Energy Economics 98 (2021) 105214

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneeco

Pass-through of oil supply shocks to domestic gasoline prices: evidence from daily data

Etsuro Shioji

Department of Economics, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan

ARTICLE INFO

Article history:
Received 25 November 2019
Received in revised form 1 February 2021
Accepted 25 February 2021
Available online 11 March 2021

Keywords:
SVAR-IV (or proxy VAR)
High frequency identification
News shock
Oil price
Pass-through
JEL codes:
E31
F41
Q41
Q43
C32

ABSTRACT

Oil prices react to various types of shocks, and their impacts could reach our lives quickly. For example, on one day, we might hear on the radio that OPEC has decided to cut their oil production for the next year, and, only a few days later, find out that our local gasoline station has just raised the price. This paper examines daily data on gasoline prices, produced by a price comparison site in Japan, to estimate how they respond to a shock that hit the world oil market. In doing so, we take seriously the possibility that an increase in oil prices might cause different reactions depending on the source of the change. This paper focuses on one particular type of shock, namely changes in expectations about future supplies of crude oil. Identification is achieved via estimating a version of the Structural VAR with External Instruments (SVAR-IV or proxy-VAR) coupled with High Frequency Identification (HFI). The result confirms that pass-through is indeed very fast: about 70% of the entire adjustment process is completed within just 18 days.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

This paper estimates how shocks to world oil prices are transmitted to domestic gasoline prices in Japan. Our experiences suggest that this process might be very fast. For example, on September 28, 2016, the OPEC member countries agreed to reduce production of oil. On October 14, a little over two weeks later, the Nikkei Newspaper of Japan was reporting that gasoline stations were raising their prices. In such an environment, a conventional approach, which relies on monthly data, may not be able to capture the entire process of the price adjustment adequately. For that reason, this paper examines a new daily data set on gasoline prices in Japan. The source of this valuable data is a popular price comparison site specialized in reporting prices of gasoline and diesel oil. Tens of thousands of registered drivers, as well as shop managers, report prices posted at gasoline stations to this web site from all over Japan, on daily basis.

In our analysis, we take seriously the claim made in the literature that the effects of oil price changes might vary depending on the nature of the shock that caused them. That is, we need to disentangle the intertwined relationship between supply and demand. This paper adopts an approach similar to that of Känzig (2021) to identify oil *supply* shocks, or more precisely, shocks to expectations regarding future supplies of oil, and estimates their effects on domestic prices.

This approach is based on the methodology called the Structural VAR with External Instruments, or SVAR-IV (also referred to as the proxy VAR), which has been developed by Mertens and Ravn (2013) and Stock and Watson (2012). It involves use of an "external instrument", whose requirement is to be correlated with the shocks of interest (oil supply shocks in our case) but uncorrelated with all the other types of shocks. Introduction of such an instrument allows us to achieve identification of the shock that we are interested in, without requiring us to achieve identification of the entire structural form of the model.

In this paper, such an external instrument is constructed in three steps. The first step is to find candidate periods during which important news about the future course of the world supply of oil might have occurred. I look for two types of events. The first set of events are the ones related to OPEC's decisions concerning its future oil production (such as an announcement on an agreement to cut future output). The second are those related to the US-led efforts to impose or to lift sanctions on crude oil imports from Iran. I utilize *Google Trends* for this task of narrowing down the candidate periods. The second step is to go through news articles around those candidate periods to determine on which date(s) there was a news about future oil supply. The third step is to measure reactions of the market to such news. They should tell us which of those news have truly been "shocks" to market expectations,

E-mail address: shioji@econ.hit-u.ac.jp.

https://doi.org/10.1016/j.eneco.2021.105214
0140-9883/© 2021 Elsevier B.V. All rights reserved.

and, if so, magnitudes of the surprises. This is done by computing daily changes in Brent futures on those days. This produces an indicator which reflects at least some parts of changes in people's expectations about the future oil supply.

This news indicator is incorporated, as an external instrument, into a VAR model with three endogenous variables, namely crude oil prices that are most relevant for the Japanese market (Dubai, denominated in the USD), the exchange rate (between the USD and the JPY) and domestic gasoline prices. Note that our news indicator might not fully capture all the fluctuations in the expected oil supply, but is (in our view) definitely correlated with the true sequence of all the supply shocks.

This paper finds the following. An expected oil supply shock has a significantly positive impact on gasoline prices in Japan. The effect is permanent and large. The long run pass-through rate is around 20% if we use gasoline prices that are not adjusted for taxes, but this number goes up to 40% when we use the tax-adjusted data. And this pass-through is fast: almost 70% of the long run pass-through occurs within the first 18 days after the shock. This demonstrates the importance of the use of daily data: a study which relies on monthly data would not be able to detect the dramatic price responses that happens within a month.

In addition to this study of the dynamics of the national average gasoline prices, we also study how response of gasoline prices might differ across groups of gasoline stations with different characteristics. For example, we ask if shops in urban areas tend to react faster than their non-urban counterparts to an oil price shock. The group-specific mean gasoline price series are constructed from a vast micro level data set which contains all the price information reported to the web site from different gas stations across the country for a period of about one and a half years.

The rest of the paper is organized as follows. Sections 2 and 3 offer an overview of the related literature. Section 4 introduces our new dataset on gasoline prices at Japanese gas stations. Section 5 reviews the idea behind the SVAR-IV approach. Section 6 explains the construction of our External Instruments. Section 7 presents the estimation results based on our daily data. In section 8, we compare those results with evidence we find by using weekly as well as monthly aggregated data. Section 9 studies how responses of gasoline prices might differ depending on shop characteristics. Section 10 concludes.

2. Background (1) news, oil prices, and the macroeconomy

This paper stands at the crossroads of three strands of literature. First is a series of macroeconomic research on the impacts of oil price changes on the domestic economy. Secondly, it is related to a group of event studies on the world crude oil markets, especially the ones that focus on OPEC announcements. The last is the literature on pass-through, especially papers that study impacts of oil prices on domestic gasoline prices. We discuss the first two here, and the last one in the next section.

2.1. Oil shocks literature and the importance of shock identification

Macroeconomists have long been interested in the effects of oil price changes on the aggregate economy. The most important progress in the literature in recent years has been a realization that “Not All Oil Price Shocks Are Alike” (Kilian (2009)). The effects of an increase in oil prices of similar magnitudes could differ substantially depending on the source of the change: for example, a price hike due to a negative supply shock could have a very different impact on the domestic macroeconomy compared to increases in prices triggered by booms in global demand for oil. Kilian (2009) (and also Kilian and Park (2009)) proposes a new methodology to overcome this issue.

