JOHANNES KEPLER
UNIVERSITY LINZ

DEPARTMENT OF ECONOMICS
JOHANNES KEPLER UNIVERSITY OF LINZ

The effects of oil supply and demand shocks on U.S. consumer Sentiment

by

Jochen H. F. GÜNTNER

Katharina LINSBAUER

Working Paper No. 1614
December 2016

Johannes Kepler University of Linz
Department of Economics
Altenberger Strasse 69
A-4040 Linz - Auhof, Austria
www.econ.jku.at

jochen.guentner@jku.at

The effects of oil supply and demand shocks on U.S. consumer sentiment

Jochen H. F. Güntner*
Johannes Kepler University Linz

Katharina Linsbauer
Johannes Kepler University Linz

December 14, 2016

Abstract

This paper investigates how the University of Michigan's Index of Consumer Sentiment (ICS) – a survey measure of U.S. households' expectations about current and future economic conditions – responds to structural oil supply and demand shocks. We find that the response to an observed increase in the real price of crude oil depends on the underlying reason. While oil supply shocks have little effect on the ICS, other oil demand shocks such as a precautionary demand shock, for example, have a statistically significant negative impact over a two-year horizon. The effect of aggregate demand shocks associated with the global business cycle is positive in the first few months and negative thereafter. Considering the responses of ICS sub-indices and more specific survey questions, we find that expectations about higher future inflation and the associated reduction of real household income as well as a deterioration of perceived vehicle and house buying conditions are the main transmission channels of aggregate demand and other oil demand shocks. Oil shocks also affect consumers' satisfaction with U.S. economic policy.

Keywords: Consumer sentiment; Oil price shocks; Structural VAR estimation; Transmission channels

JEL Classification: C32, E30, N50, Q41

*Corresponding author: Jochen Güntner, Johannes Kepler University Linz. Altenbergerstrasse 69, 4040 Linz, Austria. Telephone: (+43)(0)732 2468 8360. Email: jochen.guentner@jku.at.

1 Introduction

The unprecedented boom and bust in the price of crude oil since 2000 re-aroused the interest of economists and policy makers in the origins and effects of oil price fluctuations. To this day, there is a lack of consent about how the latter are transmitted to macroeconomic variables such as GDP, inflation, or stock returns. While introductory macroeconomic textbooks interpret oil price changes as exogenous shifts of the aggregate supply (AS) curve, Kilian (2008) does not find much evidence in favor of this cost-push shock interpretation. Instead, they might impact real economic activity through the demand side, via (perceived) changes in the purchasing power of disposable income, increased uncertainty about future economic conditions, or a reduction in consumer and investor sentiment, which induces households and firms to cut back on their consumption and investment expenditures, respectively.

Following an increase in the real price of crude oil, higher gasoline and energy prices reduce household disposable income (excluding energy) and thus the budget for other expenditures. At the same time, higher gasoline prices raise the operating cost of vehicles, making purchases less attractive. Beyond the reduction in current disposable income, gloomy expectations about future economic conditions due to oil price fluctuations might motivate precautionary savings and depress consumer spending further. Given that personal consumption expenditures account for more than two thirds of U.S. GDP, understanding the response of household expectations is a prerequisite for understanding the transmission of oil shocks to the U.S. economy.[1]

In this paper, we therefore use aggregate time series data from the University of Michigan's *Surveys of Consumers* to investigate how oil demand and supply shocks affect U.S. households' perception of their personal financial situation as well as their expectations about current and future economic conditions. Building on the structural vector autoregression (SVAR) approach in Kilian (2009) and Kilian and Park (2009), we identify mutually orthogonal oil supply, aggregate demand, and other oil demand shocks in order to be able to interpret the impulse responses of U.S. consumer sentiment. In addition to the well-known Index of Consumer Sentiment (ICS), we consider its five components as well as other interesting survey questions, such as whether consumers expect higher future inflation or whether they think it's a good time to buy a car or house, in order to understand the transmission channels of oil supply and demand shocks.

Kilian's (2009) finding that "not all oil price shocks are alike" implies that the estimation of economically interpretable impulse response functions requires use of a structural econometric model that identifies the origin of observed oil price fluctuations. Reduced-form regressions of

[1] For example, Matsusaka and Sbordone (1995) find that the null hypothesis of no Granger-causality from consumer sentiment to GNP can be rejected, even after controlling for lagged values of economic fundamentals.

macroeconomic variables on the real price of crude oil capture the average effect of oil price changes, at best. While various authors adopted and extended Kilian's (2009) identifying strategy to investigate the effects of oil supply, aggregate demand, and other oil demand shocks on real GDP growth, inflation, and stock returns (see, e.g., Kilian and Park, 2009; Güntner, 2014a), their transmission via U.S. consumer sentiment has not yet been studied. Hence, we contribute to the existing literature by providing prima facie evidence of the effects of oil shocks on household expectations about current and future economic conditions.[2] To the best of our knowledge, we are also the first to consider the structural impulse response functions of survey questions concerning expectations about unemployment and real household income as well as contentment with U.S. economic policy. By covering 1978Q1-2015Q4, our sample includes the unprecedented surge in the real price of crude oil before the Great Recession as well as the period of historically low oil prices starting in 2014.

While consumer sentiment used to track the ups and downs in oil prices closely, this was not the case during the historical surge in the early 2000s. According to Curtin (2005), consumers assigned this price increase to a temporary shortfall of supply rather than a permanent increase in demand, suggesting that, even from their perspective, it makes a difference whether oil prices rise due to reduced supply or increased demand.[3] An observed increase in the price of crude oil might therefore have very different consequences for personal consumption expenditures and real economic activity. Figure 1 plots the ICS against the real price of crude oil, illustrating the comovement between the two series. In particular, consumers tend to be more pessimistic in times of high oil prices, whereas consumer sentiment drops sharply as oil prices peak, for example in 1991 and 2008. Although we abstain from assigning a causal interpretation to this observation, it indicates substantial negative unconditional correlation between the ICS and the real price of crude oil during our sample period.

[**Figure 1 about here**]

The remainder of this paper is organized as follows. Section 2 surveys the related literature and discusses previous results concerning the effects of oil shocks on macroeconomic variables, including the behavior of consumers. Section 3 describes the data used in our empirical analysis. Section 4 presents the econometric methodology and identifying strategy. Our estimation results and the implied impulse response functions for the ICS, its five sub-indices, and selected further survey questions are discussed in Section 5. Section 6 investigates the robustness of our results, while Section 7 concludes.

[2] Edelstein and Kilian (2009) estimate the response of the ICS to a purchasing power shock arising from an increase in energy prices, whereas they do not distinguish between different oil demand and supply shocks.

[3] Richard Curtin is Chief Economist of the Surveys of Consumers at the University of Michigan.

2 Related Literature

The aim of this section is to give a brief overview of methodologies and results in the existing literature. For comprehensiveness, we discuss the effects of oil price changes on output, prices, investment, and asset markets as well as on consumption expenditure and consumer sentiment.

Hamilton (1983) is the first to report a systematic and statistically significant relationship between oil price changes and U.S. recessions during 1948-1972, while abstaining from a causal interpretation of his findings. Based on impulse response functions of several economic variables, Burbidge and Harrison (1984) conclude that the oil shock of 1973/74 had a significant impact on U.S. real GDP and prices, whereas this was not the case in 1979/80. On the contrary, Gisser and Goodwin (1986) argue that oil price changes have statistically significant real and inflationary effects of similar magnitude during the pre- and post-OPEC period. More recent studies include Engemann et al. (2011), who find that oil price changes have predictive power in forecasting recessions. Including the period preceding the Global Financial Crisis of 2007/08 in his analysis, Hamilton (2009) finds that oil price changes led to a significant decline in real GDP during past recessions mainly due to the large negative impact on consumer spending and automobile purchases. At the same time, he notes that the reasons for the surge in oil prices after 2000 were different than those in previous episodes.[4]

On the other hand, a number of studies questions the finding of a significant relationship between oil prices and real economy activity. Based on a discussion of major insights from previous studies, Barsky and Kilian (2004) conclude that the effect of oil price changes on the U.S. macroeconomy is smaller than common wisdom suggests. While admitting a possible impact of oil shocks on economic growth, they argue that theoretical explanations for such a relationship are not supported by the data or economically meaningless. Similarly, Segal (2011) does not back the view that oil prices played a major role during past recessions, arguing that the price of crude oil affected real economic activity primarily through its impact on monetary policy. It is important to note that both studies review the existing literature in the light of theoretical considerations rather than conducting own empirical analyses. Accordingly, a theoretical explanation for the empirical result of a significant relationship between oil prices and macroeconomic variables has yet to be found.

The insight that oil prices are driven by supply disruptions due to exogenous political events in the Middle East as well as fluctuations in global aggregate demand or other oil demand goes back to Barsky and Kilian (2002, 2004). In addition, Kilian (2009) points out that regressing

[4] A more recent strand of the literature investigates the effects of oil price *uncertainty* on real economic activity (see, e.g., Elder and Serletis, 2010; Jo, 2014).

macroeconomic variables on the price of crude oil presumes that the latter is exogenous to the former. While approximately correct for small economies, this assumption is less credible for the U.S. and the G20 (see also Berument et al., 2010). Using a structural VAR model in order to distinguish between different types and ensure the exogeneity of oil shocks, Kilian (2009) finds that global oil supply shocks cause a temporary decrease in U.S. real GDP that vanishes after about two years. This finding is in line with Kilian (2008), who further concludes that oil supply shocks have not played a prominent role in the U.S. business cycle since the 1970s. Moreover, Kilian (2009) finds a statistically significant negative effect of aggregate demand shocks on cumulative real GDP growth after three years, while the effect is insignificantly positive during the first year. Oil market-specific demand shocks such as a speculative demand shock, for example (see also Kilian and Lee, 2014), lead to a gradual decrease in cumulative real GDP growth. Kilian (2009) further finds that CPI inflation is mainly driven by aggregate and oil market-specific demand shocks, whereas oil supply shocks have an insignificant impact.

Building on this identifying strategy, another strand of the literature reinvestigates whether or not oil shocks drive asset markets. For example, Kilian and Park (2009) argue that, without distinguishing between oil supply and demand shocks, the failure to detect a significant relation between oil prices and stock returns in previous studies might be due to offsetting positive and negative effects of different types of oil shocks during a particular sample period. The authors find statistically significant differences in the response of cumulative real U.S. stock returns to oil supply, aggregate demand, and oil market-specific demand shocks. Kilian and Park (2009) further show that the responses of cumulative real stock returns differ across U.S. industries, calling for different portfolio adjustments depending on the nature of the underlying oil shock. Güntner (2014a) builds on the previous study by considering the impulse responses of cumulative real stock returns in six OECD countries, including the net oil exporters Canada and Norway. Besides replicating Kilian and Park's (2009) results for the U.S., Güntner (2014a) finds that positive aggregate demand shocks, which raise the real price of crude oil, have a positive effect on national stock markets in both oil-exporting and oil-importing economies, albeit with varying persistence. Other oil demand shocks such as a precautionary demand shock, for example, hurt stock markets in oil-importing economies, whereas they raise cumulative real stock returns in Norway and have no significant effect in Canada.

While the literature discussed so far mainly focuses on the relationship between oil prices and supply-side variables such as aggregate or firm-level output, inflation, and stock returns, much fewer studies investigate the response of personal consumption expenditures or consumer sentiment to oil supply and demand shocks. In light of recent empirical evidence that oil shocks are primarily transmitted through the demand side (see, e.g., Hamilton, 2009; Kilian and Park, 2009), this is an apparent shortcoming.

In this paper, we focus on consumer sentiment as a proxy for household perceptions of their personal financial situation as well as current and future economic conditions. To the extent that the ICS reflects changes in actual or perceived economic conditions and uncertainty, it also affects the transmission of oil shocks. Carroll et al. (1994) find that lags of U.S. consumer sentiment have explanatory power for changes in household spending.[5] They conjecture that increased uncertainty induces households to spend less, while habit formation postpones the adjustment of consumption spending. Matsusaka and Sbordone (1995) provide evidence that unpredictable changes in consumer sentiment "Granger-cause" changes in real GNP, contributing between 13 and 26% to its total variance. Referring to Keynes' "mass psychology of the market", Matsusaka and Sbordone (1995) attribute their results to uncertainty.

