



Working Papers

www.cesifo.org/wp

The Role of Oil Price Shocks in Causing U.S. Recessions

Lutz Kilian
Robert J. Vigfusson

CESIFO WORKING PAPER NO. 5743
CATEGORY 10: ENERGY AND CLIMATE ECONOMICS
FEBRUARY 2016

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

The Role of Oil Price Shocks in Causing U.S. Recessions

Abstract

Although oil price shocks have long been viewed as one of the leading candidates for explaining U.S. recessions, surprisingly little is known about the extent to which oil price shocks explain recessions. We provide a formal analysis of this question with special attention to the possible role of net oil price increases in amplifying the transmission of oil price shocks. We quantify the conditional recessionary effect of oil price shocks in the net oil price increase model for all episodes of net oil price increases since the mid-1970s. Compared to the linear model, the cumulative effect of oil price shocks over course of the next two years is much larger in the net oil price increase model. For example, oil price shocks explain a 3% cumulative reduction in U.S. real GDP in the late 1970s and early 1980s and a 5% cumulative reduction during the financial crisis. An obvious concern is that some of these estimates are an artifact of net oil price increases being correlated with other variables that explain recessions. We show that the explanatory power of oil price shocks largely persists even after augmenting the nonlinear model with a measure of credit supply conditions, of the monetary policy stance and of consumer confidence. There is evidence, however, that the conditional fit of the net oil price increase model is worse on average than the fit of the corresponding linear model, suggesting much smaller cumulative effects of oil price shocks for these episodes of at most 1%.

JEL-codes: E320, E370, E510, Q430.

Keywords: real GDP, nonlinearity, asymmetry, time variation, conditional response, prediction.

*Lutz Kilian**
Department of Economics
University of Michigan
238 Lorch Hall, 611 Tappan Street
USA - Ann Arbor, MI 48109-1220
lkilian@umich.edu

Robert J. Vigfusson
Board of Governors of the
Federal Reserve System
USA - Washington, D.C. 20551
robert.j.vigfusson@frb.gov

*corresponding author

June 18, 2015

The views 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 Conor Foley for excellent research assistance and Bob Tetlow for providing the updated Federal Reserve Board Staff Financial Stress Index. We thank Christiane Baumeister, Jim Hamilton, Ana Maria Herrera, Kirstin Hubrich, and Vince Kaminiski for helpful comments.

1. Introduction

There is a long tradition of associating U.S. recessions with oil price shocks on the one hand and with shocks to monetary policy or credit on the other. Both approaches have evolved over time. For example, the original focus on unexpected reductions in monetary aggregates has given way to measures of the unexpected reduction in credit availability as a result of policy changes or of financial frictions (e.g., Romer and Romer 1989; Kiyotaki and Moore 1997; Bernanke, Gertler and Gilchrist 1999; Hall 2011; Gilchrist and Zakrajsek 2012). At the same time, the literature on the transmission of oil price shocks has increasingly shifted from linear models to models involving nonlinear transformations of the price of oil (e.g., Hamilton 1996, 2003). It has proved difficult to separate the role of these two explanations of the business cycle because most recessions are preceded both by higher oil prices and by a tightening of monetary policy and/or of credit markets (e.g., Hoover and Perez 1994, Barsky and Kilian 2002).¹

One would have expected that over time a recession would have occurred that is unambiguously associated with one explanation only, resolving this identification problem. This has not been the case. The latest example is the recession of 2007-09. Hamilton (2009) and Ramey and Vine (2011), on the one hand, stress the importance of oil price increases for this economic slowdown. Stock and Watson (2012), on the other hand, concede that higher oil prices contributed to this recession, but attribute the bulk of the economic decline to other causes including the worsening of U.S. and global credit conditions. Likewise, Christiano, Eichenbaum and Trabandt (2014) argue that the vast bulk of movements in aggregate real economic activity during the Great Recession were due to financial frictions interacting with the zero lower bound. For a related view also see Del Negro, Giannoni and Schorfheide (2015).

How important the role of oil price shocks is in explaining recessions appears to depend largely on how their relationship with U.S. real GDP growth is modelled. Whereas linear models of the business cycle tend to assign low explanatory power to oil price fluctuations, some nonlinear models attribute a much larger role to oil prices. This paper takes a fresh look at this

¹ There is also a large empirical literature on the related question of endogenous monetary policy responses to oil price shocks (e.g., Bernanke, Gertler and Watson 1997; Balke, Brown and Yücel 2002; Hamilton and Herrera 2004; Herrera and Pesavento 2009; Kilian and Lewis 2011; Bodenstein, Guerrieri and Kilian 2012). We abstract from this literature in the current paper because there is no empirical evidence that the Federal Reserve has been responding directly to oil price shocks in the 1970s or more recently. Moreover, there is no evidence that such a policy response would have had large cumulative effects on inflation and growth outcomes. Finally, recent work has shown that there is no theoretical rationale from a welfare point of view for a direct policy response to oil price fluctuations beyond the usual response to fluctuations in inflation and output.

question based on new econometric tools. Our analysis builds on the recent literature dealing with potential asymmetries and other nonlinearities in the relationship between oil prices and U.S. real GDP growth. The most important representative of nonlinear models in the literature is the net oil price increase model originally proposed by Hamilton (1996, 2003), which embodies an asymmetric treatment of oil price increases and oil price decreases.

Asymmetric models allow positive oil price shocks to have much larger recessionary effects because they amplify negative real GDP responses to positive oil price shocks, while dampening positive real GDP responses to negative oil price shocks. For a review of the economic rationale of asymmetric responses the reader is referred to Kilian (2014). While asymmetric models of the transmission of oil price shocks have been popular in applied work, a number of recent studies including Kilian and Vigfusson (2011a), Herrera, Lagalo and Wada (2011), and Herrera and Karaki (2013) have shown that there is no statistically significant asymmetry in the responses of U.S. aggregate real output and unemployment to oil price shocks. Similar results also hold for aggregate output in other OECD economies (see Herrera, Lagalo and Wada 2011). These response estimates are unconditional in that they are based on averages of the responses observed since 1973.

This unconditional evidence, however, does not rule out that the response of U.S. real GDP growth to oil price shocks conditional on specific dates may have been more asymmetric or more nonlinear. Such a conditional perspective is essential when assessing the role of oil price shocks in explaining recessions. A natural conjecture is that asymmetries or other nonlinearities tend to play a larger role on the eve of recessions than during other times, undermining the power of conventional unconditional tests of symmetry and linearity. We explore this question in detail. We begin in section 2 by applying a newly developed test of the null hypothesis of conditionally symmetric response functions to the net oil price increase model proposed by Hamilton (2003). This conditional test is derived from the same net oil price increase regression framework already employed in Kilian and Vigfusson (2011a,b). We find no evidence of departures from the null hypothesis of conditional symmetry at conventional significance levels. Closer inspection of the conditional response estimates for 1975-2012 reveals that the nonlinearity captured by the net oil price increase model does not primarily reflect asymmetry in the response functions, as is usually presumed. In fact, there is typically little conditional asymmetry in the response estimates even for large oil price shocks. Rather this model captures variation over time

in the magnitude of the response of real GDP to oil price shocks. In other words, the net oil price increase model effectively behaves like a time-varying parameter vector autoregressive model. This point is important because it undermines one traditional motivation for this class of models as capturing asymmetries in the transmission of oil price shocks due to oil price uncertainty, frictional unemployment or behavioral motives. Our evidence shows that this interpretation is at odds with the evidence.

An alternative explanation is that the magnitude depends on how vulnerable the economy is to shocks. We provide evidence that the magnitude of the conditional response is correlated with a number of indicators of macroeconomic conditions such as consumer confidence, financial stress, the share of oil in GDP, and interest rate expectations that may indicate a heightened vulnerability of the economy to oil price shocks, but most of the time variation appears linked to the prior evolution of the real price of oil.² Notwithstanding the observed time variation in the conditional response estimates, however, formal statistical tests fail to reject the hypothesis that the nonlinear conditional response estimates equal the response estimates in the linear model specification. This fact does not necessarily mean that we should abandon the nonlinear framework. Classical hypothesis testing protects the null model. It requires overwhelming evidence before the null model may be rejected. A failure to reject the null hypothesis of conditional linearity therefore does not establish that the linear null model is the correct model. It may simply reflect low power in small samples.

The fact that statistical tests do not help us discriminate between these models would be of little concern, if the substantive conclusions from these models were similar, but we show that this is not the case. In section 3, we introduce a conditional structural prediction error decomposition that allows us to quantify in a systematic and rigorous manner the role played by

² The seeming time variation in the estimates from the net oil price increase model could also be potentially explained to some extent by changes in the nature of the oil price shocks over time. As shown in Kilian (2009a) and Kilian and Murphy (2014), oil price shocks reflect a variety of demand and supply shocks in oil markets, the composition of which evolves over time. As each shock is associated with different responses in the price of oil and in U.S. real GDP, changes in the composition of oil demand and oil supply shocks can give the appearance of time variation in the response of real GDP to aggregate oil price shocks, when there is none. For example, one would expect oil price shocks to be less recessionary, if they reflect shifts in the flow demand for oil than if they are driven by oil supply shock or speculative demand shocks. A model ignoring this distinction will tend to overestimate the contribution of oil price shocks at some points in time and underestimate it at others in an effort to fit the data on average. Proponents of the net oil price increase model and related nonlinear models, while recognizing this point, however, have tended to downplay the distinction between oil demand and oil supply shocks (e.g., Hamilton 2003). In this paper, we adopt the same perspective and examine the net oil price increase model on its own terms in sections 2, 3, and 4, but we return to the distinction between oil demand and oil supply shocks in section 5.

oil price shocks in explaining U.S. recessions. We examine the empirical relationship between oil prices and U.S. real GDP using full-sample estimates based on data for 1974Q2-2012Q4 with the objective of comparing how the linear model and the net oil price increase model differ in their interpretations of these recessions. Our analysis shows that in many episodes the pattern of overall prediction errors implied by these models is qualitatively similar, but nonlinear models tend to assign much more importance to oil price shocks in explaining U.S. recessions than linear models. For example, in the linear model oil price shocks explain at most a 1% decline in real GDP after 2007Q2 with the remaining decline attributed to other shocks, whereas in the net oil price increase model oil price shocks explain a decline as high as 5% of U.S. real GDP. The fact that net oil price increase models of the transmission of oil price shocks suggest a very different interpretation of the U.S. recessions than linear vector autoregressive models means that the question of which model to rely on matters, even if statistical tests are unable to discriminate between these models.