Their new approach is based on the SVAR to identify oil supply shock, oil demand shock, as well as global demand shock (i.e., shock that affects demand for not just oil but all types of commodities) simultaneously. In their framework, oil supply shocks are identified as innovations in the world production of crude oil. Global demand shocks are defined as innovations in a measure of the world economic activities, constructed from data on dry cargo freight rates, that are orthogonal to the above oil supply shocks. The part of innovations in oil prices that are orthogonal to the above two types of shocks are called oil market specific demand shocks. It is shown that different types of shocks have very different impacts on the US CPI. It is also shown that oil supply shocks are not a very important driver of the fluctuations in oil prices. Fukunaga et al. (2011) apply this methodology to the Japanese economy to investigate the effects of oil prices on sectoral output and prices. Also, the above approach has been extended by Iwaisako and Nakata (2015) to incorporate the exchange rate, and was applied to the Japanese data.

Kilian (2008b) proposes a different approach to directly estimate shocks to the world oil supply. It has been applied to G7 countries, including Japan, by Kilian (2008a). Baumeister and Hamilton (2019) show a way to extend this line of research by incorporating informative priors in a Bayesian framework.

In comparison, this paper tries to identify only one type of oil shock, namely shocks to the expectations about future world oil supply. It can be considered as a kind of an oil market specific demand shock in the terminology of Kilian (2009), which is “*designed to capture shifts in the price of oil driven by higher precautionary demand associated with market concerns about the availability of future oil supplies*” (lines 21–22, page 1053).

2.2. Events, news, and asset/commodity markets

There is a long tradition of event studies in financial economics. They examine how occurrence of events or news is reflected in either asset or commodity prices. This methodology has been adopted to study how oil-related news are incorporated into crude oil prices in the global market. Most of them have used OPEC announcements as events. Draper (1984) analyzes excess returns in the NYMEX Heating Oil futures market around the time of OPEC meetings. Loderer (1985) finds evidence that OPEC meetings affected gas-oil spot prices for the period 1981–1983 (but not for the years 1974–1980). In a more recent study, Demirer and Kutan (2010) show that announcements of production cuts from OPEC conferences have significant impacts on daily excess returns in both the futures and the spot markets for crude oil for the period 1983–2008. Lin and Tamvakis (2010), using data from the period 1982–2008, report that announcements from both formal OPEC conferences and ministerial meetings have significant impacts on various types of crude oil prices, though the effects are state-dependent. Schmidbauer and Rösch (2012) start with a set of three dummy variables, each corresponding to OPEC announcement of production cut, production increase, or unchanged production. They show that, when appropriately modified, all three have significant impacts on oil prices, though the effects are stronger for production cuts. Loutia et al. (2016), using data for the period 1991–2015, uncover a time varying nature of the effects of OPEC announcements. They also find some differences between reactions of WTI and Brent.

2.3. Utilizing events/news for shock identification

2.3.1. Events and shock identification

The idea of utilizing events or news to achieve a better shock identification is gaining importance in macroeconomic analysis. The idea goes back at least to Romer and Romer (1989) who create dummy variables for dates on which the Fed made surprising policy decisions which would not have been anticipated from its usual pattern of behaviors. In relation to fiscal policy, (Ramey and Shapiro, 1998) and Ramey (2011) set up dummies for dates on which the US government made decisions that would have meant larger military spendings in near future. The idea is to capture the timing on which the private sector's expectations about future fiscal policy changed. In that sense, they utilize those event dates to identify fiscal news shocks.

2.3.2. Financial market as the mirror of the mind

Some related work try to extract the private sector's expectations about the future from the financial market. For example, Fisher and Peters (2010) hypothesize that people's expectations about future US military spending are reflected in stock returns for large military contractors in the US, in an effort to study how the economy reacts to shocks to such expectations.

2.3.3. Event + market" based shock identification

Recent studies try to achieve a sharper shock identification by combining the above two approaches. That is, they look at behaviors of financial indicators on the dates of the news or events. In relation to the US monetary policy, Kuttner (2001) has shown that, when the Fed makes surprising policy announcements, they tend to be reflected in unanticipated changes in returns to the fed funds futures. Faust, Swanson and Wright (2004) utilize this finding to identify monetary policy shocks by measuring the change in the futures rate on the day of announced changes in the Fed's target federal funds rate. Gertler and Karadi (2015) use a similar variable not as an endogenous variable in a VAR model but as an external instrument for monetary policy shocks in an SVAR-IV model. For fiscal policy, Shioji (2018) uses changes in stock returns of construction companies on days of important policy developments in Japan as an external instrument to identify shocks to expectations concerning the future course of public investment spending.

2.3.4. SVAR-IV approach by Känzig (2021)

Känzig (2021) has been the first to introduce the idea of "event + market" based shock identification into the oil price shock literature. He proposes a novel approach to identify shocks to oil supply expectations. His study aims to identify shocks to expectations concerning future oil supply. For that purpose, he employs the SVAR-IV methodology, using the response of crude oil futures to news about future oil supplies as the external instrument.

2.3.5. Contributions of this paper

The basic identification strategy of this paper is quite similar to that of Känzig (2021).[1] It tries to identify shocks to expectations about the future course of oil supply, by using changes in crude oil futures prices on days on which oil-supply-related news arrived as an external instrument.

One feature that distinguishes this paper from much of the macro literature is the use of daily data. Most of those studies use either monthly or quarterly data, because those are the frequencies at which most macro variables are available. For example, Känzig (2021) converts his instrumental variable, which is based on daily observations, into the quarterly frequency. This is because aggregate variables of his interest, such as GDP, is quarterly data. This could lead to an information loss. Suppose, for example, that there were two events in a single quarter that moved oil prices by the same amount in the opposite direction. Then the value of the instrument would be equal to zero once aggregated. This paper avoids such a problem by focusing on daily prices of gasoline.

Methodologically speaking, this paper improves on Känzig (2021)'s identification scheme in two senses. They are both related to the way the list of news dates is formulated. First, his list consists of all the dates within the sample period on which there were formal announcements after the OPEC meeting. On one hand, this rule eliminates any room for a subjective judgement in choosing the dates. On the other hand, we believe only news that are "important enough" should be included in the list. If we start including news that are of minor importance, we face a risk of confounding the effects of the news in question and other types of news that happen to occur on the same day. This paper proposes a way to "prescreen" the candidate news before they are included in the list and filter out those that are deemed to be of lesser importance.

Second, the real news does not necessarily come on the day of an official announcement. For example, the media might report the likely outcome of an OPEC conference well before it concludes. In such a case, by the last day of the meeting, the outcome might be completely priced into the market. This paper tries to avoid this problem by scrutinizing newspaper articles before and during those meetings to see when the market perception changed (if any). It is also possible that important OPEC-related news come on occasions other than its formal meetings: for example, sometimes, an unofficial meeting between some major oil producing countries might turn out to be consequential. This paper expands the scope of search for relevant news to outside the official meeting periods. In addition, this paper looks at news other than those related to OPEC: specifically, it also studies news about the US government's policies on sanctions on oil imports from Iran.

