Matter?*

Zeina Alsalman† Ana María Herrera‡

Abstract

We investigate the effect of oil price innovations on the U.S. stock market using a model that nests symmetric and asymmetric responses to positive and negative oil price innovations. We find no evidence of asymmetry for aggregate stock returns, and only very limited evidence for 49 industry-level portfolios. Moreover, these asymmetries do not match up well with conventional views regarding energy-dependent sectors of the economy. Instead, asymmetries are more likely driven by the effect of oil price innovations on expected and/or realized demand. We inquire whether the size of the shock matters in that doubling the size of the shock more (or less) than doubles the size of the response, finding that the effect of a 2.s.d innovation is just about double the magnitude of the impact of a 1.s.d innovation. Furthermore, we find no support for the conjecture that shocks that exceed a threshold have an asymmetric effect on stock returns.

Key words: oil prices, U.S. stock returns, asymmetric responses.

*The on-line appendix is available at http://gatton.uky.edu/faculty/herrera/documents/AHappendix.pdf. We thank Jim Hamilton, Lutz Kilian, participants at the 2013 American Economic Association Meetings, two anonymous referees and the editor for helpful comments and suggestions.

†Department of Economics, 332D Elliott Hall, School of Business Administration, Oakland University, Rochester, MI 48309; e-mail: alsalman@oakland.edu

‡Corresponding Author: Department of Economics, 335Z Gatton College of Business and Economics, University of Kentucky, Lexington, KY, 40506-0034; e-mail: amherrera@uky.edu

1 Introduction

Headlines such as "U.S. stocks plunge after oil climbs $6" (New York Times, June 11, 2008) or "U.S. stocks rally after crude drops to 3-month low" (Wall Street Journal, August 8, 2008) highlight the shared belief among journalists and stock market commentators that oil price shocks have a direct effect on the U.S. stock market. Moreover, these headlines put in evidence the belief that the effect might depend on the behavior of crude oil prices in the recent history, and suggest that the relationship between oil price shocks and stock returns might be nonlinear.

For many years, researchers compiled conflicting evidence regarding the nature of the relationship between changes in crude oil prices and stock returns. On the one hand, Chen, Roll and Ross (1986) and Huang, Masulis and Stoll (1996) found no evidence of a negative relationship between prices of oil futures and stock returns. Similarly, Wei (2003) encountered that the oil price shock of 1973-74 had no impact on stock returns. On the other hand, work by Kling (1985) and Jones and Kaul (1996) pointed towards a negative effect of oil price shocks on stock returns. Yet, in recent years, a consensus appears to have emerged among academics: oil price shocks exert a negative impact on most stock returns, though the nature of the relationship depends on the underlying shock. In particular, Kilian and Park (2009) find that oil price shocks that are driven by innovations to the precautionary demand for crude oil have a negative impact on U.S. stock returns. They show that the response differs significantly depending on the source of the oil price shock (e.g., supply or demand driven). Thus, changes in the composition of oil price shocks over time help explain why, in the past, researchers failed to find evidence in favor of an effect of oil price innovations on U.S. stock returns.

An alternative explanation for these contrasting results could stem from the possibly nonlinear nature of the relationship between stock returns and oil price shocks. For instance, if people's perception of the importance of an oil price shock depends on the past history of oil prices (Hamilton 1996, 2003), or if firms' cash flows respond differently to positive and negative oil price innovations, then the effect of an oil price shock on stock returns will also depend on the size and the sign of the shock.

There are a number of reasons why oil price shocks could have an asymmetric, and possibly nonlinear, effect on stock returns. First, oil prices do not appear to have an asymmetric effect on aggregate real GDP (Kilian and Vigfusson 2011a, b) and aggregate industrial production (Herrera, Lagalo, and Wada 2011). Yet, they seem to have an asymmetric effect on some (but not all) industries that use energy intensively in their production process such as rubber and plastics, or in consumption such as transportation equipment (Herrera, Lagalo and Wada 2011). On the other hand, some of the industries that exhibit asymmetric response patterns are not energy intensive at all, contradicting conventional explanations of asymmetries. Asymmetries in the response of production could translate into an asymmetric response of profits and, thus, stock returns.

In addition, the optimal decision for a firm that pays dividends to its shareholders and seeks to maximize the expected present value of its dividends (without closing), could be to pay dividends only when its surplus exceeds a threshold (Wan 2007). Therefore, a negative (or a positive) oil price innovation could push the surplus below the cutoff required to pay dividends for an oil company (or an industry that uses energy intensively). If that is the case, the company could choose not to pay dividends and face a decline in stock prices. The negative impact that such a decision would have on stock returns is likely to be larger than the increase in stock returns that would stem from higher dividend payments due to a larger surplus.

Another possibility is that uncertainty and financial stress brought about by the oil price shock, could lead to asymmetries in the response of interest rates (Ferderer 1996; Balke, Brown and Yücel 2002) and, in turn, on the expected present discounted value of the dividends and stock returns. Such an effect would also be evident if people believed the monetary authority will respond differently to oil price increases and decreases. For instance, Ferderer (1996) and Bernanke, Gertler and Watson (1997) find that part of the decline in economic activity brought about by a positive oil price innovation can be attributed to a more restrictive monetary policy. Yet the importance of this systematic monetary policy response –on average and after the Great Moderation– has been questioned (see, for instance, Hamilton and Herrera 2004, Herrera and Pesavento 2009, Kilian and Lewis 2010, Kilian and Vigfusson 2011a). Thus, one would expect the asymmetries embedded in this transmission channel to be negligible.

These arguments merit a careful investigation of the presence of possible asymmetries in the response of stock returns to unexpected variation in crude oil prices –both at the aggregate and disaggregate level. Indeed, if the relationship between stock returns and oil price changes was nonlinear, then a linear VAR model in oil prices and stock returns would be misspecified. This would also be the case for linear models that decompose oil price innovations into oil demand and oil supply shocks (i.e., Kilian and Park 2009). Studying whether the responses of stock returns to unanticipated oil price increases and decreases are symmetric provides a benchmark for thinking about the mechanisms that are commonly thought to amplify the effect of oil price innovations. In particular, proponents of the hypothesis that oil price increases have led to recessions –and lower stock returns– in the U.S. commonly appeal to asymmetries in the transmission of oil price shocks. However, the asymmetry in the response of U.S. stock returns to oil price shocks is yet to be established.

Our contribution to the literature is twofold. First, we explore the question of asymmetry in the response of U.S. real stock returns, both at an aggregate and a disaggregate level. To do so we estimate a dynamic simultaneous equation model that nests symmetric and asymmetric responses to positive and negative oil price innovations using monthly data on aggregate US stock returns and 49 industry-level portfolios. We then employ state-of-the art techniques to directly test the null of symmetry in the response of real stock returns to real oil price innovations (see Kilian and Vigfusson 2011a).

Our estimation results suggest that the response of aggregate stock returns is well captured by a linear model. This is also the case for most of the 49 industry-level portfolios. Yet, there are a small number of portfolios (candy & soda, apparel, healthcare, textiles, aircraft, and insurance) where we find some evidence of asymmetry. Many of these industries are neither energy-intensive in consumption nor in production, thus our evidence should not be necessarily viewed as support for conventional models of asymmetry. Yet, these results imply that financial investors interested in the latter industries should consider asymmetries in the response of stock returns to oil price innovations when forming their portfolios. Similarly, for financial forecasters, innovations of the same magnitude but opposite sign should not enter their loss function in a symmetric manner.

Second, we investigate whether the response of stock returns depends nonlinearly on the size of the shock. To do that, we evaluate whether the test of symmetry leads to different results when we consider innovations of one and two standard deviations. In addition, we explore whether only shocks that exceed a threshold have an asymmetric effect on stock returns as one could conjecture that agents chose to be inattentive to small oil price changes but re-optimize when changes are large. This inattentiveness to small price changes might be justified by the fact that monitoring energy costs is costly, as well as by the presence of adjustment costs in production and consumption (see, e.g., Goldberg 1998; Davis and Kilian 2011). More specifically, we explore a model specification where small oil price changes have different effects than large oil price changes (i.e., changes that exceed one or two standard deviations).

