How to Construct Monthly VAR Proxies Based on Daily Futures Market Surprises

Lutz Kilian

Working Paper 2310 **July 2023**

Research Department

https://doi.org/10.24149/wp2310

How to Construct Monthly VAR Proxies Based on Daily Futures Market Surprises*

Lutz Kilian†

July 23, 2023

Abstract

It is common in applied work to estimate responses of macroeconomic aggregates to news shocks derived from surprise changes in daily futures prices around the date of policy announcements. This requires mapping the daily surprises into a monthly shock that may be used as an external instrument in a monthly VAR model or local projection. The standard approach has been to sum these daily surprises over the course of a given month when constructing the monthly proxy variable, ignoring the accounting relationship between daily and average monthly price data. In this paper, I provide a new approach to constructing monthly proxies from daily surprises that takes account of this link and revisit the question of how to use OPEC announcements to identify news shocks in VAR models of the global oil market. The proposed approach calls into question the interpretation of the identified shock as oil supply news and implies quantitatively and qualitatively different estimates of the macroeconomic impact of OPEC announcements.

JEL codes: C36, C51, E31, E32, E44, Q43

Keywords: Proxy VAR, IV, shock aggregation, time aggregation, identification, OPEC, supply news, storage demand, oil futures, oil price expectations.

---

*The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. I thank Miguel Acosta, Lukas Boer, Alexander Chudik, Jon Faust, Helmut Lütkepohl, Christian Matthes, and Xiaoqing Zhou for helpful comments. The author has no conflict of interest to disclose.

†Correspondence to: Lutz Kilian, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA and CEPR. E-mail: lkilian2019@gmail.com.

1. Introduction

There is a large empirical literature on the sources of oil price fluctuations and the link between global oil markets and the economy (e.g., Kilian 2008a; Kilian and Zhou 2023). Of particular interest has been the causal effect of shocks to oil price expectations on macroeconomic outcomes. Monthly structural vector autoregressive (SVAR) models based on sign restrictions indicate that positive shocks to oil price expectations – caused by a shortfall of expected oil supply relative to expected oil demand raising the expected price of oil– are followed by a large and immediate increase in the real price of oil, an increase in global oil inventories, a modest decline in oil production with a delay and changes in global real activity that are indistinguishable from zero (e.g., Kilian and Murphy 2014; Zhou 2020). This evidence supports the notion that changes in expectations about future oil supply shortfalls can have powerful effects, even if current oil production does not decline.

Based on an alternative SVAR methodology, Känzig (2021) recently concluded that shocks to oil price expectations cause a large and immediate increase in the real price of oil, a gradual fall in global oil production, a delayed substantial decline in global real activity, and rising global oil inventories. Moreover, in the U.S. economy, real activity falls, inflation and inflation expectations rise, and the dollar depreciates. A key difference from earlier work is that Känzig (2021) argues that these responses can be attributed to adverse oil supply news alone, as opposed to shifts in oil price expectations driven by oil market news more generally.

His analysis is based on a proxy VAR model of the global oil market that also includes selected U.S. macroeconomic aggregates. This model relies on news revealed by OPEC announcements as an instrument for unpredictable variation in the monthly average real price of oil.[1] The premise is that daily surprises in oil futures prices around OPEC

[1] OPEC refers to the Organization of Petroleum Exporting Countries, which currently includes Algeria, Angola, Equatorial Guinea, Gabon, Iran, Iraq, Kuwait, Libya, Nigeria, the Republic of the Congo, Saudi Arabia, the United Arab Emirates and Venezuela as its members.

announcement dates can be aggregated to monthly frequency by summing the daily surprises within each month. The resulting monthly proxy is used as an external instrument to identify exogenous variation in the real oil price shock implied by the reduced-form VAR model. This exogenous variation is then used to estimate the responses of global oil market variables and U.S. macroeconomic aggregates.

In this paper, I examine the rationale of this approach, which has been employed in a number of recent studies, and show that it ignores the accounting relationship between daily price data and average monthly price data.[2] I explain how this accounting relationship may be incorporated in the construction of the monthly proxy, and I examine the extent to which this changes the response estimates implied by the proxy VAR model. Although the paper focuses on the question of how to recover monthly news shocks from daily OPEC announcements, it should be noted that the same methodological approach is applicable in many other contexts.[3]

The paper departs from the existing literature in two dimensions. First, I show that the first six years of daily OPEC surprises constructed by Känzig (2021) rely on unsuitable data that must be discarded. Dropping these observations renders the VAR results sensitive to the estimation period, changes the response estimates, and lowers the explanatory power of the proxy for real oil price shocks. Whereas Känzig stresses that his instrument is strong with robust F-statistics safely above the threshold of 10, after discarding the first six years of futures prices this instrument appears weak, invalidating the use of conventional methods of inference. I address this problem by replacing the conventional impulse response confidence intervals reported by the weak-identification robust intervals developed in Montiel Olea et al. (2021) that remain valid when the instrument is strong.

[2] Examples of studies employing this approach include Degasperi (2021), Bruns (2021), Bruns and Lütkepohl (2023) and Gagliardone and Gertler (2023).
[3] For example, in the literature on monetary policy surprises, daily surprises in the fed funds target rate around FOMC dates are linked to the monthly average of the daily fed funds rate in the VAR model (e.g., Kuttner 2001; Faust et al. 2004). Similar techniques are also being used in the literature on the macroeconomic effects of climate change policies (e.g., Känzig 2023).

Second, I propose a more natural way of mapping daily surprises into monthly average shocks that implies a substantially different proxy variable than that used in Känzig (2021). I employ this new proxy variable both in the baseline model of Känzig (2021) and in the workhorse model of the global oil market originally proposed by Kilian and Murphy (2014). I show that the news shocks captured by the proxy VAR model do not in general represent oil supply news, but at least in part, if not entirely, capture oil demand news, echoing concerns expressed in Degasperi (2021). This point is revealed by the impulse responses generated by these shocks. In particular, I provide evidence that monthly average surprises derived from OPEC announcements tend to capture a linear combination of storage demand shocks driven by oil supply news and of flow demand shocks, with the latter dominating after 1989. This result complicates the interpretation of the response estimates implied by the proxy VAR model and calls into question the narrative that the proxy VAR model identifies the macroeconomic effect of oil supply news.

My analysis relates to a number of other strands of the literature. First, it contributes to the literature on the impact of news on oil prices (e.g. Demirer and Kutan 2010; Kilian and Vega 2011; Kilian and Hicks 2013) and the literature on modeling exogenous shifts in oil price expectations (e.g., Alquist and Kilian 2010; Kilian and Murphy 2014; Kilian and Lee 2014; Knittel and Pindyck 2016; Känzig 2021; Degasperi 2021; Bruns 2021; Cross et al. 2022; Bruns and Lütkepohl 2023; Gagliardone and Gertler 2023).[4]

Second, it complements a recent debate about whether OPEC surprises represent supply news or are contaminated by demand news (e.g., Känzig 2021; Degasperi 2021; Kilian and Zhou 2023; Bruns and Lütkepohl 2023). I provide new evidence based on impulse response analysis that OPEC surprises, especially in recent data, represent mainly oil demand

[4] The latter literature in turn built on earlier work on the role of precautionary demand shocks in oil markets that are driven by geopolitical or macroeconomic uncertainty (e.g., Kilian 2009; Alquist and Kilian 2010).

news rather than oil supply news.

Third, this paper relates to the literature on generating confidence sets that are robust to weak identification. A large literature has established that existing external instruments for the real price of oil tend to be weak (e.g., Kilian 2008a,b; Montiel Olea et al. 2021; Kilian 2022a). Statistical tests suggest that the OPEC announcement proxy is no exception, once we correct for the problems with the underlying oil futures data. Throughout the paper, I therefore employ Anderson-Rubin confidence sets developed by Montiel Olea et al. (2021) that remain valid whether the proxy in question is a weak instrument or a strong instrument.

Finally, and most importantly, the paper addresses the question of how to map shocks observed at higher frequency into a lower data frequency. This question is relevant more broadly for other applications of proxy VAR models and IV-local projection models that use high-frequency data in constructing the proxy, for example, when estimating monetary policy shocks (e.g., Gertler and Karadi 2015; Stock and Watson 2018; Caldara and Herbst 2019; Wolf 2020; Miranda-Agrippino and Ricco 2021; Amir-Ahmadi, Matthes and Wang 2023; Bauer and Swanson 2023; Gagliardone and Gertler 2023). It is also relevant for alternative SVAR approaches that involve temporally aggregated daily surprises (e.g., Faust et al. 2004; Jarociński and Karadi 2020). My analysis emphasized that the conventional approach to constructing monthly proxies in typical applications is not persuasive. The existence of temporal aggregation bias from constructing monthly proxies as the sum of daily surprises is well understood in the literature. For example, Jacobson, Matthes and Walker (2023) document substantial bias in macroeconomic response estimates implied by conventionally constructed monthly proxies relative to the responses obtained when the frequency of the shock and the outcome variable is aligned.