3. Background (2) oil shocks and pass-through

This paper is also related to the literature on pass-through. Many existing studies examine pass-through from the exchange rate to domestic prices. For example, Shioji (2012, 2014, and 2015) study time-varying nature of the exchange rate pass-through in Japan. This study is more closely related to studies of oil price pass-through to domestic prices. For example, Chen (2009) uses time-varying parameter models to oil price pass-through to domestic CPI inflation for 19 developed countries. The rest of this review section will focus on those studies that concern domestic gasoline prices.

3.1. Oil Pass-through to domestic gasoline prices

Meyler (2009) performs a comprehensive study on oil price pass-through to domestic prices of liquid oil products in the Euro area, with weekly data. Blair et al. (2017) estimate error correction models using US weekly data, and find substantial differences in the extent of oil price pass-through across regions. Yilmazkuday (2019), using US weekly data, estimates a structural VAR model and conclude that the rate of oil price pass-through to domestic gasoline price is 13% after a week, and then goes up to 37% after three months, and goes up further to 50% in the long run. That is, more than two thirds of the adjustment occurs within the first three months. Such a finding of fast pass-through is confirmed in this paper as well. On the other hand, Chudik and Georgiadis (2019) develop a new mixed-frequency methodology and apply it to study the relationship between daily oil prices and weekly US gasoline prices. They find evidence for even faster oil price pass-through. The estimated pass-through rate within the first five working days is 23%, which goes up to 48% after twenty days.

In relation to the macro econometric literature mentioned in the previous section, Kilian (2010) extends the framework of Kilian (2009) by incorporating gasoline specific supply shocks (or refining shocks) and gasoline market demand shocks to study how oil supply and demand shocks affect US domestic gasoline prices. The data frequency is monthly. He finds that, although refining shocks are the most important driver in the short run, long run fluctuations in gasoline prices are dominated by the world demand shocks and oil market specific demand shocks.

3.2. Asymmetry in pass-through

A significant part of the literature has focused on the issue of asymmetry in gasoline price responses between times of rises and falls of oil

[1] The original idea behind this paper had been developed independently, without the knowledge of the paper by Känzig. However, I fully acknowledge that he had written up the discussion paper version first.

prices, known as the "rockets and feathers" hypothesis, named after the work of Bacon (1991). Classic studies include Bacon (1991), (Balke et al., 1998), Borenstein, Cameron and Gilbert (1997), Duffy-Deno (1996), Karrenbrock (1991) and Shin (1994). Godby et al. (2000) apply a threshold autoregressive model to the Canadian data and fail to find support for the hypothesis. Bachmeier and Griffin (2003) find results that contradict the hypothesis using daily data, but for wholesale prices. Chesnes (2016) utilizes daily data on gasoline prices from different cities in the US and find very strong asymmetry in their responses to oil prices. The degree of asymmetry is found to be quite heterogeneous across cities. Recent studies such as Polemis and Tsionas (2016) have come to employ sophisticated nonparametric approaches. Deltas and Polemis (2020) offer a comprehensive overview of the literature and conduct a thorough investigation of the issue using multi-country data with different data frequencies and various empirical models: they find that research design matters to the results we obtain.

3.3. Studies on Japan

For Japan, Shioji and Uchino (2011) estimate time-varying parameter VAR models to study the extent of pass-through from world oil prices to domestic prices, including gasoline prices, using monthly data. Yanagisawa (2012) estimates oil price pass-through to domestic gasoline prices in Japan and find asymmetry in responses.

3.4. Contributions of this paper

Compared to most of the studies in this literature, this paper takes identification more seriously. Instead of using raw data on crude oil prices, which are driven by a mixture of forces of very different nature, this paper focuses on only one of such kind (but a very important one, we believe), namely shocks to expectations about oil supply.

Use of daily data, which is still rare in this literature, is also crucial. As pass-through to gasoline prices is very fast, one may not be able to fully capture their dynamic responses with either weekly or monthly data.

4. Overview of the gasoline data

4.1. Dataset on daily gasoline prices

As already discussed in the introduction, data on domestic gasoline prices is taken from a major price comparison site on the web in Japan. This web site, called gogo.gs (URL: https://gogo.gs/; unfortunately, it is written in Japanese only), collects information from two sources. The first is registered users who happen to use a certain local gas station, or just happen to drive by one of those stations and saw prices quoted on a billboard. The second is the group of gas stations registered with the web site, that are eager to attract customers. The web site publishes the nationwide average every day. The sample period for this analysis starts from January 1, 2013 and ends in January 7, 2020. I use the nationwide average series for "Regular Fuel (meaning not high-octane), Cash-only (i.e., non-member prices)".[2]

4.2. Japanese gasoline taxes

The data we collect from the above source is computed from prices that users pay at gas stations, meaning that they include payments for taxes. Japan is known for its very high tax rates related to purchases of gasoline. More importantly, most of them are levied based on the quantity of gasoline purchased, not its value. As a consequence, this large tax component is insensitive to changes in the production cost, such as the price of crude oil. This feature dampens fluctuations in the price of gasoline in the eyes of the consumer. On the other hand, from the viewpoint of economists, if we use just this statistic, we could be underestimating the true degree of pass-through of oil price fluctuations to domestic prices set by retailers.

For this reason, in this paper, I analyze two alternative series of domestic gasoline prices. The first one is the unadjusted series that have been taken straight out of the data source mentioned above. As the second series, I construct gasoline prices adjusted for taxes, by estimating and then removing the effects of those taxes from the first series. The first series will be denoted as **gasNON**, while the second one will be called **gasADJ**.

In Fig. 1, I plot each of those series, together with the three most prominent indicators of crude oil futures prices, namely Brent, Dubai and WTI. To facilitate the comparison, all the three series are normalized so that they would be equal to 100 at the beginning of January 2014. We can see that, though both of the gasoline price series are highly correlated with crude oil prices, gasNON, the unadjusted price, exhibits much smoother movements. The tax-adjusted series, gasADJ, appears to be much more tightly connected with changes in crude oil prices. Comparing the two gasoline price series suggests the extent to which the Japanese tax system helps dampen volatilities in gasoline prices that consumers face.

4.3. How well do the new indicators track official statistics?

Before proceeding, we take a moment to examine how closely our gasoline price series line up with data from official sources. Official data on gasoline prices that are available at the highest frequency in Japan is weekly data on *Petroleum Products Retail Price*, which is collected by the Oil Information Center of Japan from about 2000 gasoline stations around Japan, via telephone, FAX, or through personal computers, etc. The data is collected on Monday of each week, with a few exceptions. From this data set, I have chosen to work with "regular" gasoline price with cash payment. This variable is denoted as "gasGOV". For the sake of comparison, weekly data on gasNON and gasADJ are constructed by taking averages of their daily values. A "week" here is defined to end on Monday, to be consistent with the way gasGOV is measured.

