

Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?

Lutz Kilian
University of Michigan and CEPR

Robert J. Vigfusson
Federal Reserve Board

August 13, 2009

Abstract: How much does real GDP respond to unanticipated changes in the real price of oil? Commonly used censored VAR models suggest a substantial decline in real GDP in response to unexpected increases in the real price of oil, yet no response to unexpected declines. We show that these estimates are invalid. Based on a structural model that encompasses both symmetric and asymmetric models as special cases, correctly computed impulse responses are of roughly the same magnitude in either direction, consistent with formal tests for symmetric responses. We discuss implications for theoretical models and for policy responses to energy price shocks.

Key Words: Asymmetry; oil price; energy prices; net increase; shocks; propagation; transmission; vector autoregression.

JEL Classification: C32, E37, Q43

Acknowledgments: The corresponding author is Lutz Kilian, University of Michigan, Department of Economics, 611 Tappan Street, Ann Arbor, MI 48109-1220, USA. Email: lkilian@umich.edu. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. We thank Ana-María Herrera, James Hamilton, and seminar participants at the Federal Reserve Board, the 2009 NBER Summer Institute, the Federal Reserve Bank of New York, the 2009 Joint Statistical Meetings, the University of Michigan, Michigan State University, the London School of Economics, Ente Einaudi, Frankfurt University, George Washington University, the International Monetary Fund, Texas A&M, and the University of Houston for helpful comments.

Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?

Abstract: How much does real GDP respond to unanticipated changes in the real price of oil? Commonly used censored VAR models suggest a substantial decline in real GDP in response to unexpected increases in the real price of oil, yet no response to unexpected declines. We show that these estimates are invalid. Based on a structural model that encompasses both symmetric and asymmetric models as special cases, correctly computed impulse responses are of roughly the same magnitude in either direction, consistent with formal tests for symmetric responses. We discuss implications for theoretical models and for policy responses to energy price shocks.

Key Words: Asymmetry; oil price; energy prices; net increase; shocks; propagation; transmission; vector autoregression.

JEL Classification: C32, E37, Q43

1. Introduction

A common view in the literature is that the effects of energy price shocks on macroeconomic aggregates such as output or employment are asymmetric. In particular, energy price increases are perceived to have larger effects than energy price decreases. This perception has been bolstered by empirical evidence that energy price increases (obtained by censoring energy price changes to exclude all energy price decreases) have seemingly large effects on the macro economy, whereas uncensored percent changes in energy prices tend to have much smaller effects. Vector autoregressive (VAR) models relating energy price increases to macroeconomic aggregates in particular have shaped the discussion of the effects of energy price shocks in recent decades. In this paper, we demonstrate that the regression models and estimation methods typically used in that literature produce inconsistent estimates of the true effects of unanticipated energy price increases and are likely to have exaggerated the impact of positive energy price shocks. We show that fundamental changes are needed in how these effects are estimated in practice. In addition to addressing the problem of how to estimate asymmetric responses to energy price shocks, we develop improved tests of the null hypothesis of symmetric responses to energy price increases and decreases. Our empirical evidence suggests that there is no compelling evidence against the symmetry null.

The practical relevance of our analysis for the macroeconomic analysis of oil price shocks is best illustrated in the context of the widely studied question of how much real GDP responds to unanticipated changes in the real price of oil. Commonly used censored VAR models and impulse response estimation methods suggest a decline of 1.1 percent in real GDP within two years in response to a 15 percent unanticipated increase in the real price of oil, and a 0 percent increase in response to an unanticipated decline in the real price of oil of the same magnitude. In contrast, a linear VAR model would imply a decline in real GDP of 0.3 percent and an increase of 0.3 percent, respectively.¹ Appropriately computed impulse responses from unrestricted structural models that encompass both the linear symmetric model and the asymmetric model as special cases yield a 0.47 percent decline in real GDP and a 0.39 percent increase, respectively, roughly consistent with the estimates from the linear symmetric VAR model. We discuss implications of our findings for the theoretical literature on the transmission of energy price shocks and for the debate about policy responses to energy price shocks.

1.1. The Origins of Censoring Changes in Energy Prices

In the view of many economists, oil price shocks are perhaps the leading alternative to monetary policy as the determinant of U.S. postwar recessions.² Increases in the price of oil preceded the recessions of 1973-75, 1980-82, and 1990-91, for example. Given the striking coincidence of deteriorating macroeconomic outcomes and rising oil prices in the 1970s and early 1980s, it was natural at the time to suspect a strong link from oil price increases to recessions. Nevertheless, as discussed in Bernanke, Gertler and Watson (1997), it has proved surprisingly difficult to find an indicator of oil price shocks that produces the expected responses of macroeconomic and policy variables in a VAR setting. Finding a measure of oil price shocks that "works" in a VAR context in practice is not straightforward. Simple measures of energy price shocks (such as linearly unpredictable changes in energy prices) in particular sometimes imply "anomalous" effects on macroeconomic outcomes, relative to the conventional wisdom about the effects of oil price shocks on the economy. They also tend to

¹ This illustrative example is based on the censored VAR model in oil price increases. This allows the comparison with other approaches, while holding fixed the magnitude of the shock.

² For an early exposition of this idea, see Hamilton (1983).

have an unstable relationship with macroeconomic outcomes.³

Far from undermining the view that energy price shocks are important, these difficulties have led researchers to employ increasingly complicated specifications of the "true" relationship between oil prices and the economy. Today it is widely believed that the most appropriate specification of oil price shocks involves some measure of oil price increases, obtained by censoring oil price changes to exclude all oil price decreases. This consensus dates back to the work of Mork (1989). After the sharp oil price declines of 1985-86 failed to lead to an economic boom in oil importing economies, Mork (1989) pointed out that the effects of positive and negative oil price shocks on the economy need not be symmetric. He provided empirical evidence that positive changes in the real price of oil had far more important effects on U.S. real GDP than negative changes. This was widely interpreted as evidence that only oil price increases matter for the U.S. economy (see, e.g., Bernanke, Gertler and Watson 1997, p. 103). Given the a priori belief that oil price shocks have quantitatively important effects on macroeconomic aggregates and given researchers' inability to generate such responses from linear and symmetric models, VAR models of macroeconomic aggregates and oil price increases became accepted on the grounds that they produced "better looking" impulse responses (Bernanke, Gertler and Watson 1997, p. 104).⁴

The initial proposal to focus only on oil price increases has subsequently been refined by Hamilton (1996, 2003), who introduced the "net oil price increase". This measure distinguishes between oil price increases that establish new highs relative to recent experience and increases that simply reverse recent decreases. Specifically, in the context of monthly data, Hamilton's measure equals the maximum of (a) zero and (b) the difference between the log-level of the crude oil price for the current month and the maximum value of the logged crude oil price achieved in the previous 12 (or alternatively 36) months. Hooker (2002), for example, finds that

³ Kilian (2009a) recently has discussed some of the reasons for the apparent instability of such regressions in small samples and for the seemingly counterintuitive response estimates occasionally obtained from such regressions.

⁴ This finding reinforced results based on measures of oil supply disruptions such as the quantitative dummy variable of Hamilton (1996, 2003). It also seemed consistent with evidence using an alternative VAR methodology provided in Davis and Haltiwanger (2001, p. 509), who considered "the evidence for asymmetric responses to oil price ups and downs as well established". Related evidence includes Mork, Olsen and Mysen (1994), Ferderer (1996), Hooker (1996a,b; 2002), Hamilton (1996, 2003), Huntington (1998), and Balke, Brown and Yücel (2002), among others. For a critical perspective on this literature see Edelstein and Kilian (2007, 2009).

the net increase measure performs “well”, in the sense of having a relatively stable relationship with macroeconomic variables, and Hamilton (1996, 2003, 2009a) makes the case that this measure predicts declines in U.S. real GDP.

The most influential use of the net increase measure has been not for one-step ahead prediction from single-equation reduced form models, but in constructing estimates of the response of the economy to energy price shocks based on VAR models. Many of these VAR estimates have become accepted in academic and policy discussions of the transmission of energy price shocks (see, e.g., Davis and Haltiwanger 2002; Lee and Ni 2002; Jones, Leiby, and Paik 2004; Herrera 2008). The net increase measure also plays a central role in VAR analyses of the role of monetary policy in propagating energy price shocks (see, e.g., Bernanke, Gertler and Watson 1997, 2004; Hamilton and Herrera 2004; Herrera and Pesavento 2007).

1.2. Outline of the Paper

In this paper, we demonstrate that the estimation methods used in these VAR applications generate inconsistent estimates of the true effects of unanticipated energy price increases. We show, for example, that VAR models of the effects of energy price increases on macroeconomic outcomes will overestimate the true response to energy price increases asymptotically, when the underlying data generating process (DGP) is symmetric. Since one of the reasons that these models were adopted was precisely their ability to generate larger responses to energy price shocks than symmetric VAR models, our finding casts doubt on the results reported in this literature and, indeed, on the empirical relevance of asymmetric VAR models. Our results apply with equal force to measures of oil price increases and net oil price increases. More generally, we show that asymptotic biases in the censored VAR response estimates may arise whether the underlying DGP is linear and symmetric or not.

In section 2, we illustrate this point in the context of a stylized static model. We establish the inconsistency of estimators of conventional asymmetric models analytically, and we study the determinants of the asymptotic bias of the estimator by simulation. We show that static censored regressor models will be consistent only in very special and theoretically implausible cases. In section 3, we strengthen this result by showing that estimates of dynamic censored VAR models of the type frequently employed in the literature are fundamentally

misspecified and produce inconsistent impulse response estimates regardless of the DGP. We demonstrate that censored VAR models will distort the quantitative importance of energy price increases asymptotically, even when energy price decreases have no effect on the economy. The reason is that such asymmetric DGPs cannot be represented as censored VAR models. We discuss the construction and estimation of alternative, correctly specified asymmetric models.

An important problem in practice is that we may not know whether the DGP is symmetric or not, and whether energy price decreases should be included in the regression, if the DGP is asymmetric. In section 4, we propose a regression model that can be estimated consistently by standard methods whether the true model is symmetric or not, regardless of the precise form of the asymmetry.

In section 5, we show that in addition to the misspecification of the censored VAR models used in the literature and the resulting inconsistency of the parameter estimates, the dynamic responses of macroeconomic aggregates to unanticipated energy price increases have routinely been computed incorrectly in a way that exaggerates the quantitative importance of these shocks. We adapt existing methods of constructing nonlinear impulse responses to the structural model proposed in section 4 and demonstrate how asymmetric impulse responses may be computed correctly. Unlike existing methods of computing nonlinear impulse responses in the econometric literature, our approach is fully structural and avoids the ambiguities of defining a shock in nonlinear reduced form models.