How to address this problem has been less clear, given that in many applications the only data available are monthly average prices. In related work, Caldara and Herbst (2019, footnote 13) conclude that end-of-month price data, even if they are available, tend to be too noisy to replace monthly average price data. Gertler and Karadi (2015, footnote 11) also recognize that the timing of the daily surprises matters for their effect on the average monthly price. Their approach, however, differs from the proposal in this paper in they difference the average cumulated daily surprises. Moreover, they do not allow for the fact that a daily surprise in the current month may create a monthly surprise in the subsequent month. The approach I outline in this paper, in contrast, addresses these shortcomings. Like its predecessors, my approach falls short of explicitly solving the temporal aggregation problem involved in linking shocks in daily data to monthly data – a task that is known to be very difficult and perhaps intractable – but it provides a simple way of bringing the construction of the monthly proxy more in line with the construction of the VAR variable to be instrumented.

The remainder of the paper is organized as follows. In Section 2, I briefly review the proxy VAR model employed by Känzig (2021). In Section 3, I delve into the construction of the daily surprises and how they are mapped into monthly proxies. Section 4 revisits the findings reported in Känzig (2021), examines their sensitivity to changes in the estimation period and model specification, addresses the weak instrument concern, and sheds new light on the economic interpretation of these results. The concluding remarks are in Section 5.

2. Methodology

It has long been argued that OPEC announcements about its production target and economic outlook, which have been made at irregular intervals since 1982, may affect oil price expectations. Measuring changes in oil futures prices between the trading day preceding the OPEC announcement and the day of the announcement (or the next trading day if there is no trading on the announcement date) helps isolate the impact of the announcement on these prices. As long as the risk premium does not change on the day of the announcement, the change in the oil futures price relative to the preceding day is assumed to capture the revision in oil price expectations associated with the OPEC announcement. Rather than focusing on price changes for one futures contract of a given maturity, Känzig works with the first principal component of the surprises based on WTI crude futures contracts with maturities ranging from one month to one year (PC-IV). However, he observes that similar results would be obtained for specific maturities. He interprets these surprises as OPEC oil supply news.

This type of analysis is by no means new to the oil market literature (e.g., Demirer and Kutan 2010). What is novel about Känzig's analysis is the idea to convert these surprises to a monthly proxy that can be used to identify oil supply news shocks in a structural VAR model of the global oil market, building on an instrumental variable methodology developed by Stock and Watson (2012) and Mertens and Ravn (2013). Specifically, he treats the daily surprise as the proxy observation for the month in which the OPEC announcement occurs (or the sum of the daily surprises, if there is more than one OPEC announcement in that month). This proxy time series is then used as an instrument for the reduced-form residual of the real price of oil in a VAR model of the global oil market.[5]

The baseline VAR(12) model in Känzig (2021) includes the real price of oil, global oil production, global industrial production and global oil inventories, as well as selected U.S. macroeconomic aggregates, all expressed in log levels. While the proxy variable only becomes available starting in 1983.4, estimation of the proxy VAR model does not require the proxy variable to exist for the entire sample of 1975.1 through 2017.12. Either way, the first column of the structural impact multiplier matrix in the structural VAR model may be identified up to scale. Given a suitable normalization of the estimate of this impact response

[5] Further discussion of the econometric foundations of this approach can be found in Stock and Watson (2018, Montiel Olea, Stock and Watson (2021), and Arias, Rubio-Ramirez and Waggoner (2021).

vector, it is straightforward to compute from the estimated reduced-form VAR model all objects of interest such as the structural shock of interest or the responses of the model variables to this shock.

3. The Construction of the Proxy Variable

The construction of the proxy variable involves two steps. First, we measure revisions to oil price expectations on days of OPEC announcements, using the same timing conventions as Känzig (2021). Second, we map these daily surprises into a monthly measure of shocks to oil price expectations, so it can be used as a proxy for the monthly VAR model of the global oil market.[6]

3.1. Issues with measuring OPEC surprises in daily oil futures prices

When constructing the surprise change in oil price expectations on the day of OPEC announcements, Känzig (2021) works with daily oil futures prices for maturities between 1 and 12 months, starting from early 1983. An obvious concern is that before April 1989 trading in oil futures markets was limited to selected dates and maturities.

There are three distinct problems with the pre-1989.4 data. First, one-month WTI futures contracts started trading in March 1983, with longer maturity contracts only gradually being added in later years. For example, the 12-month contract did not start trading until April 17, 1984. This means that many daily observations early in the estimation sample are not available nor are the corresponding surprises on OPEC announcement dates. Känzig (2021) sets these unknown surprises to zero instead of dropping these events. The fact that a surprise cannot be measured based on futures prices, however, does not mean that the surprise is zero. Clearly, nothing prevented observers from adjusting their 12-month

[6] The adjustments we propose differ conceptually and practically from those used in Fast, Swanson and Wright (2004) to account for the definition of the federal funds futures contract as an average over the month (see also Kuttner 2001). The latter complication does not arise in the context of West Texas Intermediate oil futures contracts.

expectations in response to the OPEC announcement in 1983, for example. By replacing this surprise by zero, the data are censored, which calls into question the construction of the leading principal component of surprises across maturities (PC-IV) and any other regression analysis involving these surprises.

Second, even after new futures contracts were introduced, trading for several years was intermittent only, which explains why oil futures prices during this time often remained unchanged for extended periods. This reflects the practice of reporting the last available futures price in the absence of more recent transactions. For example, between April 17, 1984, and September 28, 1984, the 12-month oil futures price remained unchanged at $30.66/barrel. This means that any OPEC announcement over this period is associated with surprises of zero by construction. For example, the change in the 12-month futures prices associated with the July 11, 1984, OPEC announcement is zero, which is mechanically correct, but misleading since no transactions took place on those dates and a market price did not exist. Only starting in April 1989, futures prices start fluctuating on a daily basis.

Third, even when trading took place in the 1980s, as new contracts were introduced, the volumes at these maturities initially tended to be low, undermining the price discovery and making these futures prices potentially unrepresentative (see Alquist and Kilian (2010), Figure 2). These observations suggest that daily surprises around OPEC announcement dates can be constructed only starting in 1989.4 rather than starting in 1983.4 as in Känzig (2021). In Section 4.1, I examine how this affects the response estimates from the proxy VAR model.

3.2. How to construct the monthly proxy

The next step is the construction of a monthly proxy (also referred to as an external instrument) from the daily surprises in the log futures price caused by OPEC announcements. As has been common in the literature on high-frequency monetary policy shocks, Känzig treats the daily surprise as the proxy observation for the month in which the OPEC announcement occurs (or the sum of the daily surprises, if there is more than one OPEC announcement in that month). This proxy is then used as an instrument for the reduced-form shock to the real price of oil in the oil market VAR model.[7]

This approach would make sense, if the price of oil in the VAR model were measured as the percent change in the price of oil from the last day of the preceding month to the last day of the current month. However, in practice, the monthly price of oil in VAR models of the global oil market is invariably defined as the average of the daily prices, reflecting the lack of daily data before 1983. Thus, the VAR price shock is defined as the percent change in the monthly average price relative to the average price in the preceding month. Focusing on the cumulative change in the futures price caused by OPEC announcements effectively postulates that a surprise occurring on the first day of the month has the same effect on the average price of oil in this month as a surprise of the same magnitude occurring on the last day of this month. This premise is clearly counterintuitive.[8]

A more natural approach is to derive the implications of daily OPEC surprises for the surprise change in the monthly average price of oil, as outlined next. Recall that the daily oil futures price expressed in logs evolves according to

$$f_{t+1}^{h} - f_{t}^{h} = s_{t}^{h} + u_{t}\,,$$

where the surprise s_t^h is zero except on announcement dates and u_t is a white noise process that is independent of s_t^h. Consider a given OPEC surprise of magnitude α and suppose that a month contains T trading days. Furthermore, assume without loss of generality that the log oil futures price was zero prior to the OPEC announcement. By construction, a surprise by α

[7] It may seem odd at first to instrument a real price based on a nominal price, but the distinction between nominal and real surprises can be ignored to the extent that inflation is negligible at monthly frequency. For example, adding the log of the nominal price of oil as the first variable to the baseline VAR model (while dropping the log of the CPI to avoid the resulting singularity) and instrumenting for the nominal price of oil instead generates responses that are virtually identical to the baseline model in Känzig (2021).
[8] Cumulating the daily surprises is equivalent (up to scale) to averaging the daily surprises, as sometimes seen in applied work.

raises the daily oil futures price permanently by α on all days, starting with the announcement day. Summing the daily surprises, as is standard in the literature, is equivalent to computing $f_T^h - f_0^h$, which represents the cumulative increase in the daily log futures price caused by the OPEC announcement. In this example, $f_T^h - f_0^h = \alpha$.

There are two concerns with this approach. One concern is that the cumulative change is invariant to when the surprise(s) occur. A shock of α occurring early in the month would have exactly the same cumulative effect on the oil futures price as a shock of α occurring near the end of the month. This feature is counterintuitive when seeking to explain variation in the average price of oil in the VAR model, which clearly depends on the timing of the daily shocks. The other concern is that the implied growth rates based on the daily price is a mismatch for the growth rate of monthly average oil price used in the oil market VAR model, making it awkward to regress the latter on the former in the first-stage IV regression. One would expect this mismatch to undermine the predictive power of the proxy in the first stage of the IV analysis.