Most closely related to our study is the work by Edelstein and Kilian (2009). The authors address a number of questions concerning the nature and magnitude of the relationship between consumer expenditures and retail energy prices, interpreting changes in the latter as shocks to the purchasing power of households. Importantly, Edelstein and Kilian (2009) find no evidence of asymmetries in the response of consumption expenditures or consumer sentiment to positive and negative energy price changes, facilitating the use of a symmetric VAR model in their study and in this paper (see also Hooker, 1996). Their main finding is that total U.S. consumption decreases by 0.15% in response to a 1% increase in retail energy prices, due mainly to a 0.84% price elasticity of motor vehicle purchases. The authors attribute this result to a precautionary savings motive fueled by rising uncertainty and confirm their conjecture by considering the response of various indices from the University of Michigan's Surveys of Consumers such as expectations about unemployment, real interest rates, and real household income, for example. All survey indices except interest rate expectations drop on impact and remain negative and statistically significant for the following 18 months, supporting the hypothesis that increased pessimism or uncertainty about current and future economic conditions is an important link in the relationship between retail energy pries and personal consumption expenditures.

Using median expected inflation from the Surveys of Consumers, Wong (2015) investigates whether oil price changes have an effect on inflation expectations in the U.S. and whether this leads to higher actual inflation. Without distinguishing between oil supply and demand shocks, Wong (2015) concludes that, although household inflation expectations are sensitive to changes in the oil price, the expectations channel seems to be of minor importance for actual inflation.

[5] Ludvigson (2004) and Souleles (2004) find similar results.

3 Data Description

In our analysis, we use monthly series of world crude oil production, global real economic activity, the real price of crude oil, and various measures of consumer confidence for 1978:1-2015:12.

3.1 Oil Market Variables

Monthly data on world crude oil production in thousand barrels per day (tbpd) is available from the U.S. Energy Information Administration's (EIA) *Monthly Energy Review*. Conducting an augmented Dickey-Fuller (ADF) test, we cannot reject the null hypothesis of a unit root in petroleum production at conventional levels. In light of the technical constraints in oil extraction such as high adjustment costs, for example, we find this result plausible (see, e.g., Kilian, 2009; Güntner, 2014b). Following the preceding literature, world crude oil production enters our vector of endogenous variables in terms of annualized percentage changes, denoted by $\Delta lprod_t$.

As a monthly measure of global real economic activity, we use the real economic activity index created by Kilian (2009), which is based on single-voyage dry cargo ocean shipping freight rates and expressed in terms of the deviation of real freight rates from their long-run trend. The idea is that increased global demand for industrial commodities will raise ocean shipping freight rates due to the inelastic supply of shipping space in the short run.[6] Given that Kilian's (2009) index is a global business cycle measure and stationary by construction, it enters the structural VAR model in levels, denoted by rea_t.

As a measure for the world price of crude oil, we use the EIA's monthly refiner acquisition cost of imported crude oil in dollars per barrel. The nominal price series is deflated by the U.S. Consumer Price Index obtained from the Bureau of Labor Statistics (BLS). Despite the fact that unit root tests yield mixed results, the real price of crude oil enters the vector of endogenous variables in terms of log deviations from the mean, multiplied by 100 ($lrpo_t$). While over-differencing a stationary series might render the estimates inconsistent, not differencing an integrated series merely leads to a loss in estimation efficiency.

3.2 Measures of Consumer Confidence

To investigate the impact of oil shocks on and their transmission through consumer confidence, we use the well-known Index of Consumer Sentiment (ICS), denoted $senti_t$, its five components, and seven more detailed indices from the University of Michigan's Surveys of Consumers.[7] At

[6] For a detailed account of the construction of the index, its benefits and possible limitations, see Kilian (2009).

[7] The Michigan consumer survey is conducted monthly, asking approximately 50 questions about attitudes and expectations regarding financial, business, and buying conditions to a representative sample of U.S. households in the form of 500 telephone interviews. The various indices constructed based on the survey are widely used in the literature as an indicator of uncertainty among consumers and an accurate predictor of economic developments. For example, the Index of Consumer Expectations is included in the U.S. Bureau of Economic Analysis' (BEA)

monthly frequency, all variables are available since 1978:1.

The ICS is based on five sub-indices concerning current personal financial conditions and expected future economic developments. The corresponding survey questions are as follows:

1. "We are interested in how people are getting along financially these days. Would you say that you (and your family living there) are better off or worse off financially than you were a year ago?" ($pago_t$)

2. "Now looking ahead – do you think that a year from now you (and your family living there) will be better off financially, or worse off, or just about the same as now?" ($pexp_t$)

3. "Now turning to business conditions in the country as a whole – do you think that during the next twelve months we'll have good times financially, or bad times, or what?" ($bus12_t$)

4. "Looking ahead, which would you say is more likely – that in the country as a whole we'll have continuous good times during the next five years or so, or that we will have periods of widespread unemployment or depression, or what?" ($bus5_t$)

5. "About the big things people buy for their homes – such as furniture, a refrigerator, stove, television, and things like that. Generally speaking, do you think now is a good or bad time for people to buy major household items?" (dur_t)

Index scores are computed as 100 plus the difference between the percentages of favorable and unfavorable replies to the respective survey question. The ICS is then constructed as the sum of the five sub-indices divided by its value for the 1966 base period.[8] Accordingly, a higher value of the ICS indicates greater optimism amongst private households.

Moreover, we consider seven indices constructed from more specific survey questions. These additional indices are supposed to yield further insight into what drives fluctuations in consumer confidence by inquiring information on inflation and interest rate expectations, expected real income and unemployment as well as perceived vehicle or house buying conditions and opinions about the government's economic policy. The corresponding survey questions are as follows:[9]

6. "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?" And "By what percent do you expect prices to go up, on the average, during the next 12 months?" ($infl_t$)

Composite Index of Leading Indicators. For a detailed description of the survey see the University of Michigan's Survey Information website (`https://data.sca.isr.umich.edu/survey-info.php`).

[8] Detailed information on index calculations can be obtained from the website of the University of Michigan's Surveys of Consumers (`https://data.sca.isr.umich.edu/fetchdoc.php?docid=24770`).

[9] See the Surveys of Consumers codebook: `https://data.sca.isr.umich.edu/fetchdoc.php?docid=45121`.

7. “No one can say for sure, but what do you think will happen to interest rates for borrowing money during the next 12 months – will they go up, stay the same, or go down?” ($ratex_t$)

8. “During the next year or two – do you expect that your (family) income will go up more than prices will go up, about the same, or less than prices will go up?” ($rinc_t$)

9. “How about people out of work during the coming 12 months – do you think that there will be more unemployment than now, about the same, or less?” ($umex_t$)

10. “Speaking now of the automobile market – do you think the next 12 months or so will be a good time or a bad time to buy a vehicle, such as a car, pickup, van, or sport utility vehicle?” (veh_t)

11. “Generally speaking, do you think now is a good time or a bad time to buy a house?” ($house_t$)

12. “As to the economic policy of the government – I mean steps taken to fight inflation or unemployment – would you say the government is doing a good job, only fair, or a poor job?” ($epol_t$)

The computation of index scores is identical to those for 1-5, except for inflation expectations, which are expressed in terms of the median across households. Interest rate and unemployment expectations are defined such that a higher index corresponds to lower expected interest rates and unemployment, respectively.

Given that all indices are in terms of fluctuations around their long-run mean of 100 and thus stationary by construction, no further transformation is required. This procedure is consistent with the finding that ADF tests reject the null hypothesis of a unit root in the ICS and most of our alternative measures of consumer sentiment at the 10% level or better. In cases, where the null hypothesis of a unit root is not rejected at conventional significance levels, an analysis of the autocorrelation pattern suggests that the time series are very persistent but most likely stationary. Table 1 provides summary statistics of the 16 variables described above.

[Table 1 about here]

3.3 Unit Root Tests

Table 2 summarizes the results of the ADF test for unit roots in all variables under consideration. Under the null hypothesis that the time series contains a unit root, the coefficient on the lagged dependent variable is equal to zero. The test statistic follows a t-distribution, where the relevant critical value depends on the number of observations and whether a constant or a time trend is included. A statistically significant coefficient on the lagged dependent variable in Table 2 indicates that the null hypothesis of a unit root is rejected. This is the case for Kilian's (2009) real economic activity index and the ICS as well as for most of the other sentiment indices. The null hypothesis of a unit root in world crude oil production cannot be rejected. Moreover, the ADF test yields ambiguous results for the real price of crude oil. Hence, it is important to note that unit root tests have low power against persistent alternatives. Even though the null hypothesis of a unit root is not rejected in some cases, we use the corresponding series in levels, if visual inspection of its autocorrelation pattern suggests a persistent but stationary process.

[Table 2 about here]

4 Econometric Methodology

Our econometric methodology builds on the structural VAR approach in Kilian (2009) along the lines of Kilian and Park (2009). Thus, we account for potential endogeneity of the real price of crude oil and distinguish between different types of oil supply and demand shocks, addressing Kilian's (2009) critique of the prior literature.

4.1 VAR Specification

Following Kilian and Park (2009) and Güntner (2014a), we set up a four-variable autoregressive model in the vector $\mathbf{z}_t \equiv (\Delta lprod_t, rea_t, lrpo_t, senti_t)'$, where $\Delta lprod_t$ denotes the annualized percentage change in world crude oil production, rea_t Kilian's (2009) real economic activity index, $lrpo_t$ the real price of crude oil in log deviations from the sample mean, and $senti_t$ the University of Michigan's ICS.[10] All data are monthly, while our sample period is 1978:1-2015:12.

We are interested in estimating the following structural VAR(p) model:

$$\mathbf{A}_0\mathbf{z}_t = \alpha + \sum_{l=1}^{p} \mathbf{A}_l \mathbf{z}_{t-l} + \varepsilon_t, \qquad t = 1, \ldots, 456, \tag{1}$$

where α denotes a 4×1-vector of intercept terms, $\mathbf{A}_0$ a 4×4-matrix of contemporaneous coefficients, $\sum_{l=1}^{p} \mathbf{A}_l \mathbf{z}_{t-l}$ a lag polynomial of order p, and ε_t a 4×1-vector of mutually orthogonal structural innovations. Following the existing literature, we set $p = 24$ in order to account for possible seasonality in the endogenous variables and because the effects of oil shocks on macroeconomic variables are highest after around nine to twelve months (compare Hamilton and Herrera, 2004; Güntner, 2014a).[11]

[10] Note that we replace the equation for real stock market returns in Kilian and Park (2009) and Güntner (2014a) by an equation for different measures of consumer confidence.

[11] In the robustness checks, we replicate our analysis using 12 and 36 lags of the vector $\mathbf{z}_t$, respectively. In both cases, the results are qualitatively unchanged.

Conditional on $\mathbf{A}_0^{-1}$ being invertible, the reduced-form representation of (1) is given by

$$\begin{aligned}\mathbf{z}_t &= \underbrace{\mathbf{A}_0^{-1}\alpha} + \underbrace{\mathbf{A}_0^{-1}\sum_{l=1}^{24}\mathbf{A}_l\mathbf{z}_{t-l}} + \underbrace{\mathbf{A}_0^{-1}\varepsilon_t} \\ &= \quad \beta \quad + \quad \sum_{l=1}^{24}\mathbf{B}_l\mathbf{z}_{t-l} \quad + \quad \mathbf{e}_t, \end{aligned} \tag{2}$$

where $\mathbf{e}_t$ denotes a 4×1-vector of possibly contemporaneously correlated reduced-form innovations. Straightforward multivariate least-squares estimation of (2) yields consistent estimates of the coefficients in $\hat{\mathbf{B}} \equiv \left[\hat{\beta}, \hat{\mathbf{B}}_1, \ldots, \hat{\mathbf{B}}_{24}\right]$ and the innovations in $\hat{\mathbf{e}}_t$. In what follows, we assume a recursive ordering of the matrix of contemporaneous coefficients, $\mathbf{A}_0$, in order to back out the structural shocks and coefficients in (1) from the reduced-form estimates in (2). Although this assumption is particularly convenient, we must ensure that the implied exclusion restrictions are economically plausible.

4.2 Identifying Strategy

Building on Kilian (2009), we distinguish three structural supply and demand shocks that affect the real price of crude oil: (1) changes in world crude oil production (*oil supply shocks*); (2) changes in the current demand for crude oil due to fluctuations in global real economic activity (*aggregate demand shocks*); and (3) changes in the demand for crude oil that are unrelated to the global business cycle (*other oil demand shocks*). Any remaining innovations in consumer sentiment that cannot be attributed to the oil market are summarized in a residual category (*other shocks to consumer sentiment*), which does not have an economic interpretation. This category might contain changes in consumer sentiment due to fiscal or monetary policy shocks that are unrelated to the oil market as well as exogenous changes in consumer sentiment, also know as fluctuations in "animal spirits".