Both the regression model proposed in section 4 and the method of computing responses to energy price shocks developed in section 5 play a crucial role in designing tests of the symmetry of response functions in energy price increases and decreases. In section 6, we discuss the problem of testing the null hypothesis that the U.S. economy responds symmetrically to energy price increases and decreases. First, we propose an improved test of symmetry in the regression slopes with higher power than Mork's (1989) test. Second, we observe that statistically insignificant departures from symmetry in the slopes may cause statistically significant asymmetries in the implied impulse response functions, given the nonlinearity of these functions, while significant departures from symmetry in the slopes need not imply large asymmetries in the impulse response functions. Third, we observe that the

extent to which responses from the linear symmetric model provide a good approximation will depend on the magnitude of the energy price shock. Thus, traditional slope-based tests, while informative about the degree of asymmetry in reduced form parameters, do not shed light on the degree of asymmetry in the impulse response functions. As an alternative, we propose a direct statistical test of the symmetry of the economy's response to unanticipated energy price increases and decreases with reasonably accurate finite-sample size. In section 7, we apply these tools in examining the evidence against the symmetry null in three prominent empirical examples. Specifically, we model the relationship between quarterly U.S. real GDP and the real price of oil, between monthly U.S. unemployment and the real price of oil, and between monthly U.S. gasoline consumption and the U.S. real retail price of gasoline. We find no compelling evidence of asymmetric responses to positive and negative energy price shocks.

In section 8, we extend the analysis to VAR models involving net energy price increases motivated by the analysis in Hamilton (1996, 2003). Despite the widespread use of the net oil price increase measure in VAR models, none of the symmetry test results in the literature provides a justification for the use of such models. In fact, notwithstanding the evidence for asymmetries in the predictive relationship between real GDP growth and oil price changes provided by Hamilton (2003), no paper has adequately addressed the implications of the net increase model for impulse response analysis. In this paper, we address this question. We discuss how structural impulse responses to energy price shocks may be consistently estimated in this context, and we present two tests of symmetric impulse responses based on the net increase model, one based on the slope parameters of a modified version of the model proposed in section 4, and the other based on the impulse response functions implied by that model.⁵ Unlike Hamilton (2003) or Balke et al. (2002), we find little, if any, evidence against the null hypothesis of symmetric responses in our three empirical examples. In fact, traditional slope-based symmetry tests are shown to be misleading in general. Our results suggest that linear impulse response analysis will be adequate for many applications. This finding has

⁵ Our analysis of the net increase model clarifies, refines, and extends the earlier analysis in Balke et al., which recognized some of the problems discussed here but had no apparent impact on the empirical practice in this literature. A partial exception is Herrera (2008), who conducts a sensitivity analysis based on Balke et al.'s methodology in a not-for-publication appendix, as a complement to results from censored VAR models.

important implications for the theoretical literature on the transmission of energy price shocks and for the debate about monetary policy responses to oil price shocks. Finally, to the extent that there is evidence of asymmetries, our analysis suggests that important changes are needed in the way these asymmetries are modeled in the VAR literature. Section 9 contains the concluding remarks.

2. A Stylized Model

It is well known that censoring dependent variables causes OLS estimates of the coefficients of linear models to be biased. The problem of censoring endogenous variables in VAR models in contrast is not well understood. To build intuition we first consider a purely static model.

2.1. Asymptotic Biases from Using Censored Regressors

For expository purposes, consider the following static linear two-equation DGP:

$$\begin{aligned}x_t &= \alpha_1 + \varepsilon_{1,t} \\y_t &= \alpha_2 + x_t\beta + \varepsilon_{2,t}\end{aligned}\tag{1}$$

where $\alpha_1, \alpha_2, \beta$ are constants, $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are mean zero i.i.d. Gaussian random variables with variances σ_1^2 and σ_2^2 , and $t = 1, \dots, T$. It is straightforward to show that the OLS estimators of a and b in the regression model

$$y_t = a + x_t b + u_t\tag{2}$$

will be unbiased and consistent estimators of α and β . To illustrate the effect of replacing negative values of x_t with zero in this regression, define the variable x_t^+ as


$$x_t^+ = \begin{cases} x_t & \text{if } x_t > 0 \\ 0 & \text{if } x_t \leq 0 \end{cases}\tag{3}$$

and consider estimating the censored regression model:

$$y_t = a + x_t^+ b + u_t\tag{4}$$

rather than model (2). Censoring the explanatory variable renders the estimator of b inconsistent for β . Figure 1 illustrates the problem. The values of the dependent variable that are associated with negative values of x_t are moved to the y-axis. Given these points, the estimated intercept is negative and far below its true value of zero. As a result of the negative intercept, the value of the slope coefficient must be greater than the true value.

Figure 1: The Effect of Censoring Negative Values of the Explanatory Variable



The upward bias in the estimated effect of x_t on y_t is not a small-sample problem. For the simple case where α equals zero and where x_t has a symmetric distribution with mean zero and variance 1 and is uncorrelated with $\varepsilon_{2,t}$, we can derive the limit for \hat{b} . Observe that $E(x_t^+) = 0.5\mu$, where $\mu \equiv E(x_t | x_t > 0)$. It can be shown that

$$\hat{b} \xrightarrow{p} \beta \frac{1}{1 - 0.5\mu^2} \quad [5]$$

If the variable x_t has a standard normal distribution, for example, the effect of x_t on y_t is overestimated by almost 50 percent.

2.2. Further Illustrations of the Asymptotic Bias from Censoring

The analysis so far has focused on DGPs that are symmetric in x_t , but the same problem may arise even when the true response of y_t to x_t is asymmetric in positive and negative values. Consider the following DGP:

$$\begin{aligned} x_t &= \alpha_1 + \varepsilon_{1,t} \\ y_t &= \alpha_2 + x_t\beta + x_t^+\gamma + \varepsilon_{2,t} \end{aligned} \quad [6]$$

where γ captures the asymmetric response that is of interest to many economic researchers.

This process allows for both positive and negative values of x_t to affect y_t , but with different coefficients. Equivalently, we could have specified the second equation of model (6) as a regression in x_t^+ and x_t^- (i.e. the censored version of x_t with only negative values) with potentially different coefficients. Given this DGP, if one estimated the equation

$$y_t = a + x_t^+ b + u, \quad [7]$$

one would want the value of b to equal $\beta+\gamma$ in the limit. However, as shown in Table 1, this is not the case unless β equals zero, implying that only positive x_t have an effect on y_t . For all other values of β , the value of the estimated slope coefficient is biased upwards. Furthermore, in this example, if the slope for negative x_t is at least half of the slope for positive x_t , then the estimate of an increase in x_t on y_t is less biased when using the full sample than when using the censored sample.⁶

Table 1: Asymptotic Bias from Censoring in the Asymmetric Static Model

Population Slope Parameters			Average Estimated Slope \hat{b}	
β	γ	$\beta+\gamma$	Standard Model	Censored Model
0.00	1	1	0.50	1.00
0.25	0.75	1	0.62	1.12
0.50	0.50	1	0.75	1.23
0.60	0.40	1	0.80	1.28
0.70	0.30	1	0.85	1.32
0.80	0.20	1	0.90	1.37
1.00	0.00	1	1.00	1.47
1.25	-0.25	1	1.12	1.58
1.50	-0.50	1	1.25	1.70

Notes: Asymmetric DGP: $\alpha_1 = \alpha_2 = 0$, $\sigma_1 = \sigma_2 = 1$. Average results for 2,000 samples of 100,000 observations each.

Table 1 illustrates that, only when β is *known* to be zero, will the censored regression model consistently estimate the effect of an energy price increase in the static model. This point is important because Mork (1989) merely failed to reject the null hypothesis that energy price decreases have no effect on real GDP growth. He did not establish that $\beta=0$, and indeed was careful only to suggest that these coefficients are “perhaps zero”. Nor does economic

⁶ Additional simulation exercises (not shown to conserve space) confirm that the asymptotic biases reported in Tables 2 and 3 that arise from the misspecification of the regression model carry over to small samples. Relative to models with Gaussian errors, small sample biases may increase substantially when the errors are fat-tailed.

theory predict that $\beta=0$ (see, e.g., Edelstein and Kilian 2007, Kilian 2008a). If, in fact, both energy price increases and decreases matter for real GDP, but to a different extent, as suggested by economic theory, then the censored regressor model is bound to overestimate the effect of an energy price increase even in this simplest possible model. The only way to protect from this inconsistency is to include both energy price increases and decreases in the regression.

The static model is useful for building intuition, but extending the analysis to dynamic regression models leads to additional insights. In many cases, researchers are specifically interested in the response of the economy over time to an unexpected energy price increase:

$$I_y(h, \varepsilon_{1,t}, \Omega_{t-1}) = E(y_{t+h} | \varepsilon_{1,t}, \Omega_{t-1}) - E(y_{t+h} | \Omega_{t-1}), \quad [8]$$

where Ω_{t-1} is the information set at the time of the shock. In the next section, we show that the censored VAR models routinely employed in the literature will produce inconsistent impulse response estimates not only in the empirically plausible case of nonzero responses to both energy price increases and decreases but even when all coefficients on current and lagged x_t are zero in population. In other words, even if oil price decreases had no effect on real GDP growth in population, as postulated in the literature, censored VAR models would yield invalid impulse response estimates. This result is in sharp contrast to the static model.

3. How Empirically Relevant Is the Asymptotic Bias of VAR Models of Energy Price Increases?

It is common practice in the literature to compute responses to unanticipated energy price increases from censored VAR models. For example, Bernanke, Gertler and Watson (1997, p. 103) observe that “Knut Mork provided evidence that only positive changes in the relative price of oil have important effects on output. Accordingly, in our VARs we employ an indicator that equals the log difference of the relative price of oil when that change is positive and otherwise is zero.” Similarly, Leduc and Sill (2004, p. 790) state that “to get an empirical estimate of the output response to positive oil-price shocks, we run a VAR using ... oil-price increases [...] constructed by taking the first difference of the log of oil prices, then setting negative values to zero. Thus, only oil-price increases affect the other variables in the system.”