Fig. 2 (A) plots weekly averages of gasNON and gasADJ along with gasGOV. It can be seen that our indicator, gasNON, closely follows the official statistics, namely gasGOV. Table 1 (A) provides their summary statistics and Table 1(B) reports correlations among them, after taking

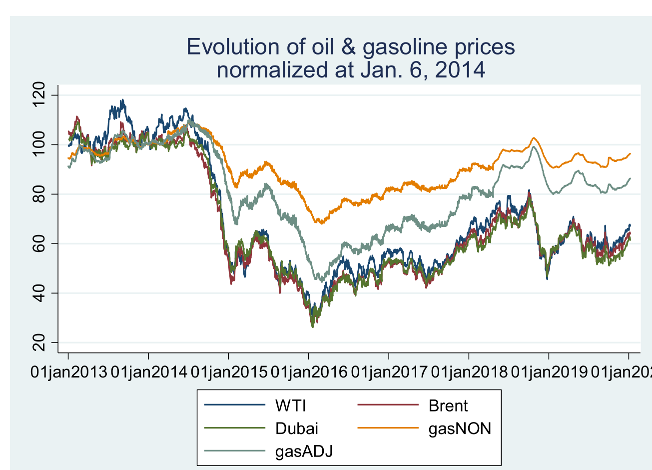

Fig. 1. Time series plot of various indicators of crude oil futures (originally in USD) vs gasNON and gasADJ (originally in JPY). Note: WTI, Brent and Dubai are crude oil futures prices. All are normalized to equal 100 at the beginning of January 2014.

[2] Between mid-July to the end of year 2018, I switch to the average of non-member and member prices, due to a technical problem. During this switching process, we lost a few days' observation by error. Those needed to be estimated through interpolation. The different series are connected using growth rates.

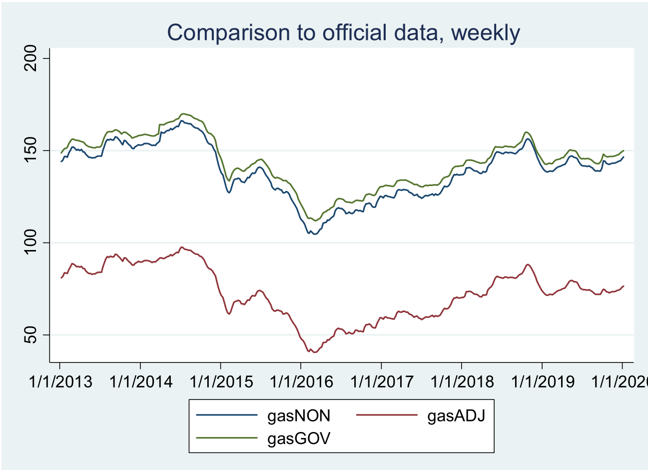

(B) Monthly data (Normalized to equal 100 in January 2014)

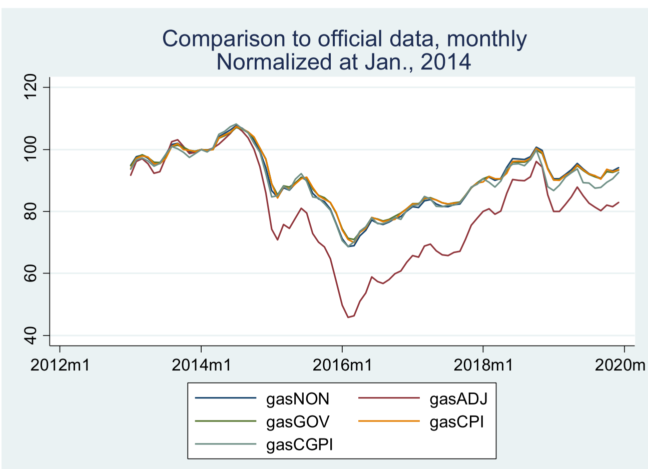

Fig. 2. National average gasoline prices: comparison to the official data.

their logs and computing differences of their values from four weeks ago. They exhibit similar characteristics except that gasADJ is more volatile, and they are highly correlated with each other.

At the monthly frequency, we can make use of the official statistics that are widely used in the empirical macro literature. In the case of Japan, gasoline price indices are available from both the Consumer Price Index (CPI) and the Corporate Goods Price Index (CGPI, a Japanese equivalent of the wholesale price index). The former will be denoted gasCPI, and the latter will be called gasCGPI. I create monthly series for gasNON and gasADJ by aggregating their daily observations up to the monthly level. I also take monthly averages of the weekly variable gasGOV.

In Fig. 2(B), the evolution of those three gasoline variables are compared to gasCPI and gasCGPI. Evidently, gasNON closely follows the movements in gasGOV, gasCPI and gasCGPI, again confirming the reliability of our daily variables. Table 2 (A) provides summary statistics for those five series, and Table 2(B) reports correlations among them, after taking their logs and computing differences of their values from one month ago. Again, their characteristics are very similar except that gasADJ is more volatile. They are also highly correlated with each other.

Table 1
Summary statistics, weekly data, log differences from four weeks ago.

(A) Mean, Standard deviation, Minimum, Maximum.

	Mean	Std.	Min	Max
gasNON	0.0000	0.0264	−0.0806	0.0586
gasADJ	−0.0010	0.0485	−0.1416	0.1194
gasGOV	−0.0001	0.0230	−0.0905	0.0473

(B) Correlation.

	gasNON	gasADJ	gasGOV
gasNON	1		
gasADJ	0.9848	1	
gasGOV	0.9756	0.9633	1

5. SVAR-IV

In this section, I review the SVAR-IV methodology using an example with just two endogenous variables, one external instrument, and one lag. For a full discussion of the methodology, the reader should refer to Stock and Watson (2018). Consider the following reduced-form VAR model with two endogenous variables denoted $y_{1,t}$ and $y_{2,t}$:

$$Y_t = AY_{t-1} + \nu_t \tag{1}$$

where.

$$Y_t \equiv \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix}, B \equiv \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix}, \nu_t \equiv \begin{bmatrix} \nu_{1,t} \\ \nu_{2,t} \end{bmatrix}$$

In the above, ν_t is a vector of reduced-form error terms which has no structural interpretation, and its two elements are in general correlated with each other. Assume there is the following linear relationship between this and a vector of structural disturbances, denoted ε_t, whose two elements are uncorrelated with each other:

$$\nu_t = B\varepsilon_t, \text{where } \varepsilon_t \equiv \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix} \tag{2}$$

Then, assuming invertibility, we can write:

$$\mathrm{Y_t = C(L)B\varepsilon_t} \text{ where } C(L) = (I—AL)^{-1} \tag{3}$$

Suppose that we are just interested in uncovering the effects of the first structural shock, $\varepsilon_{1,t}$. In such a case, there is no need to know all the elements of the entire matrix B: all we need to estimate is its first column. Suppose that we have an "external instrument" denoted Z_t, which satisfies the following two conditions:

$$(\text{Condition 1 : Relevance})\ E\varepsilon_{1,t}Z_t = \alpha \neq 0 \tag{4}$$

and

$$(\text{Condition 2 : Exogeneity})\ E\varepsilon_{2,t}Z_t = 0 \tag{5}$$

Table 2
Summary statistics, monthly data, log differences from a month ago.