From a policymakers' point of view in particular, this evidence reinforces the question of which model specification to trust more. This question comes down to a problem of model selection rather than of statistical testing. We address this question in section 4 by ranking models based on their conditional MSPE for each episode of interest. There are seven episodes of major net oil price increases in our sample, but only five recessions. We stress that the relevant question in this context is not whether recessions were preceded by net oil price increases, but rather whether net oil price increases have been followed by recessions. This distinction makes a difference. We show that in the baseline net oil price increase model the net oil price increase model improves the conditional fit on average by 22% when restricting attention to recession episodes, but worsens the conditional fit on average by 59% when considering all major net oil price increase episodes. Moreover, the improved conditional fit during recessions is driven by the recessions starting in 1979Q4 and 2007Q4, which are arguably primarily associated with a monetary policy shift under Paul Volcker and with the financial crisis, respectively. Excluding these two episodes, there is no improvement in conditional fit on average for the remaining recessions.

This result raises the question of whether net oil price increases merely proxy for other determinants of the U.S. business cycle. One concern, echoing the discussion in Hoover and Perez (1994) and Barsky and Kilian (2002) is that oil price increases and monetary or credit

contractions tend to occur at approximately the same time, making it difficult to determine their respective contributions to the U.S. business cycle. We address this question in section 4 by augmenting the baseline net oil price increase model with, respectively, the excess bond premium of Gilchrist and Zakrajsek (2012) as a measure of changes in credit supply conditions, with the term spread as an indicator of the stance of monetary policy, and with U.S. aggregate consumer sentiment. Adding consumer confidence in particular (or adding all three variables in conjunction), somewhat lowers the improvement in conditional fit obtained with the net oil price model relative to the linear model. For example, after adding consumer sentiment the average conditional fit of the net oil price model is 56% worse than for the corresponding linear model when considering all net oil price increase episodes, and 22% worse when considering recession episodes only.

Among all net oil price increase model specifications considered, in no case does the net oil price increase model have better conditional fit on average across all net oil price increase episodes than the linear model. This means that there is no compelling conditional evidence in support of the net oil price increase model, mirroring the unconditional evidence in Kilian and Vigfusson (2011a). Even after adding proxies for credit supply conditions, the monetary policy stance, and consumer sentiment, however, the conditional fit of the net oil price increase model for the 1979Q4 and 2007Q4 episodes remains superior. One possible explanation of this fact is that we failed to adequately capture the remaining determinants of the business cycle. Another possible explanation is that this result is accidental, that it is an indication of overfitting in small samples and that it would vanish if the sample were extended to include additional net oil price increase episodes. A third explanation is that there is a nonlinear effect of oil price increases after all, but that the specification of the net oil price increase model does not adequately capture this nonlinearity, causing the model to overstate the recessionary effect of oil price shocks.

An interesting question is why the recessionary effects associated with oil price shocks are so modest in the linear model. In the concluding remarks in section 5, we conjecture that a likely explanation is that oil price shocks are caused by a range of different oil demand and oil supply shocks in global markets with offsetting effects on the U.S. economy. Quantifying the role of these shocks for U.S. recessions would require a different class of econometric models, however, the development of which we defer to future research.

2. A Preliminary Look at the Relationship between the Real Price of Oil and U.S. Real GDP since 1974

It is common to stress that most recessions after 1973 have been *preceded* by oil price increases, which is often taken as evidence of recessions being caused by oil price increases. Even granting this observation, however, a more appropriate metric for the predictive power of oil prices is how often oil price increases have been *followed* by recessions. This change in focus matters.

A commonly used metric for sustained oil price increases is the net oil price increase measure of Hamilton (2003). Figure 1 plots the real price of oil measured by the U.S. refiners' acquisition cost for crude oil from 1974Q1 to 2012Q4.³ The refiners' acquisition cost for imports refers to the price of oil paid by U.S. refiners for crude oil purchased from abroad, which is a commonly used proxy for the global price of crude oil. The nominal price of oil is deflated using the U.S. consumer price index for all urban consumers. Periods in which there is a 3-year net increase in the real price of oil are marked by stars. The 3-year net oil price increase measure is defined as $\Delta r_t^{net,+3yr} \equiv \max(0, r_t - r_t^*)$, where r_t denotes the log of the real price of oil and r_t^* is the maximum log real price of oil over the most recent three years. The NBER business cycle dates have been superimposed on Figure 1 with recessions shown as shaded areas. The six NBER recessions are November 1973-March 1975, January 1980-July 1980, July 1981-November 1982, July 1990-March 1991, March 2001-November 2001, and December 2007-June 2009.

Treating the net oil price increases of 2004-06 as one episode, there have been eight distinct episodes of net oil price increases since 1974, of which only five were followed by recessions. In some cases the net oil price increase occurred well before the recession. A good example is the net oil price increase of 2000. In other cases it occurred immediately before or at the same time as the recession. Examples are 1981 and 1990. In three cases, net oil price increases were not followed by a recession at all. These episodes are 1996, 2004/05 and 2006 (the latter may be viewed as one episode), and 2011/12.

A natural conjecture would be that in the latter three episodes perhaps the net oil price increases were smaller than in those followed by recessions. This is not necessarily the case. For example, Figure 2 shows that the net oil price increase of 2011/12 is comparable in magnitude to

³ We ignore the data prior to 1974, given the evidence in Kilian and Vigfusson (2011b) of a major structural break in the relationship between oil prices and U.S. real GDP in 1973Q4. We drop the data after 20012Q4 because some of the other data we use in section 4 are not available beyond this date.

that of 1981.⁴ This evidence could be explained by the argument that a net oil price increase of this magnitude is not sufficient to cause a recession by itself. If so, the recession of 1981-82 was not caused by the net oil price increase of 1981 either and the timing is coincidence. Likewise, the net oil price increase of 2004-05 appears larger than that in 2000, and yet was not followed by a recession, which could be interpreted as evidence that the net oil price increase of 2000 was unrelated to the subsequent recession. This line of reasoning suggests that only five of eight net oil price increases have been associated with subsequent recessions, which is different from the conventional wisdom that virtually all recessions were related to net oil price increases.

Trying to explain the missing recessions in Figure 1 based on the magnitude of the net oil price increases does not help explain some episodes, however. In particular, the 2004-05 net oil price increase was as high as that of 1979 without being followed by a recession. There is little doubt that higher oil prices in 1979-80 played some role in the recession of 1980, so why did they not in 2005? Moreover, the 2004-05 net oil price increase was larger than the 2007 net oil price increase at the time when the 2007-09 recession started.

This evidence suggests that there is no mechanical link between net oil price increases and subsequent recessions, even controlling for the magnitude of the net oil price increase. One explanation of this evidence is that the recessionary impact associated with net oil price increases varies over time. In other words, the same net oil price increase sometimes is followed by deeper recessions than at other times. This perspective is new to the literature. Previous formal as well as informal studies of the relationship between net oil price increases and U.S. real GDP have typically taken an unconditional view by focusing on the average relationship across the sample (e.g., Hamilton 1996, 2003; Kilian and Vigfusson 2011a,b). In contrast, our focus in this paper is on the response of the U.S. economy conditional on the state of the economy when the oil price

⁴ The original analysis of net oil price increases in Hamilton (1996, 2003) was based on the nominal U.S. producer price index for crude petroleum instead of the real refiners' acquisition cost for crude oil imports. This choice affects the timing and magnitude of the net oil price increase measure. For example, Hamilton's definition yields a nominal net oil price increase that peaks at 28% in 1981Q1 and reaches 7% in 2012Q1. Subsequent research has considered both real and nominal net oil price increases. The use of our specific oil price definition is attractive for three reasons. One is that previous studies showed that the real U.S. refiners' acquisition cost for imports provides the strongest results in favor of asymmetric response functions and the most accurate out-of-sample forecasts of U.S. real GDP from net oil price increase models (see Kilian and Vigfusson 2011b, 2013). A second reason is that our oil price series is not subject to the distortions afflicting the corresponding domestic oil price series in the early part of the sample. Finally, the use of the real price of oil is consistent with conventional macroeconomic models of the transmission of oil price shocks (e.g., Hamilton 1988; Rotemberg and Woodford 1996; Backus and Crucini 2000).

shock occurs.

2.1. The Conditional Response of U.S. Real GDP to Oil Price Shocks

To examine more formally the relationship between net oil price increases and U.S. real GDP growth, it is useful to work with the following quarterly nonlinear dynamic structural model introduced by Kilian and Vigfusson (2011a), which may be estimated by the method of least-squares:

$$\begin{aligned}
\Delta r_t &= \alpha_1 + \sum_{i=1}^4 B_{11,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{12,i} \Delta y_{t-i} + e_{1,t} \\
\Delta y_t &= \alpha_2 + \sum_{i=0}^4 B_{21,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{22,i} \Delta y_{t-i} + \sum_{i=0}^4 \gamma_i \Delta r_{t-i}^{net,+3yr} + e_{2,t} \\
\Delta r_t^{net,+3yr} &\equiv \max(0, r_t - r_t^*)
\end{aligned} \tag{1}$$

where Δy_t denotes U.S. real GDP growth and $e_t = (e_{1,t}, e_{2,t})'$ is a vector of mutually uncorrelated i.i.d. error terms. We will refer to this model as the net oil price increase model.⁵

Net oil price increases may reflect exogenous as well as endogenous variation in the real price of oil. This observation cautions against treating the net oil price increase as exogenous with respect to U.S. real GDP in inferring the causal effects of oil price shocks. Model (1) imposes the weaker identifying assumption that the real price of oil is only predetermined with respect to U.S. real GDP. This assumption rules out instantaneous feedback from U.S. real GDP to the real price of crude oil imports, which is broadly consistent with empirical evidence in Kilian and Vega (2011). The same assumption has been maintained by related studies as far back as Rotemberg and Woodford (1996) and Hamilton (1996).