The standard approach used in the literature is based on a recursively identified VAR model, in which the energy price increase variable is ordered above the macroeconomic

aggregate of interest. The prototypical example is a linear bivariate autoregression for $(x_t^+ \ y_t)'$.⁷ Whereas impulse responses in linear models are independent of the history of the observations, impulse responses in nonlinear models such as the censored VAR model are dependent on the history of the observations and on the magnitude of the shock (see, e.g., Gallant Rossi and Tauchan 1993; Koop, Pesaran and Potter 1996). Thus, the construction of impulse responses requires Monte Carlo integration over all possible paths of the data. This point has been routinely ignored in the applied literature. In practice, researchers have presented impulse response estimates computed exactly as in linear VAR models. For now we follow that convention because we wish to illustrate the asymptotic biases in the results reported in the literature. Discussing one problem at a time will facilitate the exposition. We will return to this point in section 5, however, and show how impulse responses can be computed correctly and what difference this computation makes.

3.1. Linear and Symmetric VAR Data Generating Processes

The points we made in section 2 directly extend to censored VAR models. Consider a linear symmetric bivariate VAR(p) DGP of the form

$$\begin{aligned} x_t &= b_{10} + \sum_{i=1}^p b_{11,i} x_{t-i} + \sum_{i=1}^p b_{12,i} y_{t-i} + \varepsilon_{1t} \\ y_t &= b_{20} + \sum_{i=0}^p b_{21,i} x_{t-i} + \sum_{i=1}^p b_{22,i} y_{t-i} + \varepsilon_{2,t} \end{aligned} \tag{9}$$

where x_t denotes the percent change in energy prices, y_t denotes the percent change in the macroeconomic aggregate of interest, and $\varepsilon_t \sim (0, \Sigma)$ is uncorrelated white noise. We focus on three illustrative examples that are representative of models employed in the literature:

- (1) A quarterly VAR in the percent changes in real crude oil prices and the growth rate of U.S. real GDP. The sample period is 1973.II-2007.IV. The oil price series is based on an

⁷ Adding additional macroeconomic variables does not affect the econometric points of interest in this paper, and, indeed, is not required for consistently estimating the response of y_t to an unanticipated increase in energy prices under the maintained assumption of predetermined (or contemporaneously exogenous) energy prices (see Kilian 2008a). The identifying assumption of predetermined energy prices with respect to the macroeconomic aggregate of interest is not only standard in the literature but is consistent with empirical evidence presented in Kilian and Vega (2008).

index of U.S. refiners' acquisition cost, extrapolated as in Kilian (2009a), and deflated by the U.S. CPI. The real GDP data are from the BEA.

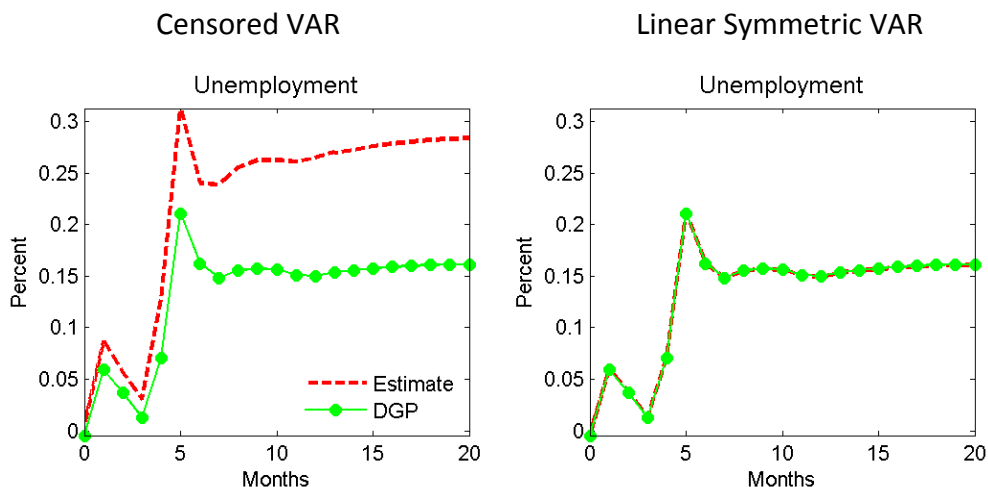
(2) A monthly VAR in the percent change in crude oil prices and the change in the U.S. unemployment rate. The sample period is 1973.2-2007.12. The unemployment rate data are from the BLS.

(3) A monthly VAR in the percent change in real gasoline prices and the percent change in real U.S. real gasoline consumption, as constructed by the BEA. The sample period is 1973.2-2007.12.

The lag order p is set to 6 for expository purposes. For each data set, we construct a DGP by replacing the model parameters by their least-squares estimates obtained from fitting this model to the data set in question and by treating the structural errors as Gaussian white noise. For each DGP, we generate a sample of length T , fit a VAR(6) model for $(x_t^+ y_t)'$, and construct the cumulative response of y_{t+h} , $h = 0, 1, \dots$ to a unit increase in energy prices. We compare that response to the response of y_{t+h} to a unit increase in energy prices in the DGP. Since we are interested in evaluating the asymptotic bias of the responses implied by the censored VAR model, all results in Figure 2 are based on $T = 1,000,000$. To conserve space we focus on the response of unemployment, but note that analogous results hold in the other two applications.

The left panel of Figure 2 quantifies the asymptotic bias induced by censoring energy price decreases. The impulse response implied by the DGP is shown as the solid line. The estimated impulse response from the censored VAR model is shown as the dashed line. As shown in section 2, the censored VAR response tends to overestimate the true response. For example, the response of unemployment after 12 months is overestimated by more than one third. The right panel of Figure 2 illustrates the point that our results are not driven by sampling uncertainty. It shows that the true and the estimated response lie exactly on top of one another if we fit a linear symmetric VAR model to the same data. The results in Figure 2 suggest strong caution in interpreting the results of censored VAR models.

**Figure 2: Inconsistency of the Estimated Effect of Energy Price Increases
Symmetric VAR DGP**



Notes: Simulations from symmetric VAR DGP based on U.S. data. T=1,000,000.

3.2. Asymmetric Data Generating Processes

Figure 2 showed that censored VAR models will be misleading when the DGP is linear and symmetric. It is equally interesting to investigate the consistency of the censored VAR model when the DGP involves asymmetric effects of energy price increases on macroeconomic aggregates. Here we follow the bulk of the empirical literature on energy price shocks and focus on the leading example of models in which only energy price increases matter for macroeconomic aggregates. For expository purposes, first consider the simplest possible dynamic model, in which energy price decreases have no effect:

$$\begin{aligned} x_t &= \alpha_1 + \rho x_{t-1} + \varepsilon_{1,t} \\ y_t &= \alpha_2 + x_t^+ \gamma_2 + \varepsilon_{2,t} \end{aligned} \tag{10}$$

where x_t^+ is defined as above. Setting the initial conditions to zero, in this system the impact response of y_t to a positive shock to x_t would be γ . In the next period, the response would be $\rho\gamma$, provided that ρ is positive.⁸ If this system is estimated, then, as expected, estimates of both γ and ρ will be unbiased. If instead a researcher estimated the censored system:


⁸ If ρ were sufficiently negative that x_t being positive implies that x_{t+1} would be negative, then the response in the second period would be zero. This result is discussed in more detail in section 5.

$$\begin{aligned}
x_t^+ &= \alpha_1 + \rho x_{t-1}^+ + \varepsilon_{1,t} \\
y_t &= \alpha_2 + x_{t-1}^+ \gamma + \varepsilon_{2,t}
\end{aligned}
\tag{11}$$

the estimate of ρ would be asymptotically biased and so would be the impulse response estimate, even though the estimate of γ is unbiased.

This example illustrates that the DGP cannot be represented as a bivariate VAR for $(x_t^+ \ y_t)'$. In fact, a censored VAR DGP with positive probability would generate realizations for x_t^+ that may be negative. It may seem that this contradiction could be avoided by censoring the realizations much like researchers have censored percent changes in actual energy prices, but in that case the same asymptotic biases would arise that we already documented for the linear symmetric model. This point is illustrated in Figure 3. Based on a censored VAR DGP, the censored VAR run on censored realizations of x_t^+ generates responses to energy price increases that are systematically higher than the pseudo-true response even in the limit.

Figure 3: Inconsistency of the Estimated Effect of Energy Price Increases Censored VAR DGP with Censored VAR Variable



Notes: Simulations from censored VAR DGP based on U.S. data. T=1,000,000.

The source of the problem in Figure 3 is that the censored VAR regression model is not a description of the DGP. This problem can be avoided only by fully specifying the underlying structural model:

$$\begin{aligned}
x_t &= \alpha_1 x_{t-1} + \alpha_2 y_{t-1} + \dots + \varepsilon_{1t} \\
y_t &= \beta_1 x_t^+ + \beta_2 x_{t-1}^+ + \beta_3 y_{t-1} + \dots + \varepsilon_{2t}
\end{aligned}
\tag{12}$$

where the structural shocks ε_{1t} and ε_{2t} are uncorrelated and where for expository purposes we have omitted the definition of x_t^+ as a function of x_t .⁹

The DGP in model (12) postulates that percent changes in energy prices evolve in an unconstrained fashion; only the feedback from energy prices to the macroeconomic aggregates is constrained. This model is easily recognizable as a generalization of model (6) with $\beta = 0$ to the VAR context. Note that in this model a negative shock to x_t may have a nonzero effect on y_{t+h} if the negative shock over time induces positive values in x_{t+h} . Also note that model (12) is not equivalent to the following model:

$$\begin{aligned} x_t &= \alpha_1 x_{t-1} + \alpha_2 y_{t-1} + \dots + \varepsilon_{1t} \\ y_t &= \beta_2 x_{t-1}^+ + \beta_3 y_{t-1} + \dots + \beta_1 \varepsilon_{1t} + \varepsilon_{2t} \end{aligned} \quad [13]$$

The key difference between models (12) and (13) is that the impact effect of a negative value of ε_{1t} is zero in model (12) and is $-\beta_1$ in model (13). Furthermore, model (12) cannot be estimated from:

$$\begin{aligned} x_t &= \alpha_1 x_{t-1} + \alpha_2 y_{t-1} + \dots + \varepsilon_{1t} \\ y_t &= \beta_2 x_{t-1}^+ + \beta_3 y_{t-1} + \dots + u_{2t} \end{aligned} \quad [14]$$


where $u_{2t} = \beta_1 \varepsilon_{1t} + \varepsilon_{2t}$, and applying a Cholesky decomposition to the variance covariance matrix of the two error terms ε_{1t} and u_{2t} . The key difference is that the Cholesky decomposition does not discriminate between positive and negative shocks.