Suppose that the decomposition of the 4×1-vector of reduced-form error terms, $\mathbf{e_t} \equiv \mathbf{A_0^{-1}}\varepsilon_\mathbf{t}$, has the following recursive representation:

$$\mathbf{e_t} \equiv \begin{pmatrix} e_t^{\Delta lprod} \\ e_t^{rea} \\ e_t^{lrpo} \\ e_t^{senti} \end{pmatrix} = \begin{pmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix} \begin{pmatrix} \varepsilon_t^{oil supply\ shock} \\ \varepsilon_t^{aggregate\ demand\ shock} \\ \varepsilon_t^{other\ oil\ demand\ shock} \\ \varepsilon_t^{other\ shocks\ to\ consumer\ sentiment} \end{pmatrix}. \tag{3}$$

Following the existing literature, the identification in (3) implies that aggregate demand shocks or other oil demand shocks as well as shocks to consumer sentiment that are unrelated to the oil market do not affect world crude oil production *within the same month*. In light of nontrivial adjustment costs in petroleum extraction and uncertainty about whether an observed change in oil demand represents a transitory shock or a permanent shift, the assumption of a vertical short-run supply curve appears highly plausible (compare Hamilton, 2009; Kilian, 2009). Güntner (2014b) shows that the short-run price elasticity of country-level crude oil supply in response to aggregate demand and other oil demand shocks is statistically indistinguishable from zero. As a consequence, we assume that only shocks to the supply side affect world crude oil production levels on impact.

The zero restriction in the second row and third column of $\mathbf{A}_0^{-1}$ implies a low short-run price elasticity of crude oil demand, consistent with the delayed response of global real economic activity after historical oil price hikes (compare Kilian, 2009). Accordingly, real economic activity is assumed to react contemporaneously to oil supply and aggregate demand shocks only. Note that the latter are the structural counterpart of innovations in rea_t, which capture exogenous fluctuations in the demand for *all kinds* of industrial commodities associated with the global business cycle.[12] In contrast, the real price of crude oil is free to respond to disruptions of the physical supply of crude oil, shifts in global aggregate demand for industrial commodities, and other oil demand shocks that are orthogonal to the global business cycle *within the same month*.

The block-recursive structure of the identification in (3) implies that the global oil market block is contemporaneously predetermined with respect to *other shocks to consumer sentiment*. Formally testing the widely used identifying assumption of predeterminedness of energy prices, Kilian and Vega (2011) find no evidence of feedback from a wide range of U.S. macroeconomic aggregates to the price of WTI crude oil at monthly frequency. For this reason, we do not allow for reverse causality from $senti_t$ to world crude oil production, the global business cycle, and the real price of crude oil within the same month. In contrast, U.S. consumer sentiment may respond immediately to oil supply, aggregate demand, and other oil demand shocks.[13]

By imposing six exclusion restrictions on the matrix of contemporaneous coefficients the structural VAR model in (1) is exactly identified. The recursive representation in (3) allows us to obtain $\mathbf{A}_0^{-1}$ by Cholesky-decomposing the sample covariance matrix of reduced-form residuals,

[12] Given that providers of shipping services hold large buffer stocks of bunker fuels – a residual product in the petroleum refining process –, the use of monthly data suggests also imposing a vertical short-run demand curve. As a consequence, oil price changes will not affect real economic activity within the same month, regardless of whether they are induced by oil supply or aggregate demand shocks. Imposing the overidentifying restriction $a_{21} = 0$ in (3) hardly affects the results in Kilian (2009). Conditional on the correctness of the other exclusion restrictions, Hansen's J-test does not reject the null hypothesis of a vertical aggregate demand curve.

[13] While high adjustment costs in petroleum extraction and the slow response of global real economic activity justify the zero restrictions in the first and second line of the fourth column, the assumption of no feedback from U.S. consumer sentiment to the real price of crude oil is at least debateable. Given that, in each period, the latter is determined by the supply of and the demand for crude oil, however, shocks that move both the oil price and consumer sentiment contemporaneously will be attributed to the oil market rather than to the residual category, justifying the zero restriction in the third row of the fourth column in $\mathbf{A}_0^{-1}$.

$\Sigma_{\hat{\mathbf{e}}_t}$, i.e.

$$\mathbf{A}_0^{-1} = chol\left(\Sigma_{\hat{\mathbf{e}}_t}\right) = chol\left(E\left[\hat{\mathbf{e}}_t\hat{\mathbf{e}}_t'\right]\right). \tag{4}$$

Our second research question concerns the channels through which oil supply and demand shocks affect the ICS. For this purpose, we replace $senti_t$ by one of its five components as well as by the index for one of seven more specific survey questions in the vector $\mathbf{z}_t$ and compute the impulse response functions after a typical oil supply shock, aggregate demand shock, and other oil demand shock, respectively. The alternative indices reflect U.S. household's perception of current personal financial conditions ($pago_t$), expected future personal financial conditions ($pexp_t$), expected business conditions in the next 12 months ($bus12_t$), expected business conditions in the next 5 years ($bus5_t$), perceived conditions for buying major household items (dur_t), motor vehicles (veh_t), and houses ($house_t$), expectations about future real household income ($rinc_t$), inflation ($infl_t$), interest rate ($ratex_t$), and unemployment expectations ($unex_t$), as well as opinions about the government's economic policies($epol_t$). The survey questions corresponding to each of these indices are listed in Section 3.

5 Empirical Results

5.1 Historical Evolution of Oil Shocks

Figure 2 plots annual averages of structural oil supply and demand shocks for 1980-2015. As in Kilian (2009), we annualize the monthly shock series to facilitate their readability. Despite the fact that the structural innovations are contemporaneously uncorrelated by construction, oil shocks never occur in isolation. At each point in time, the real price of crude oil is driven by a time-varying combination of oil supply, aggregate demand, and other oil demand shocks.

For example, the top panel of Figure 2 exhibits large negative oil supply shocks in 1980 and 1981, after Iran's invasion of Iraq marked the beginning of the First Gulf War. In 1980, this supply shock was accompanied by a positive aggregate demand shock due to the global business cycle, whereas other oil demand was relatively subdued.[14] In contrast, the oil price peak during the Second Gulf War between Iraq and Kuwait is associated with an unanticipated increase in other oil demand such as precautionary demand, for example, rather than a disruption of the physical supply of crude oil. The upper two panels of Figure 2 suggest the absence of oil supply shocks and a comparatively low aggregate demand in 1990.

[Figure 2 about here]

[14] Note that our findings do not necessarily imply that the Iran-Iraq War eased concerns about the future availability of crude oil. Instead, Kilian (2009) found a positive other oil demand shock in 1979 already, indicating that oil market-specific demand increased in anticipation of the war. Other political developments in the Middle East might have contributed to high oil market-specific demand in 1979, as well.

In 1998, the real price of crude oil reached a historical low, followed by a continuous increase over the next ten years. While other oil demand shocks are particularly prominent prior to and following the Asian Financial Crisis, we find large positive and negative aggregate demand shocks after 2000. In particular, the unprecedented surge in oil prices prior to and its drop during the Great Recession seems to be driven by unexpected fluctuations in global real economic activity, accompanied by moderate other oil demand shocks. Restrained oil production might have contributed to a faster oil price increase, as well, especially in 2001-2002 and after 2005. Consistent with the view that increased use of hydraulic fracturing, also know as "fracking", in U.S. shale oil production contributed to persistently low oil prices since the start of the Great Recession, we find a series of positive oil supply shocks after 2007. The most recent drop in the real price of crude oil in 2014-2015 is attributed to a combination of positive oil supply shocks, negative aggregate demand shocks, and increasingly negative other oil demand shocks.

5.2 Historical Variance Decomposition

The historical evolution of structural innovations in Figure 2 illustrates that, at each point in time, the real price of crude oil is determined by a combination of oil supply, aggregate demand, and other oil demand shocks. However, it does not indicate the relevance of these shocks for historical fluctuations in oil prices and consumer sentiment. In this section, we therefore conduct a historical variance decomposition (HVD) of the real price of crude oil and the University of Michigan's Index of Consumer Sentiment (ICS). In contrast to impulse response functions, which illustrate the effect of a one-time typical (e.g. one-standard-deviation) shock on the variable of interest, the HVD decomposes the historical fluctuations in the variable of interest into the cumulative effects of the identified time series of structural shocks. Hence, the HVD indicates the relative importance of each of the three shocks at a given point in time.

[**Figure 3 about here**]

Figure 3 presents the cumulative effect of oil supply, aggregate demand, and other oil demand shocks, respectively, on the real price of crude oil during 1980:1-2015:12. To a large extent, our results coincide with those found in previous studies (compare, e.g., Kilian, 2009; Kilian and Park, 2009; Güntner, 2014a). Oil supply shocks have comparatively less explanatory power for historical oil price fluctuations. While shocks to global real economic activity generate oil price fluctuations with low frequency and high amplitude from 1995 onwards, other oil demand shocks such as a precautionary demand shocks, for example, cause large and sudden swings in the real price of crude oil.

Due to the different sample period and the presence of consumer sentiment, there are also a number of discrepancies. In contrast to Güntner (2014a), for example, we find a larger cumulative contribution of oil supply shocks to fluctuations in the real price of crude oil during the First Gulf War, when disruptions of the supply of crude oil in Iran and Iraq triggered a persistent oil price hike (see also Figure 1). Moreover, we find little cumulative effect of aggregate demand shocks throughout the 1980s and negative cumulative effects of other oil demand shocks in 1981-1984 and 1986-1987, suggesting that shortfalls in oil supply are the dominant reason for high oil prices during the Iran-Iraq War. It is not until the Second Gulf War in 1990-1991 that other oil demand shocks such as a precautionary demand shock, for example, have a noticeable positive cumulative effect on the real price of crude oil.[15]

The unprecedented surge of oil prices after 2000 is driven by a series of unexpectedly high growth rates of global real economic activity starting around 2003 and a shift in oil market-specific demand, amplified by simultaneous negative oil supply shocks (see also Figure 2). The sharp fall of the real price of crude oil after the bankruptcy of Lehman Brothers in September 2008 as well as its temporary recovery is attributed to the joint effect of aggregate demand and other oil demand shocks, whereas positive oil supply shocks have a negative cumulative effect during this episode. Consistent with the historical evolution of structural shocks in Figure 2, the renewed fall of the real price of crude oil at the end of our sample occurred in an environment of ample petroleum supply due to the use of fracking in the U.S. and increased production from OPEC members such as Saudi Arabia, falling aggregate demand due to a slowing global economy, and comparatively low other oil demand.

[**Figure 4 about here**]

Figure 4 plots the cumulative effect of oil supply, aggregate demand, and other oil demand shocks to historical fluctuations in the ICS. On average over the sample period, oil supply shocks play a comparatively minor role, albeit depressing consumer sentiment during the early years of the First Gulf War. Around 2000, when restrained oil supply puts upward pressure on the real price of crude oil again, the cumulative effect on the ICS is also negative. Otherwise, the HVD contribution in the top panel is quantitatively small and erratic, suggesting that consumers are not overly concerned about oil supply shocks. This is rational insofar as disruptions of the physical supply of crude oil are expected to be temporary and compensable through inventories, i.e. oil reserves above the ground.

[15] Note that differences at the start of the sample period must be taken with a grain of salt. The computation of cumulative contributions to the HVD of oil prices and consumer sentiment is associated with a phase-in, the length of which depends on the persistence of the shock processes and their VAR propagation.

Similarly, other oil demand shocks such as a precautionary demand shock, for example, play a limited role for historical fluctuations in the ICS. While contributing positively in the second half of the 1990s, they have a negative cumulative effect on U.S. consumer sentiment after 2000 and again after the Global Financial Crisis, reflecting the corresponding shifts in the cumulative effect on oil prices in the bottom panel of Figure 3. It is not until the end of our sample, which is marked by a historically low real price of crude oil, that other oil demand shocks contribute to the recovery of the ICS. On the other hand, negative oil market-specific demand shocks have a positive cumulative effect on U.S. consumer sentiment during much of the 1990s.

In line with our findings in Figure 3, from about 1990 onwards, shocks to global real economic activity generate fluctuations with low frequency and high amplitude in the ICS. Following an initial decrease, aggregate demand shocks have an increasingly positive cumulative effect on consumer sentiment during 1997-2005, before contributing to its renewed decline until 2010. Note that U.S. consumers might have anticipated an economic downturn already before the bankruptcy of Lehman Brothers triggered the Global Financial Crisis. After the Great Recession, aggregate demand shocks are the main contributor to the ICS' recovery during 2010-2014.