The response of U.S. real GDP growth to the exogenous oil price innovation $e_{1,t}$ depends on the magnitude of $e_{1,t}$ and on the lagged data in the model, denoted by Ω_{t-1} , that we condition on in forming expectations. This conditional impulse response at horizon h is defined as

$$I_{\Delta y}(h, \delta, \Omega_{t-1}) = E(\Delta y_{t+h} | e_{1,t} = \delta, \Omega_{t-1}) - E(\Delta y_{t+h} | \Omega_{t-1}),$$

where δ denotes the magnitude of the oil price innovation in the impact period. The

⁵ Local projection methods have occasionally been discussed as an alternative to model (1), but these methods are designed for prediction and cannot be used for constructing output responses to $e_{1,t}$ without knowledge of the structural impact multiplier matrix (e.g., Kilian and Kim 2011). Moreover, iterated forecasts from model (1) have been shown to be more accurate than direct forecasts from single-equation models, an advantage that one would also expect to carry over to impulse response analysis (see Kilian and Vigfusson 2013).

corresponding impulse response function for $h = 0, 1, \dots, H$ is denoted $I_{\Delta y}(H, \delta, \Omega_{t-1})$. It is common to express δ in multiples of the standard deviation σ of e_{1t} . In practice, expression (2) may be evaluated by Monte Carlo integration conditional on Ω_{t-1} adapting the procedure described in Kilian and Vigfusson (2011).⁶ For each response function, we average 200,000 draws. The response function of the level of real GDP, denoted by $I_y(H, \delta, \Omega_{t-1})$, is obtained by cumulating the responses of the growth rate. If we are interested in the effect of an oil price innovation of magnitude $\delta = 2\sigma$ and $\delta = -2\sigma$, respectively, conditional on 1990Q2, the quarter preceding the 1990-91 recession, we construct the response of U.S. real GDP for the following $h = 0, \dots, H$ quarters as

$$I_y(h, 2\sigma, \Omega_{1990Q2}) = E(y_{1990Q2+h} | e_{1,1990Q3} = 2\sigma, \Omega_{1990Q2}) - E(y_{1990Q2+h} | \Omega_{1990Q2})$$

and

$$I_y(h, -2\sigma, \Omega_{1990Q2}) = E(y_{1990Q2+h} | e_{1,1990Q3} = -2\sigma, \Omega_{1990Q2}) - E(y_{1990Q2+h} | \Omega_{1990Q2}),$$

for example.⁷

Model (1) nests the linear structural vector autoregressive model

$$\begin{aligned} \Delta r_t &= \alpha_1 + \sum_{i=1}^4 B_{11,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{12,i} \Delta y_{t-i} + e_{1,t} \\ \Delta y_t &= \alpha_2 + \sum_{i=0}^4 B_{21,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{22,i} \Delta y_{t-i} + e_{2,t} \end{aligned} \tag{2}$$

which is obtained by setting $\gamma_i = 0 \forall i \in \{0, \dots, 4\}$. In this linear model, the conditional response function no longer depends on Ω_{t-1} and equals the unconditional response function:

$$I_{\Delta y}(H, \delta, \Omega_{t-1}) = I_{\Delta y}(H, \delta),$$

for which closed-form solutions are readily available. The only effect of changing δ in this case is to scale the entire response function up or down.

The following example illustrates the importance of the history Ω_{t-1} for the conditional response of U.S. real GDP to innovations in the real price of oil. Figure 3 contains two sets of

⁶ An appendix summarizing this procedure can be found in the working paper version of this paper.

⁷ Oil price shocks can be decomposed into oil demand and oil supply shocks, as discussed in Kilian (2009a) and Kilian and Murphy (2014). The conditional responses ignore changes in the composition of oil demand and oil supply shocks, but may be interpreted as the expected change in real GDP given an oil price innovation of average composition.

results. The upper panel focuses on positive one standard deviation oil price shocks, while the lower panel considers positive two standard deviation shocks. A one standard deviation shock is usually considered a typical shock in that about two thirds of all oil price innovations are no larger than one standard deviation of e_{1t} . It implies a persistent 15% increase in the real price of oil. Two standard deviation oil price innovations are rare in that about 95% of oil price innovations are no larger than two standard deviations. Each panel includes three conditional response estimates that differ in the choice of Ω_{t-1} . The first estimate conditions on 1990Q3, when the real price of oil had just peaked; the second estimate conditions on 2008Q4 when the real price of oil had just reached a trough; and the third estimate conditions on 1994Q4 when the real price of oil was about average. These three examples were chosen to illustrate the range of possible outcomes associated with an oil price shock in the net oil price increase model. The plot also includes the response estimate from the linear model for comparison.

The upper panel of Figure 3 shows a strongly negative effect on real GDP after 1990Q3 because the real price of oil was already high when the additional positive oil price shock occurred. In contrast, after 2008Q4, the effect of the same shock on real GDP is essentially zero because even after the additional unexpected oil price increase, the real price of oil is too low for the net oil price increase to turn positive. Finally, after 1994Q4, there is a modestly negative effect of about the same magnitude as the response of U.S. real GDP in the linear model.

Doubling the magnitude of the positive oil price innovation, as shown in the lower panel of Figure 3, amplifies the conditional response after 1994Q4 because such a large shock is sufficient to render the net oil price increase positive and to activate the nonlinear term. It does not affect the response after 2008Q4, however, because even such a large shock is not large enough to activate the nonlinear features of the model. Finally, doubling the magnitude of the oil price innovation about doubles the already large negative response of real GDP after 1990Q3.

The obvious question is how the conditional response has evolved over time. Figure 4 shows how the response of U.S. real GDP at a horizon of eight quarters to an oil price innovation of constant magnitude has evolved over time since 1975. Similar results hold at the four-quarter horizon. The upper panel of Figure 4 focuses on positive and negative one-standard deviation oil

price shocks.⁸ In some periods, an unexpected 15% increase in the real price of oil has no effect on U.S. real GDP, while in other periods U.S. real GDP declines by more than 1%. There is no evidence of asymmetry in the point estimates. Up to sign the conditional responses to positive and negative oil price innovations are virtually identical at all points in time. The lower panel of Figure 4 repeats this analysis for two-standard deviation shocks. The conditional responses to such large shocks show signs of modest asymmetry, but the pattern across time of the responses to positive and negative oil price innovations is very similar (up to sign). The conditional responses of real GDP range from near zero to 2.5%, depending on the date of the oil price shock, and the pattern of time variation resembles that in the upper panel.

2.2. Testing for Conditional Symmetry in the Real GDP Responses

While there is no discernible asymmetry in the conditional responses of U.S. real GDP in response to one standard deviation shocks, Figure 4 shows weak evidence of conditional asymmetry in response to two standard deviation shocks. Table 1 examines this evidence of asymmetry in the response functions more formally. Building on Kilian and Vigfusson (2011), we implement a Wald test of the null hypothesis that the response functions associated with positive and negative shocks of the same magnitude δ are identical up to sign. Under the null hypothesis $H_0 : I_y(H, \delta, \Omega_{t-1}) = -I_y(H, -\delta, \Omega_{t-1})$, this test has an asymptotic χ^2_{H+1} -distribution where H denotes the maximum horizon of the impulse response function. The variance-covariance matrix of the test statistic $I_y(H, \delta, \Omega_{t-1}) + I_y(H, -\delta, \Omega_{t-1})$, is computed by bootstrap simulation.⁹ To limit the risk of data mining when conducting repeated tests for conditional symmetry for different H , we focus on selected episodes at horizons of one year and two years only. The first two columns of Table 1 show that there is no statistically significant evidence at the 10% level against the null of conditionally symmetric response functions at $H = 4$ or $H = 8$. This result holds for both 1σ and 2σ oil price shocks. In fact, this result holds even conditional on 1990Q3, for which the test statistic falls just short of standard significance levels.

⁸ Although it may seem that a negative oil price shock in the net oil price increase model would have no effect on U.S. real GDP, this is not the case. This response is positive because knowing that $e_{it} = -\delta$ reduces the probability of the real price of oil increasing enough to result in a net oil price increase. Thus, $E(\Delta y_{t+h} | \Omega_{t-1}, e_{it} = -\delta)$ is more positive than $E(\Delta y_{t+h} | \Omega_{t-1})$.

⁹ We generated 100,000 bootstrap replications from the unrestricted model by resampling the regression errors with replacement. Each conditionally nonlinear response function was computed as the average over 25,000 draws.

The weak evidence in Table 1 in conjunction with the weak degree of conditional asymmetry in the point estimates casts doubt on the view that asymmetries of the type embodied in the net oil price increase model are economically relevant. In fact, the net oil price increase model appears to behave much like a time-varying parameter VAR model of the real price of oil and U.S. real GDP.

2.3. Is the Time Variation in the Conditional Real GDP Responses Related to the State of the U.S. Economy?

So far we showed that, to the extent that the implications of the net oil price increase model differ from those of the linear model, the reasons are not related to any of the explanations of asymmetries in the relationship between oil prices and U.S. real GDP customarily invoked by users of these models. One potential explanation of this result is that the U.S. economy tends to be particularly vulnerable to oil price shocks prior to recessions. If so, one would expect larger conditional responses during recessions than during expansions. The upper panel of Figure 4, which NBER recessions superimposed as shaded areas, helps address this conjecture. The conditional responses in this figure indicate not how strong the recessionary effect of the actual oil price innovations was at each point in time, but rather how strong the hypothetical effect of a positive oil price innovation would have been, had it occurred at that point. These responses tend to be large at the beginning of many U.S. recessions, but equally large in 1975/76, 1996, 2004-06, and 2011/12 during expansions. Thus, there is no systematic relationship between the magnitude of the conditional responses and the state of the U.S. business cycle.

A natural conjecture is that the magnitude of the conditional response is correlated with indicators of how vulnerable the U.S. economy is to real oil price shocks. This conjecture is explored in Table 2. One indicator of the vulnerability of the U.S. economy is the Chicago Fed National Financial Conditions Index (NFCI). Because positive values of this index indicate a tightening of financial conditions, *ex ante* one would expect a negative correlation between the NFCI and the conditional response, if the response to positive oil price shocks is magnified in times of financial stress. Such a link seems plausible given the effects of oil price shocks on the spending power of households. Table 2 shows that the correlation in question is indeed negative and highly statistically significant. We focus on the NFCI because of all indicators of U.S. financial stress this index is available for the longest time span, but Table 2 also includes results for the Federal Reserve Board Staff Financial Stress Index, which is available since 1988 (see

Hubrich and Tetlow 2015). Again, there is evidence of a negative correlation, but the correlation is weaker and only marginally statistically significant at the 10% level. Either way, however, the predictive power of financial conditions for the conditional responses, as measured by the R^2 of a regression of the conditional response on indicator of interest and a constant, is at most 6%, suggesting that this link is of lesser importance.

Another possibility is that the vulnerability of the U.S. economy increases in the share of nominal oil use in U.S. nominal GDP. The role of variation in the energy share has been stressed by Edelstein and Kilian (2009), for example. If a higher U.S oil share amplifies the effects of oil price shocks, one would expect a negative correlation with the conditional response. Table 2 shows a highly statistically significant correlation of -0.463. While the fact that an oil price surge increases both the nominal oil share and the net oil price measure increase (and hence the conditional response) at the same time is not surprising, Table 2 suggests that 21% of the variation in the conditional response can be explained based on fluctuations in the oil share.

