

Geiger, Martin; Güntner, Jochen

Working Paper

How are oil supply shocks transmitted to the U.S. economy?

Working Paper, No. 1913

Provided in Cooperation with:

Johannes Kepler University of Linz, Department of Economics

Suggested Citation: Geiger, Martin; Güntner, Jochen (2019) : How are oil supply shocks transmitted to the U.S. economy?, Working Paper, No. 1913, Johannes Kepler University of Linz, Department of Economics, Linz

This Version is available at:

<https://hdl.handle.net/10419/207126>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

How are oil supply shocks transmitted to the U.S. economy?

by

Martin GEIGER

Jochen GÜNTNER

Working Paper No. 1913

May 2019

How are oil supply shocks transmitted to the U.S. economy?*

Martin Geiger

Jochen Güntner[†]

Liechtenstein Institute

JKU Linz

May 20, 2019

Abstract

We investigate how oil supply shocks are transmitted to U.S. economic activity, consumer prices, and interest rates. Using a structural VAR approach with a combination of sign and zero restrictions, we distinguish between supply and demand channels in the transmission of exogenous changes in crude oil production. We find that the adverse effects of negative oil supply shocks are transmitted mainly through the demand side, as both output and interest rates react more strongly to oil supply shocks that shift the U.S. aggregate demand curve, while the supply side matters in transmitting oil supply shocks to consumer prices.

Keywords: Business cycles, oil supply shocks, structural VAR estimation, transmission channels

JEL classification: C32, E30, Q41, Q43

*We thank Christiane Baumeister, Maximilian Breitenlechner, and Johann Scharler for helpful comments. Our research on this project was co-financed by the Anniversary Fund of the Austrian National Bank (OeNB), Project Number 17027.

[†]Corresponding author: Jochen Güntner, Johannes Kepler University Linz. Altenberger Straße 69, 4040 Linz, Austria. Telephone: (+43)(0)732 2468 7360. Email: jochen.guentner@jku.at.

1 Introduction

The macroeconomic effects of oil supply shocks continue to receive considerable attention in academia and policy circles alike. Several candidate channels, through which oil supply shocks may be propagated, have been proposed in the literature. An oil supply shock may, for example, materialize as a cost-push shock that shifts the Phillips curve. While this is perhaps the most traditional way of thinking about the macroeconomic effects of oil supply shocks,¹ empirical evidence in favor of this interpretation remains limited (see Barsky and Kilian, 2002; Lee and Ni, 2002; Hamilton, 2009; Baumeister and Kilian, 2016a, among others). Instead, oil price shocks may affect real economic activity through the demand side, if households and firms respond to adverse oil supply shocks by cutting back on consumption and investment expenditures due to, for example, actual or perceived changes in their purchasing power induced by changes in energy prices, or uncertainty about future economic conditions. Baumeister and Kilian (2016a) provide informal evidence in favor of aggregate demand effects. Güntner and Linsbauer (2018) show that oil supply and demand shocks affect U.S. consumer sentiment mainly through expectations of future inflation and a change in real household income as well as perceived vehicle and house buying conditions, indicating a propagation through the demand side, as well.²

Although the different channels are well understood theoretically, quantitative assessments of their relative contributions to the overall effect of oil supply shocks on business cycle dynamics remain limited. In this paper we explicitly identify and evaluate two propagation channels using a small-scale structural vector-autoregressive (VAR) model that comprises both global oil market variables and U.S. macroeconomic aggregates. Following Baumeister and Peersman (2013b), we impose sign restrictions on the impulse response functions of oil market variables to identify exogenous oil supply shocks as opposed to changes in the real price of crude oil resulting from variations in oil demand (see Baumeister and Peersman, 2013a; Kilian, 2009). Following Kilian and Murphy (2012), we furthermore impose elasticity bounds on the impact multiplier matrix based on extraneous information on oil supply elasticities in order to exclude unreasonable candidate models and narrow the set of admissible models. To disentangle propagation channels, we assume that oil supply shocks affect U.S. macroeconomic aggregates by shifting either the aggregate supply or the aggregate demand curve. Hence, our main contribution is that

¹In fact, this is the interpretation commonly found in introductory textbooks, such as Blanchard (2017).

²Bloom (2009) identifies fluctuations in uncertainty based on stock market volatility and finds that spikes in his uncertainty measure are closely related to large fluctuations in oil prices.

we quantify the effects of exogenous oil supply shocks through two distinct propagation channels by combining sign and zero restrictions on global oil market variables and U.S. macroeconomic aggregates.

The restrictions that we impose on U.S. macroeconomic variables are frequently used in the empirical business cycle literature and consistent with a wide range of macroeconomic models (see, e.g., Smets and Wouters, 2005; Peersman, 2005; Smets and Wouters, 2007; Fry and Pagan, 2011). For example, consider a negative oil supply shock. We identify this shock by restricting oil production to fall and the real price of oil to increase on impact, as common in the literature. If the oil supply shock is propagated through a shift in the U.S. aggregate supply (AS) curve, according to the predictions of a standard New Keynesian DSGE model, it should ultimately result in lower real economic activity and higher consumer price inflation. The response of the interest rate is ambiguous and depends on the relative weights attached to output stabilization and inflation in the central bank's interest rate rule.³ If the shock is transmitted through a shift in the U.S. aggregate demand (AD) curve instead, then real economic activity still declines, while the reduction in aggregate demand exerts downward pressure on the price level. While a standard Taylor rule suggests that the central bank lowers the policy rate in order to counter the reduction in both real economic activity and inflation in this case, we do not impose any sign restrictions on the impulse response function of the policy rate.

Our analysis uncovers several interesting patterns in the propagation of oil supply shocks. We find that the propagation of oil supply shocks through AD is generally more important in explaining the dynamics of real economic activity and the policy rate, while the AS channel is comparatively more important in shaping the price level. Evaluating the effects of oil supply shocks on the U.S. CPI as well as CPI less energy, our analysis reveals that the relatively strong and immediate effects of AS-transmitted oil supply shocks on the price level are almost entirely driven by the direct effect of higher oil prices on energy prices that enter the consumption basket and thus the overall CPI. This result, in turn, casts doubts on the interpretation of oil supply shocks as cost-push shocks, which raise production costs across the board due to higher energy prices. Overall, AS propagation does not seem to be the dominant channel in the transmission of oil supply shocks in the U.S. during our sample period. Based on a historical decomposition, however, we detect episodes where the AS channel of oil supply shocks played some role.

³A positive interest rate response implies that monetary policy attributes a higher weight to fighting inflation. Assuming that the so-called Taylor principle is satisfied, the nominal interest rate increases by more than the inflation rate in order to ensure an increase in the real interest rate.

The rest of the paper is organized as follows. Section 2 presents the econometric methodology and our identifying strategy as well as the global oil market and U.S. macroeconomic variables. Section 3 discusses the empirical results for our baseline specification and Section 4 a selection of robustness checks. Section 5 concludes.

2 Methodology

We estimate reduced-form VAR models for which the vector of endogenous variables consists of a block of global oil market variables and a block of U.S. macroeconomic variables. Using a combination of zero and sign restrictions we then map the reduced-form VAR representation into different structural representations that allow us to separately identify distinct transmission channels of oil supply shocks to the U.S. macroeconomy — i.e. propagation through aggregate supply (AS) and propagation through aggregate demand (AD). The benefit of identifying the channels in separate models is that they need not be orthogonal by construction. Given that both channels derive from the same structural shock and may thus be active at the same time, we allow them to be correlated rather than mutually orthogonal.⁴ However, in our robustness checks, we disentangle the propagation channels within a single model by imposing mutually exclusive sign restrictions on selected impulse response functions. As a consequence, oil supply shocks propagated through AS and AD, respectively, will then be orthogonal by construction.

2.1 Reduced-form representation

The reduced-form VAR model is given by

$$x_t = c + \sum_{l=1}^p B_l x_{t-l} + e_t, \quad (1)$$

where x_t denotes a vector of endogenous variables, c a vector of intercept terms, B_l the matrix of reduced-form coefficients at lag l , and e_t a vector of possibly contemporaneously correlated residuals with covariance matrix $\Sigma_e = E(e_t e_t')$. We estimate the VAR model with $p = 12$ lags in monthly data in order to account for delays of up to one year in the transmission of oil price shocks. Our goal is to identify the domestic demand and supply channels through which oil

⁴The purpose of our identifying strategy therefore differs from Jarocinski and Karadi (2018), who use sign and exclusion restrictions in a Bayesian VAR framework to distinguish between mutually orthogonal shocks in U.S. monetary policy announcements.

supply shocks are transmitted to the U.S. macroeconomy. For this purpose, our VAR model features six variables consisting of a global oil-market block and a domestic U.S. economy block.

As a measure for the world price of crude oil, we use the logarithm of the monthly refiner acquisition cost of imported crude oil in dollars per barrel provided by the *Energy Information Administration* (EIA) since 1974:1 deflated by the U.S. CPI. By now, it is well understood that oil price *changes* are not a sufficient measure of oil price *shocks* and that the price of oil reflects a combination of oil supply and demand shocks at each point in time (see Kilian, 2009). While the focus in this paper is on disentangling the different transmission channels of oil price shocks rather than on identifying the determinants of a given oil price increase (see, e.g., Kilian and Murphy, 2014), we would like to avoid confounding shocks with expansionary and contractionary effects on the U.S. macroeconomy. For this reason, we use the monthly growth rate of world crude oil production as a second variable in the global oil market block in order to be able to identify oil supply shocks (see, e.g., Baumeister and Peersman, 2013a). In addition, we include Kilian’s (2009) index of real global economic activity based on dry-cargo ocean shipping freight rates in order to control for oil price changes that are due to fluctuations in the global business cycle rather than disruptions of the physical supply of crude oil.

As a proxy for real economic activity, we use monthly growth rates of the seasonally adjusted industrial production (IP) index provided by the *Board of Governors of the Federal Reserve System*, which measures real output for all facilities in manufacturing, mining, electric and gas utilities in the United States. As a measure of domestic price dynamics, we use monthly growth rates of the *U.S. Bureau of Labor Statistics’* seasonally adjusted consumer price index (CPI) for all urban consumers and all items. In order to capture the stance of domestic monetary policy, we add monthly averages of the *Federal Reserve Boards’* effective Federal Funds rate in levels.⁵ In 2008:11, we splice the effective Federal Funds rate with the shadow short rate constructed in Krippner (2015) in order to avoid the kink in nominal interest rates implied by an effectively binding zero lower bound in a linear VAR model. All variables are in monthly frequency and for the period 1974:2–2018:10.

In our baseline model, we use month-on-month growth rates of world crude oil production while we use the logarithm of the real oil price in levels. The chosen transformations are intended to facilitate a historical decomposition analysis, which requires that all variables in the VAR model are stationary (see Kilian and Lütkepohl, 2017).

⁵The corresponding *FRED* identifiers of the raw time series are INDPRO, CPIAUCSL, and FEDFUNDS.

2.2 Identifying strategy

Table 1 presents our identifying restrictions. Starting from identical reduced form VAR representations, we impose two distinct identification schemes in order to disentangle the propagation of oil supply shocks through aggregate demand (AD) and aggregate supply (AS), respectively. In both specifications, U.S. macroeconomic shocks are identified in addition to the (propagation of) oil supply shocks, given that the former very likely account for the bulk of domestic business cycle fluctuations.

In response to an adverse *domestic* AD shock, economic activity and inflation decline. We thus restrict IP and the U.S. CPI to decrease. According to a Taylor rule, the central bank responds to this shock by lowering the interest rate. Consequently, we restrict the policy rate to go down. Standard macroeconomic models predict that *domestic* AS shocks — i.e. exogenous shifts of the Phillips curve, such as price mark-up shocks, wage mark-up shocks, or technology shocks (see, e.g., Smets and Wouters, 2007) — move economic activity and inflation in opposite directions. Consistent with this prediction, we restrict IP to decrease and CPI inflation to increase. In addition, we restrict the response of the policy rate to be nonnegative, essentially assuming that the central bank puts relatively more weight on price stability rather than output stabilization in its objective function. Finally, *domestic* monetary policy (MP) shocks are identified by imposing a positive sign restriction on the U.S. policy rate and negative sign restrictions on IP and the U.S. CPI. These restrictions are consistent with standard macroeconomic models (see, e.g., Smets and Wouters, 2005, 2007) and commonly used in the empirical business cycle literature (Fry and Pagan, 2011). All sign restrictions are weakly imposed on impact and for six consecutive months.

While we impose sign restrictions on the *cumulated* impulse response functions in case of the month-on-month growth rates of IP and the CPI, the restrictions are imposed on the impulse responses of the U.S. policy rate *in levels*. In addition, we impose zero restrictions on the impulse responses of the global oil market variables to domestic macroeconomic shocks originating in the U.S., as it is standard in the related literature. Kilian and Vega (2011), for example, show that global energy prices, and thus the global oil market, are predetermined with respect to U.S. macroeconomic developments at a monthly frequency. By imposing the exclusion restrictions on the impulse response functions of the global oil market block, we ensure that the structural oil supply and demand shocks are orthogonal to U.S. macroeconomic developments by construction.

In addition to the U.S. domestic shocks, we identify an oil supply shock by imposing *opposite* contemporaneous sign restrictions on the real price of oil and world crude oil production, as in Baumeister and Peersman (2013a,b), for example. An exogenous change in the physical supply of crude oil shifts the oil supply curve along a given oil demand curve and moves thus oil production and the oil price in opposite directions. Moreover, we restrict the impact response of Kilian’s (2009) measure of global real economic activity to be weakly negative, as in Kilian and Murphy (2012, 2014). In order to distinguish the macroeconomic channels through which an oil supply shock is propagated, we impose additional sign restrictions implied by standard economic theory on the impulse response functions of U.S. domestic variables. Note that we vary the sign restrictions on the U.S. CPI *across VAR models* in order to identify the respective transmission channel, i.e. AD or AS, of the oil supply shock. In contrast to the restrictions for U.S. domestic shocks, the sign restrictions on the impulse response functions of U.S. IP and CPI are imposed six months after the impact period in order to allow the oil supply shocks to propagate to the U.S. economy with a certain delay. Further note that we impose the sign restrictions on the month-on-month growth rates of IP and the CPI on the cumulated impulse response functions. Thus, we require that the *average* growth rates over a six month horizon are consistent with the respective sign restrictions, while our approach does *not* impose any assumption on the speed at which oil supply shocks are propagated to the U.S. economy. In the robustness section, we show that our qualitative results are independent of the horizon for which the sign restrictions on U.S. macroeconomic aggregates are imposed.

The identification of U.S. domestic shocks and the structural oil supply shock leaves us with two additional candidate shocks. It is important to note that our identifying approach with separate propagation channels requires one residual shock category in the global oil market block that picks up the propagation channel which is left unspecified in the current VAR specification. Consider the identification of oil supply shocks that are propagated through AD. In this case, the residual shock in Table 1 captures, among other things, the effects of oil supply shocks that are propagated through AS. Although this residual shock has no structural interpretation, it is orthogonal to U.S. macroeconomic developments and to all other shocks by construction. In addition, we explicitly identify shocks to the global business cycle that are well known to affect the real price oil while, at the same time, affecting the U.S. through a number of other, potentially counteracting channels. The inclusion of Kilian’s (2009) real economic activity index and the corresponding identifying restrictions allow us to distinguish between flow supply and

flow demand shocks in the global market for crude oil, while preserving a residual category for oil-specific demand shocks such as a precautionary oil demand shock, for example, and flow oil supply shocks that are propagated through the domestic channel currently left unspecified.⁶

2.3 Econometric algorithm

We estimate the reduced-form VAR model using Bayesian techniques. Assuming that the data are generated by a multivariate Gaussian process and that the prior distribution has a flat Normal-Wishart density, the posterior distribution is also Normal-Wishart. The location parameters of the posterior distribution are summarized by the coefficient matrix $B = [B_1, \dots, B_L]'$ and the covariance matrix Σ_e .

To identify structural shocks using both zero and sign restrictions, we apply the algorithm proposed by Arias et al. (2018), which yields a set of permissible models for which the structural shocks are mutually orthogonal and all sign and zero restrictions are satisfied. The procedure is as follows: Multiply the Cholesky factor P of the reduced-form covariance matrix $\Sigma_e = PP'$ by a random orthogonal matrix Q , where $Q'Q = I$, in order to obtain the alternative decomposition $\Sigma_e = PQQ'P'$ of the covariance matrix. Pre-multiplying the reduced-form shocks by $(PQ)^{-1}$ then yields a new set of mutually orthogonal structural shocks, $\epsilon_t = (PQ)^{-1}e_t$, for which we can check whether the sign restrictions are satisfied. Note that the construction of the matrix Q guarantees that the zero restrictions in Table 1 hold.