Below we confirm that even when the data are generated from model (12), asymptotic biases may arise when estimating the response to energy price increases from a censored VAR model. We focus on the same illustrative example as in Figure 2, except that we now construct the DGP under the working hypothesis that the data are generated by the asymmetric model (12) in which there is no effect from current or lagged x_t on y_t . We treat the least-squares estimates of the slope parameters and innovation variances obtained on the actual data as the population parameters in the simulation and impose the zero correlation of the innovation variances. All results are based on $T = 1,000,000$. Figure 4 shows that even in this case the

⁹ Although the slope parameters of Model (12) can be estimated consistently by OLS, the resulting residuals will not be uncorrelated. To impose the latter restriction requires the use of a restricted maximum likelihood estimator.

response estimates implied by the censored VAR model will be inconsistent. The direction of the asymptotic bias is ambiguous in general. For the unemployment rate, after 18 months, the estimated response to an energy price increase is about 80 percent of the true response. Figure 5 illustrates the source of the problem by plotting the corresponding responses of the price of energy. Although there is no problem in consistently estimating the second equation of the system that includes the censored regressors (and, indeed, the impact response of y_t is correctly estimated in Figure 4), the fact that the censored VAR model misspecifies the first equation of the system causes both response estimates to be inconsistent as the energy price shock is propagated over time. The results in Figures 4 and 5 represent the best possible scenario in that we postulated that only energy price increases matter in the DGP. Additional asymptotic biases would arise if the asymmetric DGP allowed for nonzero effects from energy price decreases, and those biases would affect even the impact responses.¹⁰


Figure 4: Inconsistency of the Estimated Effect of Energy Price Increases on Outcome Variable Asymmetric Structural Model DGP: Fitting Censored VAR Model



Notes: Simulations based on model (12) estimated on U.S. data. T=1,000,000.

¹⁰ Although we focused on bivariate VAR models for $(x_t^+, y_t)'$, the same impulse response inconsistency problems would arise when fitting trivariate VAR models for $(x_t^+, x_t^-, y_t)'$. Similar problems also would arise if we were fitting a VAR model involving $(x_t^+, |x_t^-, y_t)$.

Figure 5: Inconsistency of the Estimate of Energy Price Increases on Energy Price Asymmetric Structural Model DGP: Fitting Censored VAR Model



Notes: Simulations based on model (12) estimated on U.S. data. T=1,000,000.

4. Eliminating the Inconsistency: A General Model of the Oil Price-Economy Link

Up until now, we have imposed a strong form of asymmetry in which energy price declines have no effect on the macro aggregate of interest. In the interest of full generality, we now relax this assumption by allowing for both energy price increases and decreases to have an effect, but to a different extent.¹¹ The first equation of the resulting model is identical to the first equation of a standard linear VAR in x_t and y_t , but the second equation now includes both x_t and x_t^+ and, as such, both energy price increases and decreases affect y_t .

$$\begin{aligned}
 x_t &= b_{10} + \sum_{i=1}^p b_{11,i} x_{t-i} + \sum_{i=1}^p b_{12,i} y_{t-i} + \varepsilon_{1t} \\
 y_t &= b_{20} + \sum_{i=0}^p b_{21,i} x_{t-i} + \sum_{i=1}^p b_{22,i} y_{t-i} + \sum_{i=0}^p g_{21,i} x_{t-i}^+ + \varepsilon_{2,t}
 \end{aligned}
 \tag{15}$$

Given estimates of these coefficients, one can calculate the dynamic responses to unanticipated positive and negative energy price changes.¹² Note that the OLS residuals of

¹¹ Theoretical models of asymmetry do not imply the strong form of asymmetry but allow for nontrivial effects of both energy price increases and decreases (see, e.g., Edelstein and Kilian 2007, 2009; Kilian 2008a).

¹² If energy prices never declined, this model would suffer from collinearity. But in the data, we observe both energy price increases and declines.

model (15) are uncorrelated, whereas the OLS residuals of model (12) may be correlated. This means that model (15) may be estimated by standard regression methods.

As demonstrated in [the working paper version], the key advantage of model (15) is that the dynamic responses are consistently estimated regardless of whether the true DGP is symmetric or asymmetric. In short, the advantage of model (15) is that it may be used without knowing the nature of the DGP. Its only limitation is that the response estimates are not efficient asymptotically and that they may be slightly biased in finite samples because of the nonlinearity of the response function.

5. Computing Responses to Energy Price Shocks in Nonlinear Models

So far we have followed the convention in the empirical literature on energy price shocks of computing impulse responses as one would for linear and symmetric VAR models. While this approach simplifies the computation of the responses from asymmetric models, it can be misleading in that the effect of a given shock in asymmetric models depends on the recent history of the series in question and on the magnitude of the shock. For example, the net effect of a negative innovation to energy prices on macroeconomic aggregates will depend on the extent to which the effect of this shock on energy prices is dampened or amplified by the cumulative effect of previous shocks. Thus, nonlinear impulse responses must be computed for a given shock as the average of impulse response draws obtained using alternative initial conditions. This point is well known (see, e.g., Gallant et al. 1993; Koop et al. 1996), but has been typically ignored in the literature on estimating the effects of energy price increases.

In this section we propose an adaption of methods for computing impulse responses from nonlinear models specifically designed for model (15).¹³ Having estimated model (15), we proceed as follows:

Step 1 Take a block of p consecutive values of x_t and y_t . This defines a history Ω^i . Note that the choice of history does not affect the coefficients of the model. For all histories, the model coefficients are fixed at their estimated values.

¹³ The same algorithm could be applied to model (12) as well, if that model were considered appropriate for applied work.

Step 2 Given Ω^i , simulate two time paths for x_{t+i} and y_{t+i} for $i = 0, 1, \dots, H$. In generating the first time path, the value of ε_{1t} is equal to a pre-specified value δ . In generating the other time path, the value of ε_{1t} is drawn from the marginal empirical distribution of ε_{1t} . The value of ε_{2t} and the values of all subsequent shocks ε_{1t+i} and ε_{2t+i} , $i = 1, \dots, H$, are drawn from their respective marginal distributions. Since the structural errors ε_{1t} and ε_{2t} are by definition uncorrelated, we treat the draws as independent in practice.

Step 3. Calculate the difference between the time paths for y_{t+i} , $i = 1, \dots, H$.

Step 4. Average this difference across $m=500$ repetitions of Steps 2 and 3.

This average is the response of y_{t+i} at horizon $i = 0, \dots, H$, to a shock of size δ conditional on Ω^i :

$$I_y(h, \delta, \Omega^i). \quad [16]$$

The unconditional response $I_y(h, \delta)$ is defined as the value of $I_y(h, \delta, \Omega^i)$ averaged across all histories:

$$I_y(h, \delta) = \int I_y(h, \delta, \Omega^i) d\Omega^i \quad [17]$$

Depending on the application, both the conditional response $I_y(h, \delta, \Omega^i)$ and the unconditional response could be of potential interest. To determine the general importance of oil shocks as a source of fluctuations, the unconditional response is of most interest. Likewise, for the evaluation of DSGE models the unconditional response is the most appropriate object to match. In contrast, for forecasting and policy purposes, the response conditional on current history is the more relevant statistic. In the remainder of the paper, we will focus on the unconditional response (17).

It is important to highlight the difference between the response in (17) and the impulse response functions typically computed in the literature on the transmission of oil price shocks. For this purpose, it is useful to introduce two alternative definitions of impulse responses to energy prices. For example,

$$I_y(h, \delta, \underline{0}) \quad [18]$$

refers to the response obtained conditioning on a hypothetical historical path involving $x_{t-i} = y_{t-i} = 0$ for all i . Another possible definition would be the response

$$\begin{aligned}
& I_y^*(h, \delta, \Omega^i) \\
& = E \left(y_{t+h} \mid \varepsilon_{1,t} = \delta, \Omega^i, \{\varepsilon_{1t+j} = 0\}_{j=1}^h, \{\varepsilon_{2t+j} = 0\}_{j=0}^h \right) \\
& - E \left(y_{t+h} \mid \Omega^i, \{\varepsilon_{1t+j} = 0\}_{j=0}^h, \{\varepsilon_{2t+j} = 0\}_{j=0}^h \right)
\end{aligned} \tag{19}$$

Upon integrating out across alternative histories, this expression simplifies to:

$$I_y^*(h, \delta) = \int I_y^*(h, \delta, \Omega^i) d\Omega^i \tag{20}$$


The response most often reported in the empirical literature on energy price shocks is $I_y^*(h, \delta, \underline{0})$ and may be viewed as a combination of (18) and (20). That response conditions neither on history nor does it allow for nonzero future shocks.

The time series literature on non-linear impulse responses has favored $I_y(h, \delta, \Omega^i)$ over $I_y^*(h, \delta, \underline{0})$ for at least three reasons. First, the process of $I_y^*(h, \delta, \underline{0})$ may not accurately characterize the true dynamics of a nonlinear system. For example, Koop, Pesaran and Potter (1996, p. 125) observe that $I_y^*(h, \delta, \underline{0})$ may not converge to zero, even when the nonlinear stochastic DGP is stationary. Second, Potter (2000) observes that the theoretical treatment of nonlinear impulse responses is facilitated by treating future shocks as random variables rather than fixed values. Third, there is no intrinsic economic interest in any impulse response that conditions on $\underline{0}$.

One of the chief drawbacks of nonlinear impulse response analysis of the type considered by Koop et al. (1996) is that shocks cannot be defined unambiguously in the context of reduced-form models. The reason is that reduced-form regression errors are inevitably mutually correlated. In practice, some researchers have resorted to presenting responses for representative draws from the joint distribution of the reduced-form errors (see, e.g., Balke et al. 2002; Herrera 2008). Such responses may be useful in characterizing the persistence of the data, but they are devoid of any economic interpretation. An attractive feature of the nonlinear impulse responses (16) and (17) that we defined above is that in the context of the structural models (12) and (15) an energy price shock is orthogonal to other shocks and uniquely defined (up to scale). This fact addresses the chief limitation of nonlinear impulse responses as discussed in Koop et al. (1996).