(A) Mean, Standard deviation, Minimum, Maximum

	Mean	Std.	Min	Max
gasNON	−0.0001	0.0258	−0.0816	0.0446
gasADJ	−0.0012	0.0473	−0.1455	0.0962
gasGOV	−0.0002	0.0233	−0.0860	0.0395
gasCGPI	−0.0001	0.0290	−0.0850	0.0487
gasCPI	0.0000	0.0255	−0.0858	0.0465

(B) Correlation

	gasNON	gasADJ	gasGOV	gasCGPI	gasCPI
gasNON	1				
gasADJ	0.9831	1			
gasGOV	0.9780	0.9618	1		
gasCGPI	0.9210	0.9039	0.8623	1	
gasCPI	0.9519	0.9353	0.9797	0.8185	1

Then we obtain the following:

$$E\nu_t Z_t = \begin{bmatrix} b_{11}\alpha \\ b_{21}\alpha \end{bmatrix} \tag{6}$$

Normalizing b_{11} to be equal to 1, we can focus on estimating the coefficient b_{21}. The actual estimation proceeds as follows. First, estimate the following equation, using Z_t as the instrument:

$$y_{2,t} = b_{21}y_{1,t} + d_1 y_{1,t-1} + d_2 y_{2,t-1} + b_{22}\varepsilon_{2,t} \tag{7}$$

This produces our estimate for the coefficient b_{21}, which will be denoted as $\widehat{b}_{21}$. Next, we estimate the reduced-form VAR model in eq. (1) to obtain the following:

$$\widehat{C}(L) = \left(I - \widehat{A}L\right)^{-1} \tag{8}$$

Combining the two results produces our estimate for the h period ahead impulse response function to the first shock, of the form:

$$IRF_h = \widehat{C}_h \begin{bmatrix} 1 \\ \widehat{b}_{21} \end{bmatrix} \tag{9}$$

6. Construction of instruments

6.1. Overview

This paper utilizes the above methodology to estimate the effects of a shock to expectations about future oil supply. For that purpose, we need an external instrument which is correlated with such a shock but is uncorrelated with the other types of shocks. In this paper, this is constructed in three steps. The first of those is the prescreening process: we select candidates for periods during which an *important* new development about the world oil supply might have occurred. In the second step, we examine news reports during those periods and determine if such an event indeed occurred, and, if it did, on which day it was *first* reported. The third step is to measure the change in the perception of the market participants caused by the arrival of the news, assuming that it is fully reflected in the price of crude oil futures.

The main idea behind this three-step approach is that fluctuations in the prices of futures on the days thus selected would be dominated by the news about future oil supply, as long as the news are sufficiently important. Also, they are unlikely to be correlated with the other types of shocks (such as the world business cycles) in systematic ways. The use of daily data is crucial here: if we use lower frequency data (such as weekly, monthly or quarterly data), it becomes more likely that other important news occur within the same time interval, and our estimated oil supply shock series are more likely to be "contaminated" by those events.

In the existing literature, Känzig (2021) constructs an external instrument based on a similar idea. He lists up all the dates within his sample period (which is 1983–2017) on which OPEC issued a formal announcement at the end of its official meeting. He then measures changes in oil futures prices on those dates. In what follows, I will explain details of the three-step procedure and point out some differences between my approach and that of Känzig (2021).

6.2. What kind of events to focus on?

Before we get into the three-step procedure, the first thing to do is to think what types of events are likely to have large impacts on the market participants' expectations about future oil supply. Note that it is not necessary for us to come up with a perfectly exhaustive list of all the supply-related news during the sample period. This is because, with the SVAR-IV methodology, all we need to construct is an indicator which is *correlated* with the true supply shocks (and is uncorrelated with the other types of shocks), not necessarily the one that *represents* the entire sequence of those shocks.

The first type of news chosen here are those related to OPEC's decisions regarding the member countries' production of crude oil. This idea is basically the same as that of Känzig (2021). It has long been argued that the power of OPEC to control the world oil price had diminished significantly since their heyday between the 1970s and the mid-1980s. However, recently, most notably since around 2014, they seem to have regained their influences, at least to some extent.

The second type of news is related to the US-led sanctions on oil exports from Iran. This type of events have not been taken up by Känzig (2021): as my sample period (from early 2013 to early 2020) is much shorter and more recent than his (i.e., 1983–2017), it is more important for this analysis to include this kind of new developments in the world oil market. As Iran is a large oil producer, when their supplies are cut out of the market, it is likely to act as a large negative supply shock. During my sample period, the US administration under President Obama negotiated and eventually succeeded in lifting the sanction. Then the Trump administration, which was critical of the predecessor's move, eventually reversed the course and re-imposed the sanction. All of those events are likely to have affected trends in world crude oil prices.

6.3. Step 1: prescreening

The first step of the three-step procedure is to narrow down the list of candidate periods within the sample during which events related to either OPEC or the Iran sanction might have occurred. We need events that can be deemed "sufficiently important". This is because we need the news to be powerful enough to dominate influences of other news that might have occurred on the same day.

Känzig (2021), as explained earlier, uses *all* the dates on which there were formal announcements by OPEC and his approach thus involves no prescreening. This approach is free from any subjective judgement: one does not need to take any stance on what is important and what is not (or we could just "let the data tell" as I do in step 3). On the other hand, including too many unimportant events might increase the risk of our instrument to be contaminated by other events that happened to occur on those dates.

Here, to introduce some element of objectivity into this prescreening process, I rely on statistics on the number of frequencies at which a word (or a phrase) was used as a key word for internet searches on each day. For that purpose, I use *Google Trends*, which can be used to produce exactly that kind of statistics. It is provided by Google, a company which provides the world's most popular internet search engine.

The question is what kind of key words serve our purposes here. After some trial and error, I decided to use "OPEC" for the OPEC-related news and a combination of "Iran" and "sanction" for the news related to the sanction on Iran.

Fig. 3 presents the results. Panel A shows evolution over time of the frequency at which the key word "OPEC" was used for internet searches in the entire world on each day. For details, refer to the footnote beneath the Figure. Panel B is for the combination of key words "Iran" and "sanction". It is notable that, although both series appear quite noisy, there are clear "spikes" in them, which signal sudden surges in general interests in the topic. We could hope to be able to find occasions on which important news about those topics arrived among or around those spikes. Here, I define "spikes" in each panel as those that exceed two standard deviations above the mean.