We iterate this algorithm a large number of times proceeding along the following four steps:

1. Draw a set of parameters from the posterior distribution of the reduced-form VAR model.
2. Pre-multiply the reduced-form parameters by a random matrix $Q^{(1)}$ and check whether the resulting transformation satisfies the sign restrictions imposed on the impulse response functions.
3. If $Q^{(1)}$ does *not* satisfy the sign restrictions, keep drawing matrices $Q^{(i)}$, $i = 2, \dots, N$, until a permissible transformation is found or a maximum number N of draws is reached.
4. In the former case, retain the candidate model and proceed with the next iteration of the algorithm. In the latter case, discard all candidate models and return directly to step 1.

⁶We purge flow oil demand shocks, as they have been shown to explain the majority of fluctuations in the real price of crude oil. The identification of flow oil demand shocks also allows us to complement the sign restrictions with plausible upper bounds on the short-run price elasticity of crude oil supply associated with this shock.

Given that our identifying strategy only provides *set identification*, there may be multiple rotations that are consistent with the sign restrictions in Table 1 for each draw from the reduced-form posterior. When generating the set of admissible models, their implicit weight arises from the posterior density of the reduced form, as we keep a maximum of one candidate model for each draw from the reduced-form posterior. Note also that our prior is only flat over the reduced-form coefficients but not necessarily over the structural coefficients, as the decomposition of the variance-covariance matrix Σ_e based on random orthogonal matrices Q (where $Q'Q = I$) incorporates an implicit prior distribution (see Baumeister and Hamilton, 2015, 2018). As shown in Giacomini and Kitagawa (2018), however, inference is less sensitive to the distribution of Q if zero restrictions are imposed in addition to sign restrictions.

2.4 Elasticity bounds

Using the sign restrictions approach to identifying oil supply shocks comes with the caveat that some candidate models may be associated with economically implausible responses of crude oil production. Following Kilian and Murphy (2012), we narrow down the set of permissible models by restricting the short-run price elasticity of oil supply in line with the consensus view in the literature that this elasticity is close to zero. Within a given month, oil production is assumed to adjust only incrementally to demand-driven price signals, consistent with empirical evidence based on country-level data in Güntner (2014) and well-level data in Anderson et al. (2018).

To ensure that we focus on models that have an economic interpretation, numerical restrictions are imposed on the impact multiplier matrix $PQ = A$. More precisely, we impose an upper bound on the ratio of the impact response of oil production to the impact response of the real price of oil to a flow oil demand shock arising from fluctuations in global economic activity. Following Kilian and Murphy (2012), we only consider candidate models, where the price elasticity of oil production does not exceed a value of 0.0258, corresponding to the short-run price elasticity of crude oil supply observed during the Persian Gulf War in August 1990, which was arguably large and should therefore represent a plausible upper bound.

3 Results

In what follows, we discuss our empirical results regarding the prevalence of propagation channels, impulse response functions, forecast error variance decompositions and historical decom-

positions for oil supply shocks propagated through the aggregate demand (“AD”) and aggregate supply (“AS”) channel, respectively, as well as a benchmark specification (“No channel”), where the impulse response functions of U.S. macroeconomic variables are left unrestricted.

3.1 Prevalence of propagation channels

As a first step, we consider the structural shock series retained from the different identification schemes of oil supply shocks propagated through AD and AS. We start by evaluating whether the two models produce similar domestic shocks in order to make sure that the same dynamics of the U.S. business cycle are captured across our identifications of oil supply shocks. Only if this is the case, we can trace the differences in propagation channels to the structural model of the global oil market. Figure 1 plots the quarterly averages of the point-wise median shock series based on 100 candidate models of the structural VAR specifications in Table 1, where oil supply shocks are propagated through the AD and AS channel, respectively. The broken line corresponds to a specification, where no sign restrictions are imposed on U.S. macroeconomic variables. The first three panels of Figure 1 illustrate that the identification of domestic AD, AS, and monetary policy (MP) shocks is robust to the identified propagation channel of oil supply shocks. The strong comovement of the (point-wise medians of the) time series of domestic AD, AS, and MP shocks shows that the macroeconomic shocks are almost perfectly correlated across our two specifications of the structural VAR model as well as the “no channel” benchmark. The correlation between the underlying monthly time series of the 100 candidate models and their quarterly averages considered here equals 0.78–0.81, suggesting that disentangling oil supply shocks propagated through AD from oil supply shocks propagated through AS has little impact on the domestic macroeconomic shocks identified within the respective model.

Figure 2 plots the (quarterly averages of the monthly) time series of oil supply shocks propagated through AD and AS as well as the “no channel” benchmark and illustrates thus the prevalence of the two propagation channels of oil supply shocks during our sample period. With a contemporaneous correlation coefficient of 0.65, the shock series based on the two identifying schemes in Table 1 comove strongly, suggesting that oil supply shocks tend to be propagated through both channels simultaneously. We also find that the shock series for the “no channel” benchmark represents a convex combination of the oil supply shocks propagated through AD and AS. The contemporaneous correlation coefficients of the “no channel” shock series with the AD and AS channel are 0.64 and 0.66, respectively, for the quarterly aggregates in Figure 2 and

0.67 and 0.68 for the underlying monthly shock series. During a number of historical episodes such as the late 1980s and the late 2000s, for example, the two shock series are less synchronized and one or the other channel seems to dominate the transmission of oil supply shocks to the U.S. macroeconomy. In what follows, we investigate the implications of these differences for impulse response functions, forecast error variances, and historical decompositions.

3.2 Impulse response analysis

Figure 3 plots the impulse responses of the endogenous variables in the vector x_t to a typical oil supply shock propagated through AD, AS or no channel for a horizon of up to 24 months. We report the point-wise medians together with the 66th and 90th percentiles of the distribution of impulse response functions (IRFs) for all admissible structural VAR models. *Cumulated* IRFs are shown for all variables that enter the VAR model in growth rates, while IRFs *in levels* are shown for the policy rate and the (logarithm of the) real price of oil. Recall that, in the case of the global oil market variables, the sign restrictions identifying a flow oil supply shock are imposed on impact, while the sign restrictions on U.S. macroeconomic variables identifying a particular propagation channel are imposed six months after the impact period.

Consider first the IRFs of the global oil market variables in Panel A of Figure 3. While the responses of world oil production and Kilian’s (2009) real economic activity index are virtually identical, the response of the real price of oil is somewhat more pronounced for the case of an oil supply shock propagated through AS, while the response based on the specification with no propagation channel specified is in between AD and AS. Given that the responses of oil production are similar across identifying schemes, this implies a lower short-run price elasticity of oil supply associated with shocks propagated through the AS channel. It is well known that the overall short-run price elasticity of oil supply has declined considerably since the second half of the 1980s (Baumeister and Peersman, 2013a,b). As oil supply shocks propagated through the AS channel seem to exhibit larger magnitudes during the second half of our sample (see Figure 2), it is conceivable that the change in the oil supply elasticity translates into a shift towards the AS channel becoming more important over time. We explore this further below.

Considering the IRFs of U.S. macroeconomic variables, note that we impose sign restrictions on the cumulative responses of the month-on-month growth rate of IP and the CPI six months after impact. Thus, our identifying scheme puts structure on the responses in the sense that the direction of change of these variables is predetermined. Nevertheless, a typical oil supply

shock propagated through AD and AS, respectively, leads to different dynamics across the two identifying schemes. We observe an immediate, yet relatively short-lived response of IP to oil supply shocks propagated through AD, while the response is slightly more deferred and comparatively weaker for an oil supply shock propagated through AS. With regard to both timing and magnitude, the “no channel” benchmark, is quantitatively in between.

Oil supply shocks imply a deferred and more persistent IRF of the price level, when propagated through AD, while the effects are stronger yet less persistent, when propagated through AS. This suggests that the pass-through of oil prices to the CPI is more direct in the case of AS propagation. On the one hand, this is plausible, given that oil price have a direct effect on the overall CPI, which includes the prices of energy goods. On the other hand, we observe that lower demand only exerts a deferred drag on the price level in the case of the AD channel. When no propagation channel is identified, the U.S. CPI initially increases significantly before falling short of its long run mean.

The last row in Figure 3 plots the IRFs of the U.S. policy rate, which are generally less distinct compared to the IRFs of IP and the CPI. To some extent, this may be due to the fact that we do *not* impose sign restrictions on the policy rate. While qualitatively very similar, a typical oil supply shock propagated through the AD channel implies a quantitatively larger and more significant response of the policy rate. The finding of a more pronounced response to oil supply shocks propagated through AD is consistent with the central bank following a standard Taylor rule. In the event of an oil supply shock propagated through AS, the central bank faces an immediate trade-off between stabilizing output and stabilizing the price level, which move in opposite directions. In the case of an oil supply shock propagated through AD, on the contrary, output and the price level comove positively, requiring thus the same policy response without any trade-off. Consistent with this view, we do *not* observe a pronounced central bank reaction to oil supply shocks propagated through AS. If anything, there is downward pressure on the (unrestricted) policy rate, as the effect of the oil supply shock on the CPI fades out while IP falls short of its long-run trend. This is also the case for the “no channel” benchmark, albeit to a somewhat larger degree.

3.3 Variance decompositions

In order to investigate the relative importance of the two propagation channels of oil supply shocks for the U.S. business cycle, we compute their contribution to the forecast error variance

(FEV) of the endogenous variables based on each of the set-identified models. Table 2 reports the median contribution of an oil supply shock propagated through AD and AS as well as the “no channel” benchmark at selected horizons from 0 to 24 months, together with the 16th and the 84th percentile of the posterior distributions.⁷

Consider first Panel A, which reports the contribution of oil supply shocks to the forecast error variance decomposition (FEVD) of the three global oil market variables. Oil supply shocks propagated through the AD channel, which are characterized by a higher short-run price elasticity of oil supply, contribute a mere 3–4% to the FEV of the real price of crude oil but up to 73% to the FEV of world oil production. In contrast, oil supply shocks propagated through the AS channel contribute up to 18% to the FEV of the real oil price after six months and slightly less, i.e. around 47–59%, to the FEV of world oil production. This is again consistent with a lower short-run price elasticity of oil supply. When no propagation channel is identified, the contribution of oil supply shocks to the real oil price and world oil production is quantitatively in between those of shocks propagated through the AD and AS channels. While the median contribution of oil supply shocks to the FEV of Kilian’s (2009) index of global real economic activity is two to three times higher when propagated through AD, the bands overlap to a large extent. Again, the “no channel” benchmark is in between.

Given that the contribution of oil supply shocks propagated through AD to the FEV of the real oil price seems relatively low, while the contribution of oil supply shocks propagated through AS seems high, it is important to note that we distinguish by the respective propagation channel of *one and the same shock*. Hence, the FEV contribution of oil supply shocks commonly reported in the literature, where no propagation channel is explicitly identified (see, e.g., Kilian and Murphy, 2012), as well as in the third column of Table 2 represents a convex combination of the first two columns.

Consider now Panel B, which reports the contribution of oil supply shocks propagated through AD, AS or no channel to the FEV of U.S. macroeconomic variables. By and large, such shocks account for only a moderate amount in the FEVD of industrial production, consumer price inflation, and the (spliced) Federal Funds rate, consistent with the consensus view in the literature that oil supply shocks are a minor contributor to business cycle fluctuations in the U.S. and other developed economies (see, e.g., Kilian, 2008a,b).

⁷Recall that each column represents a separate model. In the robustness section, we replicate the analysis in Table 2 based on a single model, where oil supply shocks propagated through AD and oil supply shocks propagated through AS are mutually orthogonal by construction.

Oil supply shocks propagated through AD seem to be relatively more important in the FEVD of industrial production, while oil supply shocks propagated through AS contribute relatively more to the FEV of consumer price inflation. In both cases, the figures with no channel are strictly in between. Note that the 16th and 84th percentiles of the posterior distributions are, however, overlapping. Moreover, the large contributions of AS-propagated oil supply shocks to the FEV of consumer price inflation can be attributed, at least partly, to the direct effect of oil price increases on the energy component in the U.S. CPI. Considering the CPI *less energy* instead, the FEV share of oil supply shocks propagated through AS drops to less than 3%, falling short of the contribution of AD-propagated oil supply shocks, which account for about 5% of the FEV of CPI less energy at forecast horizons of 6, 12, and 24 months (see Table 3).⁸

The fact that the contribution of oil supply shocks propagated through AS to the dynamics of the price level decreases once we consider the CPI less energy and thus purge the direct effects of energy prices from the overall CPI casts further doubt on the importance of the cost-push shock interpretation in the transmission of oil supply shocks. The latter interpretation crucially depends on the pass-through of higher operating costs due to an oil price increase, for example, to higher output prices. Given that the contribution of oil supply shocks to the U.S. CPI less energy falls below 3%, when we constrain the real price of crude oil and the CPI less energy to comove positively, we find little evidence in favor of this interpretation.

3.4 Historical decompositions

Another interesting question is whether and to what extent oil supply shocks propagated through the AD and AS channel, respectively, contributed to the historical fluctuations of U.S. macroeconomic variables. For this purpose, Figure 5 plots the cumulative effects of oil supply shocks propagated through either channel on the growth rate of industrial production and the U.S. CPI as well as the (spliced) effective Federal Funds rate for selected historical episodes. For comparison, we also plot the cumulative effects for the model, where no channel is identified.⁹

⁸All other figures in the FEVD remain largely unaffected when replacing the U.S. CPI by the CPI *less energy* in the VAR model in (1). Given the imprecise estimates of FEV contribution, quantitative difference in the global oil market block between Table 2 and Table 3 should be taken with a grain of salt.

⁹Given that the cumulative effects of oil supply shocks propagated through either channel are based on *separate models*, their contributions to the historical decomposition of U.S. macroeconomic variables must not be added up. Instead, the overall cumulative effects of oil supply shocks is best approximated by the “no channel” benchmark, which represents a convex combination of the AD and AS channel.

3.4.1 The early 1980s recession

Starting in January of 1980, the U.S. industrial production index started to fall, culminating in a growth rate of negative 2.4% in May before recovering swiftly to a growth rate of positive 1.6% in September of the same year. Panel A illustrates that oil supply shocks indeed contributed negatively to the drop in industrial production (IP), were neutral during the subsequent recovery, and contributed again negatively to IP growth throughout 1981. Interestingly, oil supply shocks propagated through the AD channel were quantitatively more important during this episode. Around the same time, CPI inflation climbed to a record high of 1.43% in month-on-month growth rates (i.e. 17% annualized). In contrast to the popular narrative and more in line with recent empirical evidence (see, e.g., Barsky and Kilian, 2002), oil supply shocks contributed only about 0.1% to CPI inflation in January of 1980, regardless of the propagation channel. At the same time, oil supply shocks contributed between 0.75 and 1.5% to the effective Federal Funds rate (FFR), which rose to 17.6% in April of 1980, i.e. in the midst of the recession, before becoming neutral by the end of 1981. Interestingly, the contribution of oil supply shocks propagated through AD is estimated to be about twice as high as the contribution of shocks propagated through AS during this episode, while the “no channel” benchmark represents a convex combination of the former two.

3.4.2 The Persian Gulf War of 1990/91

Between June and October of 1990, the EIA’s refiner acquisition cost of imported crude oil more than doubled from 15 to 33 dollars per barrel, as Iraq invaded Kuwait, raising concerns about a spread of the conflict to neighboring Saudi Arabia. Consistent with the interpretation that the 1990 oil price increase was due to a shift in speculative demand rather than the physical disruption of crude oil supply associated with the war (see, e.g., Kilian, 2009; Kilian and Murphy, 2014), we detect only a short-lived negative impact of oil supply shocks on U.S. IP growth in August of 1990. At the same time, oil supply shocks propagated through the AS channel are estimated to have contributed about one quarter to CPI inflation, which peaked at 0.84% in month-on-month growth rates in September of 1990. The different relative importance of propagation channels can be explained by the identifying restriction of positive conditional comovement between the real price of crude oil and CPI inflation, which was likely due to the energy component in the U.S. CPI. Between mid-1989 and 1993, the effective FFR was

on a downward trend from 9.8 to 2.9%, which is also reflected in the receding contribution of oil supply shocks of either propagation. Again, the cumulative effects of the “no channel” benchmark are in between those of AD and AS.

3.4.3 The 1990s economic boom in the United States

In the mid to late 1990s, U.S. GDP growth was high and accelerating, while the real price of crude oil fell to a historical low in November of 1998. Panel C of Figure 5 suggests that, if anything, oil supply shocks had little cumulative effect on the U.S. macroeconomic variables considered throughout 1998. In 1999, as the oil price recovered from about 10 to 24 dollars per barrel, oil supply shocks contributed negatively to IP growth and positively to CPI inflation. While the latter effect is entirely due to oil supply shocks propagated through AS, oil supply shocks propagated through AD exerted an increasingly negative cumulative effect on the (spliced) effective FFR during 1999. When no propagation channel is identified, the average cumulative effect of oil supply shocks downplays the role of the AS channel on the CPI and of the AD channel on the FFR. It is important to note that, between February of 1998 and December of 1999, world crude oil production was decreasing by almost 4% from 68.1 to 65.4 million barrels per day, consistent with our finding that it was not positive oil supply shocks that fueled the contemporaneous boom in the U.S. economy.