We conclude this section by noting that conflicts in the Middle East at the beginning of our sample period, which contribute to historical fluctuations in the real price of crude oil, have only a small cumulative effect on the ICS. While it therefore seems that U.S. consumers are not overly concerned about distant political events, adverse aggregate demand and other oil demand shocks contribute to the fall in consumer sentiment during the Second Gulf War. However, it is not unlikely that the latter reflects the contemporaneous threat of a global economic downturn rather than concerns about the future availability of crude oil. In contrast, the period after 2000 is marked by a non-trivial contribution of aggregate demand and other oil demand shocks to historical fluctuations in the ICS.

5.3 Impulse Response Functions

Based on the structural VAR model in (1) and the identifying restrictions in (3), we can compute the impulse responses of the endogenous variables in the vector $\mathbf{z}_t \equiv (\Delta lprod_t, rea_t, lrpo_t, senti_t)'$ to each of the three structural oil shocks. Given that *other shocks to consumer sentiment* do not have a structural interpretation, we refrain from plotting and discussing the impulse responses for this residual category. Figure 5 plots the impulse responses to a typical (i.e. one-standard-deviation) oil supply shock, aggregate demand shock, and other oil demand shock, respectively. Following Kilian (2009) and others, we normalize impulses, so that each structural innovations tends to raise the real price of crude oil, in order to facilitate their comparison. Accordingly, we plot the impulse responses functions for a *negative* oil supply shock against those for *posi-*

tive aggregate demand and other oil demand shocks. One- and two-standard-error confidence bands are calculated using 5,000 replications of a recursive-design wild bootstrap that accounts for potential conditional heteroskedasticity of unknown form in the VAR residuals.[16] Under the assumption that the shocks are standard-normally distributed, these bands correspond to approximate 68% and 95% confidence intervals.

[**Figure 5 about here**]

5.3.1 Oil Market Variables

Consider first the impulse responses of world crude oil production, real economic activity, and the real price of crude oil, which can be compared with the results in Güntner (2014a), Kilian (2009) and Kilian and Park (2009). On average over the sample period, a typical negative oil supply shock leads to an instantaneous reduction in world oil production of about 15% on an annual basis. The large initial effect approximately halves during the following seven months before stabilizing around -10% in the second year of the forecast horizon. There is some evidence that global real economic activity increases after a negative oil supply shock. While this finding seems to contradict previous studies, which tend to find no response or a slight decrease in global real economic activity, the corresponding impulse response in Figure 5 is only marginally statistically significant during months 5-10. Alternatively, the difference might be due to an extended sample period in this paper. The economic environment at the end of our sample is marked by ample oil supply and a deceleration of growth rates in countries like China. Since an expansionary impact of negative supply shocks is not consistent with economic theory, we are somewhat in favor of regarding this finding as an artefact of our sample. Following a negative oil supply shock, the real price of crude oil increases by up to 2.5% relative to its sample mean. The response, which is delayed by about two months and (marginally) significant for 17 months, is more pronounced and persistent than those found in previous studies, which might again be due to the different sample period in this paper.[17]

Following a positive one-standard-deviation innovation in aggregate demand for industrial commodities, world oil production increases by up to 3.5% on an annual basis. This effect is statistically significant after eight months and marginally significant for most of the first year. Kilian's (2009) real economic activity index jumps by about 6.5 points on impact and increases further to 8.4 index points in the following month before declining steadily for the rest of the

[16] For details on the procedure and its asymptotic validity, see Gonçalves and Kilian (2004).

[17] The end of our sample is characterized by negative comovement between world crude oil production, which increased due to the U.S. fracking boom and the fight for market shares among other oil-producing countries, and the real price of crude oil, which fell to a historical low in 2014-2015. As a result, the impact of oil supply shocks on the real price of crude oil is more pronounced than in previous studies, which use a shorter sample period. This explanation is also consistent with our findings in Figures 2 and 3.

forecast horizon. A typical aggregate demand shock has a persistent positive and statistically significant impact on the real price of crude oil, which increases continuously to 5% over a horizon of 24 months. These findings are consistent with the view that oil producers do not fully accommodate the demand for oil but follow their own agenda when setting production volumes (see also Güntner, 2014b).

Other oil demand shocks such as a precautionary oil demand shock, for example, do not have an impact on world crude oil production.[18] The positive short-run effect of other oil demand shocks on global real economic activity is, while in line with similar findings in Kilian (2009) and Güntner (2014a), is hard to explain. Intuitively, one expects that increased uncertainty about the future availability relative to the future demand for crude oil impairs global real economic activity. However, other oil demand shocks tend to be associated with the occurrence of natural disasters or exogenous political events such as wars and terrorist attacks, for example. Recalling that Kilian's (2009) real economic activity index is based on dry cargo ocean shipping freight rates, it is conceivable that, in response to an earthquake or armed conflict, demand for shipping space increases in order to move technical or military equipment across the globe. Moreover, increased demand for weapons and in anticipation of an armed conflict or in response to a terrorist attack will be reflected in the demand for raw materials, such as steel, explaining (part of) the observed positive response of global real economic activity. In response to a typical other oil demand shocks, the real price of crude oil increases by about 6.5% on impact, peaking at 10.5% after two months. Subsequently, the oil price response decreases to about 4.5% after 24 months, while it remains statistically significant throughout.

5.3.2 The Index of Consumer Sentiment

Consider now the effect of structural oil supply and demand shocks on the Index of Consumer Sentiment, which has not been investigated in the existing literature. The corresponding impulse response functions in Figure 5 indicate whether and how U.S. consumers' perception of their personal financial situation and expectations about future economic conditions change in response to a typical oil supply, aggregate demand, and other oil demand shock, respectively.

The response in the first row and fourth column suggests that an exogenous disruption of world production, which raises the real price of crude oil, depresses the ICS at least temporarily, albeit the effect is small and only marginally significant in months 4-6 an 17-19. Accordingly,

[18] Note that the response of world crude oil production conceals potential differences at the country level. Analyzing the response of individual oil-producing countries for a similar sample period, Güntner (2014b) finds that OPEC members tend to curtail production, whereas non-OPEC producers tend to expand production in response to speculative oil demand shocks. In the short run, i.e. within the same month, the price elasticity of crude oil supply is not statistically different from zero even at the country level.

a typical negative oil supply shock irritates U.S. consumers with a lag rather than on impact, possibly reflecting the delayed pass-through of a change in the real price of crude oil to gas and retail energy prices. In the short run, crude oil inventories can be depleted in order to attenuate the shortfall in production. In the medium run, lower output of an oil-producing country can be compensated by higher output elsewhere. As a consequence, the negative response of U.S. consumer sentiment is short-lived and only marginally significant.

In response to a positive aggregate demand shock, the ICS first increases up to 0.8 index points after four months before decreasing to -1.1 index points during the next seven months. Both the peak and the trough are statistically significant at the approximate 5% level. Accordingly, after an initial bout of optimism, U.S. households become more pessimistic about effects of an exogenous increase in global real economic activity. It is important to note that the same qualitative pattern can be found in U.S. cumulative real stock returns, which also increase initially before turning negative after about one year, in response to a positive aggregate demand shock (compare, e.g., Kilian and Park, 2009; Güntner, 2014a). This common pattern might reflect the "law of the business cycle", which dictates that every economic boom is sooner or later followed by a recession, and households' anticipation of the end of the boom and the start of the recessionary phase. More likely, the explanation can be found in higher expected inflation, which eats up real household incomes unless they are expected to adjust accordingly. As a consequence, an exogenous increase in global real economic activity quickly makes way for a more gloomy perspective on future economic conditions.

Finally, a typical positive other oil demand shock, which raises the real price of crude oil significantly and for the entire two-year horizon, also has a statistically significant negative effect on the ICS during the first six months. The maximum effect of -1.0 index points is reached after four months. While the effect is quantitatively smaller thereafter, it remains at least marginally significant for much of the forecast horizon. Accordingly, there is strong evidence that other oil demand shocks such as a precautionary oil demand shock, for example, have a persistent negative effect on U.S. consumer confidence.

The results in this section confirm that "not all oil price shocks are alike" also from the perspective of U.S. consumers. Figure 5 illustrates the quantitative and qualitative differences in the response of the ICS to structural oil supply, aggregate demand, and other oil demand shocks. However, without further analysis, we can only conjecture how the different oil shocks are transmitted to consumer sentiment. In what follows, we therefore take a closer look at the impulse response functions of the ICS' five components as well as seven more specific indices from the University of Michigan's Surveys of Consumers.

5.4 The Transmission of Oil Shocks to Consumer Sentiment

In order to investigate the channels through which structural oil demand and supply shocks affect the ICS, we repeatedly estimate the reduced-form VAR model in (2) and apply the identifying strategy in (3), each time replacing $senti_t$ in the vector $\mathbf{z_t}$ by an alternative measure of consumer sentiment. It is then straightforward to compute the corresponding impulse response functions and compare them to the impulse responses in Figure 5.

5.4.1 Personal Financial and Economic Conditions

Consider first the five ICS components, i.e. $pago_t$ (current personal financial condition), $pexp_t$ (expected future personal finances), $bus12_t$ (expected business conditions during next 12 months), $bus5_t$ (expected business conditions during next 5 years), and dur_t (current buying conditions for major household items).[19] In Figure 6, each column plots the impulse responses to a typical oil supply, aggregate demand, and other oil demand shock, respectively, against those for ICS.

[Figure 6 about here]

While we already know that oil supply shocks have a small negative effect on the ICS that is only marginally significant in months 4-6, the first column of Figure 6 suggests that this result is mainly driven by the impulse responses of $pago_t$ and dur_t, for which we find a (marginally) significant reduction after three and four months, respectively. Qualitatively similar, albeit less persistent negative effects can be found for the impulse responses of $pexp_t$, $bus12_t$, and $bus5_t$. Accordingly, when the ICS decreases in response to a disruption of world crude oil production, it mainly reflects U.S. consumers' perception that it is not a good time for buying major household items such as furniture or a refrigerator, for example, and that they are financially worse off than a year ago. We will get back to possible reasons for this result when discussing the impulse response functions of more specific survey questions below.

In response to a positive aggregate demand shock associated with the global business cycle, we found an initial increase of up to 0.8 index points followed by a persistent decrease in the ICS, both statistically significant. The second column of Figure 6 illustrates that the temporary hike comes mainly from the responses of $bus12_t$ and dur_t, which increase by up to 1.7 and 1.8 index points, respectively, while all five sub-indices contribute to the reduction in consumer sentiment over the rest of the two-year horizon. In contrast to those of the other components, the impulse responses of $pago_t$ and $pexp_t$ display a marginally significant dip in the second month after the shock before turning increasingly negative and statistically significant during the second year.

[19] For an exact definition of these sub-indices and the corresponding survey questions see Section 3.

Accordingly, households feel worse off financially despite a boom in the global business cycle. A candidate explanation might be that consumers suffer disproportionately from higher gas and retail energy prices as well as higher commodity and real estate prices. In summary, about three quarters after a positive aggregate demand shock and modest initial optimism, consumers feel worse off financially than a year ago, believe that now is a bad time for buying large household items, and hold more pessimistic expectations about future personal financial and economic conditions in the U.S. as a whole. Interestingly, we find no evidence of initial optimism about their personal financial situation. In contrast to oil supply shocks, aggregate demand shocks also affect household expectations about future economic conditions during the next twelve months ($bus12_t$) and, to a lesser degree, during the next five years ($bus5_t$). The results in Figure 6 do not yield much insight into the reasons for growing pessimism amongst consumers in response to the shock. As a consequence, we will get back to this question when considering the responses to survey questions about future expected unemployment and inflation.

The third column plots the impulse response functions of the five sub-indices to a positive one-standard-deviation other oil demand shock, which induces an immediate and persistent increase in the real price of crude oil. As we already know from Figure 5, the ICS falls permanently and its response is at least marginally significant for most of the two-year horizon. In Figure 6, we find only minor differences in the response patterns of the five sub-indices. The decrease in consumer confidence during the first year is particularly pronounced and statistically significant for $pexp_t$, which falls to -1.2 index points after two months. While somewhat less significant, $bus12_t$ falls by 2.2 index points already in the first month after the shock. In contrast, $pago_t$ and dur_t respond more gradually, reaching a minimum of -0.95 and -1.6 index points, respectively, after five months. Consistent with the impulse response function of the ICS, all five sub-indices display a (partial) recovery over the course of the first year and a renewed decrease during the second year, which is particularly pronounced and statistically significant for $pago_t$ and dur_t. Hence, other oil demand shocks such as a precautionary demand shock, for example, trigger persistent waves of pessimism about current personal financial and expected future economic conditions among U.S. consumers. It is therefore plausible to interpret these structural shocks as jumps in uncertainty arising from natural disasters or armed conflicts, including terrorist attacks, in oil-producing countries. In what follows, we investigate whether the observed fall in consumer confidence is due to concerns about higher price inflation, widespread unemployment, or economic policy.