6.4. Step 2: selecting specific dates (news report analysis)

The next step is to determine on which days (around those "spike dates") important news related to the two topics first arrived. This involves reading lots of newspaper articles and making judgements (admittedly, with this approach, it is difficult to completely eliminate all

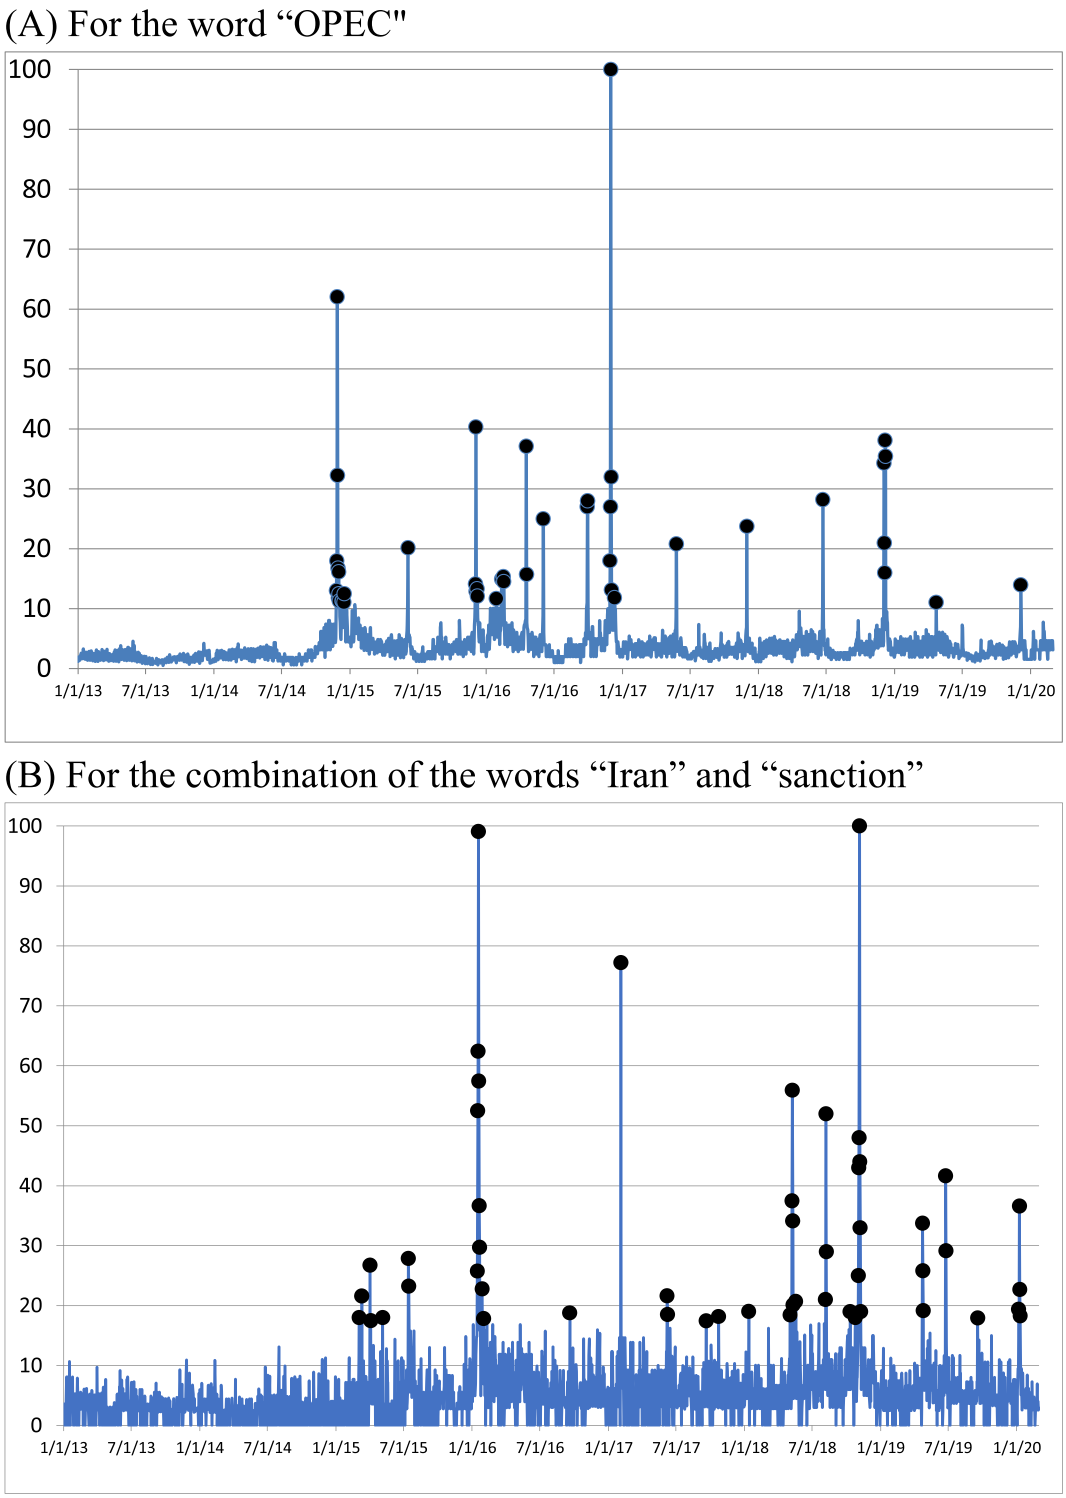

Fig. 3. Google Trends search results. Note: Based on the number of internet searches from the entire world according to Google Trends. Both series are normalized to equal 100 at the maximum. Black dots with dates are observations exceeding two standard deviations above the mean. The series are constructed in two steps. (1) We record daily search statistics for every eight months period which includes either the first or the second half of each year, such as December 2015 – July 2016 and June 2016 – January 2017. (2) Then the two adjacent series are linked at a day within the overlapping period when the value is the largest, using growth rates.

the aspects of subjective judgements). I relied mostly on Japan's *Nikkei* Newspaper (which is a major newspaper in Japan which puts emphasis on business and economic matters) and supplemented it with other sources such as Reuters and the *Financial Times*.

In contrast, Känzig (2021) includes only the dates of formal OPEC announcements. Again, this strategy has a virtue of avoiding any subjectivity. On the other hand, it sometimes happens that the content of the announcement is more or less known by the time it is made public: the real news had been conveyed by the media earlier. Also, sometimes, important decisions are made outside the formal OPEC meetings, such as an informal meeting between ministers of some OPEC countries and Russia.

Table 3
News date list (Rows in white = OPEC-related news, grey = Iran-related news).