How the impulse response is computed directly affects the magnitude of the estimated response. As shown in Figure 6, in the context of our example, $I_y(h, \delta)$ is much smaller in absolute value than $I_y^*(h, \delta, \underline{0})$. The value of $I_y^*(h, \delta, \underline{0})$ is more comparable to the response to a very large shock in excess of ten standard deviations. In other words, even if the regression model were correctly specified and hence the parameter estimates were consistent, the use of traditional impulse response functions would greatly exaggerate the effect of a positive oil price shock. This finding reinforces our earlier concern with the methods underlying the existing literature.

Figure 6: The Response of GDP to a Positive Energy Price Shock by Shock Size



Notes: $I_y(h, \sigma)$ is calculated by Monte Carlo integration over 300 histories with 10,000 paths each. All impulse responses have been scaled (as denoted in the legend) to ensure compatibility. The letter h denotes the horizon (in quarters).

To help us understand these results, below we provide a more in-depth analysis of the impact responses. We show that, as δ increases, the importance of accounting for the history Ω^i and of accounting for the variability of ε_{1t} declines. In other words,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \left(I_y(0, n\delta) \right) = I_y^*(0, \delta, \underline{0}).$$

Impact Response of x_t Because of the linear nature of the first equation in model (15), the impact effect on x_t of a shock of size δ will be a constant: $I_x(0, \delta, \Omega^i) = \delta$.

Impact Response of y_t The impact effect of $\varepsilon_{1,t}$ on y_t for a given history Ω^i is

$$I_y(0, \delta, \Omega^i) = b_{21,0} \delta + g_{21,0} \left(E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i) \right). \quad [21]$$

The term $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$ plays a central role in the construction of nonlinear impulse responses for y_{t+h} $h = 1, \dots, H$. Absent uncertainty about the value of $\varepsilon_{1,t}$, the value of $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$ would be easy to calculate. In particular, consider the value of $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \varepsilon_{1,t} = 0, \Omega^i)$ where, for ease of notation, we define

$$\tilde{x}_t \equiv E(x_t | \varepsilon_{1,t} = 0, \Omega^i) \quad [22]$$

Here \tilde{x}_t summarizes the effect of the history Ω^i on x_t . For $\delta > 0$, then

$$\begin{aligned} E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \varepsilon_{1,t} = 0, \Omega^i) &= \delta && \text{if } \tilde{x}_t > 0 \\ E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \varepsilon_{1,t} = 0, \Omega^i) &= \tilde{x}_t + \delta && \text{if } -\delta < \tilde{x}_t < 0 \\ E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \varepsilon_{1,t} = 0, \Omega^i) &= 0 && \text{if } \tilde{x}_t < -\delta \end{aligned} \quad [23]$$

Computing the value of $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$ is more of a challenge because we need to account for uncertainty about $\varepsilon_{1,t}$. With uncertainty, we have that

$$\begin{aligned} E(x_t^+ | \delta, \Omega^i) &= E \left(\max(\tilde{x}_t + \delta, 0) \mid \delta, \Omega^i \right) = \max(\tilde{x}_t + \delta, 0) \\ E(x_t^+ | \Omega^i) &= E \left(\max(\tilde{x}_t + \varepsilon_{1,t}, 0) \mid \Omega^i \right) \end{aligned} \quad [24]$$

The value of $E(x_t^+ | \Omega^i)$ depends on the variance of the shocks. Note that $E(x_t^+ | \Omega^i)$ can be positive even if \tilde{x}_t is negative. In fact, by Jensen's inequality, $E(x_t^+ | \Omega^i) \geq \tilde{x}_t$ for all values of \tilde{x}_t . In particular, if \tilde{x}_t equals zero and $\varepsilon_{1,t}$ has a standard normal distribution then $E(x_t^+ | \Omega^i)$ has a value of 0.4. Hence, when $\tilde{x}_t = 0$ and $\varepsilon_{1,t}$ has a standard normal distribution, we have that

$$\begin{aligned} E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i) &= E(x_t^+ | \delta, \tilde{x}_t = 0) - E(x_t^+ | \tilde{x}_t = 0) \\ &= \delta - 0.4. \end{aligned} \quad [25]$$


This first result implies that the larger δ , the smaller the effect of incorrectly treating $\varepsilon_{1,t}$ as equal to zero under the counterfactual path, relative to magnitude of the impulse response. In

other words, all else equal, the larger δ , the more similar the traditional incorrectly computed impulse response and the correctly computed unconditional response will be. This point is important because most energy price innovations measured at the monthly or quarterly frequency tend to be quite small (e.g. Edelstein and Kilian 2009). We conclude that traditional, incorrectly computed impulse responses will tend to exaggerate the effect of an unanticipated energy price increase. Figure 7 illustrates this point under the assumption that

$$\begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix} \right)$$


The horizontal axis shows values of \tilde{x}_t representing alternative histories. Note that the difference between the response computed by correctly accounting for the uncertainty of $\varepsilon_{1,t}$ and the incorrectly computed response obtained from treating $\varepsilon_{1,t}$ as fixed declines, as δ increases, for all possible histories \tilde{x}_t .

Figure 7: $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$



Our second point is that $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$ becomes less sensitive to \tilde{x}_t , as δ increases. Figure 8 illustrates this relationship. As δ increases, the importance of the history \tilde{x}_t declines, and the magnitude of the impulse responses becomes constant across alternative histories.

Figure 8: $E(x_t^+ | \delta, \Omega^i) - E(x_t^+ | \Omega^i)$



The reduced importance of the histories can be explained by the following limit argument:

$$\lim_{n \rightarrow \infty} \frac{1}{n} (I_y(0, n\delta)) = I_y^*(0, \delta, \underline{0}).$$

This result relies on three observations. First, note that $I_y^*(0, \delta, \underline{0}) = b_{21,0} \delta + g_{21,0} \delta$. Second,

$$\lim_{n \rightarrow \infty} \frac{1}{n} (I_y(0, n\delta)) = \lim_{n \rightarrow \infty} \frac{1}{n} \left(\int I_y(0, n\delta, \Omega^i) d\Omega^i \right)$$

In addition, observe that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \left(\int I_y(0, n\delta, \Omega^i) d\Omega^i \right) = b_{21,0} \delta + g_{21,0} \left(\frac{1}{n} \int (E(x_t^+ | n\delta, \Omega^i) - E(x_t^+ | \Omega^i)) d\Omega^i \right)$$

To complete the proof we need to show that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \int (E(x_t^+ | n\delta, \Omega^i) - E(x_t^+ | \Omega^i)) d\Omega^i = \delta.$$

Recall that by definition $\tilde{x}_t \equiv E(x_t^+ | \varepsilon_{1,t} = 0, \Omega^i)$. Therefore, $E(x_t^+ | n\delta, \Omega^i) = n\delta + \tilde{x}_t$ if $\tilde{x}_t > -n\delta$, therefore

$$\begin{aligned} & \frac{1}{n} \left(\int E(x_t^+ | n\delta, \Omega^i) - E(x_t^+ | \varepsilon_{1,t}, \Omega^i) d\Omega^i \right) \\ &= \left(\delta \mathbb{P}(\tilde{x}_t > -n\delta) + \frac{1}{n} \int_{-n\delta}^{\infty} \tilde{x}_t d\Omega^i - \frac{1}{n} \int_{-\infty}^{\infty} E(x_t^+ | \Omega^i) d\Omega^i \right). \end{aligned}$$

As long as \tilde{x}_t does not have too much mass in the left tail and the variance of $\varepsilon_{1,t}$ is small enough such that

$$\int_{-n\delta}^{\infty} \tilde{x}_t d\Omega^i - \int_{-\infty}^{\infty} E(x_t^+ | \Omega^i) d\Omega^i$$

remains finite, it follows that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \int \left(E(x_t^+ | n\delta, \Omega^i) - E(x_t^+ | \Omega^i) \right) d\Omega^i = \delta,$$

because the value of $P(\tilde{x}_t > -n\delta)$ converges toward one, as n increases, and $E(\tilde{x}_t | \tilde{x}_t < 0)$ is finite.

We conclude that the impulse response $I_y^*(h, \delta, \underline{0})$ can be thought of as the scaled version of the response to a very large shock of the form $I_y(h, n\delta, \Omega^i)/n$. Hence, for sufficiently large energy price shocks, one would expect that the importance of Ω^i diminishes such that the impulse response used in the existing VAR literature becomes a good approximation to the correctly constructed estimate, as illustrated in Figure 6. For smaller (and more typical) energy price shocks, however, such as a one standard deviation increase in $\varepsilon_{1,t}$, the interaction of the innovation in the price of energy with the history Ω^i will be quantitatively important. That is why traditional VAR impulse responses may be quite different from the correctly computed nonlinear responses in practice.

6. Testing Symmetry in Energy Price Increases and Decreases

The preceding sections have shown that the presence of asymmetries considerably complicates the econometric analysis of the transmission of positive energy price shocks. While model (15) may be used to sidestep the question of whether the model in question is asymmetric or not, this involves a cost in terms of asymptotic efficiency. Clearly, a linear symmetric VAR model would be preferable if we could convince ourselves that the DGP is symmetric. This raises the question of how to test the symmetry null hypothesis. We first discuss the case of testing for symmetry between energy price increases and decreases before adapting these tests to the problem of testing models of net energy price increases.

6.1. Slope-Based Tests

If energy price increases and decreases received exactly the same weight in regressions of y_t on lagged y_t and current and lagged x_t^- and x_t^+ , it would follow immediately that the dynamic responses to energy price shocks must be symmetric in positive and negative shocks. This line of reasoning has motivated the development of slope-based tests of symmetry. Such tests are attractive in that they do not require the complete specification of the system to be estimated nor do they require the computation of impulse responses.

The traditional approach to testing for symmetry in the transmission of energy price shocks involves tests on the symmetry of the slope coefficients in regressions of y_t on lagged x_t^- and x_t^+ (see, e.g. Mork 1989). This is equivalent to testing

$$H_0: g_{21,1} = \dots = g_{21,p} = 0$$

in model (15). This may be accomplished by a Wald test with an asymptotic χ_p^2 -distribution.

Closer inspection of model (15) reveals that this test does not exploit all restrictions implied by the null hypothesis of symmetry. Specifically, Mork's test omits the contemporaneous regressor because he works with the reduced form. A test of all symmetry restrictions on the slopes would involve the null hypothesis:

$$H_0: g_{21,0} = \dots = g_{21,p} = 0$$

Hence, our proposal is that we estimate model (15) and use a Wald test to determine whether including $\{x_{t-i}^+\}_{i=0}^p$ improves the fit of the model. This test has an asymptotic χ_{p+1}^2 -distribution. It can be shown that this modified slope-based test has similarly accurate size, but higher power than Mork's test (see the working paper version of this paper).