Does the size of the shock matter? The answer to this question is no. First, our results indicate that doubling the size of the shock doubles the size of the response function. In other words, because we cannot reject the null of symmetry, we conclude that a shift in the size of the shock shifts the response function for aggregate stock returns and almost all industry-level portfolios proportionally. In addition, we show that a transformation of the oil price change that filters out small oil price movements that do not exceed one (or two) standard deviation(s) does considerably worse in fitting the data. Our findings are consistent with Edelstein and Kilian (2009) who reject the latter model for U.S. consumption data.

This paper is organized as follows. Section 2 relates the analysis to previous literature. Section 3 describes the data on stock returns and oil prices. Section 4 explores the response of aggregate and industry-level stock returns to oil price innovations. The results of the tests of symmetry in the response to a one standard deviation innovation (hereafter 1 s.d.) are reported in section 5. The following section explores whether our findings are robust to considering larger innovations (2 s.d.) or defining the nonlinear transformation in terms of oil price changes that exceed one or two standard deviations. Section 7 concludes.

2 Relationship with the Related Literature

Although there is a broad body of literature exploring the question of asymmetry in the response of real economic activity to oil price increases and decreases, relatively few studies have addressed the issue of asymmetry for stock returns. Notable exceptions are Sadorsky (1999), and Park and Ratti (2008) who investigate possible nonlinearities in the relationship between aggregate stock returns and oil price shocks. Their approach is different from ours in three dimensions. First, neither Sadorsky (1999) nor Park and Ratti (2008) test for asymmetry in the response functions. Sadorsky (1999) addresses the question of asymmetry by evaluating the contribution to the forecast error variance of stock returns of positive and negative innovations in the level of oil prices. He also investigates the contribution of asymmetric oil price volatility shock to the variance decomposition. Park and Ratti (2008) test the null of symmetry in the slope coefficients but not in the response functions. Until recently, it was common to estimate a censored VAR model and then implement a slope based test to evaluate the presence of asymmetry in the response of economic activity (see, e.g., Park and Ratti 2008). In fact, it was customary to assume that the nonlinear transformation of oil prices was predetermined with respect to the macroeconomic aggregates. Nevertheless, Kilian and Vigfusson (2011a) show that this approach results in inconsistent estimates of the impulse responses, which invalidates their use for quantifying the degree of asymmetry. In contrast, the methodology used in this paper has two advantages: (a) it produces consistent estimates of the response coefficients in the presence of asymmetry; and (b) it enables us to test jointly the empirical implications of the different theoretical models of an asymmetric response discussed in the introduction.

Second, unlike our work, Sadorsky (1999) and Park and Ratti (2008) use only aggregate data on stock returns. Sadorsky (1999) uses data for the U.S., whereas Park and Ratti's (2008) investigation also considers aggregate stock returns for other OECD countries. We restrict ourselves to the U.S. stock market but utilize a longer sample containing data not only on aggregate stock returns but also returns for a large number of industry-level portfolios.

Third, both studies include industrial production and interest rates in their VAR model. Instead, as we will discuss in the next section, we opt for a more parsimonious bi-variate model that is better suited for our purpose of explicitly testing for symmetry in the response of real stock returns. While adding more variables might allow us to better pin down the transmission mechanism from oil price shocks to stock returns (e.g. interest rate channels versus aggregate demand channels), doing so would lower the power of the impulse response based test, stacking the odds against finding any statistical evidence of asymmetry. The reader may wonder whether by excluding these variables from the system we might incur in omitted variable biases. This is not the case. In fact, it can be shown that a higher order VAR model can be written as a bi-variate VAR in which shocks to stock returns would reflect shocks to industrial production and interest rates (see Lütkepohl 2007). Furthermore, estimating a higher dimensional model implies a very high computation cost.[1]

It is also useful to put our results in perspective relative to recent studies regarding the impact of oil price shocks on stock returns. At first sight it may seem that the central question of our paper was already addressed by Kilian and Park (2009). Although they certainly deserve credit for being the first researchers to disentangle and quantify the impact of deeper structural oil shocks on U.S. stock returns, this is not the case.[2] In fact, their approach differs from ours in two important dimensions.

First, to investigate the impact of demand and supply driven oil price shocks on the U.S. stock market, Kilian and Park (2009) estimate a structural VAR model that relates U.S. stock returns to the global oil market. Unlike our model, theirs decomposes unexpected oil price changes into oil supply, aggregate demand and oil-specific demand shocks. The structural VAR they estimate is given by

$$A_0 z_t = \alpha + \sum_{i=1}^{24} A_i z_{t-1} + \varepsilon_t \tag{1}$$

where z_t contains the percentage change in global crude oil production, a measure of real activity in global industrial commodity markets, the real price of oil, and the U.S. stock market variable of interest (e.g., aggregate stock returns or industry-level stock returns), in that order. Then, they use a standard Choleski decomposition to identify the deep structural –oil supply, aggregate demand and oil-specific demand– shocks. In other words, they impose the assumption that there is no contemporaneous feedback

[1] Estimating a higher dimensional semi-structural model while simultaneously controlling for data mining would not be feasible within a year of continuously operating all our computation resources. Note that we are already using a cluster to execute our computations.

[2] Apergis and Miller (2009) extend Kilan and Park's (2009) work to examine the effect of deep structural oil shocks on aggregate stock returns in eight OECD countries.

from stock returns to oil prices. Here we consider a simplified version of their model in the sense that we do not decompose oil price innovations into deep structural shocks. Therefore, our impulse response functions do not have the causal interpretation attached to Kilian and Park's (2009) responses; yet, they represent the expected response to an unexpected oil price shock that consists of a combination of these deep structural shocks.

Second, an important methodological difference between our work and that of Kilian and Park (2009) is that we do not impose the restriction that stock returns are a linear function of past oil price changes. Kilian and Park (2009) do not model, nor test for possible nonlinearity in the relationship between oil prices and stock returns. Instead, we allow for asymmetry in the response of stock returns to unanticipated oil price increases and decreases. Therefore, in our case, Monte Carlo integration is required to compute the impulse response functions, which implies a significantly higher computation cost.

All in all, Kilian and Park's (2009) central question is whether time-variation in the response of U.S. stock returns can be explained by changes in the composition of the deep structural oil shocks and, in turn, whether these changes account for the contrasting findings in the literature regarding the role of oil price shocks on stock returns. In contrast, we focus on the role of unanticipated positive and negative oil price shocks, which constitute an average of the deep structural shocks considered in Kilian and Park (2009). Although our impulse response function may not be given the causal interpretation of Kilian and Park (2009), they provide a benchmark for thinking about different theoretical explanations for the presence of asymmetries in the relationship between oil price shocks and stock returns. As mentioned above, our choice of a more parsimonious model is driven by the fact that introducing additional variables: (a) would lower the power of the symmetry test; and (b) would not be feasible in less than a year of continuous computation using a cluster.

Let us conclude this section by reiterating that evaluating the degree of asymmetry in the response of stock returns to unanticipated oil price changes is crucial because this asymmetry constitutes a leading explanation for the seeming instability in the response of macroeconomic aggregates to oil price shocks. Furthermore, the presence of asymmetry –if established– would have implications regarding which of the transmission mechanisms better fit the data on stock returns.

3 Data Description

We use aggregate and industry-level U.S. real stock returns spanning the period between January 1973 and July 2013. Although data on stock returns and oil prices was available starting January 1947, we restrict the sample to the period between January 1973 and July 2013. This decision is motivated by the fact that oil prices behaved very differently during the years when the Texas Railroad Commission set production limits in the U.S. In fact, it was not until 1972 when U.S. production had increased significantly that nominal oil prices stopped being fixed for long periods of time.[3] Hence, given that we lose one observation at the beginning of the sample when we take first differences, and twelve additional observations given the number of lags considered, our estimation sample starts after the structural break in late 1973.