Another possibility is that the evolution of the conditional response reflects behavioral changes among consumers of oil. This explanation was first raised by Hamilton (1996) who explicitly linked net oil price increase measures to how “unsettling an increase in the price of oil is likely to be for the spending decisions of consumers and firms” (p. 216). The role of consumer confidence in the transmission of oil price shocks has also been stressed by Edelstein and Kilian (2009), Hamilton (2009), and Ramey and Vine (2011). We investigate this possibility by correlating the conditional response estimates with measures of U.S. consumer confidence in the economy, as reported in the *Michigan Survey of Consumers*. Higher levels of consumer confidence would be expected to be associated with less negative conditional responses. Table 2 confirms the existence of a positive correlation, but the correlation is low and only marginally statistically significant with an R^2 of only 2%.

The next three entries in Table 2 focus on buying conditions for homes, large household items and vehicles, as reported in the *Michigan Survey of Consumers*. Again one would expect a positive correlation in that above average buying conditions should be associated with above average (i.e., less negative) conditional responses. Table 2 shows that all three correlations are positive, but only those for homes and vehicles are large and highly statistically significant with an R^2 near 15% each. This evidence is important because of the importance attached to the automobile and housing sector in explaining the recessionary effects of positive oil price shocks

(see Hamilton 2009).

The remaining entries in Table 2 focus not on current business conditions, but on consumer expectations about changes in unemployment, changes in business conditions, changes in the interest rate, and changes in their financial situation one year ahead. An increase in the index of expected change in unemployment, for example, signifies lower expected unemployment, and hence a positive correlation with the conditional response. Likewise an increase in the index of the change in business conditions signals expectations of good economic times, and an increase in the index of expected changes in one's financial conditions signifies an expected improvement in financial conditions, so again the expected correlation with the conditional response is positive. This is consistent with the sign of the correlation estimates in Table 2, but none of the correlations are statistically significant and the R^2 estimates are near 0% in each case. Finally, an increase in the index of expected changes in the interest rate means an expectation of lower interest rates, and hence more favorable domestic economic conditions, which again suggests a positive correlation with the conditional response. Table 2 shows a positive correlation of 0.226, which is highly statistically significant and accounts for about 5% of the variation in the conditional response.

We conclude that the apparent time variation in the conditional responses is correlated with U.S. financial conditions, the nominal share of oil in U.S. GDP, buying conditions for homes and vehicles, and the expected change in the interest rate in particular, although none of these variables explains more than 21% of the fluctuations in the conditional response and jointly they explain only 38% of the variation in the magnitude of the conditional response. The remainder appears explained by the prior evolution of the real price of oil.

2.4. Testing for Linearity in the Conditional Responses

An obvious question is whether the evidence of time-varying conditional impulse responses in Figure 4 is statistically significant or perhaps simply random variation in the estimates. This is the question we turn to next. Traditional tests for linearity in the slope parameters as discussed in Hamilton (2003) and Kilian and Vigfusson (2011a), for example, are unconditional tests in that they are based on all sample observations. In the context of model (1) such an unconditional test comes down to testing $H_0 : \gamma_i = 0 \forall i$. One way of assessing the question of conditional linearity is to conduct instead a Wald test of $H_0 : I_y(H, \delta, \Omega_{t-1}) = I_y(H, \delta)$, where $I_y(H, \delta)$ denotes the

impulse response function in the linear VAR model (2). Under the null hypothesis of no time variation in the conditional response, all conditional responses should be equal to the unconditional (linear) response. Inference involves constructing the variance-covariance matrix of the difference in the impulse response functions by bootstrap simulation, taking account of the covariance of the estimates. The null hypothesis of conditional linearity is rejected if the Wald test statistic exceeds its χ^2_{H+1} -critical value.

The last two columns of Table 1 show that only conditional on 1990Q3, which represents the peak of the 1990 oil price spike, is there any evidence of departures from conditional linearity and even then only for $H = 4$. The rejection is statistically significant at the 10% for the 1σ shock, and at the 5% level for the 2σ shock. Although this evidence is stronger than when testing for conditionally asymmetric response functions, it remains somewhat tentative. It would be easy to disregard the evidence in Figure 4 on the basis of these test results. The fact that classical hypothesis tests do not reject the linear model except in extreme cases does not necessarily mean that we should abandon the nonlinear framework, however. Classical hypothesis testing protects the null model. It requires overwhelming evidence before the null model may be rejected. A failure to reject the null hypothesis therefore does not establish that the linear null model is the correct model. It may simply reflect low power in small samples.

An alternative approach to testing the linearity hypothesis based on local projections has been suggested by Hamilton (2011). Constructing nonlinear response functions from these local projections is not straightforward, nor would these nonlinear responses functions be identical to those we derived earlier from model (1), but Hamilton's approach lends itself to a simple unconditional test of the linearity hypothesis. It involves testing for each horizon h the null hypothesis that $\gamma_{h,1} = \dots = \gamma_{h,4} = 0$ in the system of local projection equations:

$$\Delta y_{t+h} = \alpha_{2,h} + \sum_{i=1}^4 B_{21,h,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{22,h,i} \Delta y_{t-i} + \sum_{i=1}^4 \gamma_{h,i} \Delta r_{t-i}^{net,+3yr} + e_{2,t+h}, \quad h = 0, 1, \dots, 8. \quad (3)$$

Table 3 shows evidence of highly statistically significant rejections of this null of linearity at several horizons. This evidence does not, however, account for the fact that Hamilton's approach involves a sequence of tests across h , making it more likely under the null that we find rejections than if only one test is conducted (see Inoue and Kilian 2004). The last column shows the corresponding data-mining robust p -value, obtained by bootstrapping the maximum of the Wald test statistic over h . The latter test rejects at the 10% level, but not at the 5% level, providing at

best weak evidence against the hypothesis of linearity. Of course, as shown in Jorda (2005), the nonlinear response function implied by the local projections in model (3) depends only on $\gamma_{h,1}$, $h = 0, \dots, 8$, and not on $\gamma_{h,i}$, $i > 1$. Hence a more direct approach to testing the hypothesis of nonlinearities in the impulse response function would be to test exclusion restrictions only on this subset of coefficients. Table 3 shows that, despite evidence against linearity at one horizon, after accounting for data mining, there is no evidence at all at conventional significance levels. This does not mean that the impulse response functions are indeed linear and symmetric. A failure to reject linearity in the slopes may still be consistent with the implied impulse response functions being nonlinear and asymmetric, as demonstrated in Kilian and Vigfusson (2011a).

3. How Much Do Oil Price Shocks Help Explain U.S. Recessions?

Section 2 established that statistical tests are unable to discriminate between the linear model and the net oil price increase model, even after conditioning on the relevant information. This does not mean that we do not care about the choice of model because these models may suggest very different interpretation of the U.S. recessions since the 1970s than linear vector autoregressive models. In this section, we demonstrate just how different these interpretations are. Perhaps surprisingly, this question has not been examined in any depth in the existing literature. We document systematically the cumulative effects of oil price shocks on U.S. real GDP during key historical episodes.

3.1. The Importance of Oil Price Shocks in Explaining Conditional Prediction Errors in U.S. Real GDP

Our analysis relies on a conditional structural prediction error decomposition computed from the estimates of models (1) and (2) based on data for 1974Q2-2012Q4. To assess the importance of oil price shocks for U.S. real GDP in the net oil price increase model we need to condition on the information available immediately before a net oil price increase occurred. Figure 5 considers all episodes of major net real oil price increases since the late 1970s, as shown in Figure 2, whether they were succeeded by recessions or not. We condition on the information available in 1979Q1, 1980Q4, 1990Q2, 1995Q4, 1999Q4, 2004Q2, and 2007Q2. These dates correspond to the quarter that preceded major net oil price increases in Figure 2. The stars in Figure 5 denote the total prediction errors for U.S. real GDP implied by the linear model and by the net oil price increase models during each of these seven episodes.

The total prediction error of a given model at horizon at h is defined as $y_{t+h} - E(y_{t+h} | \Omega_{t-1})$. A negative prediction error, for example, means that the model over-predicts U.S. real GDP. In other words, real GDP turned out lower than predicted. One can decompose the source of these prediction errors into the component explained by unforeseen oil price shocks, $e_{1,t+h}$, and the component explained by the unforeseen residual shocks, $e_{2,t+h}$. Whereas a conditional impulse response function only captures the cumulative effect on real GDP of a one-time oil price shock at date t , the objective of constructing such a conditional structural decomposition of the prediction errors in real GDP is to capture the cumulative effect of an entire sequence of structural oil price shocks for $t, t+1, \dots, t+h$.¹⁰

Model (1) explains the evolution of U.S. real GDP based on two structural innovations: an oil price innovation and a residual innovation capturing all other shocks that move U.S. real GDP. The conditional structural real GDP prediction error decomposition $D_{i,y}, i \in \{1, 2\}$, describes how real GDP would have evolved in the absence of further oil price innovations conditional on Ω_{t-1} :

$$D_{1,y}(h, \Omega_{t-1}) = E\left(y_{t+h} | \{e_{1,t+i}\}_{i=0}^h, \Omega_{t-1}\right) - E(y_{t+h} | \Omega_{t-1})$$

and in the absence of further residual innovations conditional on Ω_{t-1} :

$$D_{2,y}(h, \Omega_{t-1}) = E\left(y_{t+h} | \{e_{2,t+i}\}_{i=0}^h, \Omega_{t-1}\right) - E(y_{t+h} | \Omega_{t-1}).$$

The expectations in question must be evaluated by Monte Carlo integration. A similar analysis can be conducted with the linear VAR model (2). Whereas in a linear model by construction

$$D_{1,y}(h, \Omega_{t-1}) + D_{2,y}(h, \Omega_{t-1}) = y_{t+h} - E(y_{t+h} | \Omega_{t-1}),$$

in a nonlinear model there may be interactions between the two structural innovations. For example, a positive residual shock may be associated with higher oil prices, causing a subsequent positive oil price innovation to turn positive the net oil price increase measure. Such interactions are negligible in our applications, however.

To the extent that unpredicted fluctuations in real GDP can be explained by innovations to the real price of oil, we may infer from Figure 5 the importance of oil price shocks for

¹⁰ The idea of a conditional structural decomposition is related to, but distinct from that of a historical decomposition for stationary structural VAR models. For related ideas the reader is also referred to Baumeister and Kilian's (2014) discussion of forecast scenarios based on structural VAR models.

explaining real GDP in each regression model. Our objective is not to address the question which model is more likely in this section, but what the substantive differences are between the linear model and the net oil price increase model when it comes to understanding the recessionary effects associated with oil price shocks. Figure 5 shows that both the linear model and the net oil price increase model would have generated similar total prediction errors, conditioning on 1979Q1, but the net oil price increase model would have attributed these errors to a much larger extent to positive oil price shocks. Depending on the quarter of interest, between one third and all of the prediction errors between early 1980 and mid-82 are attributed to oil price shocks.