From the results in Figure 6, we conclude that oil supply shocks mainly affect the ICS by making U.S. households feel worse about their current personal finances and short-run business conditions, whereas aggregate demand and other oil demand shocks have a negative effect on both perceived current and expected future personal financial and economic conditions. To learn more about the underlying reasons for the observed responses of the ICS and its five sub-indices, we investigate the response of household expectations about macroeconomic variables such as inflation, interest, and unemployment rates as well as vehicle and house buying conditions.

5.4.2 Expectations about Macroeconomic Variables and Buying Conditions

The University of Michigan's Surveys of Consumers comprise around 50 questions, albeit not all of them are available since 1978 at a monthly frequency. In what follows, we consider the impulse response functions of seven selected survey questions concerning household expectations about inflation ($infl_t$), interest ($ratex_t$), and unemployment ($umex_t$) rates during the next twelve months, real family income ($rinc_t$) during the next 24 months, current vehicle (veh_t) and house ($house_t$) buying conditions, and satisfaction with economic policy ($epol_t$).[20] Figure 7 plots the impulse responses of the corresponding indices against those of $senti_t$.

[**Figure 7 about here**]

Consider first the impulse responses to a typical negative oil supply shock in the left column. While this generates only a small, temporary, and marginally significant decline in the ICS, the effect is stronger and more persistent for certain alternative measures of consumer confidence. For example, interest rate expectations ($ratex_t$) trend upwards and are significantly above their long-run mean after one year and again from 16 months onwards, peaking at +4.1 index points after 19 months. Moreover, perceived vehicle and house buying conditions deteriorate in response to a negative oil supply shock. While the decrease in veh_t is statistically significant only during the fourth month, that of $house_t$ remains significant from five to nine months after the shock. Despite a (partially) significant increase in the real price of crude oil in Figure 5, the median expected inflation rate ($infl_t$) increases only slightly during the first few months before decreasing until the end of the first year. The corresponding impulse response function is not significantly different from zero. Note that an increase in $ratex_t$ implies *lower* expected interest rates in the near future (e.g. due to expansionary monetary policy). This is consistent with the view that a disruption of petroleum production leads to a contraction of demand, also reflected in the deterioration of perceived vehicle and house buying conditions, rather than higher inflation. U.S. consumers' expectations about unemployment ($umex_t$) and real household income ($rinc_t$) during the next one or two years are not affected significantly, despite a tendency

[20] The exact formulation of each survey question and the computation of the corresponding indices can be found in Section 3. In the case of inflation expectations, the response of the median expected inflation rate is reported.

for $rinc_t$ to be below its long-run mean in the first six month and above thereafter. While none of these indices enters the ICS directly, the drop in overall consumer sentiment after a negative oil supply shock seems to reflect a perceived deterioration of buying conditions, in general.

Consider now the effects of a positive aggregate demand shock in the second column of Figure 7. With the exception of interest rate expectations ($ratex_t$), the impulse responses are statistically significant and diverse. After 22 months, for example, U.S. consumers expect prices to rise by an additional 0.12 percentage points during the next year. The corresponding impulse response function is statistically significant from the fifth month on. After five months, consumers also expect higher interest rates in the future, indicating a contractionary monetary stance, although this effect quickly vanishes. Despite higher expected inflation and interest rates, the ICS initially increases in response to higher demand for industrial commodities associated with the global business cycle. The reason is that U.S. consumers initially expect an increase in real household income, no change in unemployment rates, and slightly better vehicle and house buying conditions on impact. Accordingly, the perceived improvement of future economic and buying conditions overcompensates concerns about higher expected oil prices and inflation in the first months after the shock. Consistent with a reversal of the responses of $bus12_t$ and $bus5_t$ in Figure 6, consumers then become increasingly pessimistic about future economic conditions. The reason is a sudden increase in expected unemployment ($umex_t$) by 1.5 index points within only two months, accompanied by a drop in expected real household income by 0.8 index points, and a substantial deterioration of expected vehicle buying conditions to -1.8 index points after ten months.[21] The fall in $rinc_t$ is consistent with the persistent increase in the real price of crude oil in Figure 5 and expected inflation in Figure 7. Note also that inflation expectations seem to be firmly anchored, remaining within a narrow band over the entire two-year horizon. While an exogenous boom in global real economic activity causes an initial bout of optimism, expectations of a future surge in inflation and unemployment rates, a deterioration of perceived vehicle and house buying conditions, and lower expected real income depress consumer sentiment in the medium run, mirroring the response of U.S. real stock returns in Kilian and Park (2009) and Güntner (2014a).

Finally, the third column of Figure 7 plots the impulse responses to a positive other oil demand shock such as a precautionary demand shock, for example. The (largely) significant drop in the ICS during the first year resonates with a pronounced increase in expected future inflation and interest rates, as $infl_t$ rises by 16 basis points after four months, while $ratex_t$ falls by 3 basis

[21] Recall that a drop by one index point means that the share of favorable minus the share of unfavorable replies to the respective survey question decreases by 1%. For example, ten months after a positive aggregate demand shock, the net percentage of households replying that they expect *higher unemployment* during the next year increases by 1.5%.

points after one month. Consistently, U.S. consumers expect a statistically significant reduction in their real household income during the first year, whereas unemployment expectations are virtually unaffected over the entire two-year horizon. On average across households, the survey responses are therefore consistent with an "income Euler equation" (see also Dräger et al., 2016). Similarly, both perceived vehicle and house buying conditions deteriorate significantly in response to a typical other oil demand shock, dropping by −2.4 and −2.0 index points after three and two month, respectively. About six month after the shock, $infl_t$ and $ratex_t$ start to converge towards their long-run mean. We also find evidence of overshooting in interest rate expectations during the second year, consistent with an expected change in the Federal Reserve's monetary stance. On their way to recovery, the impulse response functions of veh_t and $house_t$ turn significantly negative again after twelve and 13 months, respectively. Accordingly, a sudden increase in the real price of crude oil arising from higher precautionary or speculative demand due to political or military events in the Middle East rather than disruptions of the physical supply of crude oil affects the ICS mainly through its effect on inflation and interest rate expectations as well as perceived vehicle and house buying conditions. Importantly, survey participants do not expect a disruption of the U.S. job market.

The detailed analysis of the impulse response functions of alternative measures of consumer confidence in this section contributes to our understanding of how the three structural oil supply and demand shocks affect the ICS. Apparently, U.S. consumers agree with Kilian's (2009) statement that "not all oil price shocks are alike". As a consequence, oil supply, aggregate demand, and other oil demand shocks have qualitatively and quantitatively different effects on the ICS, its component indices, and the selected other indices in Figure 7.

5.4.3 Contentment with Economic Policy

Besides inquiring households' perception of their personal financial situation and expectations about future economic conditions, the University of Michigan's Surveys of Consumers also query survey participants about their satisfaction with the government's economic policy, in particular steps taken to fight inflation and unemployment. Making use of the corresponding index, we now investigate the potential effects of structural oil supply and demand shocks on U.S. consumers' contentment with economic policy.

The bottom line of Figure 7 illustrates that a negative oil supply shock has a short-lived, marginally significant effect on $epol_t$. While the impulse response function falls to −0.8 index points after three months before rising to +1.0 index points after 13 months, only the former estimate is marginally significant. Despite an increase in the real price of crude oil in Figure 5, oil supply shocks do not seem to affect consumers' opinion about economic policy.

In response to a positive aggregate demand shock associated with the global business cycle, $epol_t$ starts to fall from the third month onwards and peaks at -1.9 basis points after one year. The corresponding impulse response function is at least marginally significant over the entire two-year horizon, recovering only during the last two months. Besides depressing consumer confidence after an initial bout of optimism, an exogenous increase in the demand for industrial commodities has a persistent negative effect on households' opinion about U.S. economic policy.

Finally, consider the response to a positive other oil demand shock such as a precautionary demand shock, for example. During the first year, $epol_t$ decreases by up to -1.8 index points. The response is statistically significant for most of the first ten months before recovering. In line with our findings for inflation, real income, and interest rate expectations, the response is concentrated in the first year following the shock. After that, U.S. households on average reply that the government is doing a fair job in fighting inflation and unemployment.

When triggering higher inflation and unemployment expectations, oil supply and demand shocks also put economic policy to the test, regardless of the fact that these shocks are beyond the control of the U.S. government or central bank. It is important to note that the impulse response functions in Figures 6 and 7 are free to respond *on impact*. The fact that various indices respond to oil supply, aggregate demand, or other oil demand shocks with a lag is determined by the data rather than by our identifying restrictions in (3).

5.5 Forecast Error Variance Decomposition

We conclude our empirical study by performing a forecast error variance decomposition (FEVD) based on the structural VAR model in (1) and the identifying restrictions in (3). Table 3 reports the percent contribution of each of the three structural oil market shocks to the forecast error variance (FEV) of the Index of Consumer Sentiment, its five component indices, and the seven alternative measures of U.S. consumer confidence at four different forecast horizons.[22]

[Table 3 about here]

Consistent with the small impact responses of consumer confidence in Figures 6 and 7, the contribution of each individual oil market shock to the FEV at $h = 1$ is generally smaller than 1%, with the exception of other oil demand shocks, which contribute 4.3% to the FEV of future expected inflation ($infl_t$) in the very short run. Hence, we conclude that U.S. consumer confidence tends to respond to structural oil supply and demand shocks with a lag.

[22] By construction, the contribution of *other shocks to consumer confidence* at forecast horizon h equals 100% minus the sum of the three corresponding columns in Table 3. Recall that this residual category has no structural interpretation and is therefore not discussed in this section.

At longer forecast horizons, the contribution of oil market shocks tends to increase, albeit with varying intensity across shocks and variables. After one year, oil supply shocks contribute merely 0.7% to the FEV of unemployment expectations, whereas the contribution to the FEV of dur_t and $ratex_t$ increases to 6.2 and 7.1%, respectively. Both types of oil demand shocks contribute in particular to household inflation expectation. While the contribution of other oil demand shocks to the FEV of $infl_t$ reaches 31% at $h = 12$, (slowly) declining thereafter, that of aggregate demand shocks continuously increases to about 21.6% in the very long run.

In line with prior evidence that oil shocks are transmitted through the demand side, both aggregate demand and other oil demand shocks contribute substantially to the FEV of U.S. consumers' perceptions about buying conditions of large household items (dur_t) and vehicles (veh_t), while oil supply shocks contribute comparatively more to the FEV of perceived house buying conditions ($house_t$). As we already know from Figure 7, household satisfaction with economic policy is influenced more by oil demand than by oil supply shocks. The bottom line in Table 3 reveals that the contribution of aggregate demand shocks to the FEV of $epol_t$ reaches 30.6% in the very long run, whereas that of other oil demand shocks remains below 10% $\forall h$.

6 Robustness Checks

We conduct three robustness checks, the results of which are discussed in what follows.[23]

6.1 Real Price of Crude Oil

First, we revisit our failure to reject the null hypothesis of a unit root in the real price of crude oil in terms of log deviations from the mean. Replacing $lrpo_t$ in the vector $\mathbf{z}_t$ by its first difference, $\Delta lrpo_t$, we find that our main results are virtually unaffected, both qualitatively and quantitatively. This is true for the historical evolution of structural shock series, the historical variance decomposition, the impulse response functions, and the forecast error variance decomposition.

6.2 Different Lag Order

Second, we experiment with higher and lower lag orders rather than choosing $p = 24$ discretely. Replicating our analysis for $p = 12$ and $p = 36$, respectively, we find that the impulse response functions in Figures 5, 6, and 7 are qualitatively unaffected. Despite the smaller number of coefficients, however, some impulse response functions are statistically less significant for $p = 12$, indicating that a sufficiently high lag order is a prerequisite to capture the full effect of oil supply and demand shocks.

[23] The results for all three robustness checks are available from the corresponding author on request.

6.3 Splitting the Sample in Two

Third, we divide our sample into the subperiods 1978:1-1996:12 and 1997:1-2015:12. To account for the correspondingly smaller sample size, we reduce the lag order to $p = 12$ in this case. When comparing the first and second half of our sample, we find qualitatively similar impulse response functions for both the oil market variables and consumer sentiment. While the negative effects of aggregate demand and other oil demand shocks on U.S. consumer sentiment survive throughout, we find a statistically significant response of the ICS to oil supply shocks during the first six months only for the second half. Accordingly, the impulse responses of the five ICS component indices are qualitatively very similar across the two subsamples. While the response pattern of interest rate and unemployment expectations as well as perceived vehicle and house buying conditions is stable over time, the response of households' satisfaction with economic policy is more pronounced and statistically significant during 1997:1-2015:12.