All of the data on monthly nominal stock returns were obtained from

[3]Estimation results for the full sample are available in the on-line appendix at http://gatton.uky.edu/faculty/herrera/documents/AHappendix.pdf

Kenneth French's database available on his webpage.[4] As a measure of aggregate stock returns we use the excess return on the market, which is defined as the value-weighted return on all NYSE, AMEX, and NASDAQ stocks from the Center for Research in Security Prices (CRSP) minus the one-month Treasury bill rate. For industry level stock returns we use the returns on 49 industry portfolios provided on French's webpage. In this database each NYSE, AMEX, and NASDAQ stock is assigned to an industry portfolio based on its four-digit SIC code as reported by Compustat or, in absence of a Compustat code, by the four-digit SIC classification provided in CRSP. These portfolios include industries in agriculture, mining, construction, manufacturing, transportation and public utilities, wholesale and retail trade, finance, insurance and real estate, and services. (A complete list of the 4-digit SIC industries included in each portfolio is provided in Part B of the on-line appendix.) We then compute real stock returns by taking the log of the nominal stock returns and subtracting the CPI inflation.

Regarding the nominal oil price, we follow the literature (see, for instance Mork 1989, Lee and Ni 2002) and use the composite refiners' acquisition cost (RAC) for crude oil from January 1974 until July 2013. Then, to compute prices for the previous months, we extrapolate using the rate of growth in the producer price index (PPI) for crude petroleum, after making adjustment to account for the price controls of the 1970s. The real price of oil is then computed by deflating the price of oil by the U.S. CPI.

To assess whether oil price innovations have an asymmetric effect on U.S. stock returns, we use three different nonlinear transformations of the real oil price, o_t. The first nonlinear transformation is a modified version of Mork's

[4] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. We use the file containing 49 industry portfolios.

(1989) proposal to split percent changes in oil prices into increases and decreases to allow for an asymmetric response of stock returns to positive and negative oil price shocks. That is, we use the oil price increase, which is defined as:

$$x_t^1 = \max\left(0, \ln o_t - \ln o_{t-1}\right). \tag{2}$$

Alternatively, Hamilton (1996, 2003) suggests that agents might react in a different manner if the oil price increase constitutes a correction for a previous decline and not an increase in a previously stable environment. To account for this behavior, he proposes to use the net oil price increase as a measure of oil price shocks. Thus, as a second nonlinear transformation of oil prices we use the net oil price increase relative to the previous 12-month maximum (Hamilton 1996), which is given by:

$$x_t^{12} = \max\left(0, \ln o_t - \max\left(\ln o_{t-1}, ..., \ln o_{t-12}\right)\right) \tag{3}$$

The last measure is the net oil price increase over the previous 36-month maximum (Hamilton 2003), which is defined in a similar manner:

$$x_t^{36} = \max\left(0, \ln o_t - \max\left(\ln o_{t-1}, ..., \ln o_{t-36}\right)\right) \tag{4}$$

Although, the last two measures do not have a direct grounding on economic theory, there are behavioral explanations as to why agents might react differently in the face of a positive shock if oil prices have been stable in the near past or if they only represent a correction for a previous decline. In fact, the headlines reported in the news often suggest that analysts and stock market commentators consider the behavior of oil prices in the recent past when thinking about the impact of shocks on stock returns.

4 The Effect of Oil Price Shocks on Stock Returns

To evaluate the effect of positive and negative oil price innovations on stock returns we use a simultaneous equation model that nests both symmetric and asymmetric responses of stock returns. In addition, the nonlinear nature of this model allows for small and large oil price innovations to have different effects on the stock market. Thus, consider the data generating process for each of the stock return series, $y_{i,t}$, to be given by the following dynamic simultaneous equation model:

$$x_t = a_{10} + \sum_{j=1}^{p} a_{11,j} x_{t-j} + \sum_{j=1}^{p} a_{12,j} y_{i,t-j} + \varepsilon_{1t} \tag{5a}$$

$$y_{i,t} = a_{20} + \sum_{j=0}^{p} a_{21,j} x_{t-j} + \sum_{j=1}^{p} a_{22,j} y_{i,t-j} + \sum_{j=0}^{p} g_{21,j} x_{t-j}^{\#} + \varepsilon_{2t} \tag{5b}$$

where x_t is the log growth of the crude oil price at time t, $y_{i,t-j}$ is the return on the $i-th$ portfolio at time t, $x_t^{\#}$ is one of the nonlinear transformations of oil prices described in the previous section, and ε_{1t} and ε_{2t} are, by construction, orthogonal disturbances. That is, for identification purposes, we assume that oil price changes are predetermined with respect to U.S. stock returns. This identification strategy is common in the literature on the relationship between oil price shocks and economic activity and amounts to assuming that there is no contemporaneous feedback from U.S. stock returns to oil prices (see for instance, Kilian and Vigfusson 2011a,b; Herrera, Lagalo and Wada 2011; Kilian 2014). Moreover, such assumption is justified by the work of Kilian and Vega (2011) who show that daily U.S. macroeconomic news has no predictive content for oil prices within a month but it does for changes in stock returns. As for the number of lags included in the model, we follow Hamilton and Herrera (2004) and Kilian and Park (2009) in selecting twelve monthly lags, $p = 12$, to capture the effect of oil prices on economic activity.

Note that the inclusion of $x_t^{\#}$ in equation (5b) invalidates the computation of the impulse response functions in the usual textbook manner (see Gallant, Rossi and Tauchen 1993 and Koop, Pesaran and Potter 1996). Instead, to compute the response of stock return i to an innovation of size δ in ε_{1t} we use Monte Carlo integration as in Kilian and Vigfusson (2011a). That is, we first calculate the impulse response functions to a positive innovation, $I_y(h, \delta, \Omega_t)$, and to a negative innovation, $I_y(h, -\delta, \Omega_t)$ of size δ –conditional on the history Ω_t– for $h = 0, 1, 2, ..., 12$. We perform this computation for 1,000 different histories and then calculate the unconditional impulse response functions, $I_y(h, -\delta)$, by averaging over all the histories.[5]

Figure 1 about here

The first panel of Figure 1 illustrates the response of aggregate stock returns to positive and negative innovations of one standard deviation in the real oil price. For ease of comparison, we report the response to a positive innovation and the negative of the response to a negative innovation of size $\delta = 1$ s.d. Note that, regardless of the oil price measure, the effect of a 1 s.d. innovation in oil prices has a statistically insignificant effect on aggregate stock returns in the short-run. Using the oil price increase (the net oil price increase relative to the previous 36 months) the response of aggregate stock returns to both positive and negative innovations becomes significant at the 5% level 8 months (8 and 12 months) after the shock. In both cases, an unexpected increase in real oil prices leads to a decline in U.S. aggregate

[5] See Herrera, Lagalo and Wada (2011) for a detailed description of the computation.

stock returns of less than 1%, whereas an unexpected decrease causes an increase of about the same magnitude. At a first sight, the fact that the $IRFs$ to positive and negative innovations lie almost on top of each other suggests no asymmetry is present in the response of aggregate stock returns.

The remaining panels of Figure 1 plot the response of stock returns for a group of portfolios that are thought to be affected by oil prices (see Kilian and Park 2009). To economize space, the impulse response functions for the remaining industry-level portfolios are plotted in Figures A.1a-A.1f of the on-line appendix. Evidence of a negative relationship between positive oil price innovations and real stock returns at the industry-level, for at least two of the oil price measures, is apparent for entertainment, consumer goods, chemicals, rubber and plastic products, steel works, automobiles and trucks, aircraft, shipbuilding and railroad equipment, utilities, retail, restaurants, hotels and motels, agriculture, candy & soda, beer & liquor, tobacco products, printing & publishing, apparel, medical equipment, textiles, construction materials, construction, personal services, computer hardware, electric equipment, business supplies, shipping containers, banking, insurance, real estate, and trading. For most of these portfolios, the responses to positive and negative innovations of 1 s.d. lie on top of each other. This suggests that a negative innovation of 1 s.d. would have a positive effect on stock returns of the same magnitude –but opposite sign– than a 1 s.d. positive innovation. These results are in line with Kling (1985), and Jones and Kaul (1996), who find a negative impact of oil price shocks on stock returns. Notwithstanding important differences in the sample period and model specification, our empirical results are fully consistent with Kilian and Park's (2009) finding of a negative impact of oil prices on stock returns. Furthermore, note that we find a statistically significant effect of oil price changes, even though we do not account for the source of the shock.