A more appealing approach is to focus on the average increase in the daily log futures price caused by the surprises relative to the baseline of zero. This average increase varies with the timing of the daily surprises. For expository purposes, consider a daily surprise of α occurring on day d of the month January. Then the average increase in the daily log futures price during January is given by $\alpha(T-d+1)/T$. For example, a surprise on day 1 of the month January ($d=1$) raises the average log futures price for January by α, whereas a surprise on the last day ($d=T$) only raises the average log futures price by α/T.[9]

This is not the end of the story, however. Given the monthly information set of the

[9] This result applies to the construction of monthly proxies from a given series of daily surprises and should not be confused with the analysis in Kuttner (2001) and Faust et al. (2004) of how to compute daily surprises in the fed funds target rate from changes in daily fed funds futures prices around FOMC announcements. Their analysis is designed to deal with the fact that the fed funds futures contract relates to the average federal funds rate during the contract month rather than the value of this rate on a particular date.

VAR model, if the initial shock occurred in January, as of February, the econometrician would expect the average monthly price to be $\alpha(T-d+1)/T,$ which is below the daily futures price of α that is still prevailing in February, given the surprise in January. This implies a secondary surprise in February of $\alpha-\alpha(T-d+1)/T.$ Thus, unless the original daily surprise occurs on the first day of the month, there will be an additional shock in the following month, whose magnitude depends on the timing of the original daily surprise. By construction, starting in the third month, the expected average price and the average price in the futures market coincide, so there will be no more monthly surprises.[10] Figure 1 illustrates how the timing of the daily surprise affects the magnitude of the monthly surprise. For expository purposes, I compare a daily surprise of magnitude 2 that occurs, respectively, on January 3 or on January 18. In the latter case, the monthly proxy is much smaller in January, but larger in February, than in the former case.

In rare cases there may be a second OPEC announcement in a given month. If that second surprise is of magnitude β and occurs on date $\bar{d} > d,$ the effect on the monthly log oil futures price in January will be $\alpha(T-d+1)/T+(\beta-\alpha)(T-\bar{d}+1)/T,$ the effect in February will be $\alpha+\beta- \alpha(T-d+1)/T-(\beta-\alpha)(T-\bar{d}+1)/T,$ and the effect beyond February will be zero. This situation is illustrated in Figure 2 where it is assumed that $a=1,$ $b=2,$ $d=6,$ and $\bar{d}=10.$

As illustrated in Figure 3, the monthly average proxy for maturity 12, which is only weakly first-order autocorrelated, tends to be quite different from the PC-IV proxy variable constructed in Känzig (2021). Their contemporaneous correlation is only 42%, reflecting differences in the timing and magnitude of these monthly proxies. The difference between the

[10] It is important to stress that it is not the economic agents that are surprised twice in this setting, but it is the econometrician running a VAR model whose information set only includes monthly data, whereas agents in the economy have access to daily data as well.

series can be as large as 9.5 percentage points in absolute terms. Another way of illustrating these differences is to focus on the OPEC announcement on November 27, 2014, which occurred right before the last trading day of that month. Figure 4 shows the fitted values obtained from regressing the reduced-form residual for the real price of oil in Känzig's baseline VAR(12) model specification on either of the two proxies. The plot illustrates that, according to the monthly average proxy, the oil price shock associated with that event effectively occurred in December 2014 rather than in November, as suggested by the PC-IV proxy. There is also a noticeable difference in the magnitude of these price shocks.

3.3. The Importance of Inference that is Robust to Weak Instruments

Känzig (2021) stresses that the PC-IV proxy is a strong instrument, as indicated by a robust F-statistic above 10 in the first-stage IV regression, unlike alternative instruments in the literature such as OPEC supply shocks that have been shown to be weak instruments (e.g., Kilian 2008a,b; Montiel Olea et al. 2021; Kilian 2022a). Evidence presented in Section 4, however, shows that the PC-IV proxy becomes a weak instrument after dropping the futures price data for 1983.4-1989.3 for the reasons discussed in Section 3.1. While this problem may be overcome by replacing the PC-IV proxy with the monthly average proxy, as suggested in Section 3.2, weak instrument problems resurface when changing the VAR specification, as illustrated in Section 4. Given this evidence, throughout this paper, I report CS^{AR} impulse response confidence sets obtained by inverting the Anderson-Rubin (AR) statistic, as proposed by Montiel Olea et al. (2021). The AR confidence set coincides with the standard confidence set when the instrument is strong, but retains its validity when the external instrument is weak.[11]

[11] It should be noted that even when a proxy VAR specification passes the weak instrument test, the coverage accuracy of the conventional confidence interval may be seriously distorted due to pre-testing bias, which is why many econometricians recommend dispensing with pre-tests for weak instruments altogether and applying weak-identification robust asymptotics by default, as we do in Section 4 (e.g., Andrews et al. 2018).

4. Evidence that the Construction of the Proxy Matters for the Response Estimates

Ultimately, it is the response estimates and confidence intervals implied by the proxy VAR model that matter to applied users. Our objective in this section is to illustrate what difference the construction of the monthly average proxy makes compared to simply summing daily shocks. Before we can do so, however, we need to establish a valid baseline that corrects the problems with the daily oil futures price data documented in Section 3.1.

4.1. The corrected PC-IV proxy in the baseline VAR model of Känzig (2021)

We start by replicating the VAR(12) baseline model in Känzig (2021) based on the PC-IV instrument for the estimation period of 1975.1-2017.12 (see Figure 5a). The impact effect of the OPEC shock on the real price of oil is normalized to 10%. The only difference is that we discard the proxy data for 1983.4-1989.3 for the reasons discussed in Section 3.1. An immediate consequence of correcting the proxy in this way is that the robust F-statistic drops well below 10, so we can no longer reject the null that the proxy is a weak instrument at conventional significance levels. Thus, impulse response inference must be conducted based on weak-identification robust confidence intervals.

Figure 5a suggests that bad news is associated with a persistent decline in global oil production and an increase in the real price of oil that that dies out only after three or more years, as storage demand for oil surges and market participants build inventories in anticipation of shortages. The price increase causes a delayed, but persistent decline in world industrial production and U.S. industrial production as well as a short-lived increase in the U.S. inflation rate (which may be inferred from the response of the log CPI by differencing the data). These responses are seemingly consistent with the interpretation of the VAR shock as oil supply news. The estimates obtained in Figure 5a are qualitatively similar to those reported in Figure 3 of Känzig (2021) with one important difference. The decline in world industrial production in response is no longer statistically significant at any horizon at the

10% level, which raises the question of why we see a statistically significant decline in U.S. industrial production, but not in the rest of the world, which tends to be more dependent on imported crude oil than the United States.

There is another even more important difference, however. A key argument in support of the interpretation of the proxy VAR shock as oil supply news in Känzig (2021) is that the response estimates are "qualitatively similar" when restricting the estimation period to coincide with the dates for which the PC-IV proxy is available. Figure 5b shows that this robustness no longer holds, when estimating the model with the corrected PC-IV proxy. Not only is the response of the real price of oil much more transitory and that of the global oil inventories much more muted, but there is no more evidence of a decline in global oil production, as would be required under the oil supply news interpretation. In fact, oil production, if anything, responds positively in the first few months, consistent with producers responding to price incentives, but inconsistent with negative supply news. At the same time, there is evidence of a (statistically insignificant) increase in world industrial production at most horizons, consistent with OPEC announcements bringing news about higher oil demand rather than lower oil supply. Moreover, the large and statistically significant decline in U.S. industrial production in the original specification has been replaced by a slight and statistically insignificant decline. The response of the CPI also has shrunk.[12]

The concern here is not so much that the responses in Figure 5b are different from those in Figure 5a, which may be explained by the nature of the news evolving over time, but that the pattern of the response functions in Figure 5b is difficult to reconcile with a coherent economic narrative. Certainly, these responses are not consistent with the narrative that the

[12] This apparent instability in the responses across subsamples is also potentially consistent with evidence in Bruns and Lütkepohl (2023) of a structural change around 1990 in a similar structural oil market VAR model identified by changes in heteroskedasticity. Such instability, however, in small samples may also arise under the maintained assumption of a time invariant error covariance matrix, as the nature of the news evolves over time. For related discussion see Kilian and Park (2009).

proxy VAR model recovers responses to adverse oil supply news.

Figure 5b also provides a third important insight. The motivation in Känzig (2021) for extending the estimation period far beyond the period for which the proxy can be constructed was the belief that this increases the precision of the response estimates, while producing similar point estimates. However, Figure 5b shows that the response estimates based on the estimation period starting in 1988.4 not only differ substantially from those for the full estimation period, but that they are more tightly estimated in many cases, so these results cannot simply be disregarded.

Of course, as explained earlier, the PC-IV proxy is likely to be misleading since it ignores the accounting relationship between daily futures prices and monthly average futures prices. Thus, the results in Figure 5 are mainly relevant as a baseline for the further analysis of the implications of the construction of the monthly proxy. The key question is how these results change when using monthly average surprises as the proxy.