While the pattern of the overall prediction errors after 1980Q4 is again similar for the two models, Figure 5 shows that the net oil price increase model relies on a much larger negative contribution from the residual shock to offset the positive contribution of unexpectedly lower real oil prices from 1982 to 1983. In contrast, the linear model explains this episode largely without references to oil price shocks. This evidence is interesting because it does not match common accounts of this episode. Figure 5 illustrates that the net oil price increase model does not attribute the reduction in real GDP in 1982 and 1983 to positive oil price shocks. The entire decline is attributed to the residual shock. Put differently, the 1980 oil price shock was inconsequential, according to both the linear and the net oil price increase model.

After 1990Q2, one again sees a similar pattern in the total prediction errors for the two models, but the models imply very different accounts of how these errors came about. Whereas positive oil price shocks played a minor role in the 1990-91 recession according to the linear model, the net oil price increase attributes the bulk of the reduction in real GDP to positive oil price shocks. Regarding the evolution of U.S. real GDP after 1995Q4, the models agree that the positive prediction errors must be mainly attributed to the residual shock; the oil price shock played no important role in this episode, despite the net oil price increase in Figure 1. As to 1999Q4, again the models agree at least qualitatively, with the net oil price increase model assigning a somewhat larger role to positive oil price shocks in lowering U.S. real GDP.

In contrast, during the 2004Q2 episode, the net oil price increase model strongly underpredicted U.S. real GDP. This underprediction is explained by very large positive residual shocks after 2004Q2 in this model, which more than offset the recessionary effects of the positive oil price shocks after 2004Q2. While the linear model also mildly underpredicts U.S. real GDP, these prediction errors are largely unrelated to subsequent oil price shocks. Finally, the

results for 2007Q2 illustrate that the net oil price increase model attributes more than half of the prediction errors during the Great Recession period to oil price shocks, whereas the linear model views oil price shocks as tangential to the Great Recession. The two models agree that both the residual shock and the oil price shock lowered U.S. real GDP.

Figure 5 highlights that the net oil price increase model provides a substantially different economic interpretation of the evolution of U.S. real GDP in four of the seven episodes considered. These episodes are 1979Q1, 1990Q2, 2004Q2 and 2007Q2. Columns (1) and (2) of Tables 4a and 4b provide a succinct summary of the cumulative effects of oil price shocks on U.S. real GDP during each episode. Table 4a focuses on the net oil price increase model and Table 4b on the linear VAR model. Of particular interest are the recession episodes. The net oil price increase model, for example, suggests that the cumulative effects of oil price shocks had lowered U.S. real GDP by 3.5% eight quarters after 1979Q1. Similarly, this model attributes a decline in U.S. real GDP of 3.2% between 1990Q2 and 1992Q2 to the cumulative effects of oil price shocks. The largest cumulative effect of oil price shocks is a 5.1% decline in U.S. real GDP between 2007Q2 and 2009Q2. This is more than half of the decline in U.S. real GDP of 8.4% that could not have been predicted by the net oil price increase model as of 2007Q2. In sharp contrast, the linear model implies that the declines in U.S. real GDP that can be attributed to oil price shocks were only between -0.5% and -0.7% depending on the episode.

3.2. The Role of Credit Supply, the Monetary Policy Stance, and Consumer Confidence

The evidence presented so far raises the question of whether net oil price increases merely proxy for other determinants of the business cycle. One concern, echoing the discussion in Hoover and Perez (1994) and Barsky and Kilian (2002) is that oil price increases and credit contractions tend to occur at approximately the same time, making it difficult to determine their respective contributions to the U.S. business cycle. This is not the only alternative explanation of U.S. recessions, however. We address this concern in Table 4 by augmenting the baseline net oil price increase model (1) by one of three alternative additional variables that are widely believed to be determinants of the business cycle.

The first variable is the excess bond premium, as defined in Gilchrist and Zakrajsek (2012). The excess bond premium is a measure of variation in credit supply conditions proposed by Gilchrist and Zakrajsek (2012). The inclusion of this variable is important in particular during the financial crisis, but variation in credit supply conditions more generally has been shown to

explain fluctuations in economic activity over and above lags of real GDP.¹¹

We also consider an alternative specification of the net oil price increase model obtained by adding the term spread, defined as the spread between the ten-year treasury bond yield and the 5-year treasury bond yield. It is widely believed that fluctuations in the term spread help predict recessions (see, e.g., Wheelock and Wohar 2009). Many studies attribute the apparent predictive ability of the term spread for real GDP to actions by monetary authorities to stabilize output growth, while other attribute it to consumption smoothing in response to expected changes in short-term interest rates. The yield curve has also been interpreted as a measure of the stance of monetary policy. The use of the term spread allows us to capture changes in the stance of monetary policy without taking a stand on the identification of monetary policy shocks. This feature is important because of the changes in monetary policy regimes that have taken place over the course of our sample.

Finally, we consider a net oil price increase model augmented by the overall confidence in the U.S. economy as reported in the *Michigan Survey of Consumers*. The role of innovations to consumer confidence for explaining fluctuations in U.S. economic activity has been stressed, for example, by Barsky and Sims (2012). As shown in section 2, there is no tight relationship overall between variation in consumer confidence and net oil price increases. Our model is designed to permit consumer confidence to respond to oil price shocks, but also allows for exogenous variation in consumer confidence to affect U.S. real GDP.

As in the baseline model, the augmented structural models treat the real price of oil as predetermined with respect to the U.S. economy. Consider the example of the excess bond premium discussed in column (3). The augmented net oil price increase model is

$$\begin{aligned}
\Delta r_t &= \alpha_1 + \sum_{i=1}^4 B_{11,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{13,i} \Delta y_{t-i} + \sum_{i=1}^4 B_{12,i} EBP_{t-i} + e_{1,t} \\
\Delta y_t &= \alpha_2 + \sum_{i=0}^4 B_{21,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{22,i} \Delta y_{t-i} + \sum_{i=1}^4 B_{23,i} EBP_{t-i} + \sum_{i=0}^4 \delta_i \Delta r_{t-i}^{net,+3yr} + e_{2,t} \\
EBP_t &= \alpha_3 + \sum_{i=0}^4 B_{31,i} \Delta r_{t-i} + \sum_{i=1}^4 B_{32,i} \Delta y_{t-i} + \sum_{i=1}^4 B_{33,i} EBP_{t-i} + e_{3,t} \\
\Delta r_t^{net,+3yr} &\equiv \max(0, r_t - r_t^*)
\end{aligned} \tag{4}$$

where EBP_t denotes the monthly average of the excess bond premium. We follow Gilchrist and

¹¹ An updated version of this series was obtained from <http://people.bu.edu/sgilchri/Data/data.html>.

Zakrajsek (2012) in postulating that real GDP growth (and hence the real price of oil) does not respond within the quarter to a shock to the excess bond premium, but that the excess bond premium as a financial variable responds instantaneously to all real variables. The structural errors $e_t = (e_{1t}, e_{2t}, e_{3t})'$ are mutually uncorrelated and serially uncorrelated. Because the real price of oil is predetermined in model (4), the response of U.S. real GDP to the oil price shock is unaffected by the implicit ordering of the remaining structural shocks in model (4). It is important to stress that the augmented structural model allows the oil price shock to affect real GDP both directly and through its effect on the augmented variable.

Columns (3) and (4) in Table 4a show that augmenting the net oil price increase model by the excess bond premium slightly lowers the explanatory power of oil price shocks in the three episodes considered above, but does not change the qualitative result. Perhaps surprisingly, this is true even for the financial crisis episode. A similar result holds when augmenting the net oil price increase model with the term spread instead. As shown in columns (5) and (6), the latter change actually slightly increases the explanatory power of oil price shocks in some cases. There is no indication that the monetary policy stance soaks up the explanatory power of oil price shocks in the net oil price increase model. Similarly, columns (7) and (8) of Table 4a show that adding consumer confidence to the baseline net oil price increase model instead may reduce the explanatory power of oil price shocks, but again without overturning the qualitative results for the baseline model.¹² Finally, Table 4a considers a model that includes all three extra variables in an effort to account for possible interactions between these variables. The last two columns of Table 4a confirm that after including all three additional business cycle determinants in the net oil price increase model, the results for the baseline model remain qualitatively intact. The cumulative impact of oil price shocks at the eight-quarter horizon remains virtually unchanged at -3.5% after 1979Q1. For 1990Q2, it drops from -3.2% in the baseline model to -2.5%. For 2007Q4, it drops from -5.1% to -4.3%.

The corresponding results for the linear VAR model in Table 4b show that augmenting the baseline linear model by all three extra variables raises the explanatory power of oil price

¹² A case can be made that the augmented model including consumer confidence should also allow for direct feedback from lagged net oil price increases to consumer confidence. Table 2a reports the results for the more parsimonious specification (4) in which this direct feedback is ruled out. Additional results (not included to conserve space) show that this restriction does not change the qualitative results. The only difference is that in many cases the explanatory power of oil price shocks for U.S. real GDP is closer to the baseline model. Either way, the explanatory power of oil price shocks in the net oil price increase model is remarkably robust.

shocks from -0.5% to -0.9% after 1979Q1, from -0.1% to -0.4% after 1990Q2, and leaves almost unchanged the -0.7% decline in U.S. real GDP after 2007Q2. Taken in conjunction, the estimates in Tables 4a and 4b considerably narrow down the range of uncertainty about the ability of oil price shocks to explain recessions, even if one wishes to remain agnostic about the role of net oil price increases, but they also highlight that linear models assign much less importance to oil price shocks than net oil price increase models. The differences may be as high as 3.5 percentage points of U.S. real GDP in some cases.¹³

These results for the linear model are somewhat sensitive to the choice of the autoregressive lag order, however. We followed the conventional choice of four autoregressive lags in the linear benchmark model maintained by Hamilton (2003, 2011). It can be shown that increasing the order to 8 lags may increase the explanatory power of oil price shocks slightly to -1.2% two years after 1979Q1, -0.9% two years after 1990Q2, and -1.0% two years after 2007Q2. Even when one or more of extra variables are included in the higher-order VAR model, however, we find that the linear model explains at most about a 1% decline in U.S. real GDP after 1979Q1, after 1990Q2, and after 2007Q4, leaving a gap as high as 3 percentage points of U.S. real GDP between the explanatory power of oil price shocks in the linear model and in the net oil price increase model.