Date	News
Nov. 25, 2014	OPEC's major member countries fail to agree to cut production.
Nov. 27, 2014	OPEC meeting fails to agree to cut production.
Apr. 02, 2015	Countries agree on a framework to resolve the nuclear issue.
Jun. 05, 2015	OPEC decides against cutting production.
Jul. 14, 2015	Final agreement on the Iran nuclear deal.
Dec. 04, 2015	Market expects OPEC will fail to cut production.
Dec. 07, 2015	OPEC meeting fails to agree to cut production.
Jan. 15, 2016	Iran sanctions will be lifted.
Feb. 16, 2016	OPEC agrees to freeze production increase [disappoints the market which had anticipated a production cut].
Apr. 18, 2016	OPEC meeting fails to agree to cut production on 04/17 (Sun). (Note: for this news only, I take log difference between the closing price of Friday (04/15) and the opening price of Monday (04/18).)
Jun. 02, 2016	OPEC meeting decides not to freeze production increases.
Sep. 28 2016	OPEC's informal meeting agrees to cut production.
Nov. 30, 2016	OPEC meeting agrees to cut production.
May 25, 2017	OPEC agrees to extend coordinated production cut [disappoints the market which had anticipated a longer extension.]
Nov. 30, 2017	OPEC agrees to further extend production cut [but the market had largely anticipated this decision.]
Jan. 12, 2018	US decides not to restart sanctions on Iran.
May 4-9 2018	US will leave the Iran nuclear agreement.
Jun. 22, 2018	OPEC agrees to ease production cut [disappoints the market which had anticipated a larger production increase.]
Jun. 26, 2018	US urges others to halt imports of Iranian oil.
Nov. 02, 2018	US makes exemptions for sanctions: the list includes Iranian oil imports by most participating countries.
Dec. 07, 2018	OPEC agrees to cut production.
Apr. 22 2019	US will not extend exemptions on Iranian oil imports.
Jun. 20 2019	Iran shoots down a US drone.
Sep.14 2019	A Saudi oil facility is attacked.
Dec.6 2019	OPEC and Russia agree to raise oil supply from January.
Jan. 3-7 2020	Soleimani assassination.

Table 3 lists the dates thus selected. In the table, rows in white represent news about OPEC. As can be seen, they include, for example, an announcement that they have agreed to jointly cut oil production, or news that they have failed to reach such an agreement. Compared to the list of news used by Känzig (2021), the most notable difference is that my list includes two dates in 2016, one in February and the other in April. The one in February corresponds to a meeting held by ministers in charge of oil related matters from Saudi Arabia, Russia, Venezuela and Qatar, which ended with an agreement to halt increases in oil production. But it was conditional on that all the other major oil producing countries would join this agreement. This led to another meeting in April, this time by all the major oil producing countries, including non-OPEC members. But the meeting failed to produce an agreement. Although those two were not the formal OPEC meetings, they were perceived as important events by the market.[3]

In the same table, rows in grey correspond to news about Iran sanctions. As can be seen, they are mostly about the US government's decisions on whether to lift the sanction or to withdraw from the nuclear deal.

6.5. Step 3: measuring market reactions to the news

The final step is to measure the market's reaction to the news. To understand the importance of this last phase, suppose, for example, that OPEC has just made an announcement that it would cut the member countries' production by 10%. If this was truly a big surprise, we would expect it will show up in the market prices: as the participants would come to expect oil prices to go up in future, prices of oil futures would jump up immediately. On the other hand, suppose that the market had fully anticipated this announcement beforehand. Then, assuming market efficiency, such an expectation would have been already incorporated into futures prices. As a consequence, on the day of the announcement, we would see zero reaction to the news. As a third possibility, imagine that the market had anticipated a production cut of 20% rather than 10%. In such a case, the above announcement should be regarded as a positive shock to the expected future supply of oil. We would expect futures prices to go down, rather than up, on that day.

We measure market response to a piece of news by the log change in the closing price of crude oil futures on the news date from the previous business day's closing price. This procedure is basically the same as the one in Känzig (2021), though the choice of crude oil futures is different, as discussed later.

6.6. Choice of crude oil price indicator

In our analysis, crude oil prices play dual roles. First, they are used to construct the external instruments. Second, they are included as an endogenous variable in the VAR. I have decided to use Brent for the former purpose. I will use Dubai for the latter. The main reason for choosing Brent for the construction of the news indicator is that it is currently considered as the most representative of global oil prices. Baumeister and Kilian (2016) argue, between pages 147 and 148, that WTI ceased to be the global representative indicator after 2011. Note that the entire sample period of this study falls in that range of time. The reason for the change is that WTI is now partly driven by the local supply condition in the market for light sweet crude oil in the central United States. And the market is now under heavy influence from the US shale oil production. Kilian (2016) offers detailed accounts on how the shale revolution has disrupted the US oil market, and created a gap between the US crude prices and the world prices represented by Brent.

The second reason is related to the timing of openings and closings of each of the markets. Fig. 4 illustrates this point. It indicates trading hours of the three major crude oil futures markets, for a non-daylight-saving-hours day. The scale above indicates the GMT and the one underneath is for the JST (Japan Standard Time). The Brent market is open for much longer hours than the other two. Its trading hours cover much of the time of the day when important news from OPEC and the Middle-eastern oil producer countries (and also Russia) are likely to come. It covers much of the business hours on the eastern side of the US, where some of the news related to sanctions against Iran is likely to come from. The market closes just before most Japanese start their business activities of the day. So its closing price (of the previous day from the Japanese viewpoint) is likely to reflect most of the important developments of the (previous) day without being contaminated by news emitted from Japan on the day.

On the other hand, Dubai is widely considered as the indicator which is most closely related to the price of oil exported to East Asia, including Japan. I therefore use its closing price, denominated in the USD, as an endogenous variable in the VAR.[4]

I also tried different types of crude oil futures for the construction of the IVs, but did not find much difference in the results.

6.7. Details

As both Brent and Dubai are futures contracts, at any given point in time, there are multiple types of contracts with different times to maturity are sold. For Brent, the media usually reports the price of the front (or near) month contract, as it tends to be most heavily traded. Dubai is peculiar in that it is the far month contract (i.e., ones with 6 months ahead delivery period) which has the largest trading volume. For each of the commodities, I follow those conventions. For Dubai, I make a switch to contracts with duration of 5 months when their trading volume surpasses those of the 6 months contracts toward the end of each month.

Unlike the gasoline price series, crude oil prices cannot be observed when the markets are closed. When there was no observation, I simply set its value equal to that of the previous day.

6.8. Resulting instruments (in the baseline specification)

I create two indicators, which will be used as external instruments in the baseline SVAR-IV estimation. The first is the market reaction to OPEC-related news, denoted IV1 (OPEC). The market reaction is measured as the log difference in the closing prices of Brent crude oil futures. Suppose that there was an important announcement from OPEC on November 30, 2016, say around 5 PM GMT (or 2 AM of the following day in Japan). Then IV1 for December 1 is equal to the Brent closing price on November 30 (or the morning of December 1 in Japan) minus the closing price of the previous day. On the other hand, if there was no news on that day, IV1 is equal to zero.

The second IV has to do with Iran-sanction-related news, and will be called IV2 (Iran), which is constructed in a similar way.

Fig. 5 plots both IV1 and IV2 together with the level of Dubai futures prices, measured in US Dollars. In the figure, IV1 is in red while IV2 is in green. Clearly, some news cause much larger market reactions than some others, which are captured by differences in absolute values of those indicators across different news dates.

7. Estimation details and results

7.1. Estimation details

Each of the SVAR-IV models estimated in this paper involves three endogenous variables and one external instrument. The three variables are:

[3] On February 17, 2016, the *Financial Times* put an article with headline "Saudis and Russia agree output freeze in bid to halt oil price slide" as the top news of the day. On April 18, 2016, it reports on its front page: "Deal to freeze oil production collapses after Saudi Arabia holds out over Iran".