4.2. The monthly average proxy in the baseline VAR model of Känzig (2021)

The next set of results replaces the PC-IV proposed by Känzig (2021) by the monthly proxy proposed in Section 3.2. As before, I discard the oil futures price data for 1983.4-1989.3. Inference is based on the weak-identification robust AR confidence sets. For expository purposes I focus on the 12-month maturity (IV(12)) rather than constructing the leading principal component across all maturities. Since the correlation across maturities is high, consistent with the entire futures curve shifting up and down in response to OPEC news, this simplification is immaterial for the results.

Figure 6a shows that, when working with the full sample, this instrument is much stronger than the PC-IV proxy. The robust F-statistic increases from 6.92 to 20.46, indicating that the method of the temporal aggregation of the daily shocks has important practical implications. The implied VAR responses are broadly similar to those in Figure 5a, but often less precisely estimated. For example, the increase in world oil inventories now is delayed and barely statistically significant. Likewise, the declines in world oil production and in U.S. industrial production are barely statistically significant at the 10% level and only at horizons of two years. The decline in world industrial production is statistically insignificant at all horizons.

Despite some similarities with the responses in Figure 5a, there are also important differences. What does not quite fit the narrative of negative oil supply news is the delayed increase in world oil inventories. We know from economic theory that a storage demand shock driven by expectations of rising oil prices, all else equal, will raise inventories on impact, so adverse oil supply news should be associated with rising global oil inventories. The lack of such a response may be explained by the positive (if statistically insignificant) short-run response in global industrial production, except the sign of the impact response is at odds with the definition of a storage demand shock. Note that, all else equal, an increase in the real price of oil driven by higher storage demand should lower real activity rather than increase it. Thus, it is not clear what this proxy model identifies.

Figure 6b shows the corresponding proxy VAR estimates for the shorter estimation sample. This further raises the robust F-statistic to 33.54 compared to 8.63 in Figure 5b. There are some striking changes in the responses compared to Figure 6a. Except for a short-lived increase in the real price of oil and decline in global oil inventories, all other global responses are statistically insignificant. There is no evidence that the shock of interest is associated with a systematic decline in world oil production. If anything there is an initial statistically insignificant increase in world oil production that may represent an endogenous response to higher oil prices. There is again some statistically insignificant increase in world industrial production that suggests that OPEC news may be capturing global demand news, which would help explain the initial decline in global oil inventories, but is difficult to reconcile with adverse oil supply news. The recessionary effect on U.S. industrial production is small and statistically insignificant. Only the blip in U.S. inflation remains.

Figures 6a and 6b raise the question of what exactly is behind the large increase in the real price oil, given that the pattern of the responses does not match the narrative of Känzig (2021) or, for that matter, any other economically plausible narrative. In particular, the response of the real price of oil in Figure 6b appears too large to be consistent with the remaining oil market responses when compared to Figure 6a. One possible explanation we examine next is that these puzzling results may reflect a misspecification of the reduced-form oil market VAR model.

4.3. The monthly average proxy in the workhorse model of the global oil market

Since the baseline oil market VAR(12) specification used by Känzig (2021) is by no means standard in the literature, it is important to examine how the monthly average proxy performs when applied to more conventional oil market models such as the VAR(24) model of the global oil market introduced by Kilian and Murphy (2014), which was explicitly designed to capture shifts in oil price expectations.[13] The latter model uses variable definitions and data transformations that are more appropriate for the task at hand and allows for more lags, which has been shown to be important in modeling oil markets.[14]

In this section, I estimate this model with the state-of-the-art measure for global oil inventories developed in Kilian (2022b) that incorporates changes in China's oil inventories that are ignored by more traditional measures of global oil inventories. I also use the global

[13] Applications of this framework include Boer, Pescatori and Stuermer (2023), Cross et al. (2022), Herrera and Rangaraju (2020), Herwartz and Plödt (2016), Inoue and Kilian (2022), Kilian and Zhou (2020, 2022), and Zhou (2020), among others.
[14] For example, the U.S. price of oil used by Känzig was heavily regulated until the early 1980s and not representative for the global price of oil, which is why a more common choice for the global oil price has been the U.S. refiners' acquisition cost for imported crude oil (e.g., Mork 1989; Alquist et al. 2013). Likewise, world industrial production has been shown to be a problematic proxy for global real activity in commodity markets (e.g, Kilian and Zhou 2018). Another importance difference is that Kilian and Murphy (2014) does not express changes in inventories in percent, but in barrels.

real activity index discussed in Kilian (2009, 2019) and Kilian and Zhou (2018), which captures the fact that demand for industrial commodities tends to rise well before global industrial production since these commodities must be shipped to the producer first.

Figures 7a shows the response estimates obtained when applying the monthly average proxy to the real oil price residual implied by the Kilian-Murphy reduced-form model estimated on the full sample. The first point to note is that the robust F-statistics are low, indicating that the OPEC proxy has considerably less predictive power for real oil price shocks, once the information set is specified as in conventional global oil market models. The second point of interest is that the responses in some dimensions are similar to those reported in Känzig (2021), but in other dimensions are fundamentally different.

Specifically, there is strong evidence in Figure 7a that the increase in the real price of oil identified by the proxy VAR model is associated with a rise in expected demand, as can be seen from the statistically significant impact response of the Kilian (2009, 2019) index of global real activity derived from ocean bulk dry cargo shipping rates, which is known to be a leading indicator for global industrial production (e.g., Kilian and Zhou 2018; Funashima 2020). This evidence is consistent with OPEC announcements revealing news about expected demand for oil.[15]

In contrast, there is no evidence that the increase in the price of oil is associated with expectations of supply cuts. In fact, world oil production rises for the first year after the shock, consistent with an endogenous response to higher oil prices, before declining. The decline in global oil production after about two years is statistically significant, but does not

[15] If OPEC reacted to demand news already known to oil futures market participants, one would not expect the oil futures price to move in response to OPEC announcements. Thus, as long as we believe that the market is well informed, it makes sense to treat OPEC news as exogenous, which in turn implies that it is the news causing the global real activity index to move rather than the other way around. If participants in the oil futures market were not well informed about oil demand news and only learned about these news from the OPEC announcement, in contrast, the exogeneity of the news with respect to the global real activity measure in the VAR model would have to be questioned. For related discussion of concerns about the exogeneity of high-frequency monetary policy shocks see Bauer and Swanson (2023).

appear related to OPEC supply news, because supply cuts announced by OPEC tend to be implemented within the next few months rather than two years later. Some of this decline may be explained by a general economic slowdown caused by rising oil prices, given the simultaneous statistically significant decline in U.S. industrial production. As noted in Kilian (2009), demand-driven oil price booms carry the seeds of their own destruction. The magnitude of the decline in the response of oil production, however, argues against this being the only explanation and is suggestive of a storage demand shock driven by the anticipation of declining oil production at longer horizons.

The fact that the global real activity index responds positively on impact is consistent with expectations of rising demand for oil driving up the real price of oil contemporaneously. It is this effect that helps explain the large positive response of the real price of oil in Figure 7a for the first two years. Such demand shifts are labelled flow demand shocks in Kilian and Murphy (2014), suggesting that the proxy VAR model underlying Figure 7a identifies a linear combination of flow demand and storage demand shocks. These positive flow demand shocks mask the decline in real activity and the increase in inventories expected from a positive storage demand shock driven by negative oil supply news. This interpretation is consistent with the decline in global inventories during the first two years, when the flow demand shock dominates the response to the expectations shock. Only at longer horizons, as this demand boom dies out, we see the positive response in global inventories expected from an anticipation of reduced oil production. The evidence of shifts in the flow demand for oil also helps explain the more persistent inflationary pressures in response to this shock in Figure 7a.

We conclude that the narrative proposed by Känzig (2021) is not robust to changes in the specification of the VAR model of the global oil market. These model specifications, however, suggest an alternative interpretation of the oil price expectations shock as a linear combination of flow demand and storage demand shocks. This point is important not only for the narrative, but also because it prevents us from externally validating the estimated model by means of historical decompositions. The importance of externally validating VAR estimates of global oil market models has been stressed as far back as Kilian and Murphy (2014). External validation involves comparing historical decompositions of the data against extraneous evidence about events in the global oil market to verify that the model estimate makes economic sense.[16] For example, as discussed in Kilian and Murphy (2014), the only explanation of the lack of an oil inventory response to the invasion of Kuwait in August 1990 is that storage demand must have increased sharply in anticipation of further supply disruptions, amplifying the spike in the real price of oil, but offsetting the decline in oil inventories caused by the disruption of Iraqi and Kuwaiti supplies. Thus, a storage demand shock in a proxy VAR model could be validated by showing that a historical decomposition of the real price of oil recovers this shift in oil price expectations. To the extent that the shock identified by the proxy VAR model is a linear combination of flow demand and storage demand shocks, however, this validation exercise becomes infeasible.