7 Conclusion

Although consumers are directly affected by oil price fluctuations, most of the existing literature investigates the effect of oil shocks on real GDP, inflation, or stock prices. Motivated by the fact that consumption accounts for two thirds of U.S. GDP and by prior evidence that oil shocks are transmitted also through the demand side, we investigate the effect of structural oil supply and demand shocks on U.S. "consumer sentiment" – a barometer of private households' perception of current and expected future economic conditions and uncertainty.

Our findings for the University of Michigan's Index of Consumer Sentiment (ICS) resonate with those for real stock returns in Kilian and Park (2009) and Güntner (2014a), for example. While disruptions of the physical supply, which raise the real price of crude oil, have a limited impact on the ICS, both aggregate demand shocks associated with the global business cycle and other oil demand shocks significantly depress consumer confidence during the subsequent years. The minor role for oil supply shocks indicates that U.S. consumers perceive these shocks to be only temporary, as reduced production in one country is quickly offset by other oil producers. On the contrary, other oil demand shocks such as a precautionary demand shock, for example, have a persistent negative effect on the ICS. In line with the results in Kilian and Park (2009), positive aggregate demand shocks cause a bout of optimism among U.S. consumers, followed by a statistically significant decrease of the ICS during the following two years. On the one hand, households benefit from an exogenous increase in global real economic activity through higher income and employment. On the other hand, the expected increase in nominal household income may be "eaten up" by higher future energy and consumer prices. We find that the latter effect dominates in the medium run. Although our results are not directly comparable to those in Edelstein and Kilian (2009), who investigate the effects of purchasing power shocks arising from changes in retail energy prices in a bivariate VAR model, the impulse response pattern of the ICS is very similar to that for other oil demand shocks in this paper.[24]

Insights into the transmission of oil supply and demand shocks are gained by investigating the impulse response functions of the ICS' five component indices and the replies to seven more specific survey questions concerning inflation, unemployment, and interest rate expectations, vehicle and house buying conditions, and contentment with the government's economic policy. Disruptions of the physical supply of crude oil play a limited role for the ICS, as they have only marginally significant, short-lived effects on expected inflation and perceived vehicle buying conditions, whereas they do not affect expectations about future unemployment or real household income. Aggregate demand shocks are transmitted through higher expected inflation and a substantial increase in expected unemployment after three quarters, while perceived vehicle and house buying conditions deteriorate gradually over time. Other oil demand shocks lead to pronounced increases in expected inflation and interest rates and the corresponding decreases in expected real household income, vehicle and house buying conditions in the short run, whereas unemployment expectations are literally unaffected.

Based on these results, we conclude that different oil shocks have diverse effects on the ICS because they are transmitted via different aspects of consumers' perceptions and expectations. Importantly, we also find that aggregate demand and other oil demand shocks have significant influence on household satisfaction with economic policy to fight inflation and unemployment. Given that the effect of repeated oil shocks on consumer sentiment can be sizeable, these shocks must not be neglected by U.S. policy makers when trying to contain the economic consequences of oil price fluctuations. Our results suggest that oil supply and demand shocks potentially make U.S. consumers feel worse off and put at least part of the blame on the government.

In this study, we focus on the response of the University of Michigan's ICS and several more detailed survey measures of consumer confidence. While analyzing the response of U.S. personal consumption expenditures would be equally interesting, it is beyond the scope of this paper. Similarly, an analysis of the responses of business survey measures of investor sentiment to oil supply and demand shocks is left for future research.

[24] In Edelstein and Kilian (2009), the ICS decreases by 1.6 index points on impact and remains negative and statistically significant for 18 months in response to a typical purchasing power shock.

8 References

Barsky, R. B. and L. Kilian (2002). Do we really know that oil caused the great stagflation? A monetary alternative. In B. S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*, Volume 16, pp. 137–198. MIT Press.

Barsky, R. B. and L. Kilian (2004). Oil and the macroeconomy since the 1970s. *Journal of Economic Perspectives 18*(4), 115–134.

Berument, H., N. B. Ceylan, and N. Dogan (2010). The impact of oil price shocks on the economic growth of selected MENA countries. *Energy Journal 31*(1), 149–176.

Burbidge, J. and A. Harrison (1984). Testing for the effects of oil-price rises using vector autoregressions. *International Economic Review 25*(2), 459–484.

Carroll, C. D., J. C. Fuhrer, and D. W. Wilcox (1994). Does consumer sentiment forecast household spending? If so, why? *American Economic Review 84*(5), 1397–1408.

Curtin, R. (2005). Energy and the consumer: "What me worry?". Federal Reserve Bank of Chicago, Detroit Branch.

Dräger, L., M. J. Lamla, and D. Pfajfar (2016). Are survey expectations theory-consistent? The role of central bank communication and news. *European Economic Review 85*, 84–111.

Edelstein, P. and L. Kilian (2009). How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics 56*(6), 766–779.

Elder, J. and A. Serletis (2010). Oil price uncertainty. *Journal of Money, Credit and Banking 42*(6), 1137–1159.

Engemann, K. M., K. L. Kliesen, and M. T. Owyang (2011). Do oil shocks drive business cycles? Some U.S. and international evidence. *Macroeconomic Dynamics 15*(S3), 498–517.

Gisser, M. and T. H. Goodwin (1986). Crude oil and the macroeconomy: Tests of some popular notions: A note. *Journal of Money, Credit and Banking 18*(1), 95–103.

Gonçalves, S. and L. Kilian (2004). Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. *Journal of Econometrics 123*(1), 89–120.

Güntner, J. H. F. (2014a). How do international stock markets respond to oil demand and supply shocks? *Macroeconomic Dynamics 18*(8), 1657–1682.

Güntner, J. H. F. (2014b). How do oil producers respond to oil demand shocks? *Energy Economics 44*, 1–13.

Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *Journal of Political Economy 91*(2), 228–248.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity 40*(1), 215–283.

Hamilton, J. D. and A. M. Herrera (2004). Comment: Oil shocks and aggregate macroeconomic behavior: The role of monetary policy. *Journal of Money, Credit and Banking 36*(2), 265–286.

Hooker, M. A. (1996). What happened to the oil price-macroeconomy relationship? *Journal of Monetary Economics 38*(2), 195–213.

Jo, S. (2014). The effects of oil price uncertainty on global real economic activity. *Journal of Money, Credit and Banking 46*(6), 1113–1135.

Kilian, L. (2008). Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *Review of Economics and Statistics 90*(2), 216–240.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99*(3), 1053–1069.

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 and Finance 42*, 71–87.

Kilian, L. and C. Park (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review 50*(4), 1267–1287.

Kilian, L. and C. Vega (2011). Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics 93*(2), 660–671.

Ludvigson, S. C. (2004). Consumer confidence and consumer spending. *Journal of Economic Perspectives 18*(2), 29–50.

Matsusaka, J. G. and A. M. Sbordone (1995). Consumer confidence and economic fluctuations. *Economic Inquiry 33*(2), 296–318.

Segal, P. (2011). Oil price shocks and the macroeconomy. *Oxford Review of Economic Policy 27*(1), 169–185.

Souleles, N. S. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the Michigan Consumer Sentiment Surveys. *Journal of Money, Credit and Banking 36*(1), 39–72.

Wong, B. (2015). Do inflation expectations propagate the inflationary impact of real oil price shocks? Evidence from the Michigan Survey. *Journal of Money, Credit and Banking 47*(8), 1673–1689.

9 Tables and Figures

Table 1: Summary statistics

Time series	Mean	Std.Dev.	Min	Max
$\Delta lprod_t$	0.675	16.113	-113.882	54.319
rea_t	0.428	24.644	-61.766	66.075
$lrpo_t$	3.009	0.528	1.743	4.066
$senti_t$	85.319	12.844	51.700	112.000
$pago_t$	106.070	16.546	58.000	142.000
$pexp_t$	121.290	10.755	90.000	145.000
$bus12_t$	100.436	29.613	31.000	165.000
$bus5_t$	89.875	18.338	40.000	136.000
dur_t	144.739	19.506	77.000	182.000
$infl_t$	3.663	1.734	0.400	10.400
$ratex_t$	64.932	24.623	18.000	133.000
$rinc_t$	78.952	10.788	52.000	101.000
$umex_t$	82.048	16.532	33.000	129.000
veh_t	130.340	18.222	68.000	165.000
$house_t$	139.559	30.905	37.000	182.000
$epol_t$	90.886	20.857	48.000	143.000

Notes: The sample period is 1978:1-2015:12. $\Delta lprod_t$ denotes annualized percentage changes in world crude oil production. For illustrative purposes, the real price of crude oil is expressed in terms of log levels rather than log deviations from the mean.

Table 2: Results of augmented Dickey-Fuller (ADF) unit root tests

	Time series															
Regressors	$lprod_t$	rea_t	$lrpo_t$	$senti_t$	$pago_t$	$pexp_t$	$bus12_t$	$bus5_t$	dur_t	$infl_t$	$ratex_t$	$rinc_t$	$umex_t$	veh_t	$house_t$	$epol_t$
Lagged dependent variable	-0.039	-0.036***	-0.013	-0.041*	-0.047*	-0.048	-0.053*	-0.063**	-0.053*	-0.062***	-0.074***	-0.035	-0.097***	-0.049	-0.025	-0.034
	(-2.824)	(-2.613)	(-2.221)	(-2.692)	(-2.706)	(-2.260)	(-2.899)	(-3.128)	(-2.819)	(-4.876)	(-3.562)	(-1.832)	(-4.117)	(-2.454)	(-2.164)	(-2.476)
Constant	0.419	-	0.038	3.589	4.977	5.850	5.291	5.766	7.691	0.192	4.777	2.813	7.978	6.473	3.551	3.062
	(2.821)	-	(2.144)	(2.710)	(2.703)	(2.286)	(2.811)	(3.123)	(2.830)	(3.995)	(3.339)	(1.860)	(4.092)	(2.469)	(2.171)	(2.402)
Linear time trend	0.000	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	(3.112)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Included lagged difference terms	7	4	2	12	13	5	8	3	13	20	4	5	13	5	1	4
Critical values for significance at the 10% level	-3.13	-1.62	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57	-2.57

Note: Under the null hypothesis, the series contains a unit root and the estimated coefficient on the lagged dependent variable is zero. */**/*** denotes statistical significance at the 10/5/1% level. t-statistics are reported in parentheses.

Table 3: Forecast error variance decomposition of consumer confidence based on the identifying restrictions in (3)

	Oil supply shock				Aggregate demand shock				Other oil demand shock			
Time series	$h=1$	$h=12$	$h=24$	$h\to\infty$	$h=1$	$h=12$	$h=24$	$h\to\infty$	$h=1$	$h=12$	$h=24$	$h\to\infty$
$senti_t$	0.0200	1.5342	2.4324	1.6756	0.9212	3.5933	13.554	44.038	0.0294	6.4147	8.4207	11.224
$pago_t$	0.0003	1.8897	2.2028	1.6108	0.1127	1.6820	12.517	47.983	0.2296	4.7110	8.0951	11.110
$pexp_t$	0.5995	1.4078	2.6596	1.7818	0.1953	6.2070	15.472	46.243	0.1333	11.744	14.567	14.327
$bus12_t$	0.0021	1.1486	2.0530	1.9848	0.1675	3.2103	12.761	33.407	0.0298	5.0332	7.2690	11.520
$bus5_t$	0.3807	1.5662	1.6853	1.3922	0.8649	3.2186	8.3435	30.961	0.0042	2.9410	3.3537	6.6387
dur_t	0.0920	6.1883	10.843	8.5224	1.4006	4.5061	10.775	35.618	0.0182	7.8449	11.304	12.934
$infl_t$	0.6448	2.5503	3.3289	3.8006	0.0246	10.279	19.009	21.610	4.2736	31.347	29.098	28.354
$ratex_t$	0.0332	7.1029	19.729	21.922	0.0737	2.3066	1.8878	6.1497	1.4760	4.8087	7.5304	8.9772
$rinc_t$	1.2433	2.3854	5.1466	3.8158	0.2038	6.0423	14.016	45.348	0.3346	12.758	12.416	10.730
$umex_t$	0.0064	0.7123	1.9896	2.9239	0.4550	2.5423	11.718	17.856	0.1790	0.4281	0.6741	9.3933
veh_t	0.0103	3.2635	4.9376	4.9563	0.2056	8.1369	21.796	39.209	0.0571	20.428	20.716	19.292
$house_t$	0.0568	5.9519	6.6842	5.9642	0.2619	1.6401	6.8426	11.349	0.4465	8.9939	7.2582	5.7460
$epol_t$	0.2354	1.0383	1.5892	1.6889	0.0844	3.7866	9.0176	30.565	0.1675	9.6767	8.3847	9.4064

Note: Percent contribution of structural oil market shocks to the forecast error variance of the variable in line at a forecast horizon of $h = 1, 12, 24$, and ∞ months

Figure 1: Relationship between the real price of crude oil and consumer sentiment in the U.S.