3.4.4 The oil price surge of 2003–2008

Between January of 2003 and July of 2008, the EIA’s refiner acquisition cost of crude oil more than quadrupled from 30.3 to 127.8 dollars per barrel. The consensus in the literature is that this unprecedented oil price surge was primarily driven by a booming global economy rather than negative oil supply or speculative oil demand shocks (see, e.g., Kilian, 2009; Kilian and Murphy, 2014). Consistently, Panel D of Figure 5 indicates a largely neutral and quantitatively small effect of oil supply shocks propagated through either channel on U.S. IP growth and CPI inflation. Note that, while largely unsystematic, the cumulative effect of AS-propagated oil supply shocks on CPI inflation is estimated to be twice as large. On the contrary, oil supply shocks propagated through the AD channel initially exerted a *decreasingly positive* cumulative effect on the (spliced) effective FFR, which climbed from 2.2% in January of 2005 to 5.25% in September of 2006. Accordingly, our model attributes the monetary tightening during this period to positive flow or other oil demand shocks or factors unrelated to the global oil market.

3.4.5 The oil price drop of 2014

Following a period of relatively stable oil prices, the refiner acquisition cost of imported crude oil dropped from 100 dollars per barrel in June of 2014 to 45 dollars per barrel in January of 2015, while U.S. IP growth accelerated slightly from 0.33 in June to 0.8% in December of 2014 before dropping to -0.68% in January of 2015 and CPI inflation dropped from 0.14% in June of 2014 to -0.6% in January of 2015 (see Panel E of Figure 5). Baumeister and Kilian (2016b) and Kilian (2017) find that more than half of the oil price drop between June and December of 2014 was associated with a decline in global real economic activity and increased oil production in the United States and other countries such as Canada and Russia, that was predictable as of June 2014. The rest was associated with an unexpected decline in inventory demand in July of 2014 and an unexpected weakening of the global economy in December of 2014, respectively (see Kilian, 2017). Consistently, we find an increasingly positive cumulative effect of oil supply shocks on IP growth and a pronounced negative effect on CPI inflation during the second half of 2014, both of which are propagated equally through the AD and the AS channel. Note also that the (spliced) FFR picks up from low levels, as indicated by the upward trend in its HD.¹⁰ The fact that the cumulative effect of oil supply shocks propagated through AS tend to be above those of oil supply shocks propagated through AD suggests that the AS channel was somewhat more important in the transmission of oil supply shocks to the (spliced) FFR during this episode.

4 Robustness

In this section, we investigate the robustness of our results with regard to two crucial aspects of the identifying strategy and our measure of domestic economic activity in the vector x_t in (1). Based on these robustness checks, we conclude that our findings of a relatively larger role of oil supply shocks propagated through AD for real economic activity and AS for CPI inflation, respectively, and a minor role of oil supply shocks for U.S. business cycle fluctuations overall are qualitatively robust in these important dimensions.

¹⁰The recovery of the (spliced) FFR over this period is more evident in Figure A.1 in the Online Appendix, which plots the cumulative effects of oil supply shocks propagated through AD and AS over the entire sample.

4.1 Identification within one model

Given that we are interested in the propagation of one and the same structural shock — namely a typical oil supply shock identified by negative instantaneous comovement of world oil production and the real price of crude oil — through one of two transmission channels, our benchmark econometric approach identifies these channels in two separate models in order to allow for contemporaneous correlation. Across models, the two transmission channels are only separated by a sign restriction on the impulse response function of the U.S. CPI.

In recent work, Elbourne and Ji (2019) show that opposing sign restrictions on a single variable might not be sufficient to distinguish different *structural shocks from separate models*.¹¹ Although Elbourne and Ji (2019) re-examine a common way of identifying unconventional monetary policy shocks and are interested in the distinction of different structural shocks rather than different transmission channels of one and the same structural shocks, pundits may nevertheless question our approach of identifying the two transmission channels in separate models.

For this reason, we replicate our main results, while identifying oil supply shocks propagated through AD and AS, respectively, in *one and the same model*. Note that, in this alternative setting, the two channels are orthogonal by construction. On average over the sample period, oil supply shocks propagated through AD will thus be contemporaneously uncorrelated with oil supply shocks propagated through AS. By comparing the results across identifying approaches, we investigate whether opposing sign restrictions on the impulse response functions of the U.S. price level, as in Table 1, are sufficiently informative to disentangle the AD from the AS channel.

Table 4 summarizes the sign restrictions on the impulse response functions for identification within a single VAR model with six variables. The sign restrictions for domestic macroeconomic shocks are the same as in Table 1. In contrast to our baseline identifying scheme, however, we now identify two oil supply shocks within the same VAR model that differ only by the sign restriction imposed on the impulse response function of the U.S. CPI in order to be consistent with a propagation through AD and AS, respectively.¹² In addition, we summarize all kinds of

¹¹In particular, Elbourne and Ji (2019) show that sign restrictions on the central bank’s balance sheet are uninformative, when distinguishing unconventional monetary policy from other financial shocks, and that an otherwise identical VAR model with a series of random numbers in place of central bank assets yields qualitatively and quantitatively identical unconventional monetary policy shock series.

¹²Jarocinski and Karadi (2018) use a similar approach to disentangle conventional monetary policy shocks from news shocks associated with central bank announcements. Note that, also in this case, the authors are interested in identifying separate structural shocks rather than different transmission channels of one and the same shock.

oil demand shocks in a residual category without an economic interpretation.¹³

4.1.1 Prevalence of propagation channels

Figure 6 plots quarterly averages of the point-wise medians of oil supply shock series transmitted through AD and AS, respectively. In contrast to Figure 2, the two shock series are now orthogonal by construction. Interestingly, we find that oil supply shocks propagated through the AD channel were especially prevalent in the late 1970s and 1980s, whereas oil supply shocks propagated through the AS channel dominated in the 1990s and after 2000, in particular.

4.1.2 Impulse response analysis

The fact that we literally identify two different structural shocks rather than two transmission channels of one and the same structural shock implies that the shock series in Figure 6 may yield entirely different impulse response functions. Figure 7 plots the impulse responses of the global oil market and U.S. domestic variables to each of the two oil supply shocks. Consider first Panel A. The impulse responses in the first line indicate that oil supply shocks propagated through AD have a quantitatively smaller effect on the real price of crude oil than oil supply shocks propagated through AS. While the probability masses covered by the error bands for oil production in the second line are very similar, the pointwise median response plotted as a solid black line indicates a clustering of models at different ends of the impulse response distribution. In light of the relative prevalence of AD-propagated (AS-propagated) oil supply shocks in the 1980s (1990s and 2000s) in Figure 6, the impulse responses of oil production and prices in Figure 7 are consistent with the finding in Baumeister and Peersman (2013a) that the price elasticities of crude oil supply and demand decreased after mid-1980, indicating a stronger response of crude oil prices relative to production.

Despite the differences in the impulse response functions of the global oil market variables, we continue to find very similar responses for U.S. domestic variables in Panel B. The impulse response functions of U.S. industrial production and the CPI are in line with the sign restrictions in Table 4 by construction. The response of industrial production is somewhat more pronounced for oil supply shocks propagated through the AD channel, whereas the (cumulated) response of the CPI (inflation) is much more pronounced and immediate for oil supply shocks propagated

¹³With a single model, it is *not* possible to distinguish global aggregate demand shocks from other oil demand shocks without increasing the number of variables. At the same, it is less crucial, as the two transmission channels of oil supply shocks are now modeled explicitly.

through the AS channel. While the former difference is marginal, the latter difference is statistically significant by construction. In both cases, the (spliced) FFR decreases (marginally) significantly after three to six months.

4.1.3 Variance decompositions

Table 5 reports the shares of either oil supply shock to the forecast error variance (FEV) of the global oil market and U.S. macroeconomic variables. Similar to the FEVD for two separate models in Table 2, we find that oil supply shocks propagated through AD contribute relatively more to the FEV of U.S. industrial production, especially during the first year, while oil supply shocks propagated through AS contribute two to ten times more to the FEV of the U.S. CPI.

4.2 Horizon of sign restrictions

While imposing sign restrictions on the impact response of crude oil production and the real price of oil in order to identify oil supply shocks is standard in the literature, we need to ensure the robustness of our results with respect to the horizon for which we impose sign restrictions on the responses of U.S. macroeconomic variables to an oil supply shock propagated through aggregate demand (AD) and aggregate supply (AS), respectively.

Figures A.2 and A.3 replicate the impulse response functions (IRFs) in Figure 3, where the sign restrictions on the cumulated IRFs of (the month-on-month growth rate of) U.S. industrial production (IP) and CPI are imposed after three and twelve months, respectively, rather than after six months. The corresponding forecast error variance decompositions (FEVDs) are reported in Tables A.1 and A.2.

Consider first the IRFs in Figures A.2 and A.3. It is reassuring that a change in the timing of the identifying restrictions in Table 1 affects neither qualitatively nor quantitatively the IRFs of the oil market variables in Panel A. With regard to the IRFs of U.S. macroeconomics variables, two effects are worth mentioning. Relative to our baseline specification with $h = 6$, imposing the sign restrictions on U.S. IP and CPI after three months implies a sharper distinction of propagation channels, as both the drop in the CPI following an AD-propagated oil supply shocks occurs immediately, while the hump-shaped increase in the CPI following an AS-propagated oil supply shock fades more quickly. Imposing the sign restrictions after twelve months instead, the economic outlook becomes more pessimistic. In Panel B of Figure A.2, the cumulated IRFs of the growth rate of IP fall by more and take longer to recover relative to our baseline specification

with $h = 6$.¹⁴

Consider now the FEVDs in Tables A.1 and A.2. Imposing the sign restrictions on U.S. variables after three (twelve) months, lowers (raises) the contribution of oil supply shocks to the FEV of the real price of oil and raises (lowers) the contribution to the FEV of global real economic activity for either propagation channel at all forecast horizons (see Panel A). The differences in Panel B appear less systematic. With regard to the overlap of the 16th and 84th percentiles across specifications, the horizon for which we impose sign restrictions on the responses of U.S. macroeconomic variables does not seem to affect our qualitative results.

4.3 Alternative economic activity measures

Given that it misses services and measures output rather than value added, industrial production may be criticized for being a poor proxy for real economic activity. For this reason, we replicate our main results for two alternative measures of real economic activity, namely the U.S. civilian unemployment rate and the Chicago Fed National Activity Index (CFNAI).¹⁵ The CFNAI corresponds to the index proposed by Stock and Watson (1999), which has been found to provide a useful gauge on current and future U.S. economic activity and inflation.

Figures A.4 and A.5 replicate the IRFs in Figure 3, where U.S. IP has been replaced by the CFNAI and the civilian unemployment rate, respectively, and the sign restrictions in Table 1 have been adjusted accordingly. The corresponding FEVDs are reported in Tables A.3 and A.4.

Consider first the IRFs in Figures A.4 and A.5. Substituting the CFNAI or the civilian unemployment rate for economic activity has virtually no effect on the IRFs of the oil market variables in Panel A, while there are only minor effects on the U.S. CPI and the policy rate in Panel B. Despite the fact that the sign restrictions on all U.S. variables are imposed after six months only, the CFNAI decreases on impact in response to an oil supply shock propagated through both AD and AS, where the latter effect is quantitatively stronger. The civilian unemployment rate increases in a hump-shaped fashion for either propagation channel. Again, the response is quantitatively more pronounced for AD-propagated oil supply shocks. This is not surprising, given that the CFNAI contains a *forward-looking* component, whereas the unemployment rate is *slow-moving*. For all specifications, the “no channel” benchmark represent a convex combination of the impulse responses to AD- and AS-propagated oil supply shocks.

¹⁴Varying h has no effect on the “no channel” benchmark, where no identifying restrictions are imposed on U.S. macroeconomic variables.

¹⁵The corresponding *FRED Economic Data* identifiers are UNRATE and CFNAI.

Consider now the FEVDs in Tables A.3 and A.4. Panel A suggests somewhat larger (smaller) contribution of AS-propagated oil supply shocks to the FEV of the real oil price (oil production). Given that the latter finding carries over to the “no channel” benchmark, it seems to be driven by the reduced form of the VAR model rather than by our identifying restrictions. In Panel B, the contribution of oil supply shocks propagated through either channel to the FEV of the CFNAI is in the ballpark of that for ΔIP , whereas the contribution to the FEV of the unemployment rate is still lower. Varying the measure of real economic activity has no discernible impact on the FEVDs of the CPI and the (spliced) effective FFR, albeit the contribution of AS-propagated oil supply shocks to the FEV of ΔCPI is somewhat higher for CFNAI, while the contribution of AD-propagated oil supply shocks to the FEV of FFR is somewhat lower for both CFNAI and the unemployment rate than for the specification with ΔIP .

5 Conclusion

We investigate how oil supply shocks are transmitted to the U.S. economy. Imposing a combination of sign and zero restrictions (Arias et al., 2018) as well as elasticity bounds (Kilian and Murphy, 2012, 2014), we identify an oil supply shock propagated through aggregate demand (AD) and aggregate supply (AS), respectively, which i. raises the real price of crude oil and lowers world petroleum production contemporaneously and ii. reproduces the impulse response patterns of a domestic AD and AS shock in U.S. macroeconomic aggregates with a delay.

Although the oil supply shocks identified for either channel are correlated by construction, we find interesting differences in their contribution to the forecast error variances (FEVs) of U.S. industrial production and consumer price inflation. While the contribution of oil supply shocks is generally moderate, in line with prior informal evidence, shocks propagated through the AD channel explain between 1.4 and 5.7% of the FEV of the growth rate of industrial production, whereas shocks propagated through the AS channel contribute between 2 and 6.8% to the FEV of CPI inflation. The latter finding can be attributed to the direct effect of oil prices on the energy component of the U.S. CPI and largely disappears once we investigate the *CPI less energy* instead. Moreover, we find quantitative and qualitative differences in the cumulative effects of oil supply shocks propagated through either channel during selected historical episodes such as the early 1980s recession and the oil price surge of 2003–2008, for example.

In our benchmark specification, we do not decompose fluctuations in oil supply into *mutually*

orthogonal “channels” (a.k.a. “shocks”), as these channels may operate simultaneously most of the time. Instead, we estimate separate models for each channel in order to allow for a nonzero correlation between oil supply shocks propagated through AD and oil supply shocks propagated through AS. Accordingly, we focus on the relative contribution of oil supply shocks propagated through AD and AS, respectively, to the FEV of U.S. macroeconomic aggregates.

When identifying two different oil supply shocks in one and the same model as a robustness check, our main results go through, although the resulting shock series are mutually orthogonal by construction. We interpret this as evidence that our separating sign restrictions on U.S. CPI are sufficiently informative to distinguish oil supply shocks propagated through AD from oil supply shocks propagated through AS (see Elbourne and Ji, 2019).

References

- Anderson, S. T., Kellogg, R., Salant, S. W., 2018. Hotelling under Pressure. *Journal of Political Economy* 126, 984–1026.
- Arias, J. E., Rubio-Ramírez, J. F., Waggoner, D. F., 2018. Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications. *Econometrica* 86, 685–720.
- Barsky, R. B., Kilian, L., 2002. Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative. *NBER Macroeconomics Annual* 16, 137–198.
- Baumeister, C., Hamilton, J. D., 2015. Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica* 83, 1963–1999.
- Baumeister, C., Hamilton, J. D., 2018. Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations. *Journal of Monetary Economics* 100, 48–65.
- Baumeister, C., Kilian, L., 2016a. Lower Oil Prices and the U.S. Economy: Is This Time Different? *Brookings Papers on Economic Activity* 47, 287–357.
- Baumeister, C., Kilian, L., 2016b. Understanding the Decline in the Price of Oil since June 2014. *Journal of the Association of Environmental and Resource Economists* 3, 131–158.

- Baumeister, C., Peersman, G., 2013a. The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market. *Journal of Applied Econometrics* 28, 1087–1109.
- Baumeister, C., Peersman, G., 2013b. Time-Varying Effects of Oil Supply Shocks on the US Economy. *American Economic Journal: Macroeconomics* 5, 1–28.
- Blanchard, O., 2017. *Macroeconomics*, 7th Edition. MIT Press.
- Bloom, N., 2009. The Impact of Uncertainty Shocks. *Econometrica* 77, 623–685.
- Elbourne, A., Ji, K., 2019. Do zero and sign restricted SVARs identify unconventional monetary policy shocks in the euro area? CPB Discussion Paper.
- Fry, R., Pagan, A., 2011. Sign Restrictions in Structural Vector Autoregressions: A Critical Review. *Journal of Economic Literature* 49, 938–960.
- Giacomini, R., Kitagawa, T., 2018. Robust Inference for Set-Identified Models. Mimeo.
- Güntner, J., 2014. How Do Oil Producers Respond to Oil Demand Shocks? *Energy Economics* 44, 1–13.
- Güntner, J., Linsbauer, K., 2018. The Effects of Oil Supply and Demand Shocks on U.S. Consumer Sentiment. *Journal of Money, Credit and Banking* 50, 1617–1644.
- Hamilton, J. D., 2009. Causes and Consequences of the Oil Shock of 2007-08. *Brookings Papers on Economic Activity* 40, 215–283.
- Jarocinski, M., Karadi, P., 2018. Deconstructing Monetary Policy Surprises — The Role of Information Shocks. CEPR Discussion Paper 12765.
- Kilian, L., 2008a. A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries. *Journal of the European Economic Association* 6, 78–121.
- Kilian, L., 2008b. Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *Review of Economics and Statistics* 90, 216–240.
- Kilian, L., 2009. Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* 99, 1053–1069.