5 Does the Sign of the Shock Matter?

Recent research into the question of asymmetry in the response of economic activity to positive and negative oil price innovations suggests that the magnitude of the effect of a positive innovation is not larger (in absolute terms) than the magnitude of the effect of a negative innovation. Is this also the case for the response of U.S. stock returns? We address this question by implementing Kilian and Vigfusson's (2011a) impulse response based test. That is, we use the impulse response functions computed in the previous section to construct a Wald test of the null hypothesis:

$$I_y(h, \delta) = -I_y(h, -\delta) \text{ for } h = 0, 1, 2, ..., 12.$$

Note that this test jointly evaluates whether the response of stock returns (for a particular portfolio) to a positive shock of size δ equals the negative of the response to a negative shock of the same size, $-\delta$, for horizons $h = 0, 1, 2, ..., 12$. Our motivation for focusing on a one-year horizon is twofold. First, the extant literature on the effect of oil price shocks has found that the largest and most significant impact on economic activity takes place around a year after the shock (see, for instance, Hamilton and Herrera 2004). Therefore, one could conjecture a similar lag in the transmission of oil price shocks to dividends, and thus to stock returns. But, even in the case where financial investors rapidly incorporate the information regarding oil price changes in their expected dividends, since the Wald test is a joint test for horizons $h = 0, 1, 2, ..., 12$, we take into account the response at shorter horizons.

Second, by focusing on the 12-months horizon we mitigate issues of data mining related to repeating the test over a different number of horizons. That is, if we were to repeat the impulse response based test with a 5% size say for 6 different horizons H, then the probability of finding at least one rejection would exceed 5% under the null.

Having addressed the possible issue of data mining across horizons by focusing on $H = 12$, we still have to tackle data mining concerns related to repeating the impulse response based test over 49 different portfolios. To avoid this potential problem, we compute data-mining robust critical values by simulating the distribution of the supremum of the bootstrap test statistic, under the null, across all portfolios for each of the oil price transformations.[6] To compute the data mining robust critical values we generate 100 pseudo-series using the estimated coefficient for the 49 portfolios in model (5). We then use 100 histories to get the conditional impulse response functions for each pseudo-series and compute the $IRFs$ by Monte Carlo integration. We repeat this procedure 100 times to obtain the empirical distribution of the test statistic.

Before we review the test results it is worth noting that the χ^2_{H+1} distribution becomes a less accurate approximation under the null as the number of restrictions increases. For instance, Kilian and Vigfusson (2011a) results suggest that at $H = 12$ the actual size of the symmetry test is about twice the nominal size. Thus, it is not surprising that the χ^2_{H+1} critical values exceed in a number of instances the bootstrapped critical values. For this reason, and because the bootstrap critical values control for data mining, we rely heavily on the data mining robust critical values to conduct our

[6] See Inoue and Kilian(2004) and Kilian and Vega (2011) for the effect of data mining and solutions to the problem of data mining in the related context of tests of predictability.

inference.

The left panel of Table 1 reports the $p-values$ for the test of symmetry in the response to positive and negative innovations of 1 s.d. in the real oil price. In addition, we denote significance at the 5% and 10% level, after controlling for data mining, by ** and *, respectively. As the 'eyeball metric' would have suggested when looking at Figure 1, there is no evidence of asymmetry in the response of aggregate stock returns. Regardless of the oil price transformation ($x_t^{\#} = x_t^1, x_t^{12}, x_t^{36}$), we are unable to reject the null at a 5% level. As for the industry-level portfolios, we find some evidence of asymmetry when we use the net oil price increase relative to the previous 36-month maximum, x_t^{36}. In particular, we reject the null at a 5% significance level for candy & soda, healthcare, and aircraft, and we reject the null at a 10% level for apparel, and insurance. Interestingly, we fail to reject the null for all industry-level portfolios when we use the oil price increase, x_t^1, or net oil price increase with respect to the previous 12-month maximum, x_t^{12}.

TABLE 1 about here

Finding asymmetries in the response of aircraft, or apparel might not be surprising to the reader, as the use of transportation equipment requires considerable amounts of refined products and apparel is somewhat energy intensive in production (see Table 2). Thus, a-priori, one could anticipate the demand for these goods to contract more in response to positive oil price innovations than it would expand when faced by negative innovations. After all, firms might postpone the purchases of planes when hit by an unexpected oil price surge, but they might not increase their demand when faced by an unexpected price drop. As a consequence, one would expect the response of profits, and thus stock returns, to be asymmetric. Yet, by the same token, we would expect to find asymmetries in automobiles and truck, but none is evident. On the contrary, evidence of asymmetry in candy & soda as well as in insurance might be more puzzling as the total (direct and indirect) cost of crude petroleum and natural gas used to produce a dollar of output in these industries is less than 4 cents (see Table 2). A possible explanation for this finding could be that consumers increase precautionary savings when faced with a positive shock (Edelstein and Kilian 2009), reduce the demand for these goods, and this shortfall in demand leads to lower expected dividends and stock returns. However, Edelstein and Kilian (2009) found no evidence of asymmetry in aggregate consumer expenditures.

TABLE 2 about here

It is interesting to compare our results with those obtained by Herrera, Lagalo and Wada (2011) who study the question of asymmetry in the response of industrial production, as such a comparison could shed some light on the source of the asymmetry in stock returns. Using data mining robust critical values, they fail to reject the null of symmetry for $H = 12$ for the total industrial production index, as well as for all the industry-level indices, when using x_t^1 and x_t^{12}. Instead, they find evidence of asymmetry in transit equipment, petroleum and coal, plastics and rubber, and machinery, when using x_t^{36}. In brief, there is no correspondence between our results and those for industrial production, which suggests that asymmetries in the response of industry-level stock returns are not driven by asymmetries in the response of production. Instead, other transmission mechanisms are at play. In particular, the story driving the asymmetries in stock returns appears to be one in which increases in oil prices have a greater contractionary effect in the expected and/or realized demand for aircraft, apparel, candy & soda, healthcare, and insurance, than the expansionary effect of an oil price decline. Hence the asymmetry in the response of these stock returns. Furthermore, our inability to reject the null of symmetry in aggregate returns rules out the monetary transmission story. Recall that any individual stock return is a function of expected discounted dividends where the discount rate is an average of interest rates overtime; thus, changes in the level and term structure of interest rates affect all stock returns. Therefore, any nonlinearity in the response of interest rates to oil price shocks should be translated into a nonlinear response of individual and aggregate stock returns. Yet, that is not the case.

Does the sign of the shock matter? For aggregate stock returns, the answer is only to the extent that the response has the opposite sign; yet, positive and negative innovations have a symmetric effect. For most industry-level portfolio returns, we find no evidence that positive innovations have a larger impact than negative innovations up to a year after the shock. Yet, there are a few industries where the sign of the shock matters in that the response of real stock returns is asymmetric.

6 Does the Size of the Shock Matter?

In a linear model, the magnitude of the response to a 2 s.d. shock is simply twice that of the response to a 1 s.d. shock. Nevertheless, in a nonlinear model such as that in (5) the magnitude of the response depends on the size of the shock, and on the history of oil price changes and stock returns. Thus, we estimate the $IRFs$ to 2 s.d. innovations and test for symmetry in the response to positive and negative innovations of this magnitude, as we did in the previous section. Figure 2 plots the $IRFs$ to a positive 2 s.d. and the negative of the $IRFs$ to a negative innovation of the same size.[7] The second panel of Table 1 reports the $p-values$ for the test of symmetry in the response to a 2 s.d. innovation.

Figure 2 about here

At first glance, it would appear that doubling the size of the innovation increases the evidence in favor of asymmetry (see Figure 2 and the second panel of Table 1). Note how there are more $p - values$ below 5%, which are marked in bold, for a 2 s.d. innovation than for a 1 s.d. innovation. Yet, this is not the case when we control for data mining. In fact, using x_t^1 and x_t^{36} we are unable to reject the null for the aggregate and all of the industry-level portfolios. For x_t^{12} we find evidence of asymmetry in textiles. The difference between the test results before and after controlling for data mining is indicative of the higher degree of uncertainty associated with the estimation of the $IRFs$ to a 2 s.d. innovation. Moreover, since our data mining robust critical values are computed using the supremum of the bootstrap test statistic across all industry-level portfolios, it would suffice for the $IRFs$ to be estimated with a higher degree of uncertainty for one portfolio in order to get larger critical values.