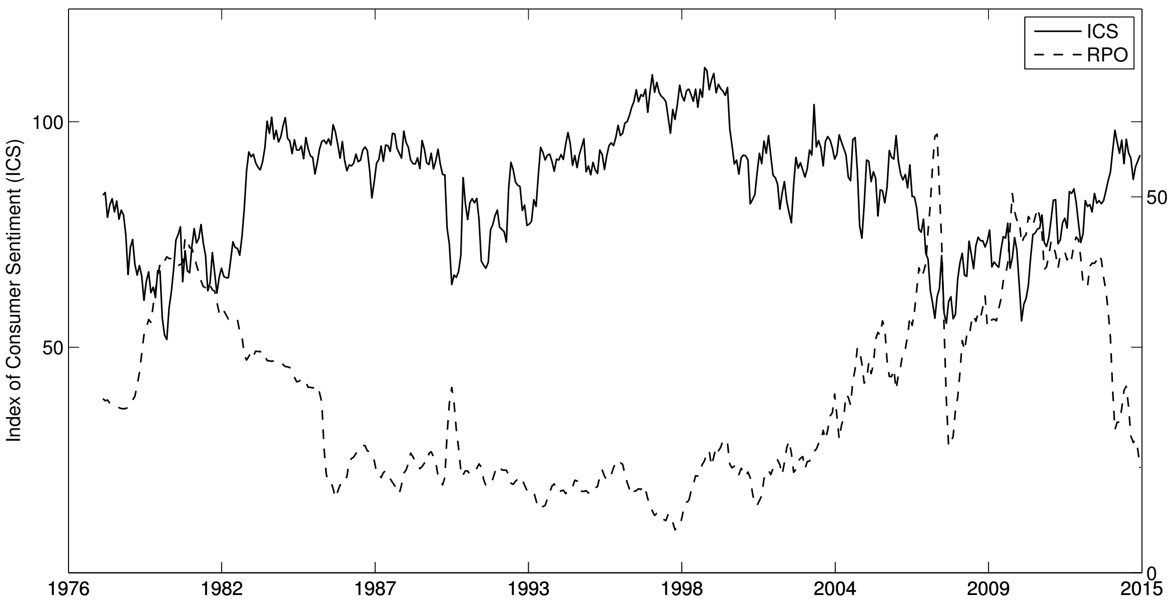

Note: The Index of Consumer Sentiment is expressed in index points, the real price of crude oil in real dollars per barrel. The unconditional contemporaneous correlation equals −0.64.

Figure 2: Historical evolution of structural oil supply and demand shocks for 1980-2015

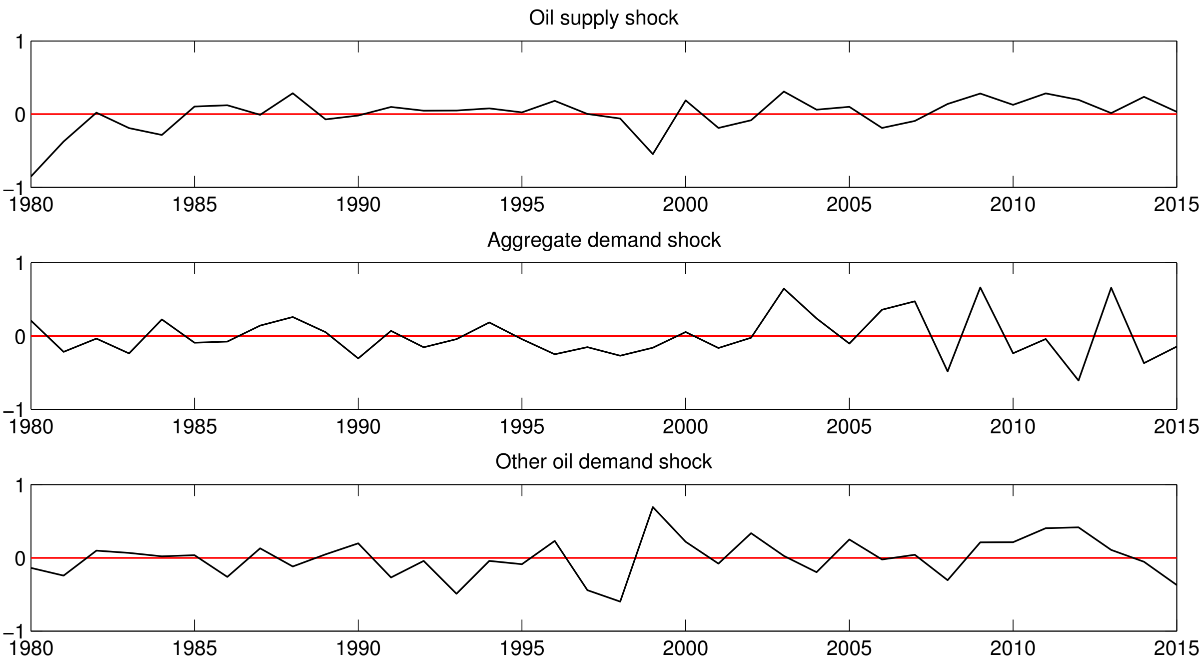

Note: Annual averages of monthly structural innovations based on the identifying strategy in (3)

Figure 3: Decomposition of historical fluctuations in the real price of crude oil for 1980:1-2015:12

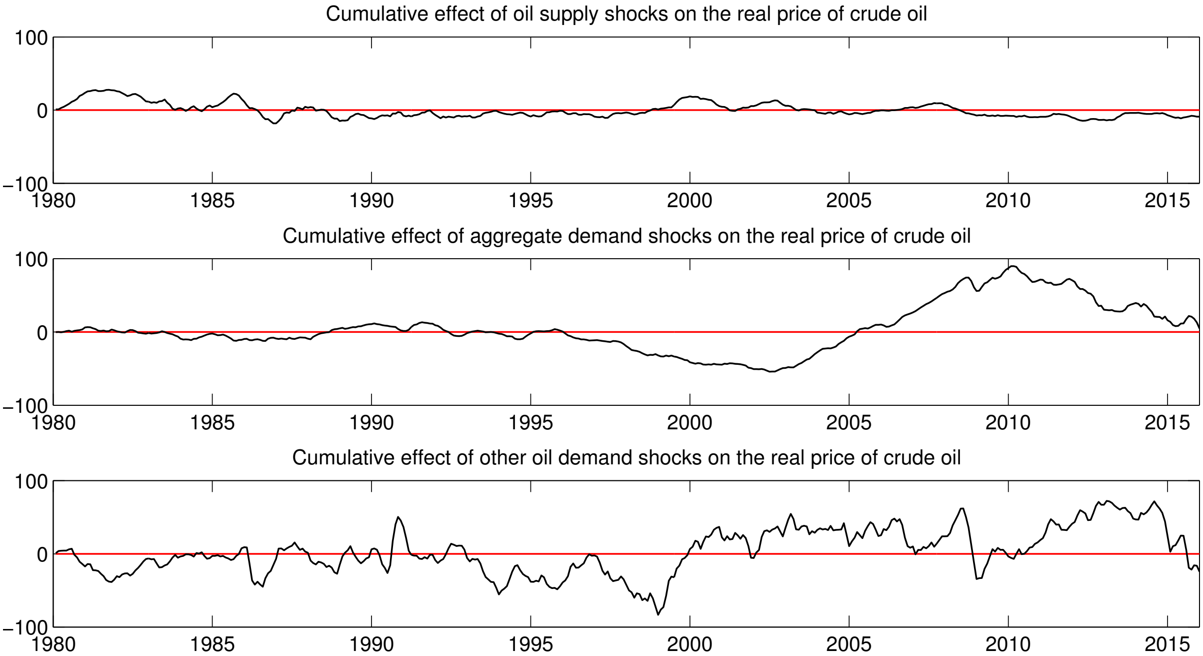

Figure 4: Decomposition of historical fluctuations in the Index of Consumer Sentiment for 1980:1-2015:12

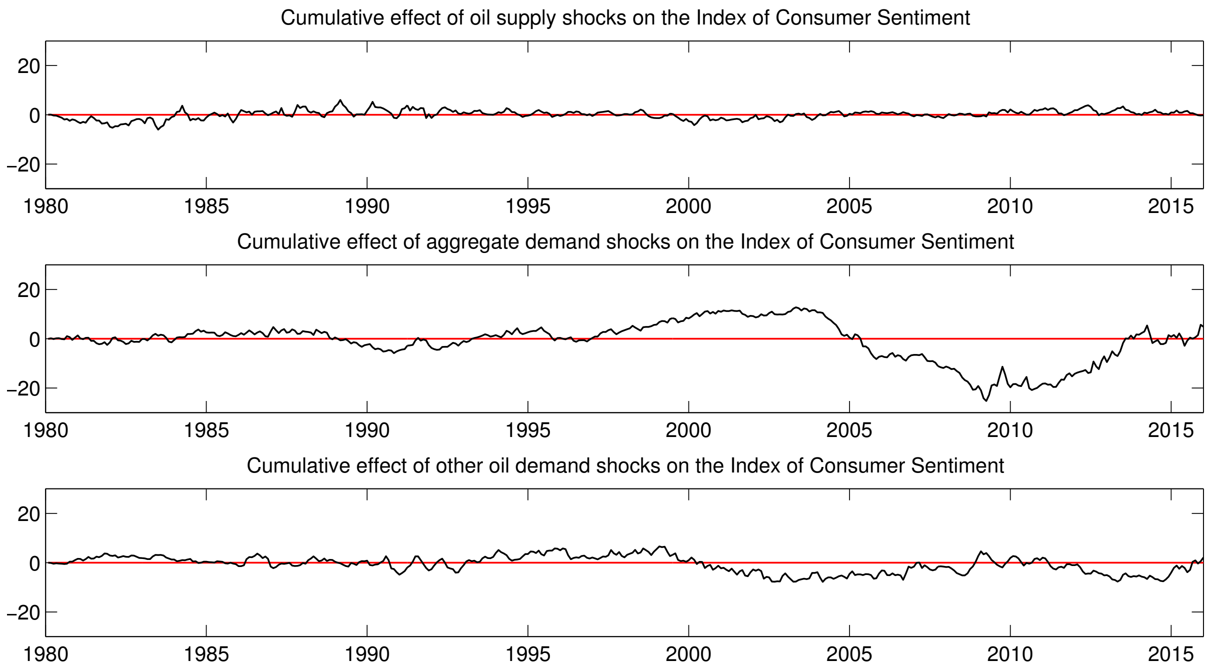

Figure 5: Impulse response functions to one-standard-deviation structural oil market shocks

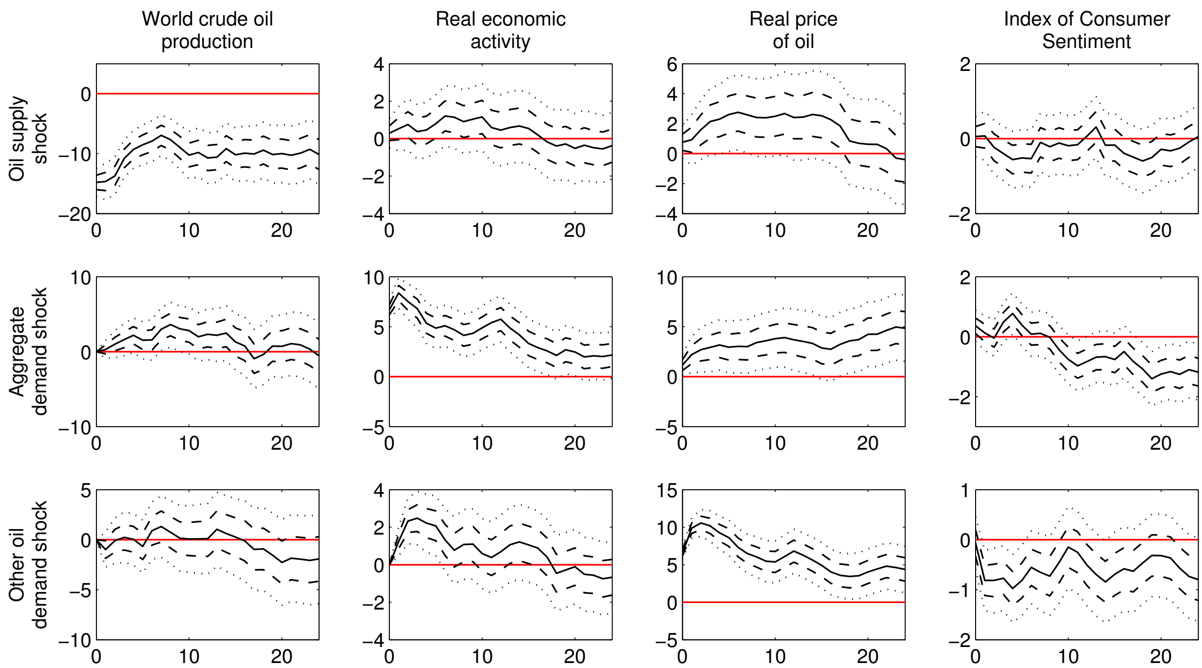

Note: Point estimates with one- and two-standard-error confidence intervals based on 5,000 replications of a recursive-design wild bootstrap (see Gonçalves and Kilian, 2004)