- Kilian, L., 2017. The Impact of the Fracking Boom on Arab Oil Producers. *Energy Journal* 38, 137–160.
- Kilian, L., Lütkepohl, H., 2017. *Structural Vector Autoregressive Analysis. Themes in Modern Econometrics*. Cambridge University Press.
- Kilian, L., Murphy, D. P., 2012. Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. *Journal of the European Economic Association* 10, 1166–1188.
- Kilian, L., Murphy, D. P., 2014. The Role of Inventories and Speculative Trading in the Global Market for Crude Oil. *Journal of Applied Econometrics* 29, 454–478.
- Kilian, L., Vega, C., 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, 660–671.
- Krippner, L., 2015. *Zero Lower Bound Term Structure Modeling: A Practitioner’s Guide*. Palgrave Macmillan.
- Lee, K., Ni, S., 2002. On the Dynamic Effects of Oil Price Shocks: A Study Using Industry Level Data. *Journal of Monetary Economics* 49, 823–852.
- Peersman, G., 2005. What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions. *Journal of Applied Econometrics* 20, 185–207.
- Smets, F., Wouters, R., 2005. Comparing Shocks and Frictions in US and Euro Area Business Cycles: A Bayesian DSGE Approach. *Journal of Applied Econometrics* 20, 161–183.
- Smets, F., Wouters, R., 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review* 97, 586–606.
- Stock, J. H., Watson, M. W., 1999. Forecasting Inflation. *Journal of Monetary Economics* 44, 293–335.

Figures and Tables

Table 1: Sign restrictions on impulse response functions

	AD	AS	MP	Propagation AD, AS	global BC shock	R
Δ oil production	0	0	0	$\downarrow(0)$	$\uparrow(0)$	
real oil price	0	0	0	$\uparrow(0)$	$\uparrow(0)$	$\uparrow(0)$
real econ. activity	0	0	0	$\downarrow(0)$	$\uparrow(0)$	
Δ IP	$\downarrow(0-6)$	$\downarrow(0-6)$	$\downarrow(0-6)$	$\downarrow, \downarrow(6)$		
Δ CPI	$\downarrow(0-6)$	$\uparrow(0-6)$	$\downarrow(0-6)$	$\downarrow, \uparrow(6)$		
spliced FFR	$\downarrow(0-6)$	$\uparrow(0-6)$	$\uparrow(0-6)$			

Notes: Oil production, the industrial production index, and the CPI enter in monthly growth rates. Sign restrictions on the responses of the industrial production index and the CPI are imposed on cumulated impulse response functions. The period for which we impose sign restriction is indicated in parenthesis. “Propagation AD, AS” indicates separate structural models and “ \downarrow, \uparrow ” the corresponding sign restrictions. For the “No channel” benchmark, no sign restrictions are imposed on U.S. macroeconomic variables in case of the oil supply shock.

Figure 1: Quarterly averages of monthly U.S. macroeconomic shock series associated with each propagation channel or no channel

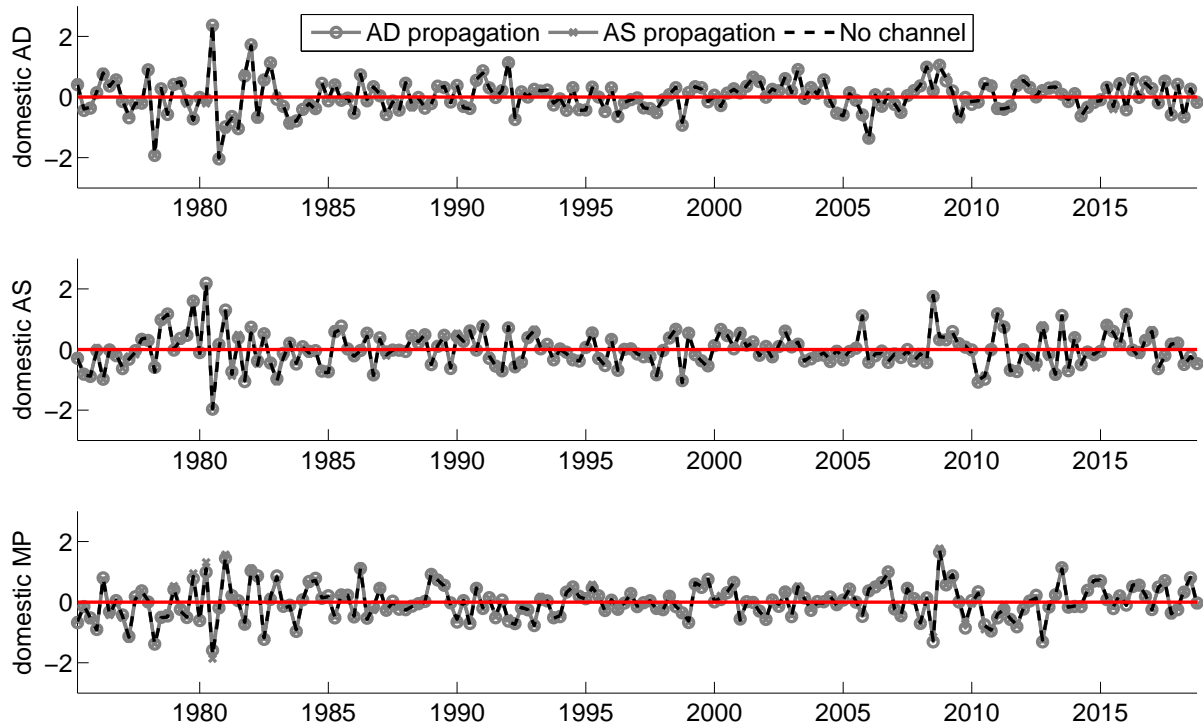


Figure 2: Quarterly averages of monthly oil supply shock series associated with each propagation channel or no channel

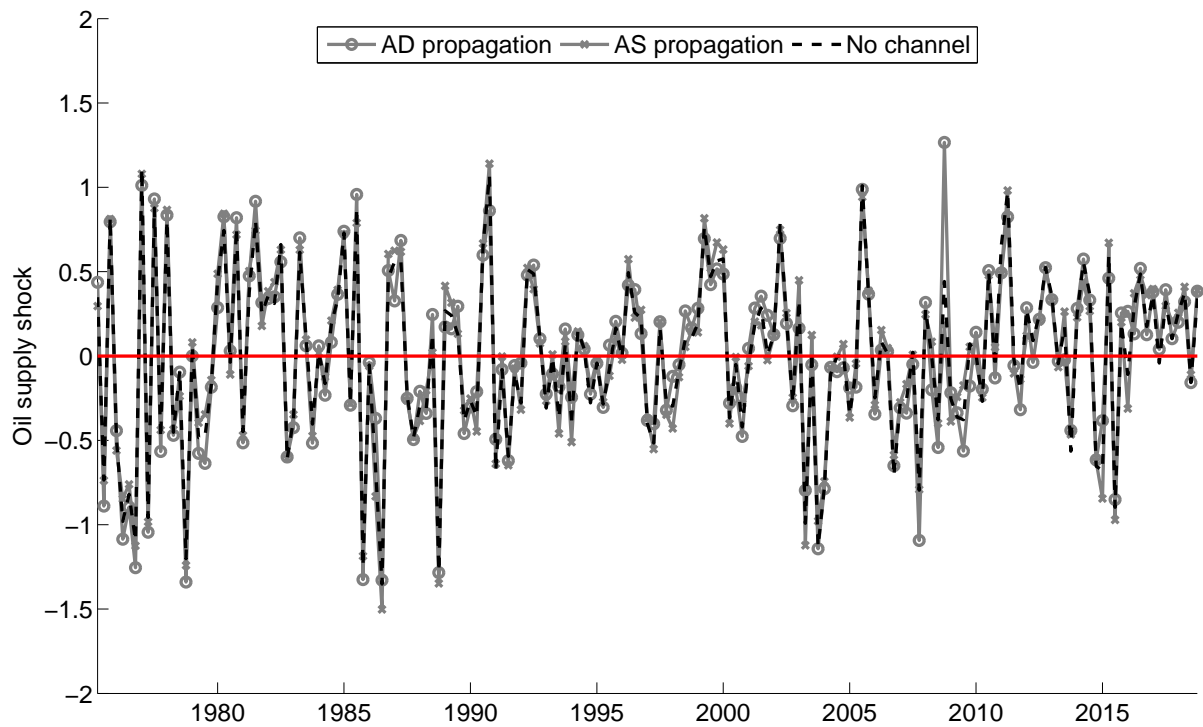
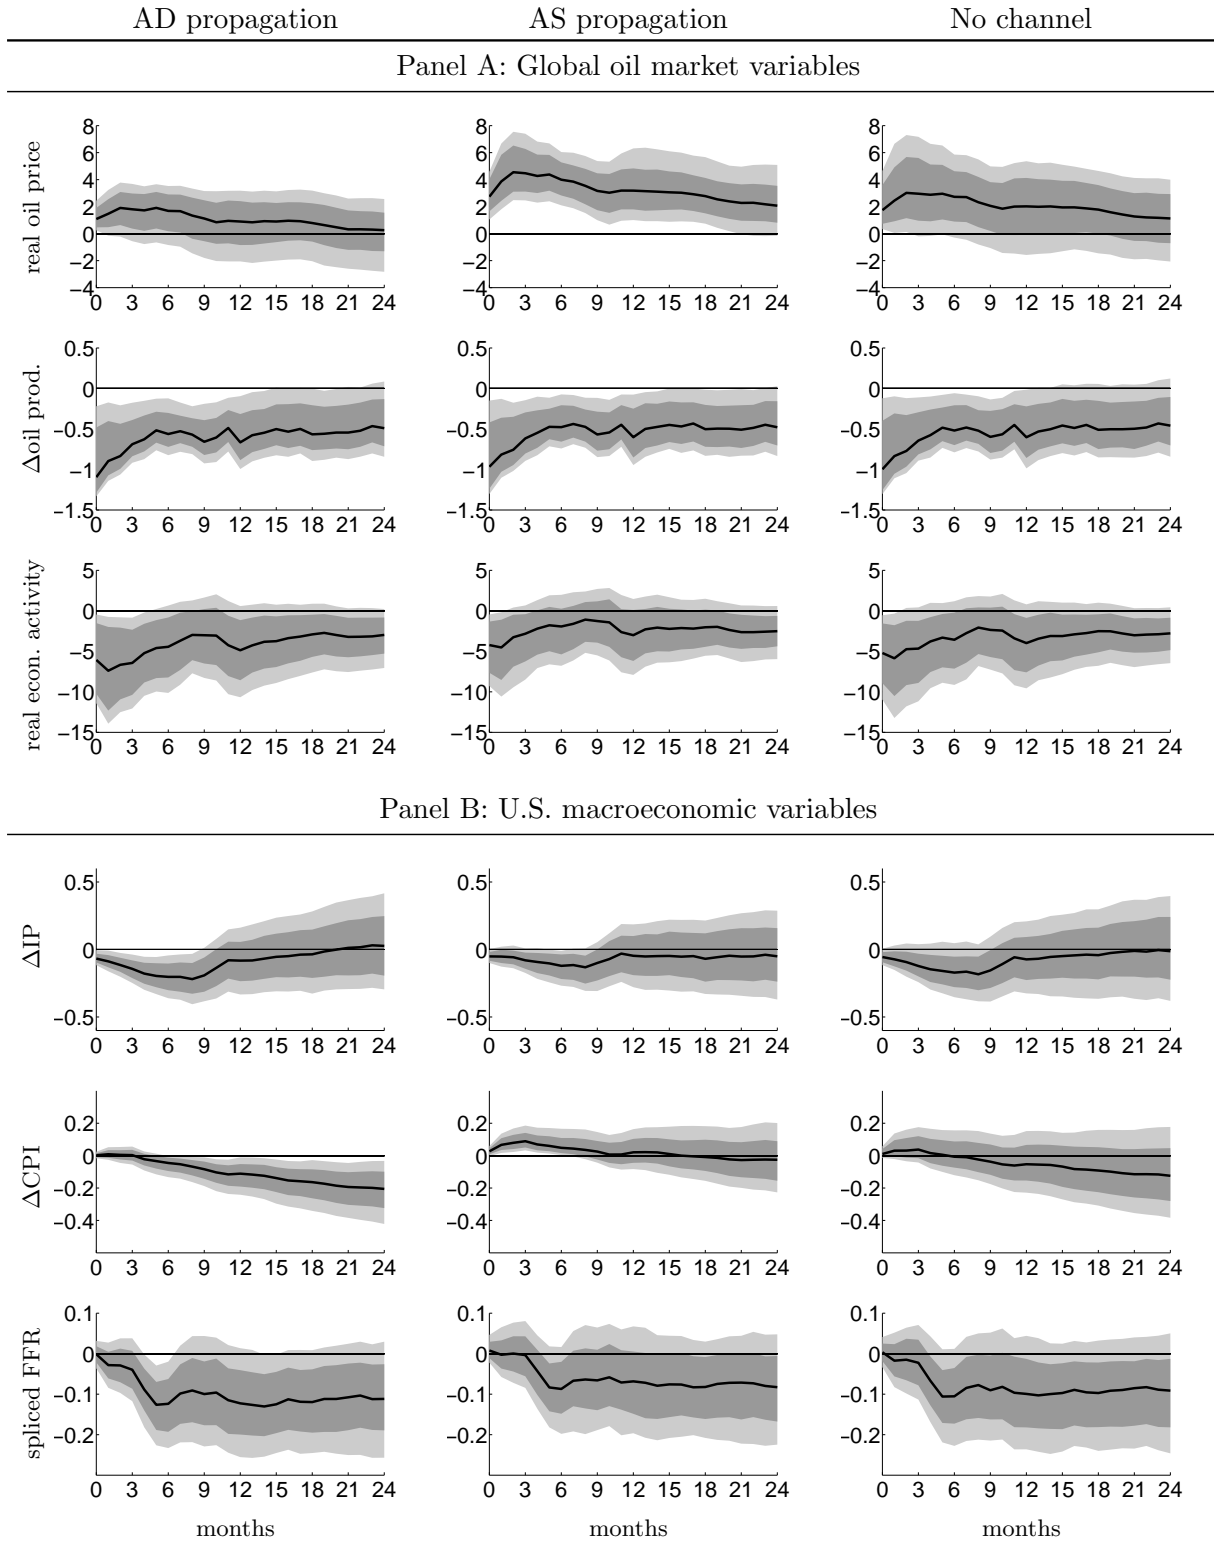


Figure 3: Impulse responses to oil supply shocks propagated through AD, AS, or no channel