To further evaluate whether the size of the oil price shock matters, we employ a different oil price transformation along the lines of Edelstein and Kilian (2007). Consider a situation in which firms and individuals only respond to shocks that exceed a certain threshold. Such behavior could be observed if there are adjustment costs that prevent agents from optimizing

[7] $IRFs$ for the remaining portfolios are illustrated in Figures A.2a-A.2f of the on-line appendix.

when the change in the price of an input or a consumption good is small, or if dividends are paid only if the surplus exceeds a threshold.

Let us define

$$x_t^{sd} = \left\{ \begin{array}{ll} 0 & \text{if} \quad |x_t| \leq \delta \\ x_t & \text{if} \quad |x_t| > \delta \end{array} \right\} \tag{6}$$

where x_t is the percentage change in the oil price, and δ equals one (6.83%) or two (13.66%) standard deviations of the oil price change.

The fourth column in the left and right panels of Table 1 report the $p - values$ for the test of symmetry computed using this alternative transformation of the oil price change. Clearly, there is no evidence of asymmetry in the response to 1 s.d. or 2 s.d. innovations when we use x_t^{sd}. In fact, our estimates suggest that x_t^{sd} does a very bad job at capturing possible asymmetries in the response of U.S. stock returns.

All in all, the response of aggregate stock returns to innovations in the real oil price, as well as that of most industry-level portfolios, is well captured by a linear model. Hence, the impact of innovations that differ only in size should differ only in the same scale. Yet, for a number of industries such as candy&soda, healthcare, aircraft, apparel, and insurance, the magnitude of the shock matters as the response is a nonlinear function of the innovation.

7 Conclusions

We started our study by inquiring whether the size and the sign of an oil price shock matter for the response of U.S. real stock returns. To answer these questions we estimated a simultaneous equation model that nests symmetric and asymmetric responses to positive and negative innovations in the price of crude oil. We found that positive oil price innovations depress aggregate stock returns, as well as the returns of about 60% of the industry-level portfolios.

We explored the question of asymmetry in the response of real stock returns by implementing Kilian and Vigfusson's (2011a) impulse response based test. To avoid issues of data mining related to the repetition of the test over all the portfolios, we bootstrapped the distribution of the supremum of the Wald test across all portfolios. Estimation results suggested that a linear model fits the data well for aggregate returns, as well as for most industry-level portfolios. Notable exceptions are candy & soda, healthcare, and aircraft for which we found evidence of asymmetry in the response to a 1 s.d. innovation using the net oil price increase relative to the previous 36-month maximum, x_t^{36}. No evidence of asymmetry is found when we use the oil price increase, x_t^1, or the net oil price increase relative to the 12-month maximum, x_t^{12}. Consistent with these findings, we concluded that, for aggregate stock returns and for most industry-level portfolios, the sign of the shock mattered only in that it determined the sign of the response. Yet, the absolute magnitude of the responses coincided.

To investigate whether the size of the shock matters we explored the question of symmetry in the response to a 2 s.d. innovation. For this larger shock, evidence of asymmetry was absent for all portfolios but textiles when we used x_t^{12}. We then explored the conjecture that only oil price innovations that exceed a threshold (1 s.d. or 2 s.d. of the percentage change in the real oil price) have an asymmetric effect on real stock returns. Our estimation results lead us to strongly reject such a model. We thus concluded that the size of an innovation in real oil prices only matters in that it determines the scale of the effect. That is, consistent with our finding of symmetry, a doubling in the size of the innovation in real oil prices leads to a doubling (no more, no less) in the response of almost all analyzed stock returns.

Comparing the test results across different oil price measures reveals only a few rejections when we use the net oil price increase relative to the previous 36-month maximum, x_t^{36}, and a typical shock. Yet, it appears that the net oil price increase relative to the previous 12-month maximum, x_t^{12}, works best at capturing possible asymmetries in the oil price-stock returns relationship for large shocks. These results are suggestive of a behavioral driven story as a possible theoretical explanation for the asymmetric effect of positive and negative oil price innovations. On the one hand, our finding of symmetry in the response of aggregate stock returns rules out monetary policy as the source of the asymmetry. In addition, the fact that some of the industries that exhibit asymmetric response patterns are not energy intensive at all suggests that mechanisms other than the transmission through a production channel are at play. On the other hand, our finding of asymmetries in aircraft, apparel, candy & soda, and insurance suggests that an asymmetric response of the demand for these goods might be the source of these patterns. Now, whether changes in demand are driven because consumers perceive oil prices to be higher than in the recent past (i.e., the behavioral explanation) or whether the decline in demand is driven by increased precautionary savings or shifts in demand, is an issue that should be explored by future research.

Do sign and size matter? The answer to this question appears to be that sign only matters in that it determines the direction of the effect on stock returns but –for the aggregate and most industry-level returns– it has no impact on the magnitude of the response. As for the size of the shock, given that in most cases the responses are symmetric, the effect of a two standard deviation shock is just twice the effect of a one standard deviation shock. In brief, our results suggest a linear model provides a good approximation to the response of real stock returns to real oil price innovations.

References

[1] Apergis, N. and S.M. Miller (2009). "Do Structural Oil-Market Shocks Affect Stock Prices? *Energy Economics* 31(4): 569-575.

[2] Balke, N., S.P.A. Brown and M.K. Yücel (2002). "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" *Energy Journal* 23(3): 27-52.

[3] Bernanke, B. S., M. Gertler, and M. Watson (1997). "Systematic Monetary Policy and the Effects of Oil Price Shocks." *Brookings Papers on Economic Activity* 1: 91-157.

[4] Chen, N-F., R. Roll and S.A. Ross (1986). "Economic Forces and the Stock Market." *The Journal of Business* 59(3): 383-403.

[5] Davis, L.W. and L.Kilian (2011). "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 16: 1187-1214.

[6] Edelstein, P. and L. Kilian (2007). "The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses about the Transmission of Energy Price Shocks." *The B.E. Journal of Macroeconomics* 7(1) (Contributions). Available at: http://www.bepress.com/bejm/vol7/iss1/art35.

[7] Edelstein, P. and L. Kilian (2009). "How Sensitive are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics* 56(6): 766-779.

[8] Ferderer, J.P. (1996). "Oil Price Volatility and the Macroeconomy: A Solution to the Asymmetry Puzzle." *Journal of Macroeconomics* 18: 1-16.

[9] Gallant, A. Ronald, P.E. Rossi, and G. Tauchen (1993). "Nonlinear Dynamic Structures." *Econometrica* 61, 871-908.

[10] Goldberg, P.K. (1998). "The Effects of the Corporate Average Fuel Economy Standards in the U.S." *Journal of Industrial Economics* 46(1): 1-33.

[11] Hamilton, J.D. (1996). "This is What Happened to the Oil Price-Macroeconomy Relationship." *Journal of Monetary Economics* 38(2): 215-220.

[12] Hamilton, J.D. (2003). "What Is an Oil Shock?" *Journal of Econometrics* 113(2): 363-398.

[13] Hamilton, J.D. and A.M. Herrera (2004). "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy." *Journal of Money, Credit, and Banking:* 36(2).

[14] Herrera, A.M. and E. Pesavento (2009). "Oil Price Shocks, Systematic Monetary Policy, and the 'Great Moderation'." *Macroeconomic Dynamics* 13 (1): 107-137.

[15] Herrera, A.M., L.G. Lagalo, and T.Wada (2011). "Oil Price Shocks and Industrial Production: Is the Relationship Linear?" *Macroeconomic Dynamics* 15 (S3): 472-497.

[16] Huang, R., R. Masulis, and H. Stoll (1996). "Energy shocks and financial markets", *Journal of Futures Markets,* 16, 1–27.

[17] Inoue, A. and L. Kilian (2004). "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?" *Econometric Reviews* 23(4): 371-402.

[18] Jones, C.M. and G.Kaul (1996). "Oil and Stock Markets." *Journal of Finance* 51: 463–491.

[19] Kilian, L. (2014). "Oil Price Shocks: Causes and Consequences." *Annual Review of Resource Economics* (forthcoming).