Figure 6: Impulse response functions of the Index of Consumer Sentiment and its five component indices to one-standard-deviation structural oil market shocks

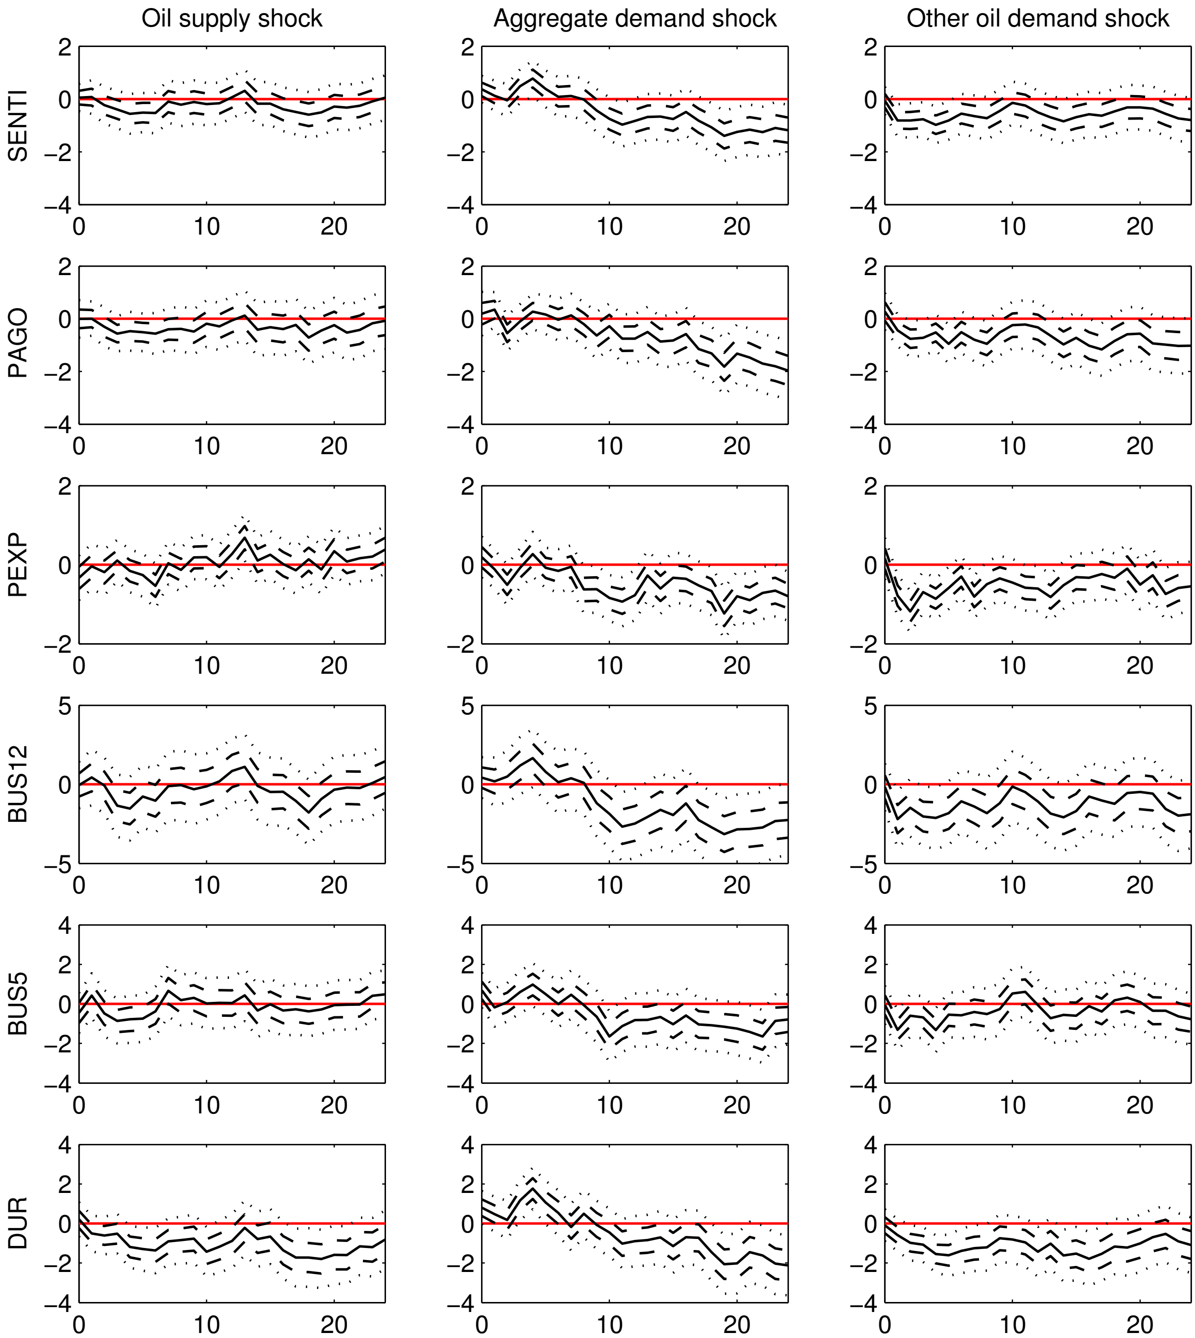

Note: Point estimates with one- and two-standard-error confidence intervals based on 5,000 replications of a recursive-design wild bootstrap (see Gonçalves and Kilian, 2004)

Figure 7: Impulse response functions of the Index of Consumer Sentiment and seven alternative measures of consumer confidence to one-standard-deviation structural oil market shocks

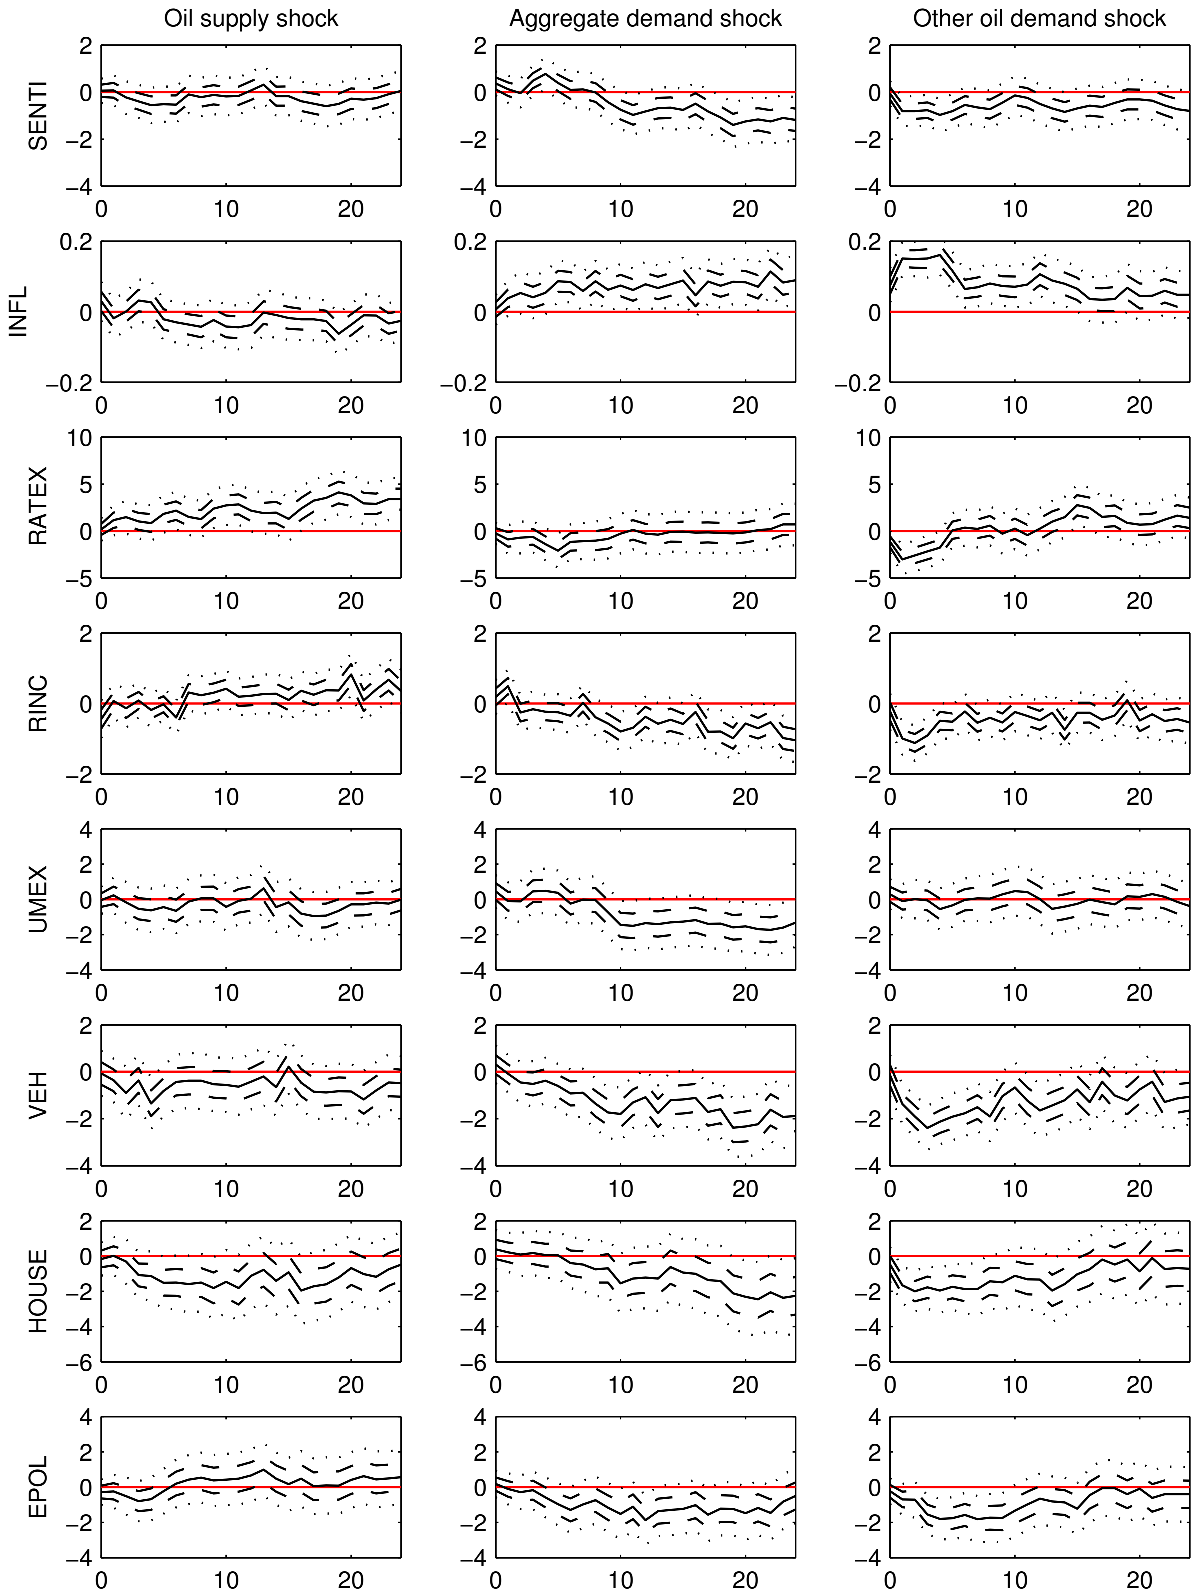

Note: Point estimates with one- and two-standard-error confidence intervals based on 5,000 replications of a recursive-design wild bootstrap (see Gonçalves and Kilian, 2004)