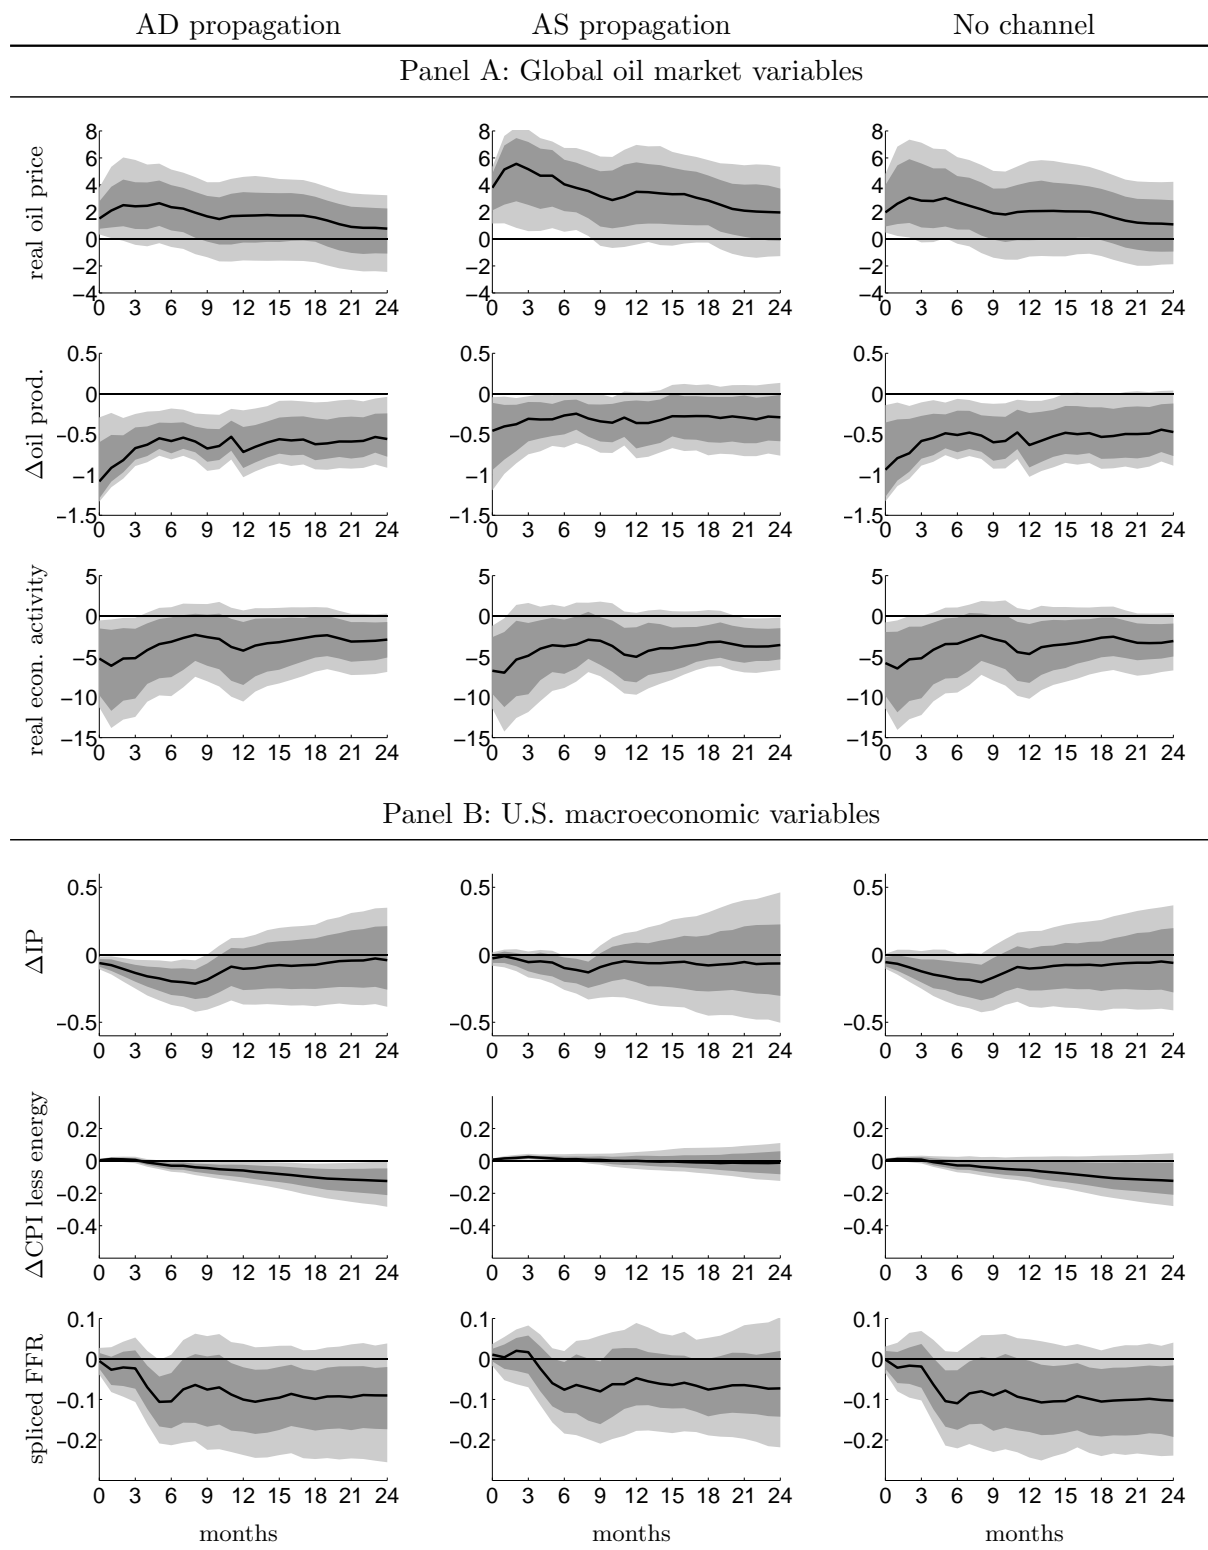
Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model. For the sign restrictions, see Table 1.

Table 2: Forecast error variance decomposition (FEVD): Baseline specification (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	3.37 (0.62, 10.27)	21.96 (7.73, 45.99)	8.37 (1.55, 37.39)
	6	3.70 (0.81, 9.52)	20.69 (10.81, 40.21)	9.72 (1.54, 32.17)
	12	3.76 (1.18, 9.67)	20.11 (10.26, 36.66)	9.52 (1.89, 29.01)
	24	3.61 (1.18, 9.04)	18.61 (8.32, 33.53)	9.01 (2.05, 25.54)
Δ oil production	0	69.27 (14.20, 97.13)	55.70 (10.16, 90.39)	57.70 (9.30, 93.65)
	6	62.15 (13.44, 86.40)	49.48 (10.08, 80.80)	52.61 (9.42, 83.45)
	12	57.22 (12.69, 79.08)	46.01 (9.59, 74.43)	48.11 (9.54, 76.65)
	24	55.39 (12.90, 75.73)	44.43 (9.65, 71.13)	46.59 (9.49, 73.97)
real econ. activity	0	24.61 (1.60, 72.35)	12.19 (1.19, 40.71)	19.06 (1.59, 54.29)
	6	20.19 (2.07, 54.69)	5.98 (0.96, 24.26)	11.82 (1.47, 39.55)
	12	18.11 (2.25, 52.58)	5.36 (1.22, 22.62)	10.33 (1.67, 37.21)
	24	18.19 (2.39, 53.69)	6.96 (1.55, 25.89)	11.18 (2.12, 39.30)
Panel B: U.S. macroeconomic variables				
Δ IP	0	1.38 (0.31, 2.80)	0.81 (0.10, 2.10)	0.96 (0.14, 2.68)
	6	2.84 (1.59, 4.47)	1.94 (0.98, 3.20)	2.50 (1.31, 4.31)
	12	5.01 (3.56, 6.93)	3.73 (2.57, 5.35)	4.55 (2.89, 6.71)
	24	5.67 (4.05, 7.61)	4.17 (2.96, 5.84)	5.05 (3.35, 7.34)
Δ CPI	0	0.17 (0.02, 0.77)	1.98 (0.23, 5.67)	0.54 (0.03, 4.48)
	6	2.79 (1.37, 4.95)	5.29 (2.41, 11.56)	3.69 (1.72, 8.78)
	12	4.37 (2.60, 7.14)	6.30 (3.37, 12.28)	5.41 (2.61, 10.07)
	24	5.36 (3.31, 8.41)	6.78 (3.89, 12.28)	6.25 (3.36, 10.85)
spliced FFR	0	0.09 (0.01, 0.43)	0.11 (0.01, 0.52)	0.11 (0.01, 0.50)
	6	2.16 (0.70, 4.80)	1.12 (0.35, 2.77)	1.55 (0.51, 4.08)
	12	3.29 (1.01, 7.32)	1.68 (0.54, 4.46)	2.48 (0.74, 6.41)
	24	4.77 (1.34, 10.05)	2.24 (0.65, 6.85)	3.29 (0.94, 8.95)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model. For the sign restrictions, see Table 1.

Figure 4: Impulse responses to oil supply shocks propagated through AD, AS or no channel: Specification with CPI less energy



Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model. For the sign restrictions, see Table 1.

Table 3: Forecast error variance decomposition (FEVD): Specification with CPI less energy (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	6.32 (1.56, 21.67)	38.90 (11.43, 65.31)	10.57 (2.30, 43.44)
	6	7.22 (1.40, 20.33)	30.05 (7.93, 56.02)	9.84 (1.64, 34.95)
	12	7.21 (1.80, 19.31)	25.89 (7.68, 50.69)	9.38 (1.91, 31.28)
	24	6.57 (1.86, 18.26)	24.09 (6.04, 45.00)	8.98 (2.10, 26.64)
Δ oil production	0	67.27 (20.18, 96.39)	12.04 (0.67, 51.51)	51.39 (7.36, 92.85)
	6	60.73 (19.00, 86.95)	11.69 (2.16, 46.85)	47.05 (7.57, 83.90)
	12	56.27 (17.84, 80.29)	11.47 (2.83, 43.92)	43.38 (7.46, 77.99)
	24	53.82 (17.71, 77.09)	11.62 (3.36, 42.32)	41.22 (7.84, 75.22)
real econ. activity	0	18.56 (1.57, 63.38)	30.08 (4.77, 62.51)	21.54 (2.50, 64.81)
	6	12.47 (1.31, 49.20)	14.82 (2.08, 43.36)	13.26 (1.48, 50.76)
	12	11.89 (1.60, 46.35)	15.22 (2.92, 42.96)	13.13 (1.65, 47.63)
	24	12.32 (1.79, 47.65)	18.81 (3.41, 43.92)	14.47 (2.14, 48.66)
Panel B: U.S. macroeconomic variables				
Δ IP	0	1.23 (0.34, 2.47)	0.24 (0.02, 1.18)	0.95 (0.14, 2.33)
	6	2.61 (1.48, 4.30)	1.82 (1.02, 2.94)	2.50 (1.27, 4.39)
	12	4.77 (3.26, 6.74)	3.40 (2.38, 5.29)	4.39 (2.85, 6.57)
	24	5.30 (3.65, 7.39)	4.11 (2.84, 6.13)	5.00 (3.29, 7.36)
Δ CPI less energy	0	0.19 (0.02, 0.69)	0.54 (0.14, 1.26)	0.24 (0.03, 0.87)
	6	4.18 (2.23, 6.65)	1.91 (0.98, 3.38)	3.91 (1.86, 6.42)
	12	4.49 (2.88, 6.77)	2.49 (1.57, 3.75)	4.28 (2.39, 7.05)
	24	5.21 (3.09, 8.85)	2.69 (1.70, 4.18)	4.94 (2.58, 9.05)
spliced FFR	0	0.10 (0.01, 0.44)	0.10 (0.01, 0.37)	0.10 (0.01, 0.46)
	6	1.65 (0.44, 4.01)	0.83 (0.34, 1.96)	1.64 (0.47, 3.97)
	12	2.34 (0.62, 6.22)	1.43 (0.50, 4.17)	2.40 (0.68, 6.56)
	24	3.53 (1.02, 9.15)	2.41 (0.86, 5.71)	3.85 (0.99, 9.71)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model. For the sign restrictions, see Table 1.

Figure 5: Historical decomposition of U.S. macroeconomic variables: Cumulative effects of oil supply shocks propagated through AD, AS or no channel

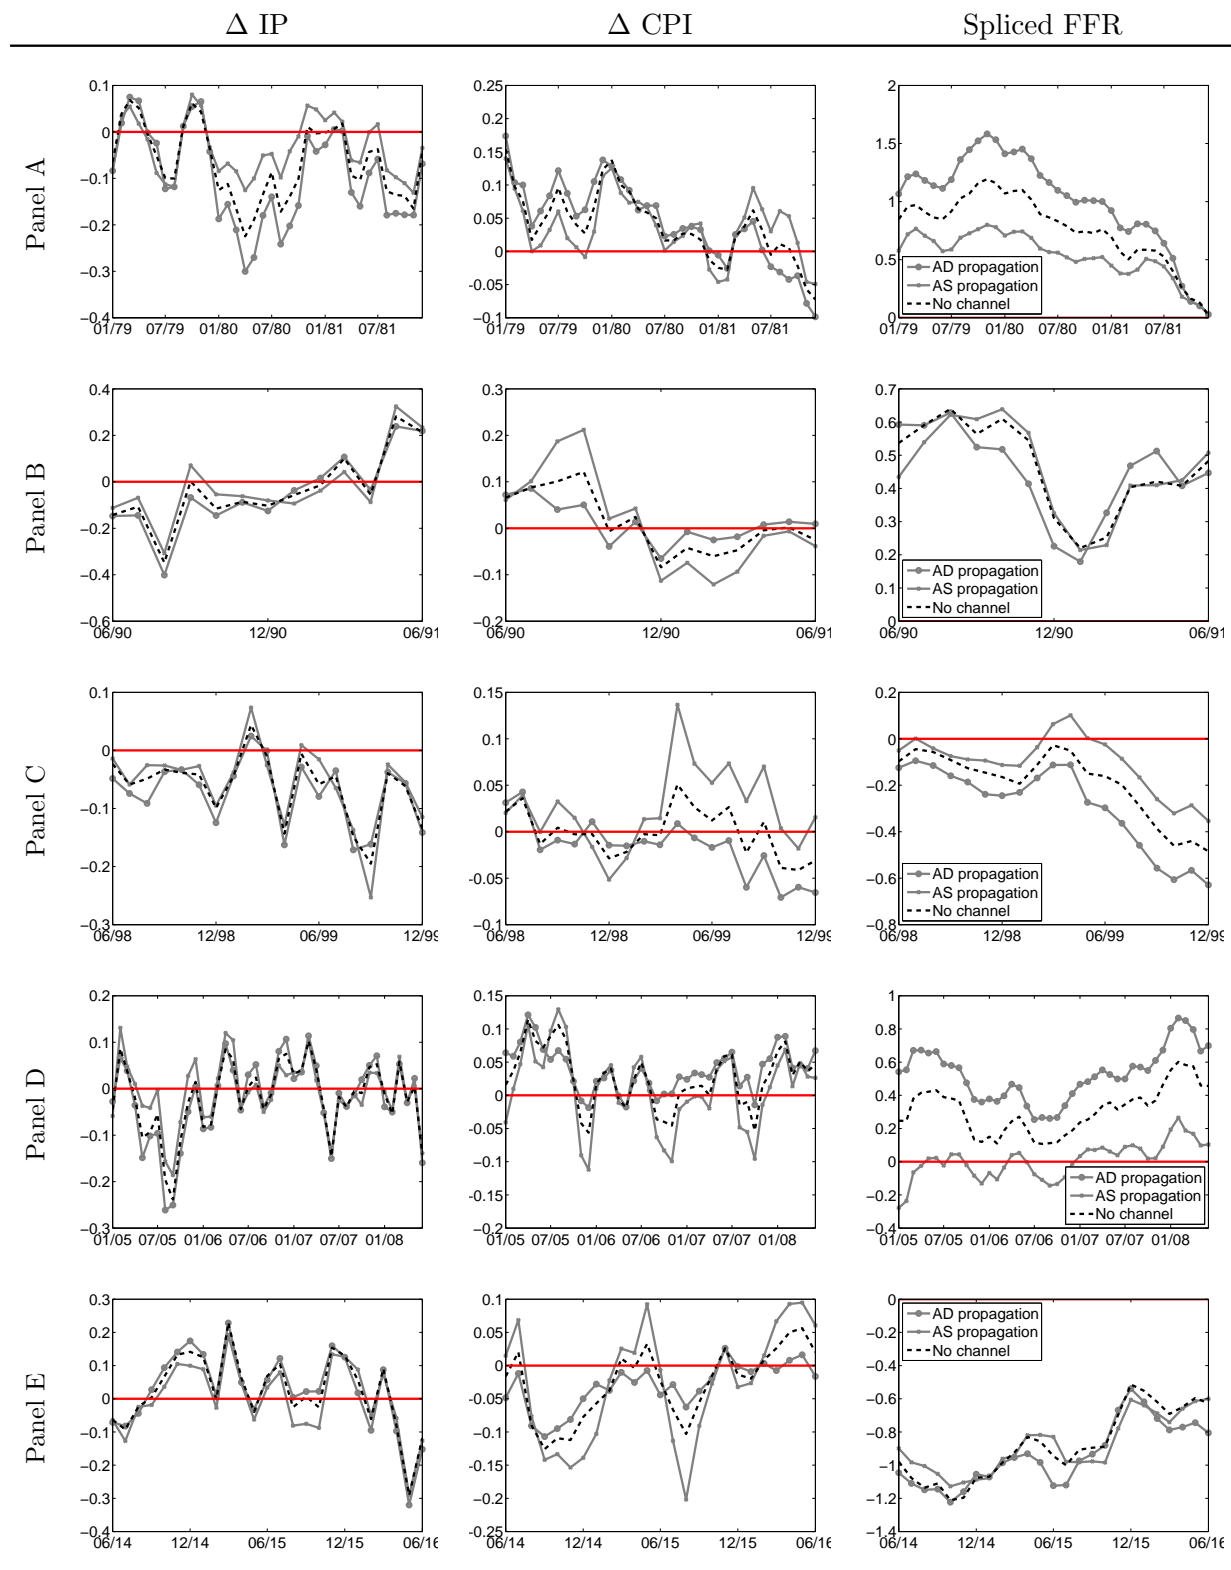


Table 4: Sign restrictions on impulse response functions: Single VAR model

	AD	AS	MP	Prop. AD	Prop. AS	oil demand
Δ Oil Production	0	0	0	$\downarrow(0)$	$\downarrow(0)$	$\uparrow(0)$
$\log(\text{realOilP})$	0	0	0	$\uparrow(0)$	$\uparrow(0)$	$\uparrow(0)$
Kilian WD	0	0	0	$\downarrow(0)$	$\downarrow(0)$	
Δ IP	$\downarrow(0-6)$	$\downarrow(0-6)$	$\downarrow(0-6)$	$\downarrow(6)$	$\downarrow(6)$	
Δ CPI	$\downarrow(0-6)$	$\uparrow(0-6)$	$\downarrow(0-6)$	$\downarrow(6)$	$\uparrow(6)$	
spliced FFR	$\downarrow(0-6)$	$\uparrow(0-6)$	$\uparrow(0-6)$			

Notes: Oil production, the industrial production index, and the CPI enter in monthly growth rates. Sign restrictions on the responses of the industrial production index and the CPI are imposed on cumulated impulse response functions. The period for which we impose sign restriction is indicated in parenthesis. “Prop. AD” and “Prop. AS” indicate mutually orthogonal shocks within a single structural VAR model and “ \downarrow, \uparrow ” the corresponding sign restrictions.