Figure 7b shows the corresponding findings for the shorter estimation period, which are even stronger. The key difference is that there is no evidence that the oil price expectations shock in the proxy VAR model captures oil supply news at all. Figure 7b shows that the expectations shock identified by the proxy VAR model is for all intents and purposes a flow demand shock, with storage demand, playing a negligible role. The positive response of the real price of oil is much less persistent than in Figure 7a, as is that of global real activity to a lesser extent, but both are statistically significant at the 10% level for several months. The response of world oil production is positive during the first 18 months, at times

[16] External validation can be viewed as an antecedent of the narrative restrictions for VAR models popularized by Antolin-Diaz and Rubio-Ramirez (2018).

significantly so, consistent with oil producers responding to the higher real price of oil. There is no evidence of a statistically significant decline in oil production at longer horizons. The delayed decline in U.S. industrial production is less persistent than in Figure 7a, as is the response of U.S. inflation.

There has been some debate over whether Känzig's proxy exclusively captures OPEC oil supply news. OPEC announcements not only include OPEC production plans, but they also reveal information about OPEC's economic outlook that tends to be closely scrutinized by the market. In response to this concern, Känzig (2021) also considered an alternative PC-IV proxy that controls for revisions in OPEC's global demand forecasts and reports substantively identical results, leading him to conclude that OPEC announcements mainly reveal information about OPEC production plans. Degasperi (2021) proposed a further refinement of Känzig's proxy designed to separate OPEC oil supply and oil demand news and reached a different conclusion. Neither focuses on the monthly average surprise, however. The analysis in this paper provides additional evidence that OPEC surprises without further refinements represent at least in part, if not entirely, news about the global economy. It also suggests that in thinking about the role of OPEC announcements on oil markets today, the estimates obtained from the sample starting in the late 1980s are more likely to be representative than the estimates based on the full sample.

5. Concluding Remarks

It is common in applied macroeconomics to estimate responses of macroeconomic aggregates to news shocks derived from surprise changes in daily futures prices around the date of policy or other announcements. This requires mapping the daily surprises into a monthly shock that may be used as an instrument in a monthly proxy VAR model or local projection (e.g., Kuttner 2001; Faust, Swanson and Wright 2004; Gertler and Karadi 2015; Känzig 2021; Gagliardone and Gertler 2023).

The conventional approach has been to sum these daily surprises over the course of a given month when constructing the monthly proxy variable. For example, in the case of only one surprise in a given month this amounts to equating daily and monthly shocks, regardless of when this shock occurs within the month. This approach is counterintuitive, if the price variable to be instrumented is defined as an average over the month, because in that case the effect of a daily shock on the monthly average price clearly depends how early or late in the month this surprise occurs. The central contribution of this paper has been to provide a new approach to recovering monthly proxies from a daily surprise series that preserves the accounting identities linking daily futures prices to monthly averages of these prices. This approach also implies that the effects of a daily surprise may extend to the subsequent month.

I illustrated this methodology by analyzing the question of how to identify OPEC news shocks in VAR models of the global oil market. This application has received considerable interest in applied work more recently (e.g., Känzig 2021; Bruns and Lütkepohl 2022; Gagliardone and Gertler 2023). The paper reexamined the problem of constructing surprises in oil futures prices associated with the OPEC announcements and explored the sensitivity of the proxy VAR estimates reported in the literature to the construction of the monthly proxy.

The paper drew attention to the fact that the first six years of daily oil futures price data used in Känzig (2021) and subsequent studies are not suitable for constructing OPEC surprises. Discarding these data not only changes the response estimates, but renders the original PC-IV proxy a weak instrument. I showed that this problem is overcome by constructing the monthly average surprise, as proposed in this paper. However, the resulting proxy VAR estimates are not robust across estimation periods and do not line up well with plausible economic narratives. Alternative oil market VAR specifications based on the same monthly average proxy produced more economically plausible results. These models suggest that the proxy VAR identifies a linear combination of storage demand and flow demand shocks, with the latter dominating in more recent data.

This evidence cautions against the narrative that OPEC announcements may be used to identify anticipated oil supply disruptions, complementing related work by Degasperi (2021). Nor is there support for the notion that this shock is a special case of the storage demand shock constructed in Kilian and Murphy (2014), which complicates the interpretation of the responses recovered by the proxy VAR model and prevents the user from externally validating the model estimates. Finally, the analysis in this paper suggested fitting proxy VAR models on a longer estimation period than the instrument is available for, may substantially change the response estimates. Restricting the estimation period to the subsample for which daily surprises can be measured is likely to provide a more accurate representation of the effects of news shocks in today's oil market.

While OPEC announcements are one prominent example of applications of the proxy VAR methodology, similar techniques are widely used in other contexts including the analysis of FOMC announcements. A question of obvious interest would be to apply the methodology for constructing monthly average surprises from high-frequency measures of surprises developed in this paper to the problem of estimating monetary policy shocks.

References

Amir-Ahmadi, P., Matthes, C., and M.-C. Wang (2023), "Understanding Instruments in Macroeconomics: A Study of High Frequency Identification in Macroeconomics," manuscript, Indiana University.

Alquist, R., and L. Kilian (2010), "What Do We Learn from the Price of Crude Oil Futures?." *Journal of Applied Econometrics*, 25, 539-573.

Alquist, R., Kilian, L., and R..J. Vigfusson (2013), "Forecasting the Price of Oil," in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, 427-507.

Andrews, I., Stock, J.H., and L. Sun (2019), "Weak Instruments in IV Regression: Theory and Practice," Annual Review of Economics, 11, 727-753.

Antolin-Diaz, J., and J.F. Rubio-Ramirez (2018), "Narrative sign restrictions for SVARs," *American Economic Review*, 108, 2802-2839.

Arias, J., Rubio-Ramirez, J.F., and D.F. Waggoner (2021), "Inference in Bayesian Proxy-SVARs," *Journal of Econometrics*, 225, 88-106.

Bauer, M.D., and E. Swanson (2023), "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," NBER Macroeconomics Annual, 37, 87-155.

Boer, L., Pescatori, A., and M. Stuermer (2023), "Not All Energy Transitions Are Alike: Disentangling the Effects of Demand- and Supply-Side Policies on Future Oil Prices," manuscript, International Monetary Fund.

Bruns, M. (2021), "Proxy Vector Autoregressions in a Data-Rich Environment," Journal of Economic Dynamics and Control, 123, 104046.

Bruns, M., and H. Lütkepohl (2023), "Have the Effects of Shocks to Oil Price Expectations Changed? Evidence from Heteroskedastic Proxy Vector Autoregressions," DIW Discussion Paper No. 2036.

Caldara, D., and E. Herbst (2019), "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs," *American Economic Journal: Macroeconomics*, 11, 157-192.

Cross, J.L., Nguyen, B.H., Tran, T.D. (2022). "The Role of Precautionary and Speculative Demand in the Global Market for Crude Oil. *Journal of Applied Econometrics*, 37, 882-895.

Degasperi, R. (2021), "Identification of Expectational Shocks in the Oil Market using OPEC Announcements," Technical Report, University of Warwick.

Demirer, R., and A.M. Kutan (2010), “The Behavior of Crude Oil Spot and Futures Prices around OPEC and SPR Announcements: An Event Study Perspective,” *Energy Economics*, 32, 1467-1476.

Faust, J., Swanson, E.T., and J.H. Wright (2004), “Identifying VARs Based on High-Frequency Data,” *Journal of Monetary Economics*, 51, 1107-1131.

Funashima, Y. (2020), “Global Economic Activity Indexes Revisited,” *Economics Letters*, 193, 109269.

Gagliardone, L., and M. Gertler (2023), “Oil Prices, Monetary Policy and Inflation Surges,” NBER Working Paper No. 31263.

Gertler, M., and P. Karadi (2015), “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44-76.

Herrera, A.M., and S.K. Rangaraju (2020), “The Effect of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?,” *Journal of Applied Econometrics*, 35, 141-159.

Herwartz, H., and M. Plödt (2016), “The Macroeconomic Effects of Oil Price Shocks: Evidence from a Statistical Identification Approach,” *Journal of International Money and Finance*, 61, 30-44.

Inoue, A., and L. Kilian (2022), “Joint Bayesian Inference about Impulse Responses in VAR Models,” *Journal of Econometrics*, 231, 457-476.

Jacobson, M.M., Matthes, C., and T.B. Walker (2023), “Temporal Aggregation Bias and Monetary Policy Transmission,” manuscript, Indiana University.

Jarociński, M., and P. Karadi (2020), “Deconstructing Monetary Policy Surprises – the Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 12, 10-43.

Känzig, D.R. (2021), “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements,” *American Economic Review*, 111, 1092-1125.

Känzig, D.R. (2023), “The Unequal Economic Consequences of Carbon Pricing,” manuscript, Northwestern University.

Kilian, L. (2008a), “The Economic Effects of Energy Price Shocks, *Journal of Economic Literature*, 46, 871-909.

Kilian, L. (2008b), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *Review of Economics and Statistics*, 90, 216-240.

Kilian, L. (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. (2022a), “Understanding the Estimation of Oil Demand and Oil Supply Elasticities, *Energy Economics,* 107, 105844.

Kilian, L. (2022b), “Facts and Fiction in Oil Market Modeling,” *Energy Economics,* 110, 105973.

Kilian, L, and B. Hicks (2013), “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?” *Journal of Forecasting*, 32, 385-394.