While slope based tests are useful in assessing the symmetry of the slope parameters of single-equation reduced form models, they are not obviously informative about the degree of symmetry of the impulse response function obtained from a fully specified model. There are two possible outcomes when conducting slope-based tests. If the test rejects symmetry, that is sufficient for concluding that the impulse responses are asymmetric, but it does not tell us whether the departures from symmetry are economically or statistically significant. Given that impulse response functions are highly nonlinear functions of the slope parameters and innovation variances it is quite conceivable that the degree of asymmetry in the impulse

responses to positive and negative energy price shocks could be quite small, making responses based on the linear model a good approximation, despite the statistical rejection. Moreover, the quality of the linear approximation will differ depending on the magnitude of the shock. For that reason, the applied user will want to plot the point estimates of the impulse response functions and inspect them.

If the test fails to reject symmetry, on the other hand, we again learn little because statistically insignificant departures from symmetry in the slopes may cause large and statistically significant asymmetries in the implied impulse response functions, given the nonlinearity of these functions. This observation suggests that a more useful approach would be to test the symmetry of the economy's dynamic responses to unanticipated energy price increases and decreases directly based on the impulse response functions (also see Edelstein and Kilian 2007, 2009). This alternative approach to testing symmetry is discussed next.

Note that what is at issue in conducting this impulse-response-based test is not the existence of asymmetries in the reduced form parameters, but the question of whether possible asymmetries in the reduced form imply significant asymmetries in the impulse response function. While any asymmetry in the reduced form representation (whether statistically significant or not) will imply some degree of asymmetry in the impulse response function, the question is whether the impulse responses constructed from linear symmetric VAR models still provide a good approximation.

6.2. Impulse-Response Based Tests

The proposal is that we estimate the unrestricted model (15) and calculate the unconditional impulse responses to both positive and negative energy price shocks. Then we construct a Wald test of the joint null hypothesis of symmetric responses to positive and negative energy price shocks up to a pre-specified horizon H . Symmetry means that

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

or equivalently

$$I_y(h, \delta) + I_y(h, -\delta) = 0 \text{ for } h = 0, 1, 2, \dots, H$$

The variance covariance matrix of the vector sum of response coefficients can be estimated by bootstrap simulation. Given the asymptotic normality of the parameter estimators of model (15), the test has an asymptotic χ^2_{H+1} -distribution.

Unlike the slope-based test, this test depends on the magnitude of δ , so the evidence against symmetry will depend on the magnitude of the shock considered. For small shocks, a symmetric model will provide a better approximation than for large shocks. How accurate and powerful the impulse-response based test is relative to the slope-based test is an empirical question. It can be shown that, in our three applications, this test has acceptable size properties, despite a slight tendency to over reject, as the horizon increases (see working paper version).

7. Empirical Tests of Symmetry in Energy Price Increases and Decreases

In this section, we use the two statistical tests discussed in section 6 to check for asymmetry in the empirical responses to energy price shocks. The p -values in Table 2 suggest no evidence against symmetry at the 10 percent significance level in monthly U.S. unemployment rates, in quarterly U.S. real GDP, and for U.S. gasoline. The same qualitative result, but with higher p -values, are obtained using Mork's test.¹⁴

Table 2: Empirical Symmetry Tests: Baseline Model

Variable	The Modified Test of Symmetric Slope Coefficients	Marginal Significance Level	Mork's Test of Symmetric Slope Coefficients	Marginal Significance Level
Unemployment	7.7224	0.358	3.1317	0.792
Gas Consumption	11.3755	0.123	9.2366	0.161
Real GDP	10.4722	0.163	9.7565	0.135

Table 3 reports the corresponding tests of the symmetry of the impulse response functions for the baseline model with six lags. Neither for U.S. real GDP nor for unemployment

¹⁴ An important question is whether these empirical results are sensitive to the choice of lag order. Our baseline results rely on six lags. As the lag order of the model exceeds six lags, we find rejections of the symmetry null hypothesis at the 5% level for gasoline consumption, but not for the other examples. The implied impulse response functions, however, are strikingly symmetric even in that case, illustrating the drawbacks of slope-based tests.

is there statistically significant evidence against the symmetry of the functions.¹⁵ For gasoline consumption, the results are mixed. Whereas there is no evidence against symmetry based on two-standard deviation shocks, based on one standard-deviation shocks the test rejects the null hypothesis of symmetry at the 5 percent level at one horizon and at the 10 percent level at several additional horizons. The evidence against symmetry appears stronger than that based on the slope-based test. However, as shown in Figure 9, which reports $I_y(h, \sigma)$ and $-I_y(h, -\sigma)$, the actual difference between these two responses seems fairly small, and one would be hard pressed to make the case for using the asymmetric model on economic grounds. In particular, the response to a negative shock is clearly not zero (as implied by the standard censored VAR model), but about as large in absolute terms as the response to a positive shock.


Table 3: Testing the Symmetry of the Response $I_y(h, \delta) = -I_y(h, -\delta)$ for $h = 0, 1, 2, \dots, H$

H	Gas Consumption		GDP		Unemployment	
	1 Std. Deviation Shock	2 Std. Deviation Shock	1 Std. Deviation Shock	2 Std. Deviation Shock	1 Std. Deviation Shock	2 Std. Deviation Shock
0	0.45	0.47	0.40	0.47	0.36	0.43
1	0.13	0.28	0.44	0.54	0.65	0.73
2	0.05	0.15	0.59	0.69	0.83	0.88
3	0.09	0.25	0.56	0.68	0.92	0.96
4	0.07	0.21	0.66	0.78	0.97	0.98
5	0.04	0.15	0.78	0.87	0.99	1.00
6	0.06	0.18	0.48	0.59	1.00	1.00
7	0.09	0.26	0.58	0.69	1.00	1.00

Notes: Based on 20,000 simulations of model (15). p -values are based on the χ^2_{H+1} distribution. actual difference between these two responses seems fairly small, and one would be hard pressed to make the case for using the asymmetric model on economic grounds.

¹⁵ These results are qualitatively consistent with the findings in Edelstein and Kilian (2007, 2009) based on a somewhat different methodology.

**Figure 9: The Response of Gas Consumption to a One Standard Deviation Energy Price Shock
Baseline Model with 6 Lags**



8. Testing Symmetry Based on Models of Net Energy Price Increases

As noted in the introduction, much of the recent empirical work in the transmission of oil price shocks has focused on the net increase in the price of oil as defined in Hamilton (1996, 2003).¹⁶ Given its widespread use in applied VAR work, it is important to assess the empirical support of that model. Hamilton (2003) proposed a formal test based on the conditional expectation function of the linear symmetric model against the alternative of a model in which only net increases matter. Rather than specifying the structural model of interest, Hamilton focused on single-equation predictive relationships in the data. His statistical evidence against the null hypothesis of linearity in the one-step-ahead predictive relationship between real GDP growth and changes in oil prices by construction implies that the underlying DGP must be asymmetric as well. Hamilton concluded that among several candidate models a conditional expectations model based on the three-year net increase measure had the most predictive power.

¹⁶ For example, Lee and Ni (2002, p. 834) note that the “oil price variable [in their VAR] is Hamilton’s (1996) ‘net oil price increase’, defined as the percentage change of oil price over the maximum value of the preceding year if the price of the current month exceeds the previous year’s maximum, and zero otherwise.” Likewise, Bernanke et al.’s (1997, p. 104) VAR analysis relies as the main measure of oil price shocks on Hamilton’s measure which “equals the maximum of (a) zero and (b) the difference between the log-level of the crude oil price for the current month and the maximum value of the logged crude oil price achieved in the previous twelve months”. Similar net oil price increase measures have also been used by Davis and Haltiwanger (2001), Lee and Ni (2002), and Hamilton and Herrera (2004), among others.

It is important to keep in mind, first, that the model Hamilton investigated differs from the VAR models estimated in the literature on estimating responses to energy price shocks such as Bernanke et al. (1997). Even if we take Hamilton's results about the existence of a nonlinear reduced-form relationship involving net oil price increases at face value, these results do not justify the use of censored VAR models of the form

$$x_t^{+,net} = b_{10} + \sum_{i=1}^p b_{11,i} x_{t-i}^{+,net} + \sum_{i=1}^p b_{12,i} y_{t-i} + \varepsilon_{1t}$$

$$y_t = b_{20} + \sum_{i=1}^p b_{22,i} y_{t-i} + \sum_{i=1}^p g_{21,i} x_{t-i}^{+,net} + a\varepsilon_{1t} + \varepsilon_{2,t}$$

where a is estimated by the Cholesky factorization of the variance covariance matrix of the reduced-form residuals. For the reasons discussed earlier, the structural models that would give rise to such asymmetries in the reduced form cannot be represented as VAR models

Second, Hamilton's work leaves unanswered the question of how much the response of real GDP to an exogenous oil price innovation is affected by the nonlinearity of the DGP relative to the linear case. Even if there is an asymmetry in the slope parameters of the reduced form, that asymmetry need not have large effects on the implied impulse response function. We will illustrate this point below. Moreover, the extent to which responses from a linear symmetric VAR model provide a good approximation will be a function of the size of the energy price shock. Answering that question requires a fully specified multivariate structural model.¹⁷

8.1. Slope-Based Symmetry Tests for the Net Increase Model

In this section, we outline two tests of the net increase model, building on the analysis in section 4. Rather than testing the null hypothesis of symmetry between net oil price decreases and net oil price increases, as in Edelstein and Kilian (2009), we nest the net increase model in the standard linear symmetric VAR model. In essence, we ask whether there is incremental

¹⁷ Hamilton (2003, Figure 14) provides an estimate of the impulse response function eight quarters ahead based on a single-equation nonlinear reduced-form model designed for one-step-ahead prediction (see p. 392). His approach is encompassed by our analysis. It suffers from three problems. First, in constructing the impulse responses Hamilton imposes that the net energy price increases are strictly exogenous. That assumption is unnecessarily restrictive and has been shown to be economically and empirically implausible (see Kilian 2008a,b). Second, Hamilton's model omits the contemporaneous regressor in constructing the impulse response function. Third, Hamilton ignores the nonlinear nature of the model in computing the impulse response function.

explanatory power in including net oil price increases in the baseline model. This results in a model structure similar to model (15) with x_t^+ replaced by $x_t^{+,net}$, where $x_t^{+,net} = \max [0, x_t - x_t^*]$ and x_t^* is the maximum of x_t over the preceding year (or three years, alternatively), following Hamilton (1996, 2003). We follow Kilian (2008b) in specifying the net increase in the real price of oil rather than the nominal price as in Hamilton (1996, 2003) because the real price is the economically relevant measure of the price of oil. This does affect the time path of the net increase variable, as documented in Kilian (2008b), but is more consistent with the recent empirical literature.