4. Evaluating the Conditional Fit of the Linear and Net Oil Price Increase Models after Net Oil Price Increases

The evidence presented so far suggests that we may not know whether the linear or the nonlinear model is the correct model, but that we clearly care about the choice of model. Even in the absence of statistically significant evidence against the linear and symmetric model, one reason to be concerned with the possibility that the correct model is nonlinear is indirect evidence in favor of a nonlinear relationship between oil prices and U.S. real GDP growth from predictive regressions. While correct models may fail to predict well, it is much harder to explain how an incorrect model could have predictive power. Predictive regressions have played an important

¹³ The low explanatory power of oil price shocks in the linear model actually is not unexpected given the changes over time in the composition of oil price shocks first documented in Kilian (2009a). Oil price increases associated with increased flow demand for crude oil driven by demand from abroad, for example, are consistent with extended periods of robust growth, as observed during 2003-07 (see Kilian and Hicks 2013). To the extent that different types of positive oil demand and negative oil supply shocks may produce responses in U.S. real GDP of opposite sign, the average effect of oil price shocks on real GDP will be dampened. Indeed, it can be shown that the response of U.S. real GDP to oil price shocks has declined in recent years, when demand shocks dominated the evolution of the real price of oil.

role in shaping the profession's perceptions of asymmetries and nonlinearities in the relationship between oil prices and the economy, dating back to the work of Mork (1989). Most importantly, Hamilton (2003, 2011) provided evidence of a nonlinear predictive relationship in population between net oil price increases and U.S. real GDP growth. Hamilton's evidence was based on measures of in-sample fit. More recently, the focus has shifted to simulated out-of-sample analysis. For example, Kilian and Vigfusson (2013) and Ravazzolo and Rothman (2013) confirmed that at longer horizons there is evidence that including nonlinear transformations of the price of oil in autoregressive forecasting models of real GDP growth may increase forecast accuracy, although this evidence appears driven by inclusion of data from the financial crisis period and vanishes when excluding data after 2007Q4.¹⁴

From a policymakers' point of view, this evidence reinforces the question of which model specification to trust more. Answering this question involves solving a problem of model selection rather than of statistical testing. We address this question by ranking alternative models based on their ability to anticipate U.S. real GDP fluctuations following net oil price increases. Several studies have compared the unconditional fit of the net oil price increase model and the linear model for U.S. real GDP growth (e.g., Hamilton 2003; Kilian and Vigfusson 2011a,b). Our objective in this subsection, in contrast, is to assess the conditional ability of the net oil price increase model for U.S. real GDP. Our approach differs from that in section 3 in that we examine the extent to which the net oil price increase model provides a better fit of U.S. real GDP conditional on the net oil price increase already having occurred. The information set matters here. Even if the net oil price increase model is the preferred model, one would not expect this model to fit the data better than the linear model before a sustained oil price increase has taken place. Only conditional on such an oil price increase already having occurred, one would expect the net oil price increase model to reveal its superior conditional fit.

Surprisingly, the conditional performance of these models has remained largely unexplored in the literature. The work most closely related to our analysis is Hamilton (2009) who first explicitly recognized the importance of conditional analysis. There are several

¹⁴ Kilian and Vigfusson (2013) also documented that the net oil price increase model proposed by Hamilton (2003) is not the most successful forecasting specification in this context, mirroring our evidence in this paper for conditional impulse responses. A net oil price change specification which imposes symmetry in positive and negative oil price changes, but allows for nonlinearities, for example, is a more accurate and more reliable forecasting model. This finding reinforces our evidence in section 2 that conditional impulse responses differ from responses estimated from linear models not based on their asymmetry so much, but based on their variation over time.

important differences, however: (1) Hamilton’s analysis was not based on a structural econometric model, but based on predictive models. (2) It combined data from before the 1973Q4 structural break in the relationship between the real price of oil and U.S. real GDP growth with data since then. (3) Hamilton provided little detail beyond averages of real GDP growth predicted with and without conditioning on the path of the price of oil. Our discussion below provides the basis for a more comprehensive and systematic analysis by episode with special attention to the timing of net oil price increases.

A natural measure of the conditional fit of the net oil price increase model is the *conditional MSPE by episode* which is defined as

$$MSPE_t(8, \Omega_{t-1}) = \sum_{h=0}^8 \left(y_{t+h} - E(y_{t+h} | \Omega_{t-1}) \right)^2,$$

where $E(y_{t+h} | \Omega_{t-1})$ is estimated by Monte Carlo integration conditional on the model estimated on the full sample. The same definition also applies to the linear model. For a given episode, the conditional MSPE is estimated by averaging the squared cumulative prediction errors of real GDP growth shown in Figure 5 across horizons $h = 0, \dots, 8$.

4.1. The Baseline Net Oil Price Increase Model

Table 5 reports the ratio of the conditional MSPE of the net oil price increase model relative to that of the linear model. A ratio below unity indicates that the net oil price increase model is more successful at explaining the variation of U.S. real GDP during the recession in question. The objective of this exercise is to evaluate the conditional fit after an oil price increase has taken place. The starting point of the first five episodes of interest corresponds to the recession dates determined by the NBER business cycle dating committee.¹⁵ Table 5 shows the conditional MSPE ratio by episode. The last two rows report the mean of the conditional MSPE ratios first for the net oil price increases followed by a recession, and, second, for all net oil price increases regardless of whether they were followed by a recession or not.

Column (1) in Table 5 shows that the net oil price increase model yields substantial improvements in fit during the recession starting in late 1979 and during the recession starting in late 2007, but it yields only small improvements during the recession starting in 1981 and 1990.

¹⁵ As a rule, these recessions were preceded by net oil price increases. The 1990Q2 recession date is an exception. Because this date precedes the 1990Q3 oil price spike triggered by the invasion of Kuwait, we adjust this date by one quarter.

During the 2001 recession, in contrast, the linear model fits as well as or better than the net oil price increase model. The average conditional MSPE ratio across all five recessions is 0.78. One interpretation of this evidence is that the net oil price increase model indeed helps explain U.S. recessions. An alternative interpretation would discount the evidence for two of the five recessions. A good case can be made that the fit of the net oil price increase model after 1979Q4 is spurious and confounds the effects of Paul Volker's tightening of monetary policy with that of unexpectedly higher oil prices (see Kilian and Lewis 2011). A similar concern arises regarding the recession after 2007Q4. The concern is that the net oil price increase model confounds the effects of the financial crisis on U.S. real GDP with the effects of unexpectedly higher oil prices (see Kilian and Vigfusson 2013). If we discard the first and the last episode, as shown in column (2) of Table 5, the mean conditional MSPE ratio increases to 1.01, indicating that the nonlinear dynamics associated with net oil price increases do not help understand the origin of U.S. recessions.

There is another concern as well. The recession episodes in Table 5 all were preceded by major oil price increases. Focusing on these recession episodes only is subject to the critique that we select on the outcome of a recession rather than considering all major oil price increases in the data and verifying that conditioning on these oil price increases improves the explanatory power of the model for subsequent changes in U.S. real GDP. The latter approach seems more relevant from a policymaker's point of view. As Figure 2 shows, one notable episode of persistently rising oil prices that was not followed by a recession was 2004Q4. Another one was 1996Q3. The conditional MSPE ratios for these episodes in Table 5 reflect the fact that the net oil price increase model underpredicted real GDP following these net oil price increases. Including these two episodes in computing the mean conditional MSPE ratio increases the mean from 0.78 to 1.59 when including all recessions and from 1.01 to 2.06 when excluding the first and last recession. This result is important because, as discussed in the introduction, we are interested not in whether recessions were preceded by net oil price increases so much, but in whether net oil price increases are associated with subsequent recessions. Table 5 shows that being precise about the question asked may change the answer.

The evidence in Table 5 suggests that even when taking a conditional perspective, it is difficult to make the case for the net oil price increase model based on the model's conditional fit. The conditional perspective also highlights that the evidence is limited to a handful of

episodes only. In fact, depending on one's interpretation of the 1979Q4 and 2007Q4 episodes, the case can be made that including net oil price increases in the model actually worsens the conditional fit relative to the linear model. The exclusion of these two episodes is not arbitrary. In fact, these are episodes which many economists associate with explanations not related to oil prices. The improvement in fit during U.S. recessions vanishes altogether when including episodes of net oil price increases not followed by recessions, revealing a tendency of the net oil price increase model to signal stronger declines in U.S. real GDP during net oil price increases than supported by the data.

4.2. Sensitivity Analysis for the Conditional Fit

The remaining columns of Table 5 examine the question of how much of the improvement in conditional fit during some episodes merely reflects a spurious correlation with other determinants of the business cycle. As in our earlier analysis, the analysis is based on a net oil price increase model augmented by, respectively, the excess bond premium, the term spread and U.S. consumer confidence. Column (3) in Table 5 shows that adding the excess bond premium slightly worsens the conditional MSPE ratio of the nonlinear model after 1979Q4 and 2007Q4. Overall the results are qualitatively similar to the baseline model except for 2001Q1. For the latter episode adding the excess bond premium substantially improves the conditional fit of the net oil price increase model. Based on all seven episodes, the average conditional MSPE ratio nevertheless is 1.17, confirming the earlier findings in favor of the linear model, whereas conditioning on recessionary outcomes the average ratio is 0.64. We conclude that the omission of the excess bond premium does not explain the 1979Q4 and 2007Q4 results for our baseline model in column (1).

Column (4) considers a net oil price increase model augmented with the term spread instead. This yields an average conditional MSPE ratio of 1.944 overall and of 0.742 when conditioning on recession outcomes. Again, this is qualitatively similar to the baseline. There is no indication that including monetary determinants of real GDP growth removes the explanatory power of models based on net oil price increases.

Column (5) considers an augmented model with U.S. consumer confidence added. This addition worsens the conditional fit of net oil price increase model considerably. The average conditional MSPE ratio exceeds 1 by a wide margin even when conditioning on recessionary outcomes. The average loss of conditional fit is 23% when focusing on recessions only. The

overall average loss of fit is 56%. We conclude that based on the average conditional fit of the net oil price increase model, after allowing for recessionary effects of shifts in consumer confidence not related to real oil price shocks, there is no systematic evidence that nonlinear dynamics associated with net oil price increases help us understand recessions in general.

An obvious question is how the conditional fit would be affected if we included all three additional variables at the same time. Column (6) in Table 5 shows that the results are similar to those in columns (3) and (4) in that the average conditional MSPE ratio is 1.19 overall, once again indicating that the linear model fits the data better. This compares with an average ratio of 0.83 when focusing on recessions only. We conclude that for each of our specifications the mean of the conditional MSPE ratio exceeds 1 when considering all major net oil price increase episodes. Thus a policymaker facing another net oil price increase in the future would have no reason to rely on the net oil price increase model based on historical evidence.