[20] Kilian, L. and C. Park (2009). "The Impact of Oil Price Shocks on the U.S. Stock Market." *International Economic Review* 50(4): 1267-1287.

[21] Kilian, L. and C. Vega (2011). "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices." *Review of Economics and Statistics* 93(2): 660-671.

[22] Kilian, L. and L.T. Lewis (2010). "Does the Fed Respond to Oil Price Shocks?" *Economic Journal* 121(555): 1047-1072.

[23] Kilian, L. and R. J. Vigfusson (2011a). "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?" *Quantitative Economics* 2(3): 419-453.

[24] Kilian, L. and R.J. Vigfusson (2011b). "Nonlinearities in the Oil Price-Output Relationship." *Macroeconomic Dynamics* 15(S3): 337-363.

[25] Kling, J.L. (1985). "Oil Price Shocks and Stock-Market Behavior," *Journal of Portfolio Management*, 12, 34-39.

[26] Koop, G., M.H. Pesaran, and S.M. Potter (1996). "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74: 19-147.

[27] Lee, K., and S. Ni (2002). "On the Dynamic Effects of Oil Price Shocks: A Study Using Industry Level Data." *Journal of Monetary Economics* 49: 823-852.

[28] Lütkepohl, H. (2007). *New Introduction to Multiple Time Series Analysis*, Springer, Heidelberg. pp. 435-438.

[29] Mork, K.A. (1989). "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results", *Journal of Political Economy* 97(3): 740-744.

[30] New York Times. "U.S. Stocks Plunge after Oil Climbs $6", June 11, 2008.

[31] Park, J. and R.A. Ratti (2008). "Oil Price Shocks and Stock Markets in the U.S. and 13 European Countries." *Energy Economics* 30(5): 2587-2608.

[32] Sadorsky, P. (1999). "Oil Price Shocks and Stock Market Activity." Energy Economics 21(5): 449-469.

[33] Wall Street Journal. "U.S. Stocks Rally after Crude Drops to 3-month Low", August 8, 2008.

[34] Wan, N. (2007). "Dividend payments with a threshold strategy in the compound Poisson risk model perturbed by diffusion." *Insurance: Mathematics and Economics* 40(3): 509-523.

[35] Wei, C. (2003). "Energy, the Stock Market, and the Putty-Clay Investment Model." *American Economic Review* 93: 311–23.

Table 1. Test of symmetry in the response to positive and negative innovations in the real oil price for H=12

	1 s.d.				2 s.d.			
Sector	$x_t^{\#} = x_t^1$	$x_t^{\#} = x_t^{12}$	$x_t^{\#} = x_t^{36}$	$x_t^{\#} = x_t^{sd}$	$x_t^{\#} = x_t^1$	$x_t^{\#} = x_t^{12}$	$x_t^{\#} = x_t^{36}$	$x_t^{\#} = x_t^{sd}$
Aggregate	0.62	0.94	0.87	1.00	0.35	0.42	0.62	1.00
Agriculture	0.88	0.96	0.90	1.00	0.86	0.66	0.94	1.00
Food Products	0.93	0.67	0.94	0.84	0.90	0.23	0.88	1.00
Candy & Soda	0.47	0.69	0.25**	1.00	0.29	**0.01**	0.71	1.00
Beer & Liquor	0.69	0.77	0.97	0.88	0.27	0.15	0.93	1.00
Tobacco Products	0.58	0.93	0.91	0.99	0.42	0.20	0.76	1.00
Recreation	0.69	0.91	0.83	0.97	0.55	0.52	0.90	1.00
Entertainment	0.46	0.55	0.74	0.99	**0.03**	**0.01**	0.47	1.00
Printing and Publishing	0.86	0.90	0.75	1.00	0.56	0.43	0.64	1.00
Consumer Goods	0.62	0.75	0.84	0.80	0.41	*0.06*	0.90	1.00
Apparel	0.30	0.61	0.53*	0.94	**0.01**	**0.02**	0.53	1.00
Healthcare	0.96	0.83	0.36**	0.67	0.91	0.12	0.41	1.00
Medical equipment	0.42	0.76	0.92	1.00	0.48	*0.09*	0.84	1.00
Pharmaceutical products	0.73	0.71	0.95	0.92	0.57	0.15	0.97	1.00
Chemicals	0.73	0.77	0.66	1.00	0.48	*0.09*	0.84	1.00
Rubber and plastic products	0.75	0.76	0.78	1.00	0.50	0.34	0.51	1.00
Textiles	0.25	0.31	0.75	1.00	**0.04**	**0.00****	0.42	1.00
Construction materials	0.53	0.89	0.85	0.99	0.36	0.23	0.93	1.00
Construction	0.97	0.98	0.80	1.00	0.94	0.66	0.46	1.00
Stee works etc.	0.79	0.98	0.84	1.00	0.38	0.45	0.74	1.00
Fabricated products	0.85	0.87	0.85	1.00	0.79	0.25	0.85	1.00
Machinery	0.58	0.93	0.88	1.00	0.41	0.13	0.69	1.00
Electrical equipment	0.77	0.93	0.73	0.99	0.71	0.29	0.73	1.00
Automobiles and trucks	0.31	0.92	0.82	0.97	**0.02**	0.14	0.47	1.00
Aircraft	0.64	0.85	0.35**	0.99	0.35	0.17	0.42	1.00
Shipbuilding, railroad equip.	0.91	0.97	0.69	1.00	0.88	0.84	0.83	1.00
Defense	0.74	0.66	0.85	0.93	0.33	**0.02**	0.61	1.00
Precious metals	0.97	0.74	0.65	0.81	0.96	**0.04**	0.70	1.00
Mines	0.94	0.99	0.98	1.00	0.91	0.83	0.93	1.00
Coal	0.95	0.91	0.68	1.00	0.88	0.18	0.44	1.00
Petroleum and natural gas	0.98	0.97	0.82	1.00	0.97	0.36	0.55	1.00
Utilities	0.87	0.91	0.73	1.00	0.70	0.27	0.22	1.00
Communication	0.52	0.64	0.96	0.96	0.15	*0.07*	0.72	1.00
Personal services	0.49	0.89	0.63	1.00	0.70	0.27	0.22	1.00
Business services	0.85	0.85	0.98	1.00	0.64	0.36	0.96	1.00
Computer hardware	0.15	0.91	0.80	1.00	**0.03**	*0.09*	0.63	1.00
Computer software	0.53	0.95	0.93	0.97	0.34	0.26	0.95	1.00
Electronic equipment	0.29	0.82	0.87	0.99	0.22	0.14	0.86	1.00
Measuring and control equip.	0.72	0.96	0.97	1.00	0.69	0.50	0.97	1.00
Business supplies	0.86	0.84	0.66	0.98	0.69	0.32	0.67	1.00
Shipping containers	0.91	0.77	0.80	1.00	0.89	*0.07*	0.80	1.00
Transportation	0.70	0.74	0.60	0.98	0.36	0.11	0.61	1.00
Wholesale	0.92	0.88	0.75	1.00	0.84	0.44	0.80	1.00
Retail	0.23	0.67	0.64	1.00	**0.00**	0.12	0.75	1.00
Restaurants, hotels, motels	0.53	0.71	0.77	0.75	0.11	*0.08*	0.71	1.00
Banking	0.45	0.51	0.68	0.85	*0.08*	**0.01**	0.24	1.00
Insurance	0.83	0.53	0.43*	0.90	0.47	**0.01**	0.29	1.00
Real estate	0.69	0.54	0.70	1.00	0.22	**0.01**	0.33	1.00
Trading	0.63	0.87	0.96	0.99	0.22	*0.08*	0.87	1.00
Other	0.39	0.71	0.77	0.94	*0.08*	**0.02**	0.54	1.00

Notes: based on 1000 simulations of model (5). p-values are based on the χ^2_{H+1}. Bold and italics denote significance at 5% and 10% level, respectively. ** and * denote significance at the 5% and 10%, respectively, after accounting for data mining.