Kilian, K., and D.P. Murphy (2014), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” *Journal of Applied Econometrics*, 29, 454-478.

Kilian, L., and 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 H. Lütkepohl (2017), *Structural Vector Autoregressive Analysis*, Cambridge University Press, New York.

Kilian, L., and C. Park (2009), “The Impact of Oil Price Shocks on the U.S. Stock Market *International Economic Review*, 50, 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, 660-671.

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

Kilian, L., and X. Zhou (2018), "Modeling Fluctuations in the Global Demand for Commodities," *Journal of International Money and Finance*, 88, 54-78,

Kilian, L., and X. Zhou (2020), "Does Drawing Down the U.S. Strategic Petroleum Reserve Help Stabilize Oil Prices?," *Journal of Applied Econometrics*, 35, 673-691.

Kilian, L., and X. Zhou (2022), "Oil Prices, Exchange Rates and Interest Rates," *Journal of International Money and Finance,* 126, 102679.

Kilian, L., and X. Zhou (2023), "The Econometrics of Oil Market VAR Models," *Advances in Econometrics*, 45B, 65-95.

Knittel, C.R., and R.S. Pindyck (2016), "The Simple Economics of Commodity Price Speculation," *American Economic Journal: Macroeconomics*, 8, 85-110.

Kuttner, K. (2001), "Monetary Policy Surprises and Interest Rates," *Journal of Monetary Economics*, 47, 523-544.

Mertens, K., and M.O. Ravn (2013), "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review*, 103, 1212–1247.

Miranda-Agrippino, S., and G. Ricco (2021), "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13, 74-107.

Montiel-Olea, J.L., Stock, J.H., and M.W. Watson (2021), "Inference in SVARs Identified with External Instruments." *Journal of Econometrics*, 225, 74-87.

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

Stock, J.H., and M.W. Watson (2012), "Disentangling the Channels of the 2007–09

Recession," *Brookings Papers on Economic Activity,* 81–135.

Stock, J.H., and M.W. Watson (2018), "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *Economic Journal*, 128, 917–48.

Wolf, C. (2020), "SVAR (Mis)Identification and the Real Effects of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 12, 1-32.

Zhou, X. (2020), "Refining the Workhorse Oil Market Model," *Journal of Applied Econometrics*, 35, 130-140.

Figure 1: Why the timing of daily shocks to the log futures price matters

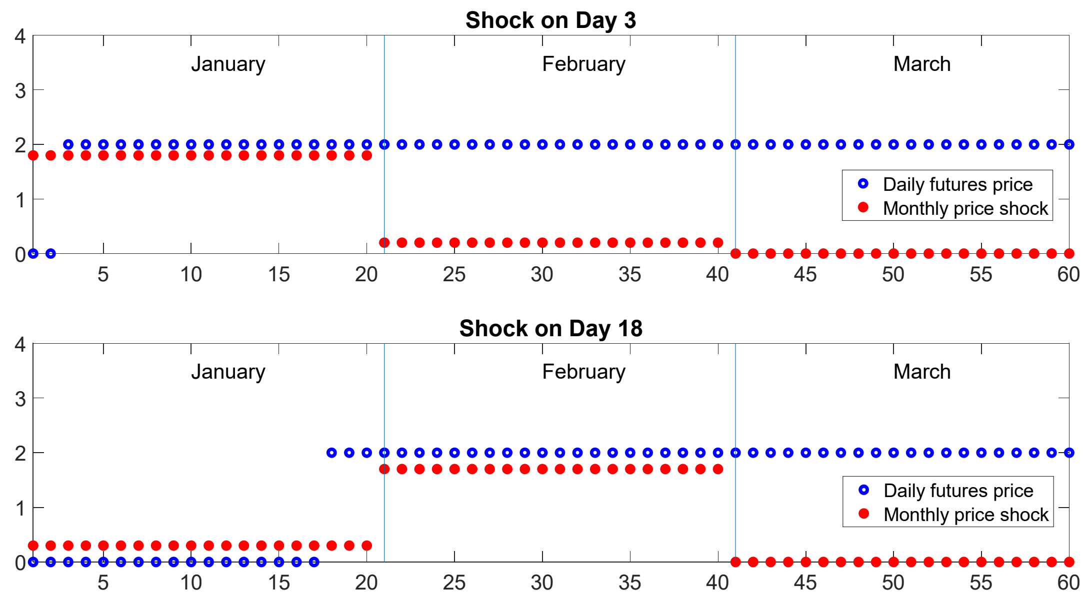

NOTES: The premise is that there are 20 business day per month and no holidays. The initial futures price is normalized to zero without loss of generality. The surprise occurs on January 3 or January 18, respectively, and amounts to 2. How much this surprise raises the average price in January depends on the timing, as does the implied monthly surprise in February, defined as the difference between the average daily price in February and the average price increase in January. By March, the average daily price coincides with the monthly average price in February and the monthly shock is zero.

Figure 2: Example of construction of monthly average price shock with two daily surprises

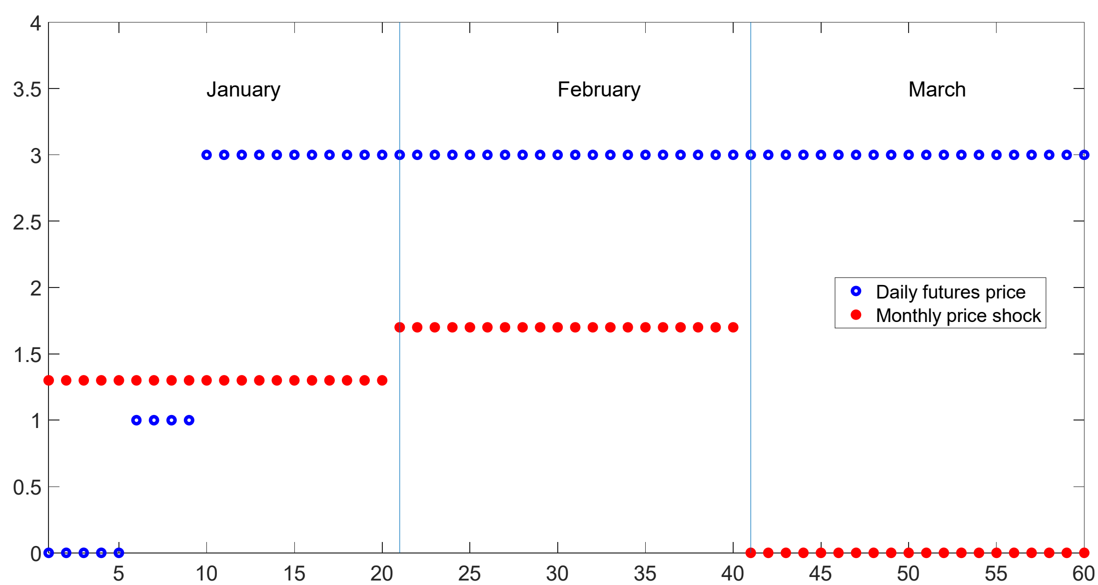

NOTES: The first surprise occurs on January 6 and amounts to 1. The second surprise occurs on January 10 and amounts to 2, raising the daily futures price to 3 cumulatively. This implies a monthly average surprise of 1.3 in January, and a further surprise of 1.7 in February, given an expectation of 1.3 for the average price going into that month and a realization of 3. There is no surprise in March, as the realized price of 3 coincides with the average futures price observed in February.

Figure 3: Alternative monthly proxies

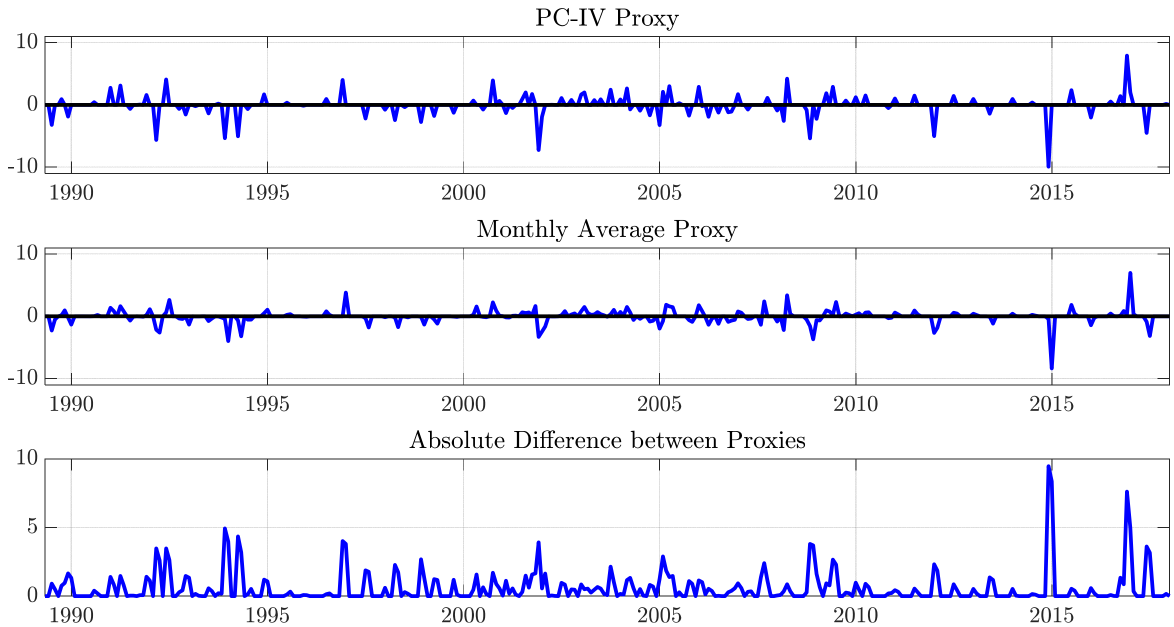

NOTES: Both proxies have been computed based on oil futures price data starting in 1989.4. The contemporaneous correlation of the proxies is 42%.