Figure 6: Quarterly averages of monthly oil supply shock series based on a single VAR model

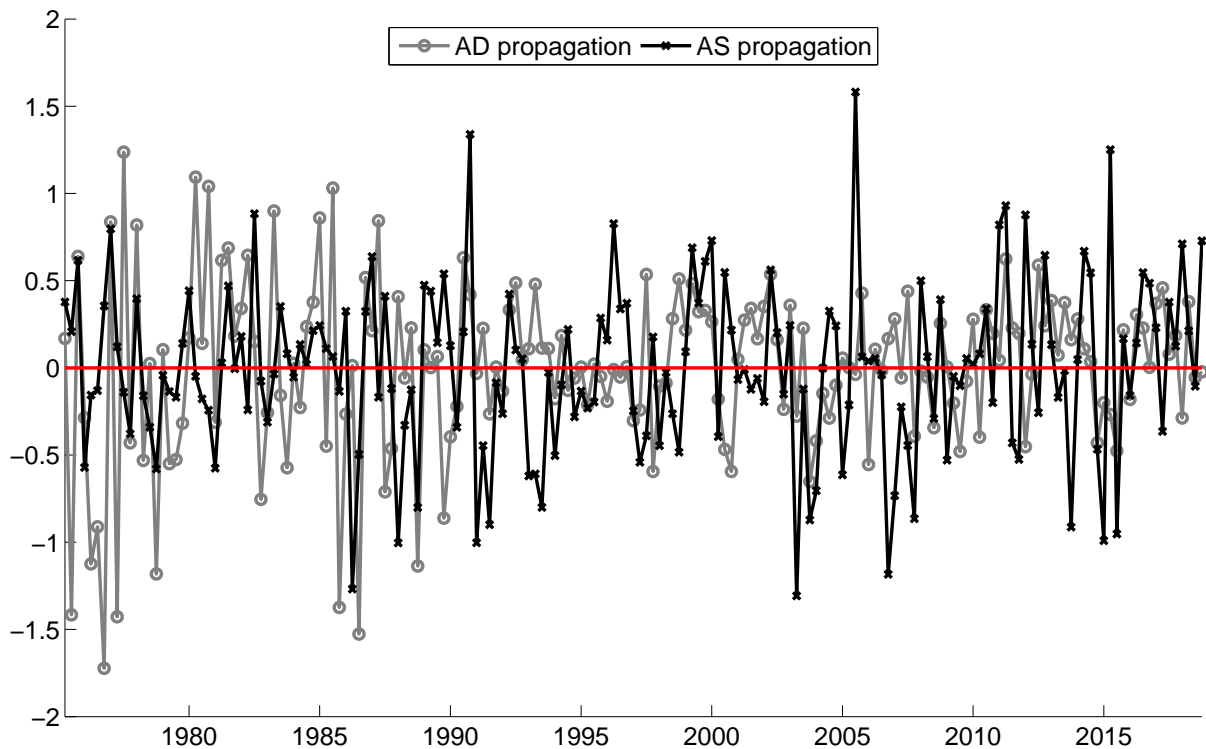
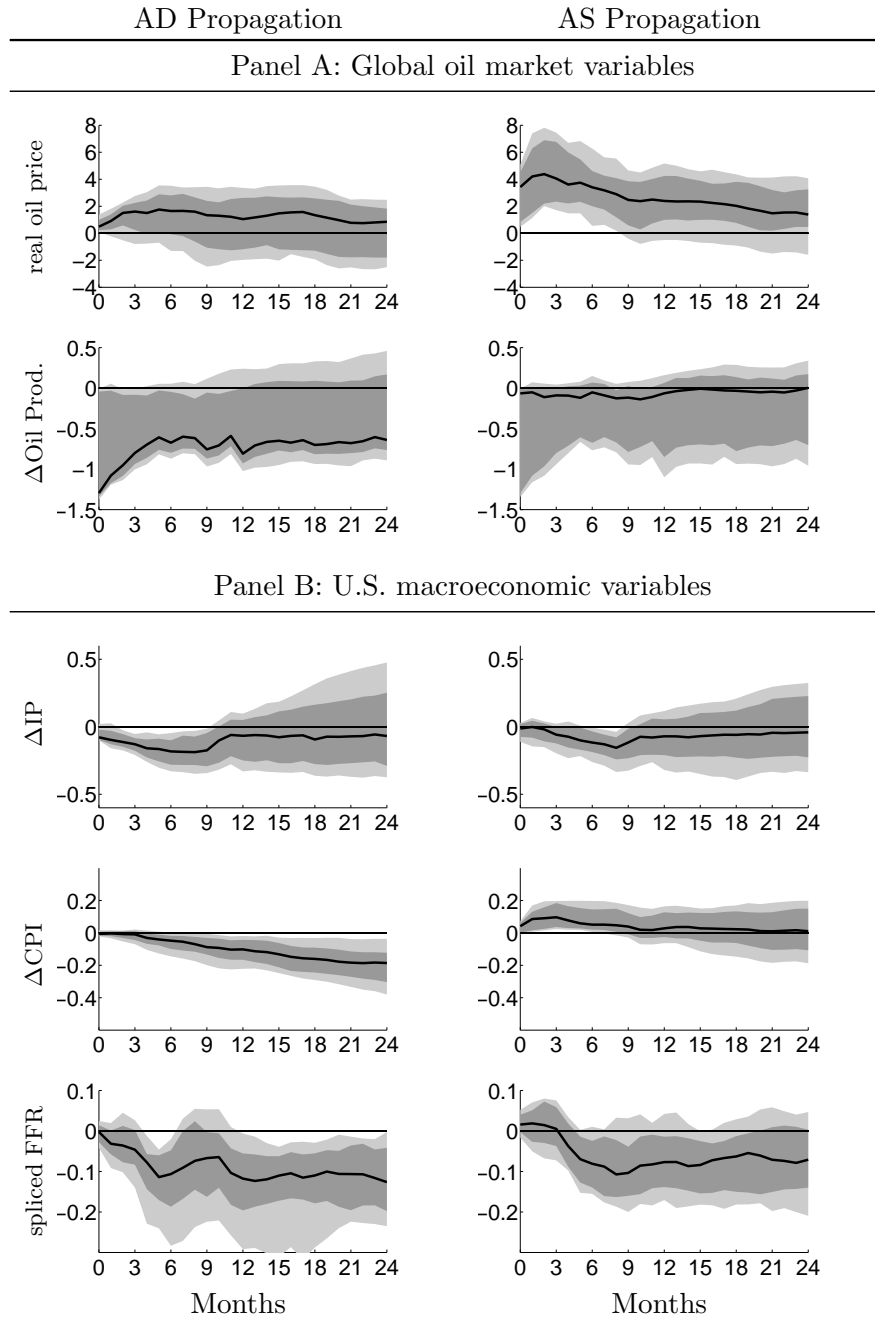


Figure 7: Impulse responses to oil supply shocks propagated through AD and AS: Single VAR model



Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a mutually orthogonal shock based on a single structural VAR model. For the sign restrictions, see Table 4.

Table 5: Forecast error variance decomposition (FEVD): Single VAR model (in percent)

	h	AD propagation	AS propagation
Panel A: global oil market variables			
real oil price	0	0.67 (0.14, 2.61)	33.41 (1.86, 61.19)
	6	3.14 (0.77, 5.63)	16.87 (6.35, 45.34)
	12	4.28 (1.20, 7.84)	14.58 (6.85, 34.48)
	24	4.87 (1.72, 10.00)	14.58 (5.21, 23.96)
Δ oil production	0	99.67 (0.11, 99.93)	0.26 (0.01, 99.37)
	6	88.21 (1.14, 91.16)	1.59 (0.86, 89.68)
	12	81.23 (2.58, 84.19)	2.66 (1.62, 81.99)
	24	78.05 (3.37, 81.21)	3.19 (1.90, 79.04)
real econ. activity	0	0.10 (0.01, 85.24)	32.40 (0.11, 62.38)
	6	0.86 (0.40, 62.88)	14.15 (0.56, 38.03)
	12	1.32 (0.59, 58.90)	13.91 (1.34, 40.36)
	24	1.78 (0.91, 57.66)	15.26 (1.58, 44.59)
Panel B: U.S. macroeconomic variables			
Δ IP	0	1.87 (0.28, 2.79)	0.14 (0.04, 1.58)
	6	2.95 (1.41, 4.24)	1.95 (0.90, 3.56)
	12	4.76 (3.36, 6.34)	3.88 (2.45, 5.52)
	24	4.82 (3.83, 7.33)	4.34 (2.98, 5.79)
Δ CPI	0	0.34 (0.04, 0.81)	4.78 (0.15, 10.08)
	6	1.88 (0.91, 3.66)	7.52 (1.50, 14.85)
	12	2.79 (1.74, 5.01)	8.72 (2.13, 14.13)
	24	4.08 (2.46, 6.61)	8.54 (2.98, 13.57)
spliced FFR	0	0.10 (0.01, 0.52)	0.17 (0.02, 0.80)
	6	1.71 (0.73, 4.15)	1.21 (0.40, 1.99)
	12	2.80 (0.98, 4.83)	2.15 (0.98, 4.50)
	24	4.01 (1.45, 8.34)	2.66 (1.04, 5.22)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a mutually orthogonal shock based on a single structural VAR model. For the sign restrictions, see Table 4.

Online Appendix Further Figures and Tables

Figure A.1: Historical decomposition of spliced FFR: Cumulative effects of oil supply shocks propagated through AD, AS or no channel

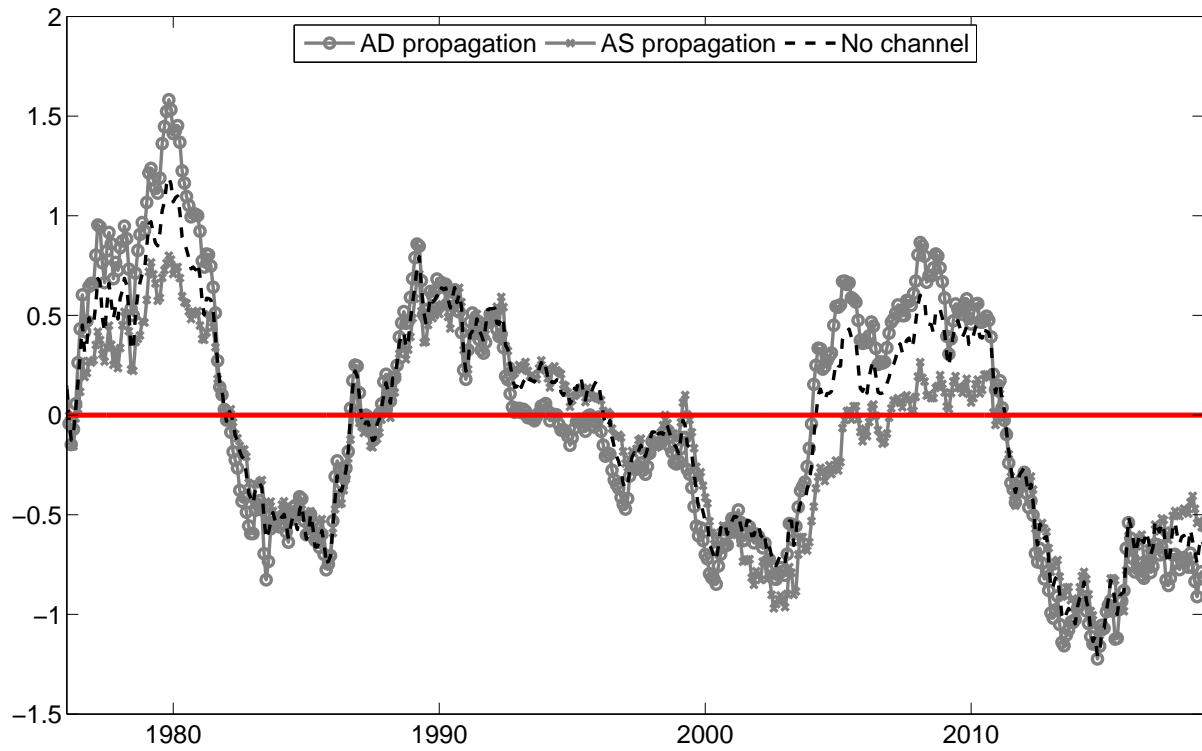
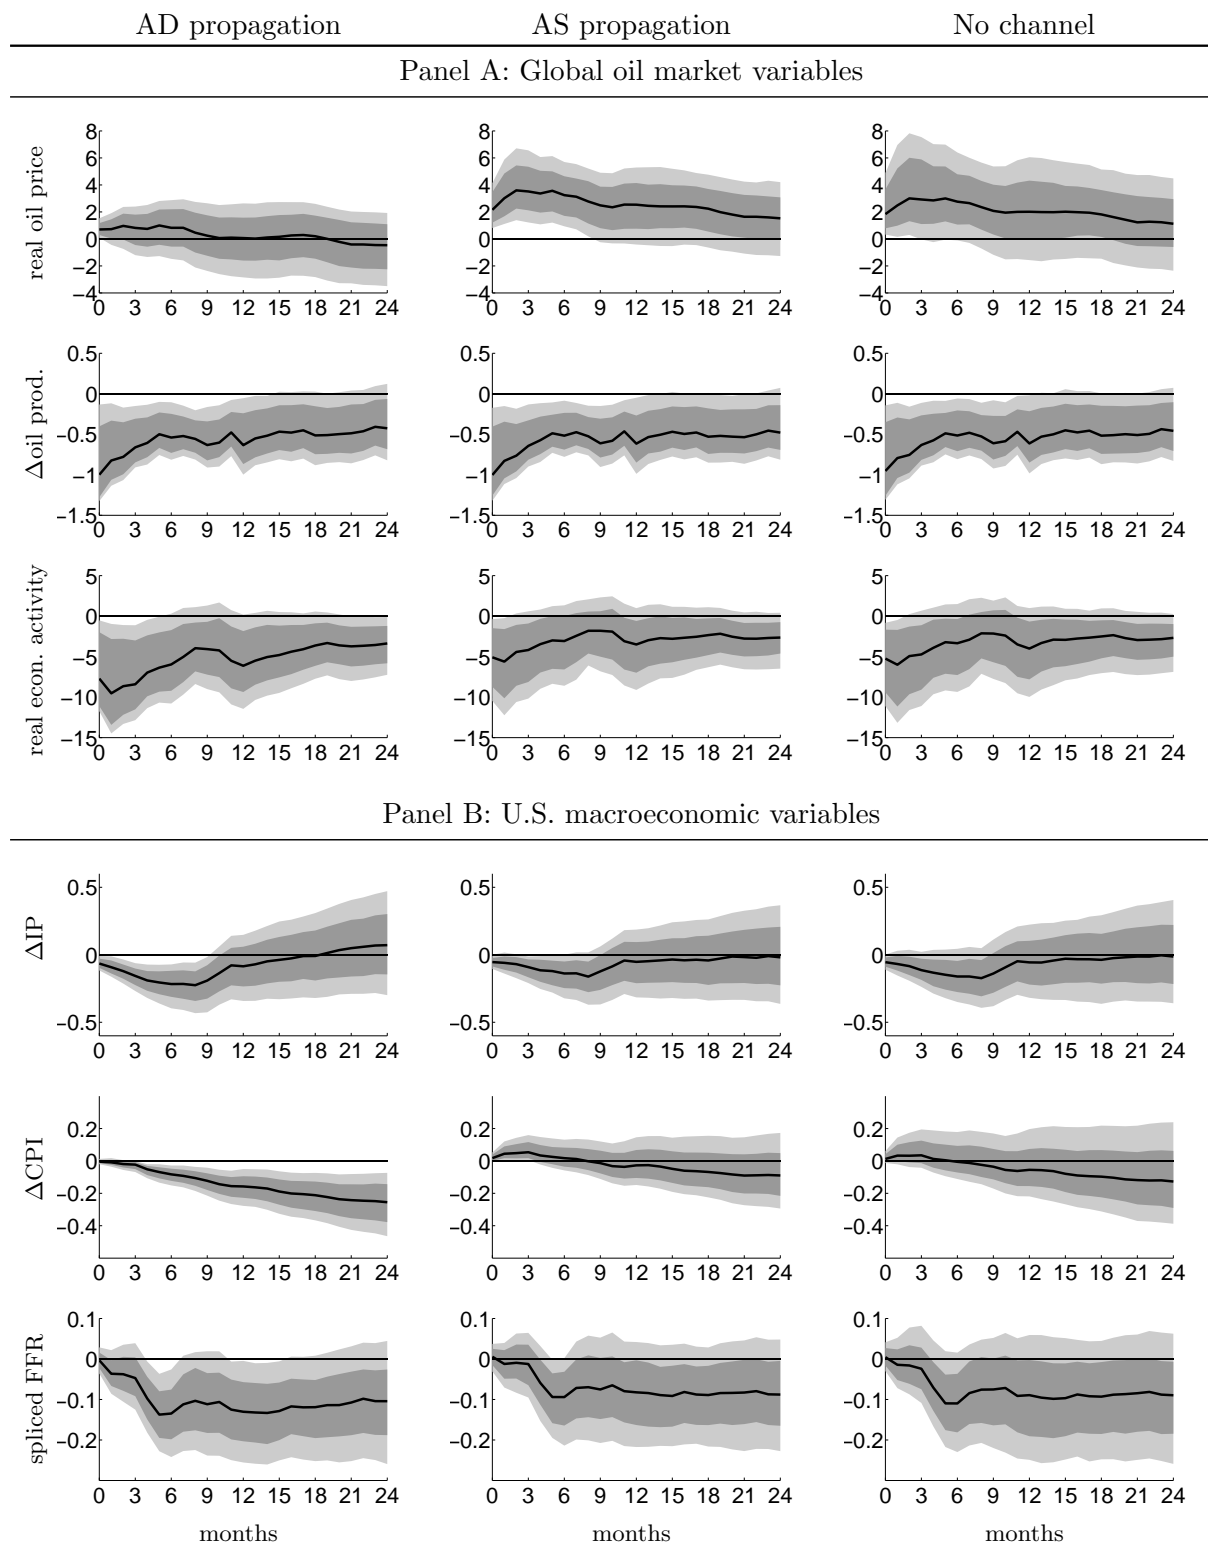