The problems with the use of net oil price increase measures in VAR models are fundamentally the same as with the use of oil price increase measures and can be addressed along similar lines. By analogy to the discussion in section 4 the structural model

$$\begin{aligned}
 x_t &= b_{10} + \sum_{i=1}^p b_{11,i} x_{t-i} + \sum_{i=1}^p b_{12,i} y_{t-i} + \varepsilon_{1t} \\
 y_t &= b_{20} + \sum_{i=0}^p b_{21,i} x_{t-i} + \sum_{i=1}^p b_{22,i} y_{t-i} + \sum_{i=0}^p g_{21,i} x_{t-i}^{+,net} + \varepsilon_{2,t}
 \end{aligned} \tag{26}$$

may be estimated consistently by least squares. Note that model (26) also allows us to compute impulse response functions taking account of the magnitude and direction of the innovation ε_{1t} as well as the history of observations, whereas the shock in the commonly used censored VAR model is not well defined.

In assessing the evidence for this structural net increase model, a natural starting point is the slope-based test:

$$H_0: g_{21,0} = \dots = g_{21,p} = 0$$

based on (26). This test relates to a similar test conducted in Balke et al. (2002) as the slope-based test in section 6.1 relates to Mork's test of symmetry. The only difference is the additional inclusion of contemporaneous regressors. Table 4 suggests that there is no evidence of asymmetries using the one-year net increase measure, but for the three-year net increase measure the symmetry test rejects at the 5 percent level for gas consumption and real GDP.

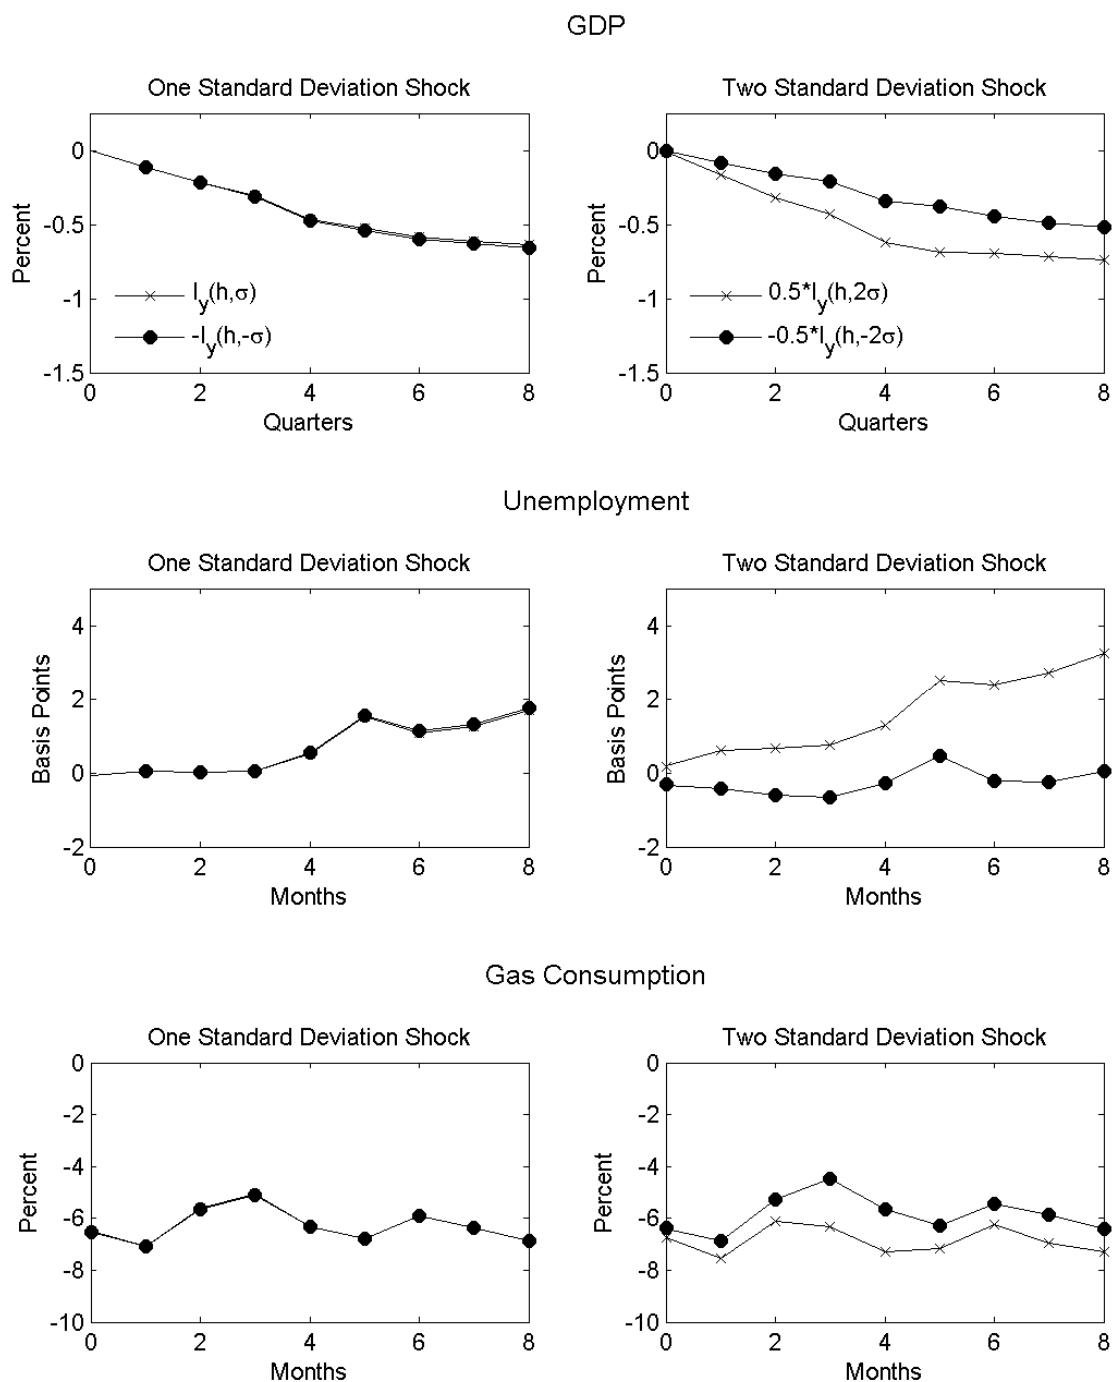
Table 4: Slope-Based Test of the Linear Symmetric VAR Model against the Net Increase VAR Model: Baseline Model with 6 Lags

Variable	1-Year Net Increase		3-Year Net Increase	
	Test of Linear Symmetric Model	Marginal Significance Level	Test of Linear Symmetric Model	Marginal Significance Level
Unemployment	10.5099	0.162	9.6332	0.210
Gas Consumption	9.8879	0.195	14.5307	0.043
Real GDP	7.2617	0.402	14.2965	0.046

Figure 10 once again illustrates that slope-based tests are of limited use in that they are not informative about the degree of asymmetry in the impulse response functions. We focus on the three-year net increase. Results for the one-year net increase are very similar. Figure 10 shows that correctly computed impulse responses from model (26) for one standard deviation shocks are almost perfectly symmetric in all three models.¹⁸ Thus, despite the partial rejection of the hypothesis of symmetric slopes in Table 4, there is no compelling reason to depart from the linear symmetric model when dealing with shocks of typical magnitude. Broadly similar results hold for two standard deviation shocks with the exception of the unemployment rate model. Ironically, the unemployment rate model was the one model that passed all slope-based tests of symmetry, highlighting the importance of actually computing the impulse response functions. In contrast, the other point estimates look fairly symmetric. Although the response of real GDP to a positive two standard deviation shock is somewhat larger in absolute terms than the response to a negative shock of this magnitude, both responses are clearly negative and have a similar pattern. In the gasoline consumption model, the symmetry of the two response functions is even more pronounced.

¹⁸ By construction, a one standard deviation shock is a typical shock in that about two-thirds of energy price shocks in historical data are no larger than one standard deviation.

Figure 10: Empirical Responses to One and Two Standard Deviation Positive and Negative Energy Price Shocks in Baseline Model with 6 Lags: 3-Year Net Increase



Notes: The responses to negative shocks are shown as mirror images to facilitate the comparison. Some of the responses in the left panel are nearly invisible because the responses are almost perfectly symmetric.

8.2. An Impulse-Response Based Symmetry Test for the Net Increase Model

It may be tempting to decide the question of symmetry based on the estimates of the impulse response functions in Figure 10. Figure 10 underscores that there is no reason to question the symmetry assumption for shocks of typical magnitude. For two standard deviation shocks the evidence is less clear, however, especially in the unemployment example. Because the point estimates in Figure 10 are subject to considerable sampling uncertainty, especially when considering large energy price shocks, it is useful to conduct a formal test of the linear symmetric model based on the impulse response functions implied by model (26). As in section 6.2, the test is based on $I_y(h, \delta) = -I_y(h, -\delta)$ for $h = 0, 1, 2, \dots, H$.¹⁹ Table 5 shows that, as expected, the p-values decline with the magnitude of the shock. At conventional significance levels, there is no evidence against the symmetry null hypothesis in any of the three empirical examples in response to a one standard deviation shock, whether we focus on the one-year or the three-year net changes. The same results hold in response to a two standard deviation shock when using the three-year net changes. Similar results also hold for the one-year net changes with the partial exception of the unemployment rate at very short horizons. We conclude that there is very little, if any, evidence of asymmetric responses to energy price increases and decreases. In particular, there is no such evidence for U.S. real GDP.

The evidence in Table 5 is based on a model that rules out responses to net decreases, but there is no compelling a priori reason that the economy is not responding to net decreases in energy prices as well, especially since the motivation for this model is behavioral rather than being grounded in economic theory. We explore this possibility by augmenting model (26) with lags of net decreases in energy prices. The corresponding results are not shown to conserve space but can be summarized as follows: For gasoline consumption there is no evidence against the symmetry null hypothesis using the three-year window, whether the energy price shock is small or large. Using the one-year window, the test fails to reject symmetry in response to

¹⁹ An alternative and asymptotically equivalent approach would have been to test the equality of the impulse responses obtained from the linear model on the one hand and either the response to an energy price increase or a decrease on the other. We do not pursue that approach because it is not clear how to use the bootstrap to evaluate the variance of the Wald test statistic in that case. In contrast, the symmetry test is straightforward to implement.

small shocks but rejects in response to large shocks. Inspection of the point estimates, however, suggests that this evidence, while statistically significant, is not economically significant. The response estimates are very similar in magnitude.