Table 5 also shows, however, that in the 1979Q4 and 2007Q4 episodes including net oil price increases in the regression model improves the conditional fit for every specification we considered. One interpretation of this evidence is that we failed to utilize an accurate measure of the determinants of the Volcker recession and of the financial crisis. This is certainly possible, although it is surprising that even the best measures currently available do little to undermine the predictive power of net oil price increases. An alternative interpretation is that this evidence reflects overfitting in small samples and would vanish if the sample could be increased. A third interpretation is that there is a nonlinear effect of oil price increases after all, but that the specification of the net oil price increase model does not adequately capture this nonlinearity resulting in false positives for several episodes. Resolving these questions is beyond the scope of this study.

5. Concluding Remarks

Much of the previous work on the ability of oil price shocks to explain recessions has focused on linear or pseudo-linear models in which the data are transformed, so linear models may be applied (e.g., Edelstein and Kilian 2009; Kilian and Lewis 2011; Ramey and Vine 2011). If the link between oil prices and the U.S. economy is inherently nonlinear, however, it becomes important to condition on the state of the economy prior to the oil price shock. Earlier studies examining the evidence for nonlinearities have tended to focus on the unconditional (or average) response of U.S. real GDP to oil price shocks across the sample rather than the conditional

response of U.S. real GDP during recessions (e.g., Kilian and Vigfusson 2011a,b; Hamilton 2003, 2011).

In contrast, the current paper focused on the conditional impact of oil price shocks during and following episodes of net oil price increases. We first documented the history of net oil price increases since 1974. It is widely known that most recessions since the 1970s were preceded by net oil price increases. We made the case that the relevant question is not whether recessions were preceded by net oil price increases, but whether net oil price increases were followed by recessions. Our findings are that by no means all net oil price increases in our sample appear to have been followed by recessions. The lack of a mechanical relationship between net oil price increases and recessions is consistent with the hypothesis that the recessionary effects of oil price shocks are time-varying. Quantifying these effects requires a structural net oil price increase model of the type proposed in Kilian and Vigfusson (2011a).

We provided for the first time a comprehensive analysis of the conditional impact of oil price shocks on U.S. real GDP during episodes of major net oil price increases. We demonstrated that asymmetries play no important role in generating these conditional response estimates, but rather variation over time in the magnitude of the conditional response of real output to oil price shocks. The latter result contradicts the usual premise among macroeconomists that asymmetries are required to generate large recessions. A central question then is why oil price shocks would have larger recessionary effects at some times than at others. We investigated whether this phenomenon might be linked to the state of the U.S. economy. Although there was some evidence of such time variation, most fluctuations in the magnitude of the response of real output appear linked to the prior evolution of the real price of oil. Further analysis showed that there is no evidence of statistically significant departures from conditional linearity.

We nevertheless proceeded to quantify the recessionary effects of oil price shocks for selected episodes of interest. We examined every NBER recession since the late 1970s. We also examined episodes in which recessions failed to occur after net oil price increases. We concluded that there is substantial disagreement between linear models and the net oil price increase on the quantitative importance of oil price shocks in explaining recessions. Only the net oil price increase model implies that oil innovations tend to be associated with a major drag on the U.S. economy. For example, oil price shocks in the baseline net oil price increase model cumulatively lowered U.S. real GDP by 3% two years after the Iranian Revolution, while explaining a 3%

reduction in U.S. real GDP two years after the invasion of Kuwait in 1990 and a 5% reduction two years after mid-2007. In the linear model, the corresponding cumulative reduction is under 1%, in contrast. The net oil price increase model does not always behave as expected, however. After the outbreak of the Iran-Iraq War in late 1980, according to the net oil price increase model, oil price shocks cumulatively raised U.S. real GDP by 1.6% after two years, while the linear model showed no cumulative effect of oil price shocks.

The fact that net oil price increase models of the transmission of oil price shocks suggest a very different interpretation of the U.S. recessions than linear vector autoregressive models means that the question of which model to rely on matters, even if statistical tests are unable to discriminate between these models. We also considered the possibility that the explanatory power of oil price shocks arises merely because of a correlation between net oil price increases and measures of credit supply conditions, of the monetary policy stance and of consumer confidence. We were able to diminish the explanatory power of oil price shocks for recessions after augmenting the net oil price increase model with U.S. consumer confidence in particular, but were unable to lower it to levels consistent with linear models. We found that oil price shocks in the net oil price increase model augmented by all three extra variables typically explain reductions in U.S. real GDP between 2% and 4% at a horizon of two years into net oil price increase episodes compared with less than 2% and sometimes less than 1% in linear models.

Given that linear models and the net oil price increase model assign a very different importance to oil price shocks in explaining these recessions, the fact that statistical tests cannot discriminate between the linear model and the net oil price increase model does not mean that the choice of model does not matter to policymakers. We addressed this question by comparing the conditional fit of these models for the eight quarters following net oil price increases in the data. This model selection exercise suggests that the linear model fits the data better on average after net oil price increases than the net oil price increase model.

This is not to deny that the net oil price increase model nevertheless appears to provide a superior fit for the 1979Q4 and 2007Q4 recessions in particular. One explanation of this fact may be that we failed to find a good measure of the determinants of the Volcker recession and of the financial crisis. This is certainly possible, although it is surprising that even the best measures currently available do little to undermine the predictive power of the net oil price increase model in these episodes. An alternative interpretation is that this evidence reflects overfitting in small

samples and would vanish if the sample could be increased. A third interpretation is that there is a nonlinear effect of oil price increases after all, but that the specification of the net oil price increase model does not adequately capture this nonlinearity resulting in false positives for several episodes. The latter interpretation suggests that further work on nonlinear transmission channels may be required.

If one accepts our evidence in favor of the linear model, the question arises why the recessionary effects associated with oil price shocks are so modest in the linear model. This result may seem at odds with the long-standing perception that real oil price shocks are clearly recessionary. One explanation is that oil price shocks are associated with a range of oil demand and oil supply shocks, some of which stimulate the U.S. economy in the short run and some of which slow down U.S. growth (see Kilian 2009a). How recessionary the response to an oil price shock is thus depends on the average composition of oil demand and oil supply shocks over the sample period. Moreover, there is a tendency for the positive and negative effects on growth associated with oil price shocks to average out, which helps explain the muted responses we observed. A good case can therefore be made that rather than studying the response to an average oil price shock, one ought to consider the response of U.S. real GDP to each type of oil demand and oil supply shock. This would require a different class of econometric models that explicitly incorporates quarterly U.S. real GDP growth into a structural model of the global oil market. Such an extended model, building on the work of Kilian and Murphy (2014) and Kilian, would help in understanding the role of the deeper determinants of the real price of oil in explaining recessions, while avoiding the generated regressor problem and loss of degrees of freedom associated with alternative econometric approaches. Developing such a framework is nontrivial, however, because many of the identifying assumptions currently in use in structural oil market models are specific to monthly data, and hence is left for future research.

References

- Backus, D.K., and M.J. Crucini (2000), "Oil Prices and the Terms of Trade," *Journal of International Economics*, 50, 185-213.
- 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.
- Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual*, 16, 137-183.

- Barsky, R.B., and E.R. Sims (2012), “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence,” *American Economic Review*, 102, 1343-1377.
- Baumeister, C., and L. Kilian (2014), “Real-Time Analysis of Oil Price Risks using Forecast Scenarios,” *IMF Economic Review*, 62, 119-145.
- Bernanke, B.S., Gertler, M., and S. Gilchrist (1999), “The Financial Accelerator in a Quantitative Business Cycle Framework,” in: Taylor, J.B., and M. Woodford (eds.), *Handbook of Macroeconomics*, Amsterdam: Elsevier Science, 1341-1393.
- Bernanke, B.S., Gertler, M., and M.W. Watson (1997), “Systematic Monetary Policy and the Effects of Oil Price Shocks,” *Brookings Papers on Economic Activity*, 1, 91–142.
- Bodenstein, M., Guerrieri, L., and L. Kilian (2012), “Monetary Policy Responses to Oil Price Fluctuations,” *IMF Economic Review*, 60, 470-504.
- Christiano, L.J, Eichenbaum, M.S., and M. Trabandt (2014), “Understanding the Great Recession,” forthcoming: *American Economic Journal: Macroeconomics*.
- Del Negro, M., Giannoni, M.P., and F Schorfheide (2015), “Inflation in the Great Recession and New Keynesian models,” *American Economic Journal: Macroeconomics*, 7, 168-196.
- Edelstein, P., and L. Kilian (2009), “How Sensitive Are Consumer Expenditures to Retail Energy Prices?” *Journal of Monetary Economics*, 56, 766-779.
- Gilchrist, S., and E. Zakrajsek (2012), “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102, 1692-1720.
- Hall, R.E. (2011), “The High Sensitivity of Economic Activity to Financial Frictions,” *Economic Journal*, 121, 351-378.
- 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. (2009), “Causes and Consequences of the Oil Shock of 2007-08,” *Brookings Papers on Economic Activity*, 1, Spring, 215-261.
- Hamilton, J.D. (2011), “Nonlinearities and the Macroeconomic Effects of Oil Prices,” *Macroeconomic Dynamics*, 15, 472-497.
- Hamilton, J.D., and A.M. Herrera (2004), “Oil Shocks and Aggregate Economic Behavior: The

- Role of Monetary Policy,” *Journal of Money, Credit & Banking*, 36, 265-86.
- Herrera, A.M., and M.B. Karaki (2013), “The Effects of Oil Price Shocks on Job Reallocation,” mimeo, University of Kentucky.
- Herrera A.M., Lagalo, L.G., and T. Wada (2011), “Oil Price Shocks and Industrial Production: Is the Relationship Linear? *Macroeconomic Dynamics*, 15, 472-497.
- Herrera, A.M., and E. Pesavento (2009), “Oil Price Shocks, Systematic Monetary Policy, and the Great Moderation,” *Macroeconomic Dynamics*, 13, 107-37.
- Hubrich, K., and R.J. Tetlow (2015), “Financial Stress and Economic Dynamics: The Transmission of Crises,” *Journal of Monetary Economics*, 70, 100-115.
- Hoover, K.D., and S.J. Perez (1994), “Post Hoc Ergo Propter Once More: An evaluation of ‘Does monetary policy matter?’ in the spirit of James Tobin,” *Journal of Monetary Economics*, 34, 47-73.
- Jorda, O. (2005), “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161-182.
- Kilian, L. (2001), “Impulse Response Analysis in Vector Autoregressions with Unknown Lag Order”, *Journal of Forecasting*, 20, 161-179.
- 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), “Comment on ‘Causes and Consequences of the Oil Shock of 2007-08’ by James D. Hamilton,” *Brookings Papers on Economic Activity*, 1, 267-278.
- Kilian, L. (2014), “Oil Price Shocks: Causes and Consequences,” *Annual Review of Resource Economics*, 6, 133-154.
- Kilian, L., and B. Hicks (2013), “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?” *Journal of Forecasting*, 32, 385-394.
- Kilian, L., and Y.J. Kim (2011), “How Reliable Are Local Projection Estimators of Impulse Responses?” *Review of Economics & Statistics*, 93, 1460-1466.
- Kilian, L., and T.K. Lee (2014), “Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories,” *Journal of International Money & Finance*, 42, 71-87.
- Kilian, L., and L.T. Lewis (2011), “Does the Fed Respond to Oil Price Shocks?” *Economic Journal*, 121, 1047-1072.