Figure A.2: Impulse responses to oil supply shocks propagated through AD, AS or no channel: Sign restrictions on $h = 3$



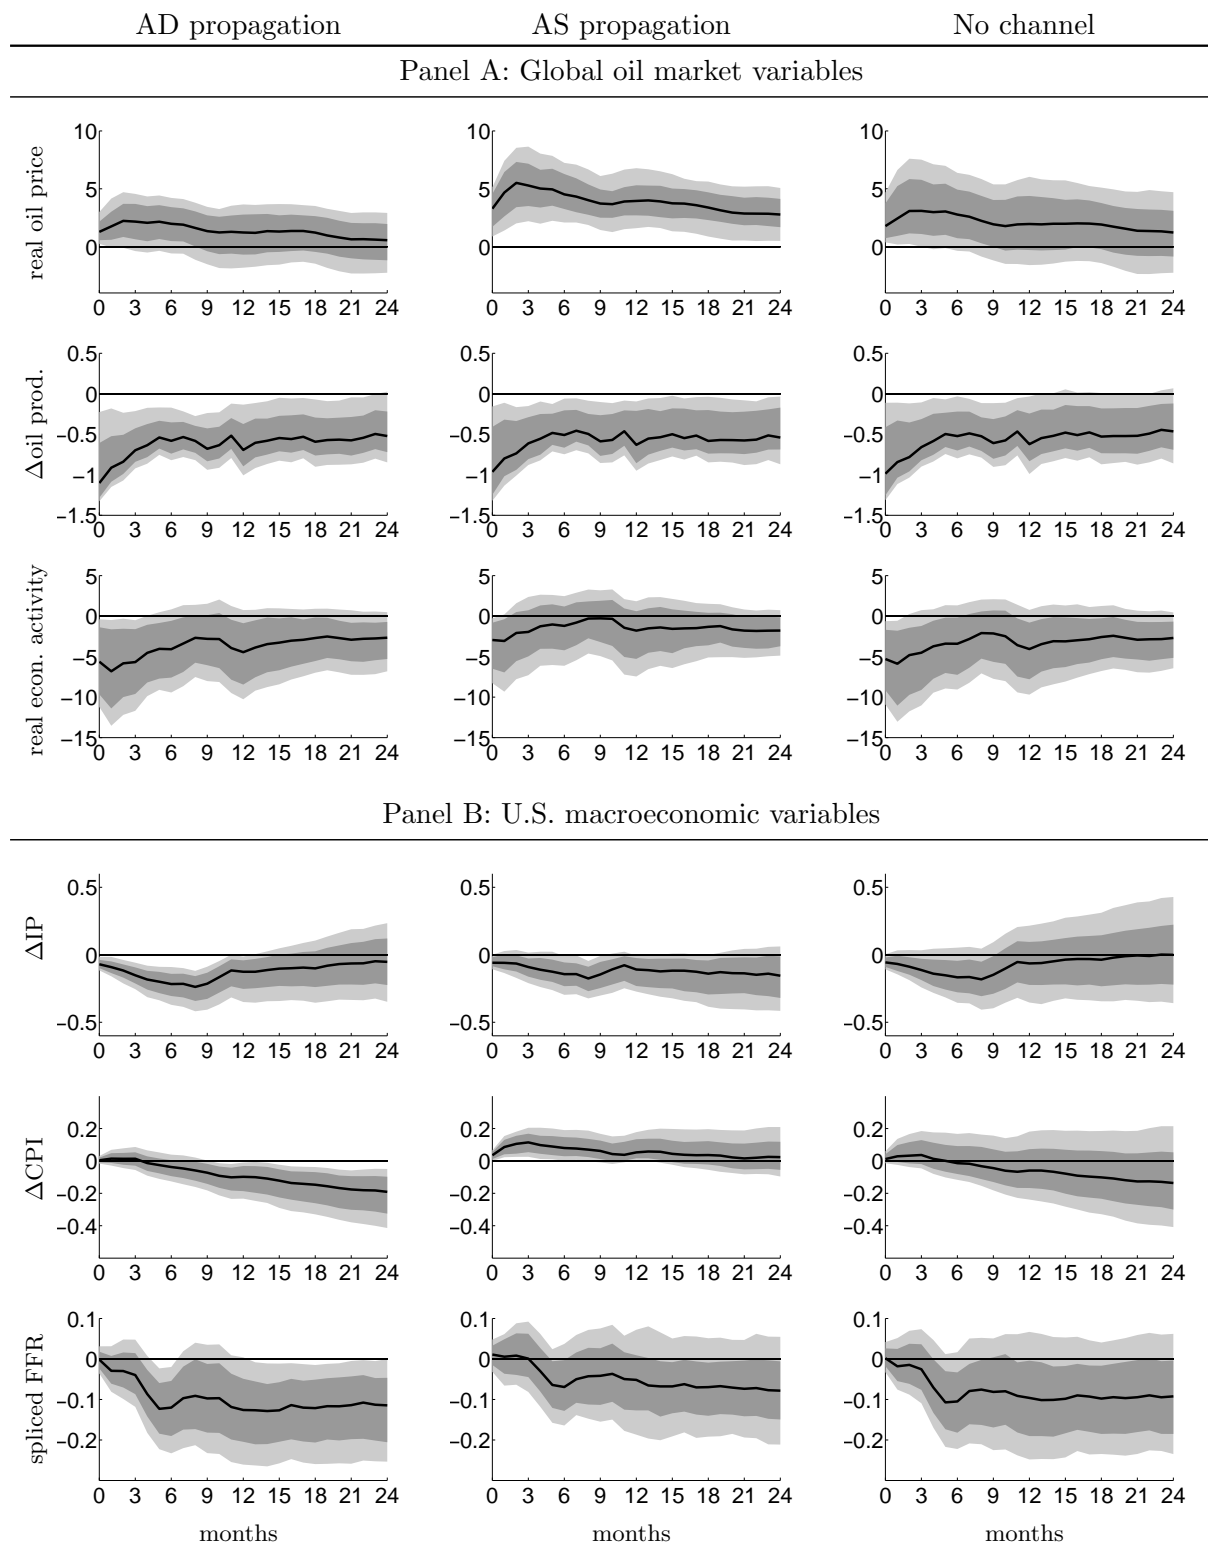
Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model.

Table A.1: Forecast error variance decomposition (FEVD): Sign restrictions on $h = 3$ (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	1.39 (0.24, 3.92)	13.40 (4.13, 33.75)	9.64 (1.81, 38.27)
	6	1.44 (0.31, 4.20)	13.33 (5.39, 29.06)	9.63 (2.01, 34.89)
	12	1.90 (0.65, 5.47)	13.09 (4.54, 27.38)	9.23 (2.59, 31.52)
	24	2.42 (0.84, 6.96)	11.36 (3.85, 25.44)	8.35 (2.44, 27.00)
Δ oil production	0	58.80 (9.18, 96.26)	58.27 (9.90, 93.05)	52.02 (6.89, 92.52)
	6	52.01 (9.17, 86.18)	52.42 (10.64, 83.11)	47.57 (7.52, 82.54)
	12	48.02 (9.53, 78.90)	48.16 (9.85, 76.53)	44.09 (7.88, 75.67)
	24	46.30 (9.62, 75.77)	46.60 (10.13, 73.68)	42.57 (7.96, 72.60)
real econ. activity	0	39.79 (2.64, 83.51)	17.48 (1.48, 53.02)	18.58 (1.95, 60.57)
	6	34.21 (3.43, 67.49)	10.25 (1.19, 35.27)	11.91 (1.75, 44.25)
	12	30.20 (3.28, 63.04)	9.17 (1.51, 33.92)	10.80 (1.98, 41.09)
	24	29.15 (3.35, 61.54)	9.72 (1.90, 36.03)	11.49 (2.50, 42.93)
Panel B: U.S. macroeconomic variables				
Δ IP	0	1.28 (0.25, 2.83)	0.88 (0.18, 2.22)	0.91 (0.15, 2.38)
	6	3.09 (1.74, 4.81)	2.03 (1.14, 3.50)	2.41 (1.21, 4.07)
	12	5.45 (3.77, 7.43)	4.16 (2.64, 5.79)	4.42 (2.86, 6.48)
	24	6.19 (4.40, 8.48)	4.55 (3.13, 6.53)	4.94 (3.40, 7.22)
Δ CPI	0	0.15 (0.01, 0.66)	0.82 (0.06, 4.17)	0.56 (0.05, 4.57)
	6	3.09 (1.62, 5.07)	3.84 (1.98, 8.45)	3.95 (1.73, 9.80)
	12	4.77 (2.72, 7.33)	5.35 (2.89, 9.59)	5.58 (3.04, 10.87)
	24	5.77 (3.55, 9.01)	6.03 (3.58, 9.98)	6.38 (3.72, 11.60)
spliced FFR	0	0.09 (0.01, 0.41)	0.11 (0.01, 0.45)	0.13 (0.01, 0.48)
	6	2.58 (0.87, 5.38)	1.22 (0.42, 3.26)	1.67 (0.54, 4.12)
	12	3.73 (1.26, 8.25)	1.77 (0.59, 5.17)	2.39 (0.72, 6.44)
	24	4.87 (1.72, 10.71)	2.87 (0.76, 7.20)	3.32 (0.90, 9.77)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model.

Figure A.3: Impulse responses to oil supply shocks propagated through AD, AS or no channel: Sign restrictions on $h = 12$



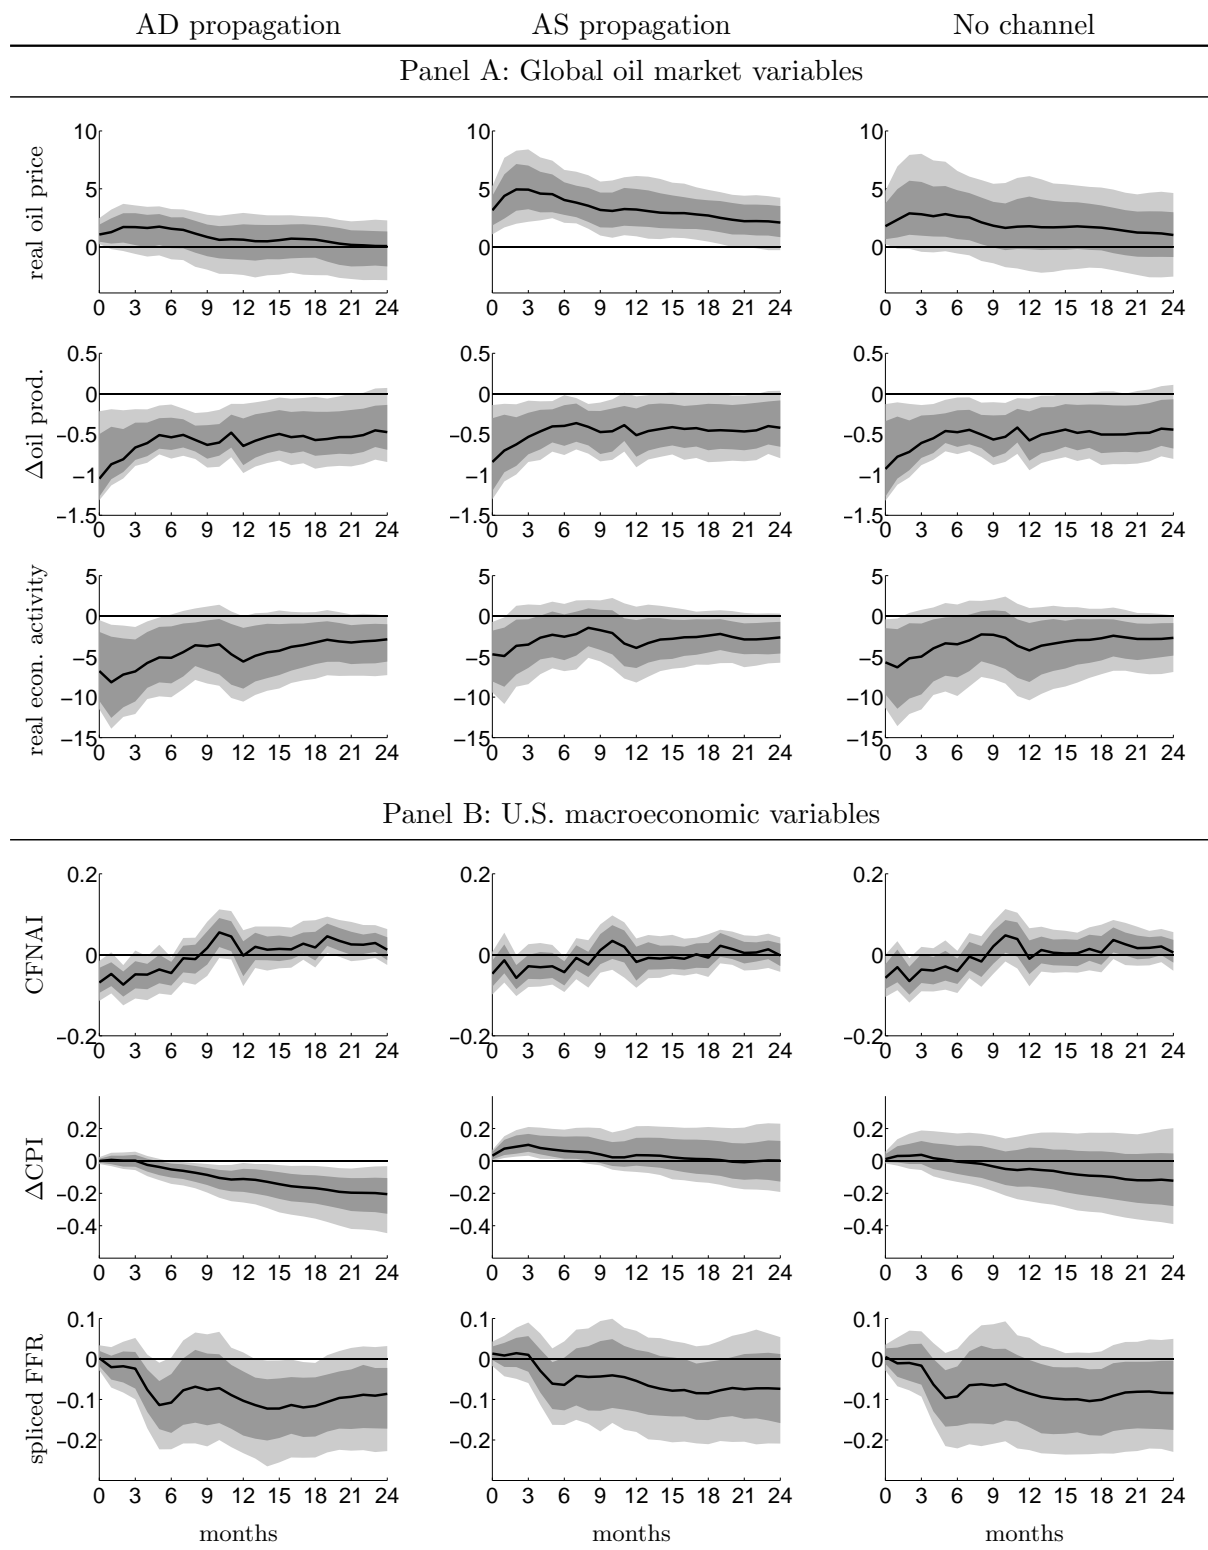
Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model.