Table 2. Direct and total requirements of crude petroleum and natural gas

	Direct cost		Total cost	
Industry	1977	1999	1977	1999
Food and kindred products	0.000	0.000	0.036	0.023
Tobacco products	0.000	0.000	0.025	0.006
Apparel	0.000	0.000	0.033	0.020
Textiles	0.000	0.000	0.057	0.033
Paper and allied products, except containers (Business supplies)	0.000	0.000	0.062	0.031
Chemicals and selected chemical products	0.025	0.097	0.187	0.268
Petroleum refining and related industries	0.607	0.484	0.720	0.785
Plastics and synthetic materials	0.016	0.006	0.118	0.075
Rubber and miscellaneous plastics products	0.001	0.000	0.059	0.030
Primary iron and steel manufacturing	0.000	0.000	0.047	0.031
Lumber and wood products	0.000	0.000	0.064	0.017
Computer and office equipment	0.000	0.000	0.021	0.010
Household appliances	0.000	0.000	0.028	0.020
Motor vehicles and equipment	0.000	0.000	0.027	0.034
Transportation and warehousing	0.001	0.003	0.067	0.177
Wholesale and retail trade	0.000	0.000	0.020	0.018
Finance and insurance	0.000	0.000	0.010	0.009
Private electric, gas, water, and sanitary services (Utilities)	0.125	0.466	0.237	0.876
Amusements	0.000	0.000	0.022	0.007

This table reports, as a measure of energy-intensity, total and direct costs of crude petroleum and natural gas required to produce a dollar of output of the particular industry in 1977 and 1999. These requirements are computed using the 1977 and 1999 annual Input-Output tables published by the BEA.

Figure 1a: Response to one standard deviation positive and negative innovation in the real oil price change

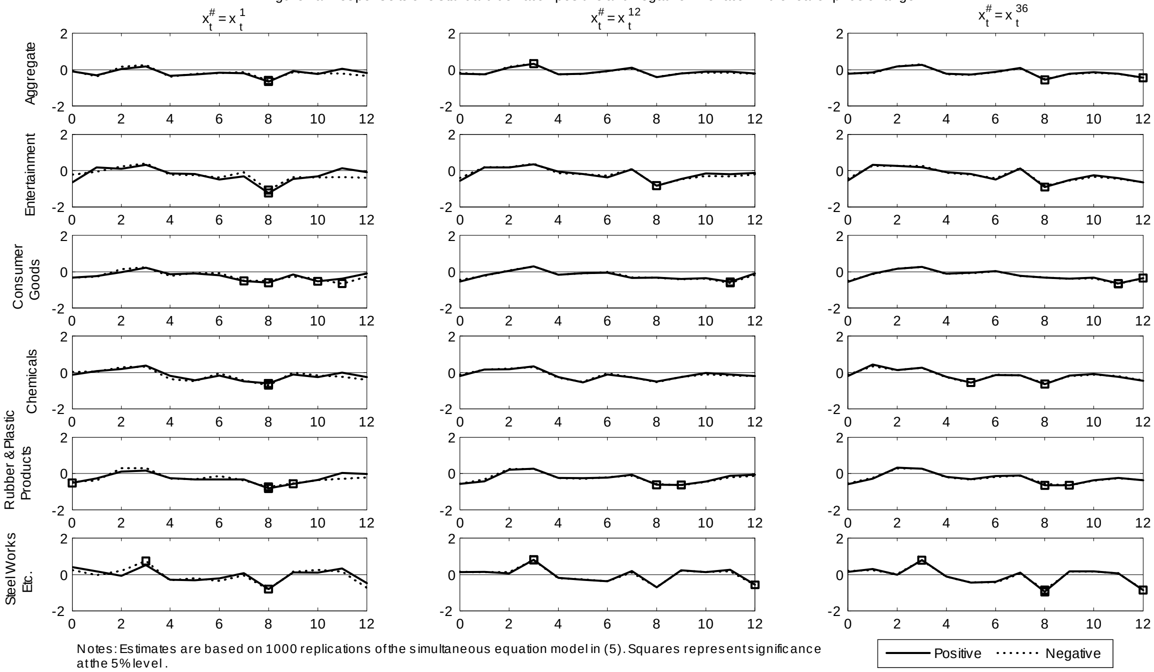

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.

Figure 1b: Response to one standard deviation positive and negative innovation in the real oil price change

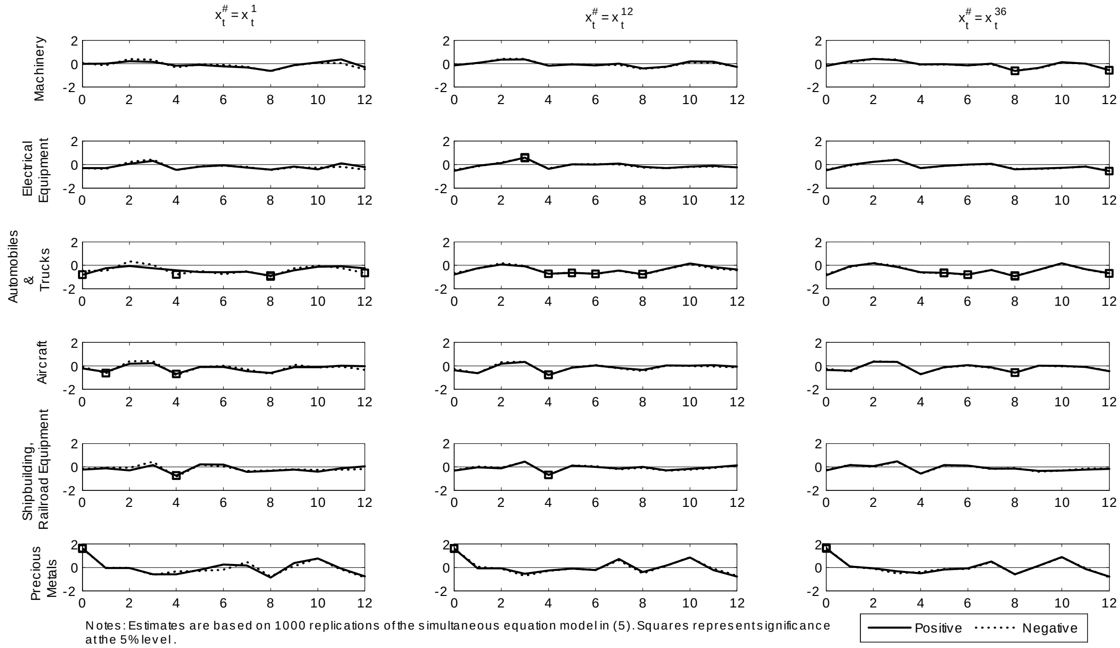

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.

Figure 1c: Response to one standard deviation positive and negative innovation in the real oil price change

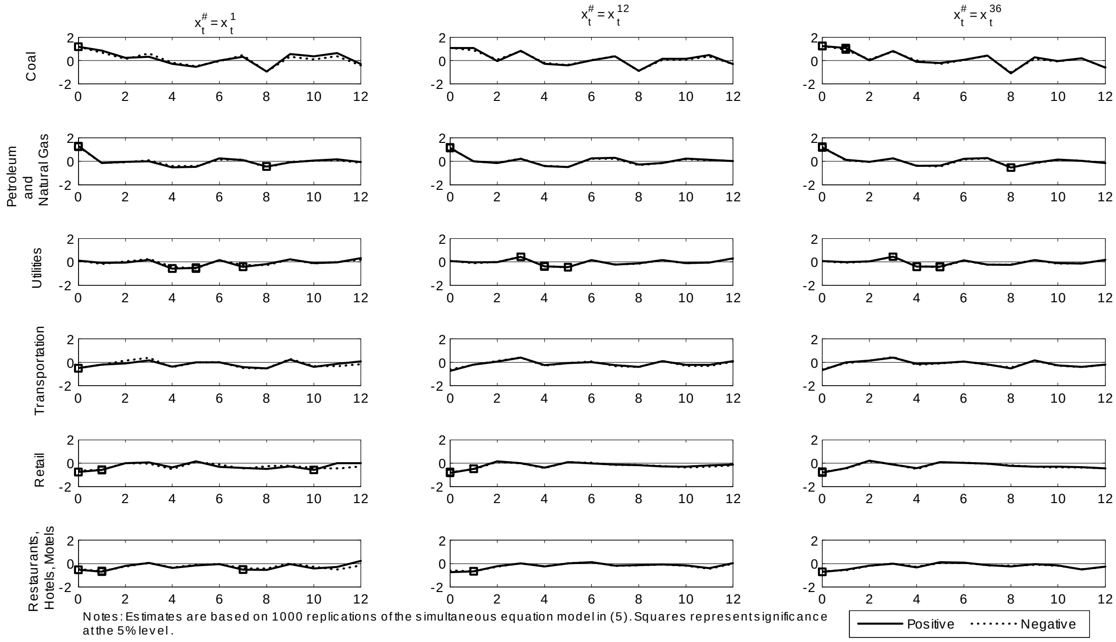

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.