Table 5: p -Values of Test of $H_0: I_y(h, \delta) = -I_y(h, -\delta)$ for $h = 0, 1, 2, \dots, H$

H	Gasoline Consumption		GDP		Unemployment	
	1 Std. Deviation Shock	2 Std. Deviation Shock	1 Std. Deviation Shock	2 Std. Deviation Shock	1 Std. Deviation Shock	2 Std. Deviation Shock
a. 1-Year Net Increase						
0	0.95	0.82	0.96	0.50	0.99	0.01
1	1.00	0.91	0.96	0.72	0.83	0.03
2	1.00	0.34	0.98	0.58	0.94	0.07
3	0.99	0.11	0.99	0.65	0.98	0.12
4	0.98	0.18	0.99	0.64	0.99	0.20
5	0.98	0.26	0.99	0.74	1.00	0.21
6	0.99	0.34	0.99	0.60	1.00	0.19
7	1.00	0.43	1.00	0.66	1.00	0.24
b. 3-Year Net Increase						
0	0.94	0.43	0.98	0.79	0.95	0.13
1	0.98	0.56	0.99	0.30	0.99	0.24
2	1.00	0.66	1.00	0.26	1.00	0.40
3	1.00	0.23	1.00	0.38	1.00	0.57
4	1.00	0.32	1.00	0.41	1.00	0.71
5	1.00	0.44	1.00	0.53	1.00	0.69
6	1.00	0.55	1.00	0.57	1.00	0.69
7	1.00	0.59	1.00	0.67	1.00	0.77

Notes: Based on 20,000 simulations of model (26). p -values are based on the χ_{H+1}^2 distribution.

For unemployment, the evidence is mixed. There is no evidence against symmetry using either window as long as energy price shocks are small. For large shocks, the symmetry test rejects at one horizon only using the three-year window, but it rejects systematically using the one-year window. Indeed, the response estimates look fairly asymmetric in the latter case, suggesting caution in imposing symmetry on unemployment responses when considering large energy price shocks. In contrast, for real GDP, which is perhaps the most interesting example from the point of view of macroeconomists, there is not a single rejection of symmetry for any combination of shock size and window.

These empirical results are important in light of the consensus view, exemplified by Davis and Haltiwanger (2001, p. 509), that the evidence for asymmetric responses to oil price ups and downs is well established. Our analysis suggests that the evidence against the symmetry hypothesis has been overstated. It is of course possible that impulse-response-based tests lack the power to detect asymmetries in the data, especially if those asymmetries are relatively weak for shocks of typical size, but the rejections found in some of our empirical examples suggest that lack of power is not a concern.

We certainly would not want to rule out the existence of asymmetries in all possible applications on the basis of our empirical evidence. Part of our objective has been to provide tools for detecting asymmetries and dealing with asymmetries on a case-by-case basis. What our empirical evidence does suggest, however, is that asymmetry in the responses to energy price shocks is clearly not a pervasive and robust feature of the U.S. data. This point is important, as a large literature has developed aiming to explain the perceived asymmetry of responses to energy price shocks from a theoretical point of view (see, e.g., Bernanke 1983; Hamilton 1988; Pindyck 1991). Our evidence casts doubt on the empirical relevance of these theoretical models.

9. Conclusions

Our empirical results have important implications for studies of the transmission of energy price shocks. First, one reason that researchers had been eager to accept the apparent finding of asymmetry in the 1990s was that it seemed consistent with theoretical models of the transmission of energy price shocks that emphasized asymmetries through shifts in uncertainty

or frictions to the reallocation of factors of production within and across sectors (see, e.g., Bernanke 1983; Hamilton 1988; Pindyck 1991). The latter models were required to rationalize large effects from oil price shocks that are difficult to obtain in conventional models based on cost shocks or aggregate demand shocks. Our evidence provides no support for theoretical models with built-in asymmetries. If such asymmetric effects exist, they appear to be too weak to be detected in aggregate data.

Second, in the absence of asymmetries, the responses of the U.S. economy to positive energy price shocks appear much more modest, which is fully consistent with conventional macroeconomic models of the transmission of energy price shocks that do not predict large fluctuations in U.S. output in response to energy price shocks (see, e.g., Kilian 2008b). Thus, the absence of larger effects is not a puzzle. We conclude that oil price shocks are only one of many factors contributing to recessions, not a key factor. At the same time, we conclude that the response to unexpected energy price declines is not negligible, a finding that matters especially for the interpretation of the data since mid-2008. In addition, our findings lend credence to recent linear models of how the oil demand and oil supply shocks that drive oil price shocks affect the U.S. economy (see, e.g., Kilian 2009a) and support the common practice of linearizing DSGE models with an oil sector about their steady state.

Finally, our analysis calls into question several empirical findings reported in the literature about the channels of the transmission of energy price shocks. To the extent that these studies used censored VAR models and/or computed impulse responses to energy price shocks incorrectly, they are invalid. For example, much of the consensus on how monetary policy responds to oil price shocks is based on the censored VAR model introduced by Bernanke, Gertler and Watson (1997). That study and subsequent papers using the same type of model will have to be reexamined in light of our findings. For example, Kilian (2009b) shows that without censoring there is little evidence that the Federal Reserve caused large effects on real output or inflation by responding to oil price shocks. Similarly, influential studies of sectoral responses to oil price shocks such as Lee and Ni (2002) or employment responses at the plant level as in Davis and Haltiwanger (2001) will have to be reexamined.

References:

- Balke, N.S., Brown, S.P.A., and M.K. Yücel (2002), "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" *Energy Journal*, 23, 27-52.
- Bernanke, B.S. (1983), "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, 98, 85-106.
- 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, 91-142.
- Bernanke, B.S., M. Gertler, and M. Watson (2004), "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: A Reply," *Journal of Money, Credit, and Banking*, 36, 286-291.
- Davis, S.J., and J. Haltiwanger (2001), "Sectoral Job Creation and Destruction Responses to Oil Price Changes," *Journal of Monetary Economics*, 48, 465-512.
- Edelstein, P., and L. Kilian (2007), "The Response of Business Fixed Investment to Energy Price Changes: A Test of Some Hypotheses about the Transmission of Energy Price Shocks," *B.E. Journal of Macroeconomics*, 7(1) (Contributions).
- Edelstein, P., and L. Kilian (2009), "How Sensitive are Consumer Expenditures to Retail Energy Prices," forthcoming: *Journal of Monetary Economics*.
- Ferderer, J.P. (1996), "Oil Price Volatility and the Macroeconomy: A Solution to the Asymmetry Puzzle," *Journal of Macroeconomics*, 18, 1-16.
- Gallant, A.R., P. Rossi, and G. Tauchen (1993). "Nonlinear dynamic structures." *Econometrica*, 61, 871-907.
- Herrera, A.M., and E. Pesavento (2007), "Oil Price Shocks, Systematic Monetary Policy and the 'Great Moderation'", forthcoming: *Macroeconomic Dynamics*.
- Herrera, A.M. (2008), "Oil Price Shocks, Inventories, and Macroeconomic Dynamics," mimeo, Department of Economics, Michigan State University.
- Hamilton, J.D. (1983), "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, 91, 229-248.
- Hamilton, J.D. (1988) "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96, 593-617.
- Hamilton, J.D. (1996), "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics*, 38, 215-220.
- Hamilton, J.D. (2003), "What is an Oil Shock?" *Journal of Econometrics*, 113, 363-398.
- Hamilton, J.D., and A.M. Herrera (2004), "Oil Shocks and Aggregate Economic Behavior: The Role of Monetary Policy: Comment," *Journal of Money, Credit, and Banking*, 36, 265-286.
- Hamilton, J.D. (2009a), "Causes and Consequences of the Oil Shock of 2007-08," forthcoming: *Brookings Papers on Economic Activity*.
- Hamilton, J.D. (2009b), "The Oil Shock of 2008," Presentation at the AEA Meetings, San Francisco, January 4, 2009.
- Hooker, M.A. (1996a), "What Happened to the Oil Price-Macroeconomy Relationship?" *Journal of Monetary Economics*, 38, 195-213.
- Hooker, M.A. (1996b), "This Is What Happened to the Oil Price-Macroeconomy Relationship: Reply," *Journal of Monetary Economics*, 38, 221-222.

- Hooker, M.A. (2002), "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications Versus Changes in Regime," *Journal of Money, Credit, and Banking*, 34, 540-561.
- Huntington, H.G. (1998), "Crude Oil Prices and U.S. Economic Performance: Where Does the Asymmetry Reside?" *Energy Journal*, 19, 107-148.
- Jones, D.W., P.N. Leiby, and I.K. Paik (2004), "Oil Price Shocks and the Macroeconomy: What Has Been Learned since 1996?" *Energy Journal*, 25, 1-32.
- Kilian, L. (2008a), "The Economic Effects of Energy Price Shocks" *Journal of Economic Literature*, 46, 871-909.
- Kilian, L. (2008b), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do they Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90, 216-240.
- Kilian, L. (2009a), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053-1069.
- Kilian, L. (2009b), "Oil Price Shocks, Monetary Policy, and Stagflation," mimeo, University of Michigan.
- Kilian, L., and C. Vega (2008), "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices," mimeo, University of Michigan.
- Koop, G., Pesaran M.H., S.M. Potter (1996), "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 74, 119-147
- Leduc, S., and K. Sill (2004), "A Quantitative Analysis of Oil Price Shocks, Systematic Monetary Policy and Downturns," *Journal of Monetary Economics*, 51, 781-808.
- 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.
- Lee, K., S. Ni, and R.A. Ratti (1995), "Oil Shocks and the Macroeconomy: The Role of Price Variability," *Energy Journal*, 16, 39-56.
- 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, 740-744.
- Mork, K.A., O. Olsen, and H.T. Mysen (1994), "Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries." *Energy Journal*, 15, 19-35.
- Pindyck, R.S. (1991), "Irreversibility, Uncertainty and Investment," *Journal of Economic Literature*, 29, 1110-1148.