Table A.2: Forecast error variance decomposition (FEVD): Sign restrictions on $h = 12$ (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	4.66 (0.89, 13.05)	30.76 (8.78, 58.68)	9.40 (1.52, 39.23)
	6	5.25 (0.95, 13.34)	30.45 (13.27, 51.64)	10.00 (1.55, 33.81)
	12	4.82 (1.41, 13.35)	28.45 (12.52, 45.19)	9.17 (1.73, 30.86)
	24	4.30 (1.53, 13.05)	25.93 (12.25, 42.65)	8.73 (2.10, 29.32)
Δ oil production	0	70.88 (21.38, 97.07)	55.73 (10.01, 90.43)	55.81 (9.87, 94.39)
	6	63.03 (19.68, 87.01)	49.89 (10.05, 80.86)	50.58 (9.71, 83.25)
	12	57.16 (18.44, 79.39)	46.58 (9.26, 73.94)	46.09 (9.93, 76.26)
	24	55.08 (18.12, 76.29)	43.85 (9.24, 71.15)	44.35 (10.27, 73.34)
real econ. activity	0	21.36 (1.29, 62.50)	5.91 (0.46, 28.06)	19.01 (1.91, 57.91)
	6	15.98 (1.24, 48.57)	2.92 (0.67, 15.06)	12.19 (1.61, 41.27)
	12	14.60 (1.44, 45.66)	3.47 (1.04, 14.74)	11.38 (1.79, 38.75)
	24	15.27 (1.70, 45.25)	4.01 (1.32, 16.91)	12.39 (1.94, 40.01)
Panel B: U.S. macroeconomic variables				
Δ IP	0	1.49 (0.45, 2.95)	1.04 (0.10, 2.17)	0.97 (0.14, 2.44)
	6	3.05 (1.73, 4.70)	2.13 (1.23, 3.44)	2.44 (1.27, 4.06)
	12	4.88 (3.25, 6.79)	3.78 (2.46, 5.36)	4.48 (2.88, 6.54)
	24	5.31 (3.82, 7.48)	4.19 (3.10, 5.75)	4.93 (3.32, 7.11)
Δ CPI	0	0.26 (0.02, 1.30)	3.32 (0.37, 8.33)	0.59 (0.05, 4.81)
	6	3.16 (1.73, 5.33)	8.28 (2.83, 15.96)	3.98 (1.81, 9.84)
	12	4.55 (2.78, 7.13)	8.24 (3.38, 15.18)	5.35 (3.06, 10.66)
	24	5.55 (3.54, 8.78)	8.53 (3.89, 14.31)	6.40 (3.73, 11.21)
spliced FFR	0	0.10 (0.01, 0.40)	0.18 (0.02, 0.63)	0.11 (0.01, 0.50)
	6	1.98 (0.57, 4.65)	0.93 (0.32, 1.91)	1.65 (0.54, 4.18)
	12	3.11 (0.79, 7.59)	1.28 (0.44, 3.06)	2.33 (0.72, 6.71)
	24	4.55 (1.35, 10.85)	1.80 (0.71, 4.71)	3.49 (0.92, 9.10)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model.

Figure A.4: Impulse responses to oil supply shocks propagated through AD, AS or no channel: Specification with CFNAI



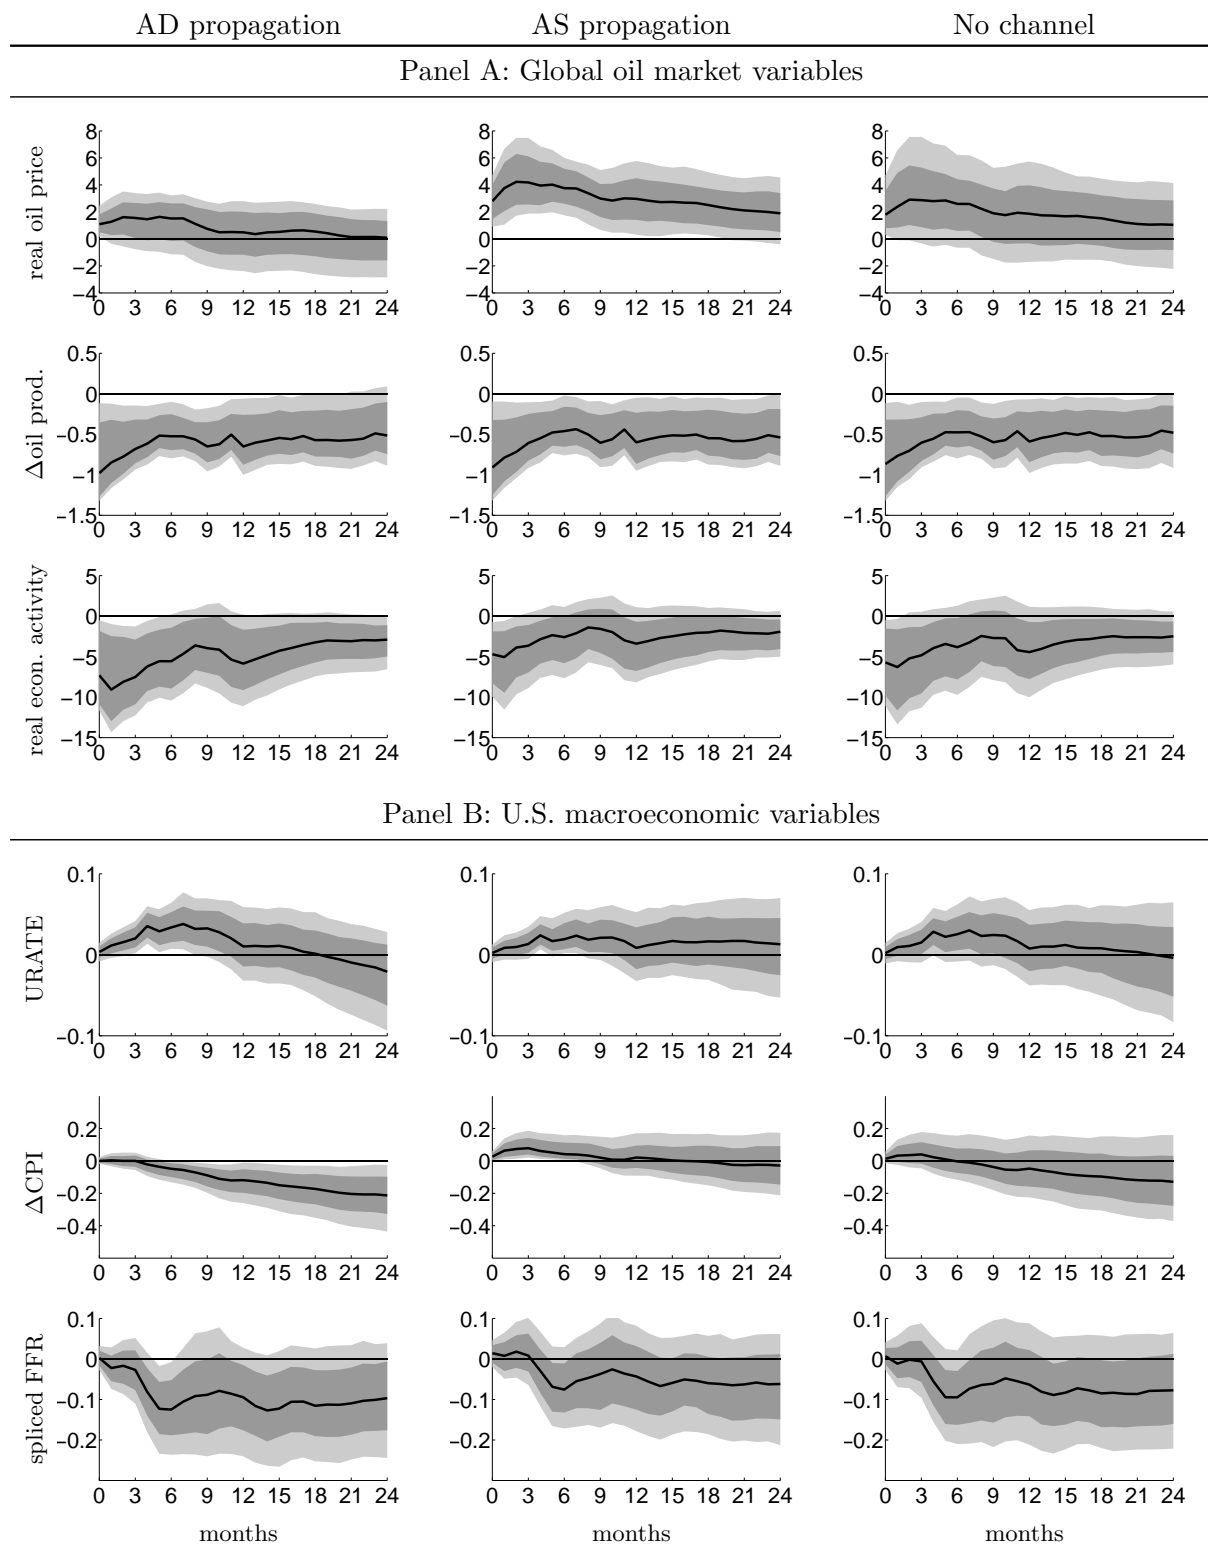
Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model.

Table A.3: Forecast error variance decomposition (FEVD): Specification with CFNAI (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	3.22 (0.51, 10.21)	29.01 (9.53, 55.41)	8.68 (1.22, 39.33)
	6	3.19 (0.52, 8.72)	24.03 (11.19, 47.97)	8.32 (1.59, 31.91)
	12	3.29 (1.01, 8.20)	21.09 (10.09, 42.58)	8.54 (2.09, 28.39)
	24	3.17 (1.17, 8.30)	17.80 (8.19, 36.44)	7.96 (2.03, 26.38)
Δ oil production	0	64.43 (14.80, 95.12)	41.31 (5.73, 84.13)	49.34 (6.31, 93.45)
	6	57.90 (14.39, 84.27)	37.60 (6.43, 75.19)	44.16 (6.78, 84.03)
	12	52.40 (13.29, 77.86)	34.86 (6.71, 68.95)	41.06 (7.31, 76.53)
	24	51.09 (13.45, 74.87)	33.66 (6.68, 66.65)	39.60 (7.57, 73.35)
real econ. activity	0	30.95 (2.56, 72.18)	14.60 (2.12, 42.82)	21.52 (1.42, 62.93)
	6	24.44 (3.16, 57.33)	7.70 (1.38, 25.06)	12.37 (1.38, 45.39)
	12	22.31 (2.84, 54.36)	7.37 (1.61, 23.42)	11.88 (1.61, 44.29)
	24	22.95 (3.05, 54.26)	9.13 (1.78, 27.18)	12.60 (1.85, 44.36)
Panel B: U.S. macroeconomic variables				
CFNAI	0	1.28 (0.29, 2.41)	0.62 (0.08, 1.66)	0.87 (0.17, 2.02)
	6	3.61 (1.99, 6.00)	2.22 (1.14, 3.87)	2.83 (1.32, 5.09)
	12	4.77 (3.20, 7.25)	3.22 (1.96, 4.95)	4.15 (2.34, 6.36)
	24	5.88 (3.60, 9.74)	3.80 (2.43, 6.00)	4.93 (2.85, 7.83)
Δ CPI	0	0.20 (0.02, 0.88)	2.76 (0.35, 8.11)	0.58 (0.04, 4.64)
	6	2.99 (1.68, 5.12)	6.76 (2.76, 14.82)	4.05 (1.68, 9.48)
	12	4.25 (2.61, 7.24)	7.28 (3.53, 14.32)	5.62 (2.80, 10.40)
	24	5.12 (3.20, 8.71)	7.75 (3.93, 13.54)	6.30 (3.49, 11.24)
spliced FFR	0	0.09 (0.01, 0.39)	0.13 (0.01, 0.53)	0.10 (0.01, 0.46)
	6	1.81 (0.60, 4.68)	0.90 (0.29, 2.34)	1.45 (0.48, 4.20)
	12	2.54 (0.80, 7.10)	1.30 (0.51, 3.99)	2.34 (0.73, 6.48)
	24	4.08 (1.25, 9.70)	2.08 (0.59, 5.74)	3.48 (0.92, 8.18)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model.

Figure A.5: Impulse responses to oil supply shocks propagated through AD, AS or no channel: Specification with unemployment rate



Notes: The solid lines represent pointwise median responses. The dark- and light-shaded areas represent pointwise 16th and 84th and 5th and 95th percentiles, respectively, of the posterior distributions of impulse response functions. Each column represents a different structural VAR model.

Table A.4: Forecast error variance decomposition (FEVD): Specification with unemployment rate (in percent)

	h	AD propagation	AS propagation	No channel
Panel A: global oil market variables				
real oil price	0	3.35 (0.72, 9.33)	22.52 (6.18, 45.94)	9.22 (1.81, 36.50)
	6	3.20 (0.61, 7.73)	18.42 (7.81, 37.80)	9.16 (1.47, 29.55)
	12	3.38 (1.04, 8.34)	17.86 (8.45, 34.38)	9.48 (1.98, 27.43)
	24	3.43 (1.23, 8.64)	15.98 (6.64, 31.51)	8.20 (2.05, 23.34)
Δ oil production	0	56.86 (7.54, 96.78)	48.10 (6.26, 92.63)	44.53 (5.95, 93.98)
	6	51.28 (7.50, 86.58)	43.33 (6.94, 83.37)	41.08 (6.63, 84.83)
	12	47.34 (7.63, 78.70)	40.28 (7.45, 76.54)	37.00 (7.13, 76.74)
	24	45.10 (7.91, 75.20)	38.21 (7.74, 73.18)	35.74 (7.39, 73.45)
real econ. activity	0	35.97 (2.24, 78.88)	14.94 (2.43, 43.96)	22.09 (1.61, 64.07)
	6	29.70 (2.99, 60.57)	7.52 (1.32, 27.07)	13.32 (1.41, 48.37)
	12	26.93 (3.03, 56.60)	6.99 (1.77, 25.54)	12.74 (1.52, 44.79)
	24	27.31 (3.35, 58.20)	8.02 (1.85, 27.66)	13.65 (2.04, 43.99)
Panel B: U.S. macroeconomic variables				
URATE	0	0.13 (0.01, 0.58)	0.11 (0.01, 0.44)	0.14 (0.01, 0.56)
	6	2.21 (0.73, 4.90)	0.87 (0.29, 2.63)	1.43 (0.40, 3.87)
	12	2.29 (0.74, 5.77)	1.10 (0.43, 3.19)	1.62 (0.48, 5.03)
	24	2.63 (0.99, 5.85)	1.72 (0.58, 4.64)	2.37 (0.70, 5.97)
Δ CPI	0	0.19 (0.01, 0.85)	1.83 (0.14, 5.90)	0.63 (0.05, 4.15)
	6	2.76 (1.33, 4.92)	4.82 (1.80, 11.43)	3.76 (1.70, 8.75)
	12	4.44 (2.43, 7.20)	5.91 (2.69, 11.96)	5.36 (2.71, 10.26)
	24	5.45 (3.06, 8.49)	6.43 (3.23, 11.41)	6.14 (3.41, 11.13)
spliced FFR	0	0.08 (0.01, 0.33)	0.17 (0.02, 0.61)	0.11 (0.01, 0.47)
	6	2.11 (0.55, 4.79)	1.02 (0.38, 2.36)	1.52 (0.47, 3.57)
	12	2.99 (0.88, 7.95)	1.38 (0.53, 3.82)	2.02 (0.72, 5.56)
	24	3.99 (1.11, 9.90)	1.73 (0.62, 5.10)	2.77 (0.87, 7.51)

Notes: Median contributions to FEVD across all admissible candidate models for a forecast horizon of h months. 16th and the 84th percentiles of the posterior distributions of FEVD contributions are reported in parentheses. Each column represents a different structural VAR model.