Figure 2a: Response to two standard deviation positive and negative innovation in the real oil price change

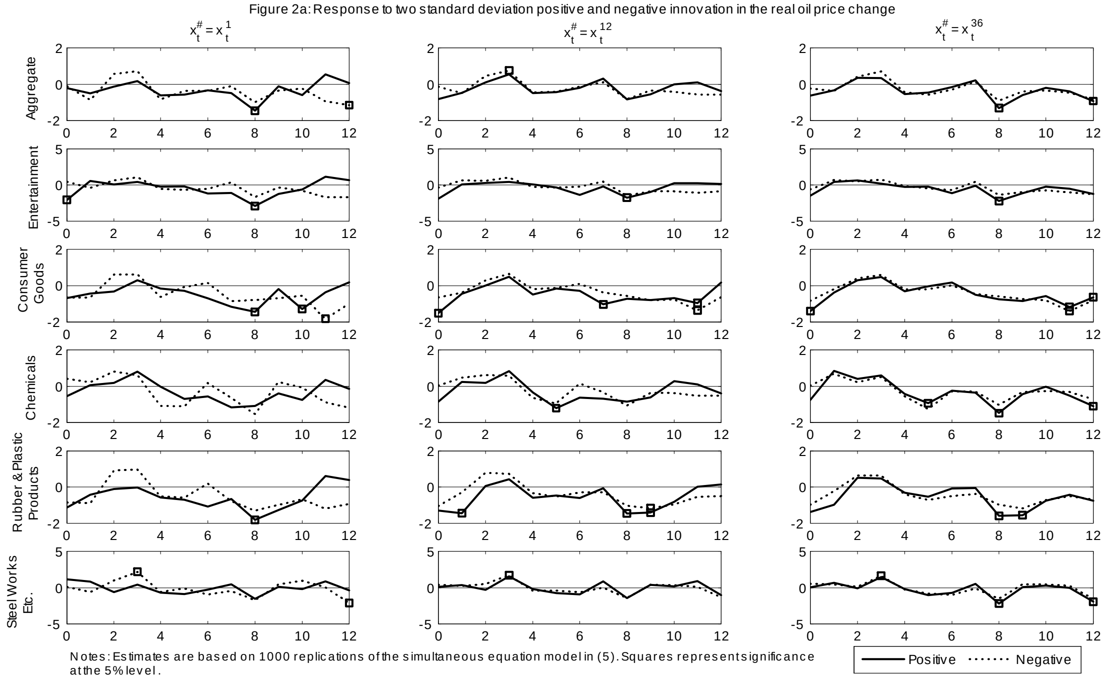

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.

Figure 2b: Response to two standard deviation positive and negative innovation in the real oil price change

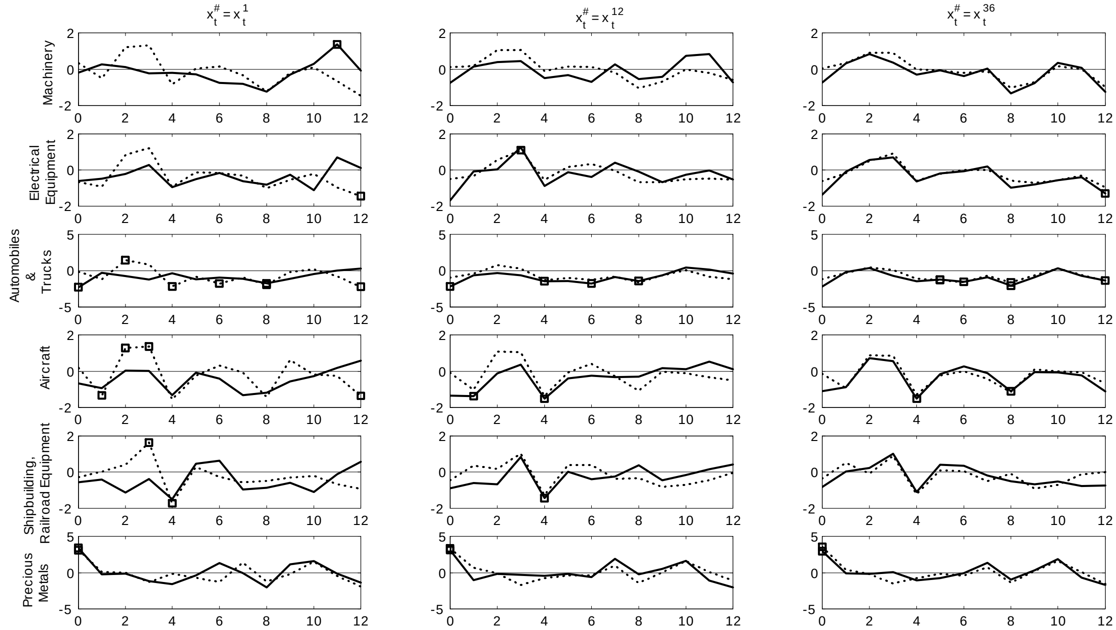

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.

Figure 2c: Response to two standard deviation positive and negative innovation in the real oil price change

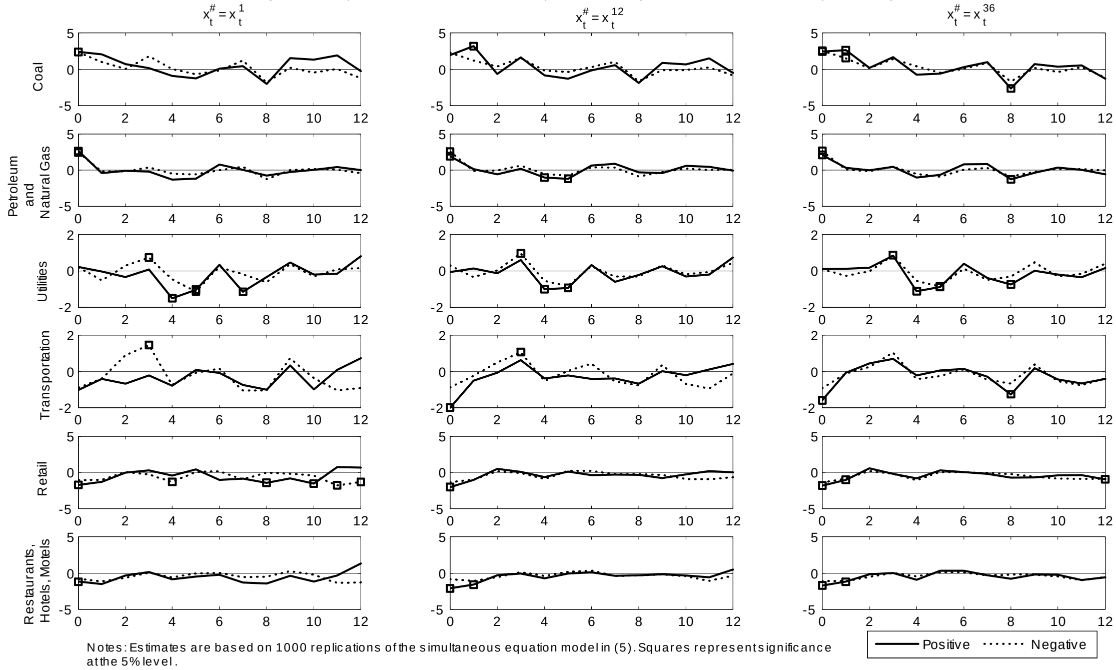

Notes: Estimates are based on 1000 replications of the simultaneous equation model in (5). Squares represent significance at the 5% level.