Figure 4: Exogenous variation in oil price predicted by November 2014 OPEC announcement

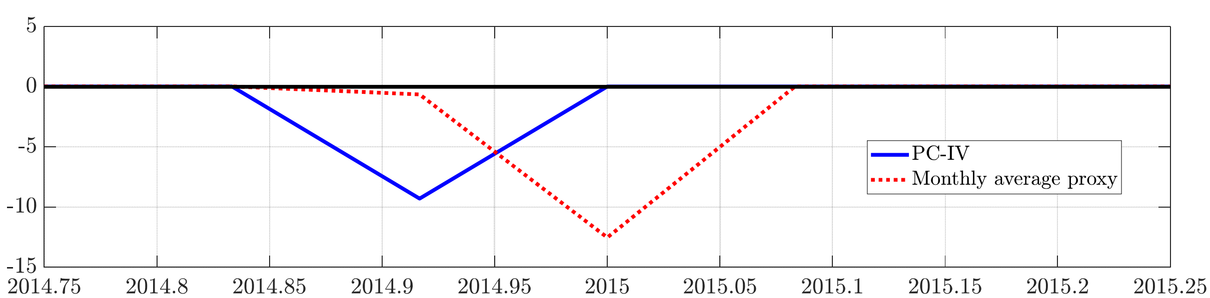

NOTES: Fitted value from first-stage IV regression based on baseline VAR(12) model specification in Känzig (2021). All proxies have been computed based on oil futures price data starting in 1989.4. The monthly average proxy is for the 12-month maturity. The OPEC announcement occurred on November 27, 2014, right before the last trading day of that month.

Figure 5a: Känzig baseline VAR(12) specification with corrected PC-IV, 1975.1-2017.12, CS^{AR}

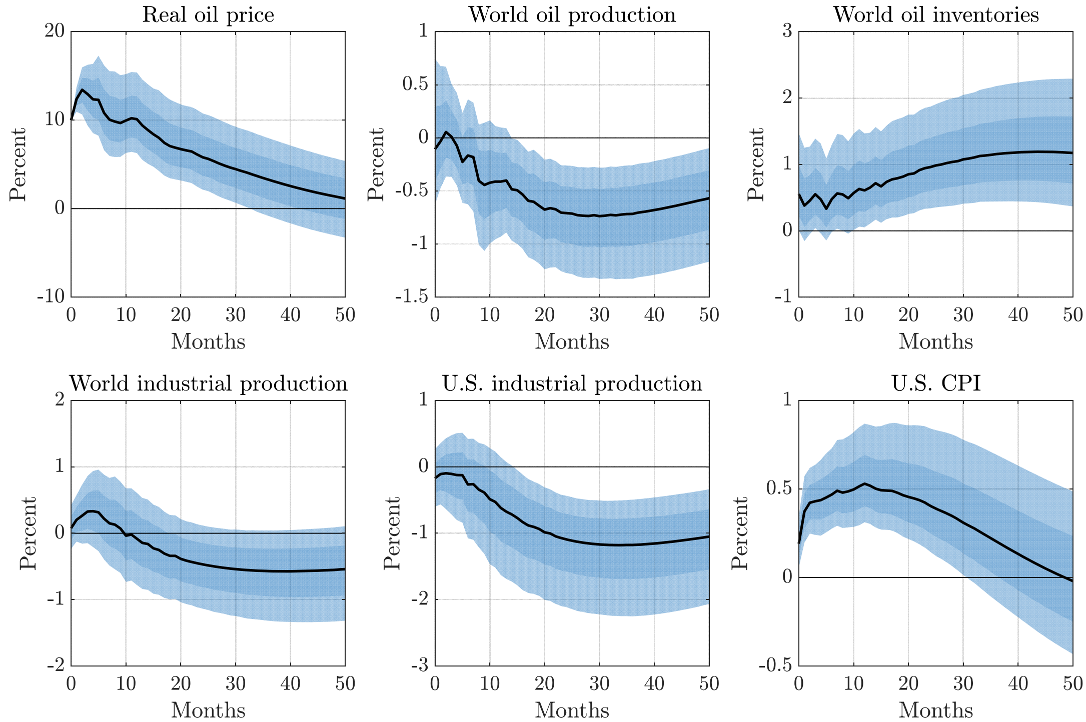

First stage regression robust F: 6.92

Figure 5b: Känzig baseline VAR(12) specification with corrected PC-IV, 1988.4-2017.12, CS^{AR}

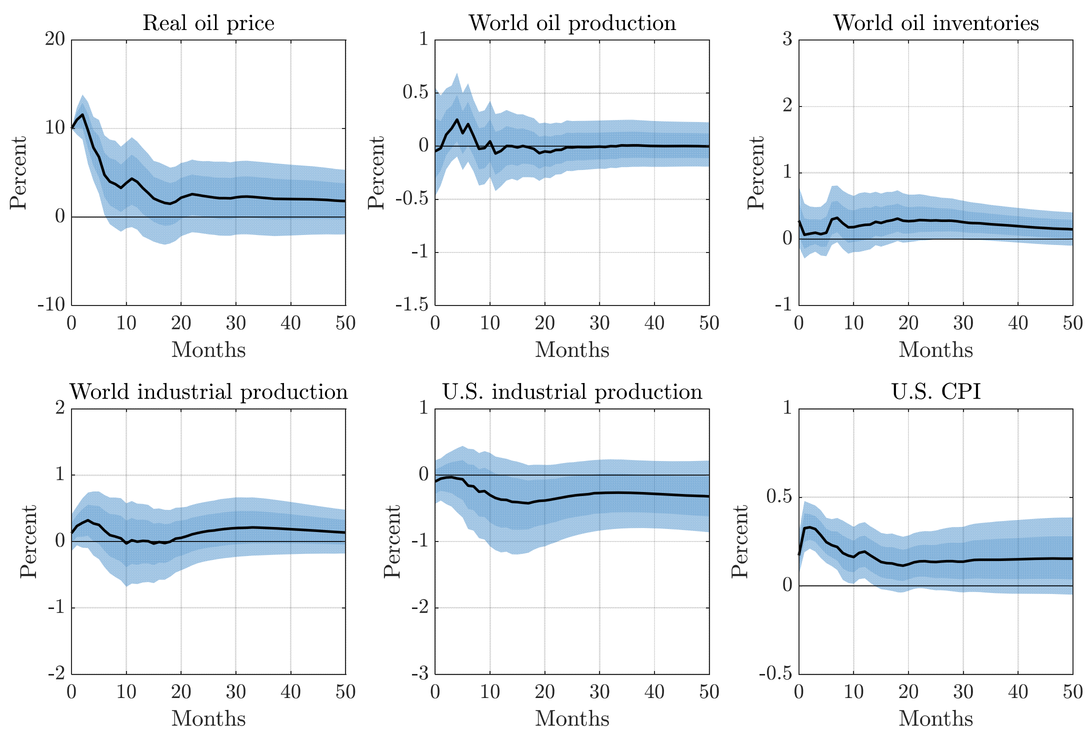

First stage regression robust F: 8.63

Figure 6a: Känzig baseline VAR(12) specification with monthly average IV(12), 1975.1-2017.12, CS^{AR}