- Kilian, L., and D.P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454-478.
- 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 & Statistics*, 93, 660-671.
- 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, 419-453.
- Kilian, L., and R.J. Vigfusson (2011b), "Nonlinearities in the Oil Price-Output Relationship," *Macroeconomic Dynamics*, 15, 337-363.
- Kilian, L., and R.J. Vigfusson (2013), "Do Oil Prices Help Forecast U.S. Real GDP? The Role of Nonlinearities and Asymmetries," *Journal of Business & Economic Statistics*, 31, 78-93.
- Kiyotaki, N., and J. Moore (1997), "Credit Cycles," *Journal of Political Economy*, 105, 211-248.
- 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.
- Ramey, V.A., and D.J. Vine (2011), "Oil, Automobiles and the U.S. Economy: How Much Have Things Really Changed?," *NBER Macroeconomics Annual*, 25, 333-367.
- Ravazzolo, F., and P. Rothman (2013), "Oil and U.S. GDP: A Real Time Out-of-Sample Examination," *Journal of Money, Credit & Banking*, 45, 449-463.
- Romer, C.D. and D.H. Romer (1989), "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in: Blanchard, O.J. and S. Fischer (eds.), *NBER Macroeconomics Annual*, 4, 121-184.
- Rotemberg, J., and M. Woodford (1996), "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit & Banking*, 28, 549-577.
- Stock, J.H., and M.W. Watson (2012), "Disentangling the Channels of the 2007-2009 Recession," *Brookings Papers on Economic Activity*, Spring, 81-135.
- Wheelock, D.C., and M.E. Wohar (2009), "Can the Term Spread Predict Output Growth? A Survey of the Literature," *Federal Reserve Bank of St. Louis Review*, 91, 419-440.

Figure 1: The Real U.S. Refiners' Acquisition Cost and NBER Recessions




Figure 2: 3-Year Net Increase in Real U.S. Refiners' Acquisition Cost for Oil Imports




Figure 3: Conditional Responses of U.S. Real GDP




Figure 4: Evolution of Conditional U.S. Real GDP Responses at 8-Quarter Horizon in the Net Oil Price Increase Model




Figure 5: Structural Prediction Error Decompositions of U.S. Real GDP during Net Oil Price Increase Episodes




Table 1: Marginal Significance Levels of Tests of Conditionally Symmetric and Conditionally Linear Response Functions

Ω_{t-1}	H	Symmetry		Linearity	
		1σ	2σ	1σ	2σ
1990Q3	4	0.892	0.104	0.063**	0.043*
	8	0.991	0.419	0.260	0.232
1994Q3	4	0.991	0.630	0.990	0.964
	8	0.997	0.794	0.583	0.698
2008Q4	4	0.999	0.942	0.687	0.691
	8	1.000	0.994	0.952	0.951

NOTES: The symmetry test is a joint test of the symmetry of the response function in the net oil price increase model at horizons $0, 1, \dots, H-1, H$ with respect to positive and negative oil price innovations of the same magnitude. The linearity test tests whether the nonlinear response function from model (1) at horizons $0, 1, \dots, H-1, H$ jointly equals the impulse response function obtained from model (2). Both tests have an asymptotic χ^2_{H+1} -distribution. Boldface indicates statistically significant rejections; * and ** statistical significance at the 5% and 10% level.

Table 2: Correlations of the Conditional Responses in the Net Oil Price Increase Model with Macroeconomic Conditions

Explanatory Variable	Contemporaneous Correlation	R^2
Index of Financial Conditions (NFCI)	-0.244*	0.06
Index of Financial Stress (FRB)	-0.171**	0.03
Nominal Oil Share in U.S. GDP	0.463*	0.21
Index of Consumer Sentiment	0.137**	0.02
Buying Conditions for Homes	0.404*	0.16
Buying Conditions for Large Household Items	0.081	0.01
Buying Conditions for Vehicles	0.391*	0.15
Expected Change in Unemployment Next Year	0.030	0.00
Expected Change in Business Conditions Next Year	-0.020	0.00
Expected Change in Interest Rate Next Year	-0.226*	0.05
Expected Change in Financial Situation Next Year	-0.058	0.00

NOTES: Boldface indicates statistically significant rejections; * and ** statistical significance at the 5% and 10% level. The R^2 is based on a regression of the conditional responses on an intercept and the contemporaneous value of the explanatory variable.³

Table 3: Bootstrap p-values of Linearity Tests based on Local Projections

$H_0 :$	Horizons h									Robust
	0	1	2	3	4	5	6	7	8	
$\gamma_{h,1} = \dots = \gamma_{h,4} = 0$	0.04*	0.07**	0.11	0.32	0.17	0.17	0.03*	0.02*	0.16	0.08**
$\gamma_{h,1} = 0$	0.10	0.01*	0.10**	0.17	0.10	0.12	0.46	0.69	0.33	0.16

NOTES: All results are based on model (3). The t- and Wald test statistics are based on Newey-West robust standard errors. The bootstrap p-values for each equation were computed by bootstrapping the system of nine equations using a fixed-design residual-based wild bootstrap under the null that $\gamma_{h,1} = \dots = \gamma_{h,4} = 0$, preserving the correlations of the regression errors across equations. The bootstrap innovations were constructed from the unrestricted model to improve the power of the test. The data-mining robust p-values were computed analogously by tabulating the distribution of the maximum of the bootstrap Wald and the bootstrap t -test statistics, respectively, across the nine equations. This approach protects against spurious rejections of the null resulting from the repeated application of the same test across horizons. Statistical significance at the 5% (10%) level is marked * (**).

Table 4a: Cumulative Percent Change in U.S. Real GDP Explained by Oil Price Shock in the Net Oil Price Increase Model during Net Oil Price Increase Episodes

Starting date Ω_{t-1}	Baseline Model		Excess bond premium added		Term spread added		Consumer confidence added		All three variables Added	
	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$
1979Q1	-1.74	-3.51	-1.85	-3.29	-1.78	-3.60	-1.41	-3.52	-1.67	-3.48
1980Q4	0.50	1.58	0.63	1.79	0.40	1.45	0.51	1.51	0.43	1.25
1990Q2	-2.41	-3.21	-2.41	-3.02	-2.33	-3.17	-1.96	-2.53	-1.95	-2.47
1995Q4	0.26	-0.02	0.27	0.12	0.24	-0.33	0.19	-0.15	0.16	-0.22
1999Q4	-0.58	-0.75	-0.64	-0.82	-0.61	-0.72	-0.40	-0.35	-0.47	-0.44
2004Q2	-0.91	-2.68	-0.84	-2.28	-0.77	-2.09	-0.64	-2.17	-0.40	-1.25
2007Q2	-1.10	-5.13	-1.21	-4.75	-1.09	-5.30	-0.88	-4.45	-1.10	-4.25

NOTES: The starting dates were selected based on Figure 2. The expanded net oil price increase models include the net oil price increase only in the real GDP growth equation.

Table 4b: Cumulative Percent Change in U.S. Real GDP Explained by Oil Price Shock in the Linear VAR Model during Net Oil Price Increase Episodes

Starting date Ω_{t-1}	Baseline Model		Excess bond premium added		Term spread added		Consumer confidence added		All three variables added	
	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$	$h = 4$	$h = 8$
1979Q1	0.25	-0.46	0.18	-0.64	0.22	-0.60	0.24	-0.57	0.10	-0.91
1980Q4	-0.06	0.00	-0.04	0.03	-0.06	0.03	-0.05	0.04	-0.04	0.05
1990Q2	-0.15	-0.14	-0.19	-0.26	-0.12	-0.28	-0.15	-0.14	-0.19	-0.40
1995Q4	0.10	-0.35	0.08	-0.37	0.10	-0.42	0.09	-0.35	0.05	-0.40
1999Q4	-0.01	-0.39	-0.01	-0.46	-0.02	-0.38	-0.05	-0.32	-0.08	-0.35
2004Q2	0.16	-0.20	0.11	-0.26	0.09	-0.24	0.14	-0.20	0.04	-0.23
2007Q2	0.32	-0.68	0.26	-0.65	0.30	-0.78	0.31	-0.69	0.20	-0.74

NOTES: The starting dates were selected based on Figure 2.

Table 5: Conditional MSPE Ratios of Net Oil Price Increase Model Relative to Linear Model following Net Oil Price Increases

Episodes	Starting date	(1)	(2)	(3)	(4)	(5)	(6)
		Real GDP, real oil price	Real GDP, real oil price	Excess bond premium added	Term spread added	Consumer confidence added	All three extra variables added
U.S. Recession	1979Q4	0.276		0.327	0.400	0.372	0.492
	1981Q2	0.929	0.929	0.892	1.051	0.870	0.823
	1990Q3	0.800	0.800	0.702	0.588	2.323	1.390
	2001Q1	1.304	1.304	0.568	1.069	1.852	0.703
	2007Q4	0.603		0.689	0.613	0.707	0.724
No U.S. recession	1996Q3	2.378	2.378	2.121	1.981	2.560	2.211
	2004Q4	4.871	4.871	2.912	7.915	2.265	1.964
Mean for net oil price increase episodes before recessions only		0.782	1.011	0.636	0.742	1.225	0.826
Mean for all net oil price increase episodes		1.594	2.056	1.173	1.944	1.564	1.187

NOTES: For the recession starting in 1990Q2 we focus on 1990Q3, given that the start of the recession as dated by the NBER predates the oil price shock associated with the invasion of Kuwait by one quarter. The conditional MSPE is based on the 8 quarters following the starting date. The excess bond premium is from Gilchrist and Zakrajsek (2012). The term spread refers to the spread between the ten-year treasury bond yield and the 5-year treasury bond yield, as reported in the FRED database. Consumer confidence refers to the *Michigan Survey of Consumers* measure of overall consumer confidence in the U.S. economy.