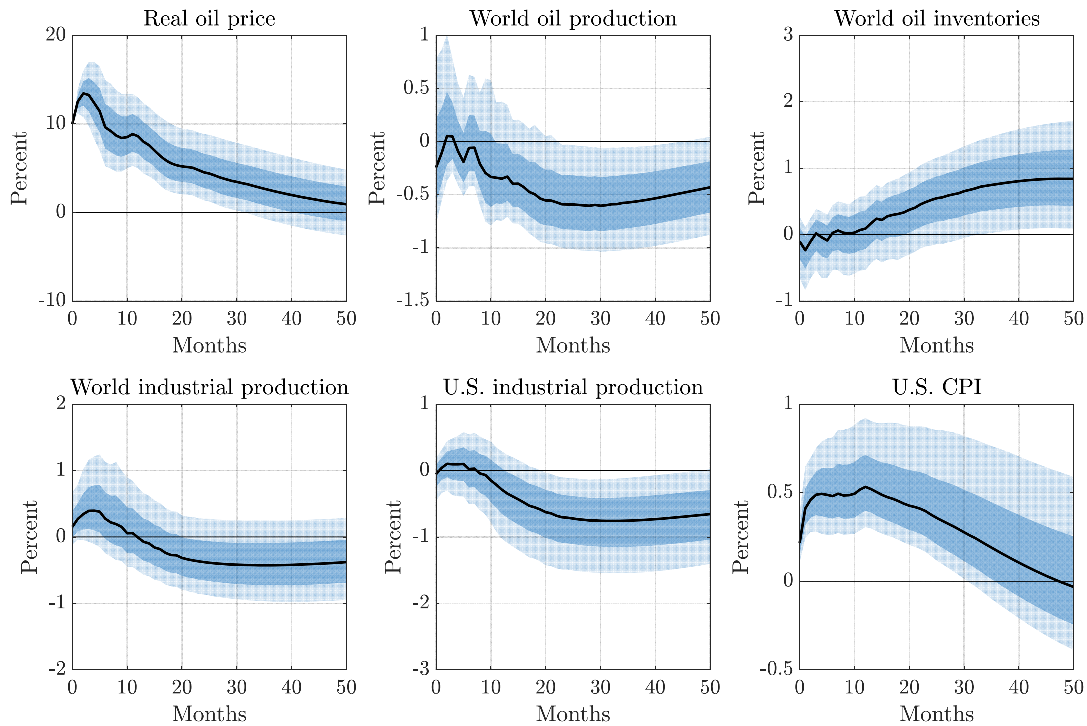

First stage regression robust F: 20.46

Figure 6b: Känzig baseline VAR(12) specification with monthly average IV(12), 1989.4-2017.12, CS^{AR}

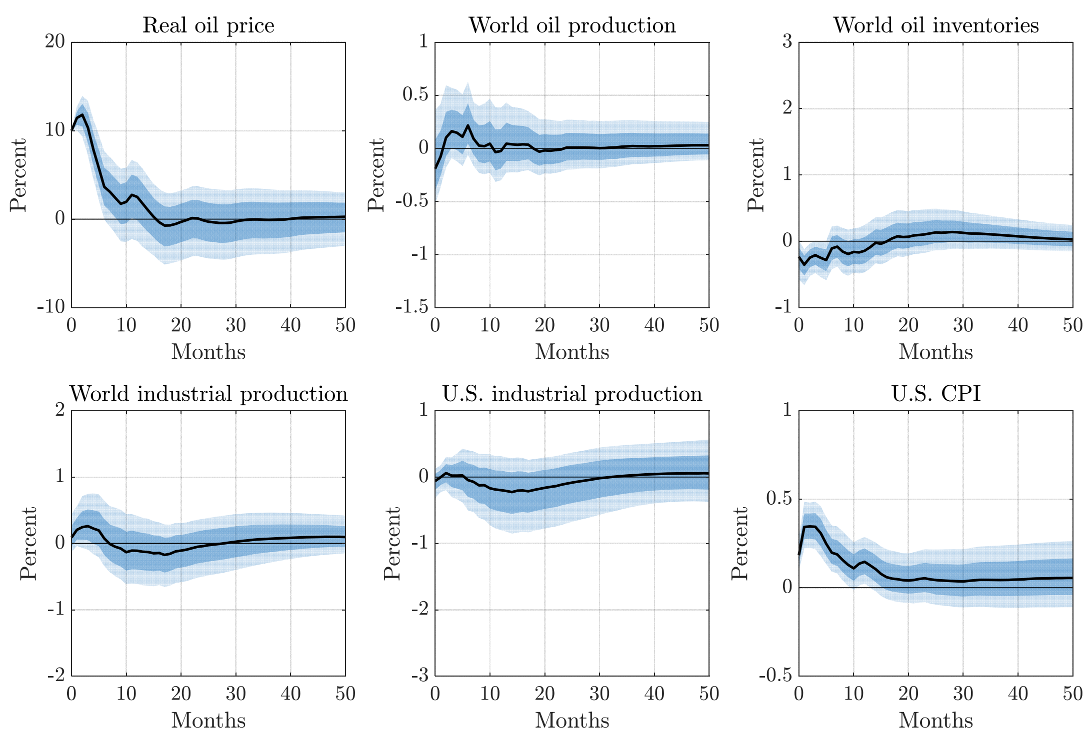

First stage regression robust F: 33.54

Figure 7a: Kilian-Murphy VAR(24) specification with monthly average IV(12), 1976.1-2017.12, CS^{AR}

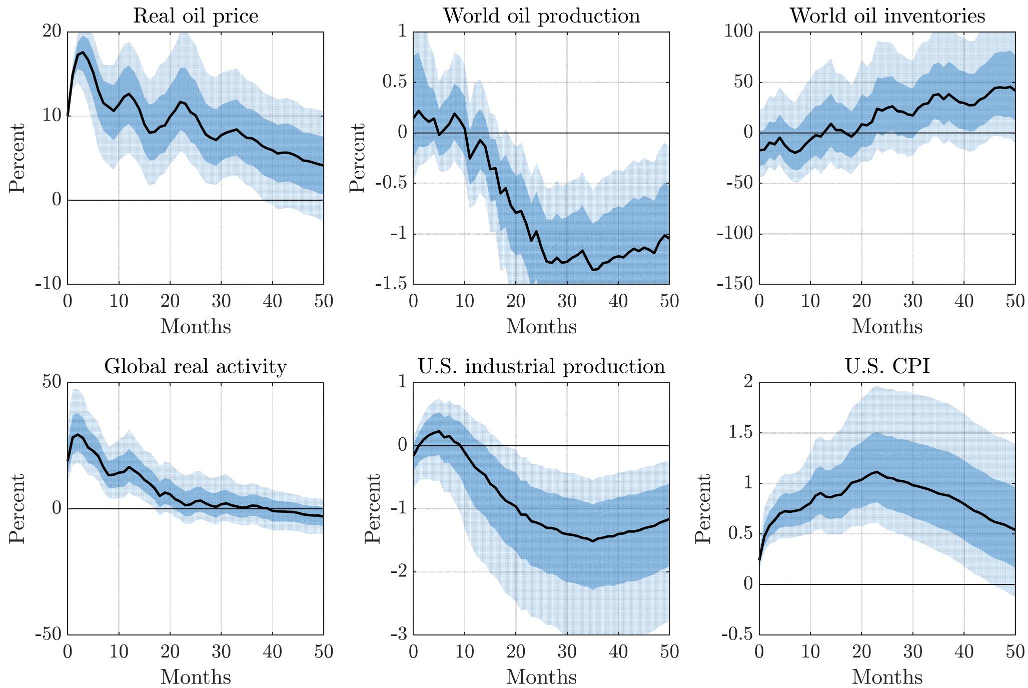

First stage regression robust F: 11.73

Figure 7b: Kilian-Murphy VAR(24) specification with monthly average IV(12), 1989.4-2017.12, CS^{AR}

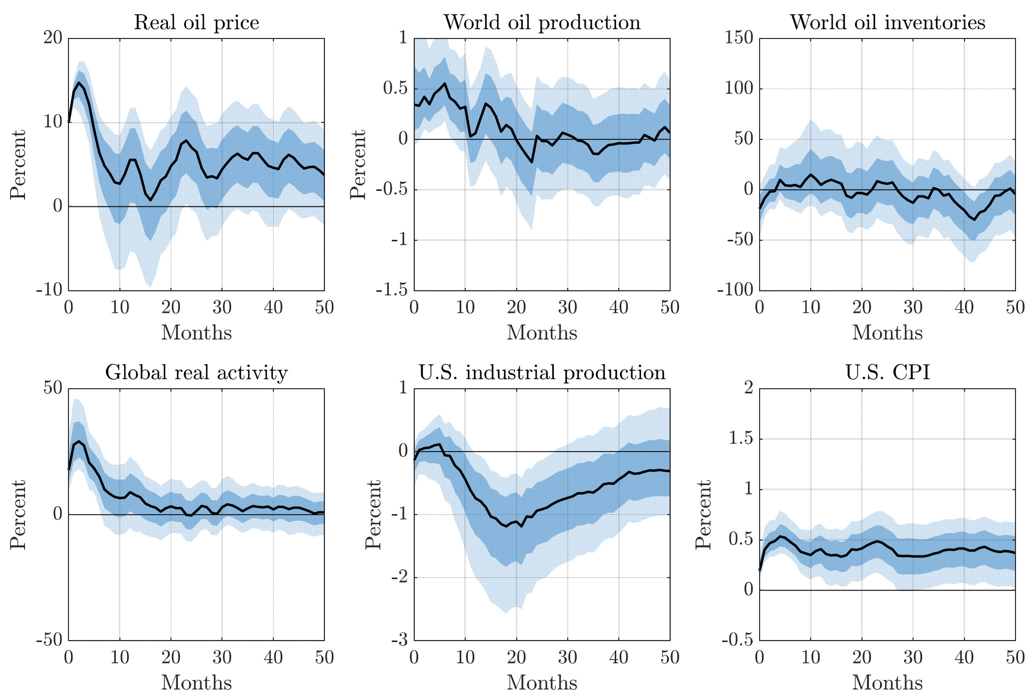

First stage regression robust F: 7.85