[4] I use the Dubai futures price as reported by Tokyo Commodity Exchange. One could argue that it is more desirable to use the Dubai spot price, but I could not find a series that is available for a long enough period on a daily basis. The movement of the spot price is said to follow that of the futures prices.

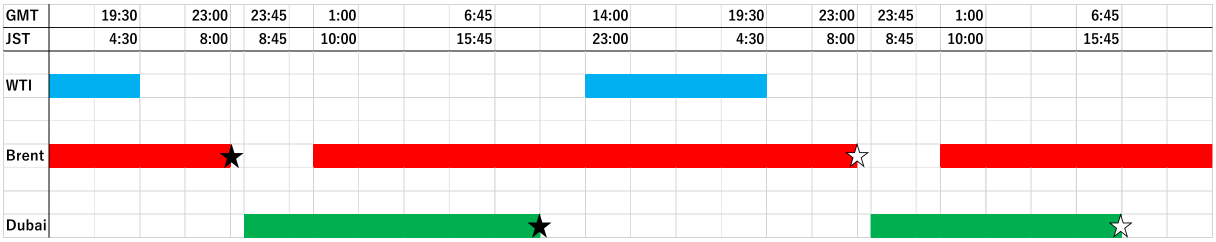

Fig. 4. Trading hours of the three major crude oil futures markets. (Note) "JST" stands for the Japan Standard Time. For WTI and Brent, hours are for the non-daylight-saving-time period.

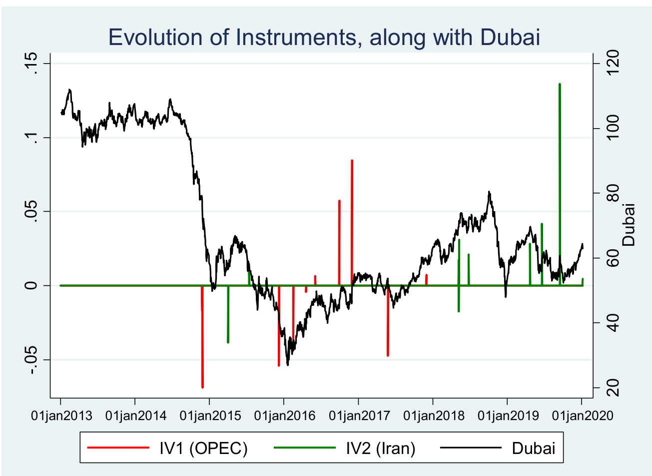

Fig. 5. Time series plot for External Instruments (along with Dubai). Note: "Dubai" = Dubai crude oil futures price (in US Dollars, normalized to equal 100 at the beginning of January 2014), "IV1 (OPEC)" = External Instrument 1 (related to news about OPEC), "IV2 (Iran)" = External Instrument 2 (related to news about the US-led coalitions' sanction on oil exports from Iran).

(1) World crude oil price = **Dubai**: log difference of the closing price in the Dubai crude oil futures market, denoted in USD, from that of the previous day.
(2) Exchange rate = **USDJPY**: the value of 1 US Dollar expressed in the units of the Japanese Yen (hence an increase in its value means a depreciation of the JPY), day's closing rate (in Tokyo).
(3) Gasoline price in Japan = **gasNON** or **gasADJ**, as explained in Section 4. Note that they are measured in the Japanese Yen.

We take log first differences of all three endogenous variables.[5] As for the External Instrumental Variable, I use either one of the following three:

[1] **IV1** (OPEC): Reaction of Brent crude oil futures prices to OPEC-related news.
[2] **IV2** (Iran) = Similar index for the news related to Iranian sanctions.
[3] **IV1** ± **IV2** = the above two are added up to create a single instrument.

In total, the baseline estimation consists of six different specifications (as we have two alternative gasoline variables and three alternative instruments).

The dataset is daily. It includes weekends as well as holidays when the markets for Brent or Dubai (or both) or the Tokyo foreign exchange market are closed. On those days when there was no observation for either Brent, Dubai or USDJPY, their values were simply set to be equal to those of the previous day. The Japanese gasoline data contains observations on all days, including weekends and Japanese holidays. The data starts from the beginning of January 2013 and ends on January 7, 2020. The number of lags is set at 14 days based on AIC.

The estimation is carried out by modifying matlab codes for Mertens and Ravn (2019), made public through Mertens' web site (https://karelmertens.com/research/).

I first compute the first-stage F-statistics with the standard error correction by (Newey and West, 1987), with the lag length chosen to be 8 following the standard recommendation. When the third endogenous variable in the VAR is gasNON, it was 84.03 with IV1, 104.43 with IV2, and 104.10 with IV1 + IV2. When gasADJ is used instead, it was 84.85, 104.61, and 104.25, respectively. As they far exceed the conventional threshold of 10, I conclude that weak instruments are unlikely to be a big issue in this study. I now turn to the estimated impulse responses to an identified oil supply shock, presented in Figs. 6–8.

7.2. Estimation results with IV1 (OPEC)

I start with a case in which IV1, the OPEC-related index, is used as the external instrument. Fig. 6 reports the impulse responses. Note that, although I take log first differences of all the endogenous variables, the reported impulse responses are cumulative ones. Hence, they can be interpreted as the responses of the levels of those variables. In each of the panels, the solid line represents the point estimates, while the dashed lines are the 95% confidence bands. The bands are computed by a moving block bootstrap method proposed by Jentsch and Lunsford (2019), which is a variant of a method developed by Brüggemann et al. (2016). They are based on 5000 draws. All the responses are normalized so that the initial response of Dubai is equal to 1. As a consequence, the underlying shock should be interpreted as a negative shock to oil supply. Responses of up to 84 days, which is equal to 12 weeks, or about three months, are shown.

Panel A shows the response of Dubai to the identified oil supply shock. It is based on the VAR model with gasNON as the third endogenous variable, but the result is very similar when gasADJ is used instead. The shock has an immediate and permanent effect on Dubai. In fact, the response immediately reaches the peak and stays flat thereafter. In panel B, the same shock also has an immediate positive impact on USDJPY, the exchange rate. Note that this variable is defined in such a way that an increase in its value signifies a depreciation of the Japanese Yen. As Japan imports almost all petroleum it needs from abroad, it is understandable that a negative shock to oil supply weakens its currency. However, the response is insignificantly different from zero.

Panel C shows the response of gasNON. Note that, as the response of Dubai shown in Panel A is basically flat at 1, this response can be considered as the "pass-through rate" (defined as the ratio between the percentage increase in a domestic price variable and that of a variable that represents a foreign cost), approximately. The response is practically zero at the beginning, but turns significantly positive in just

[5] The Johansen test rejected presence of a cointegrating relationship at any conventional level of statistical significance. In the appendix, we estimate the SVAR-IV model in the levels specification to check robustness of the main results.

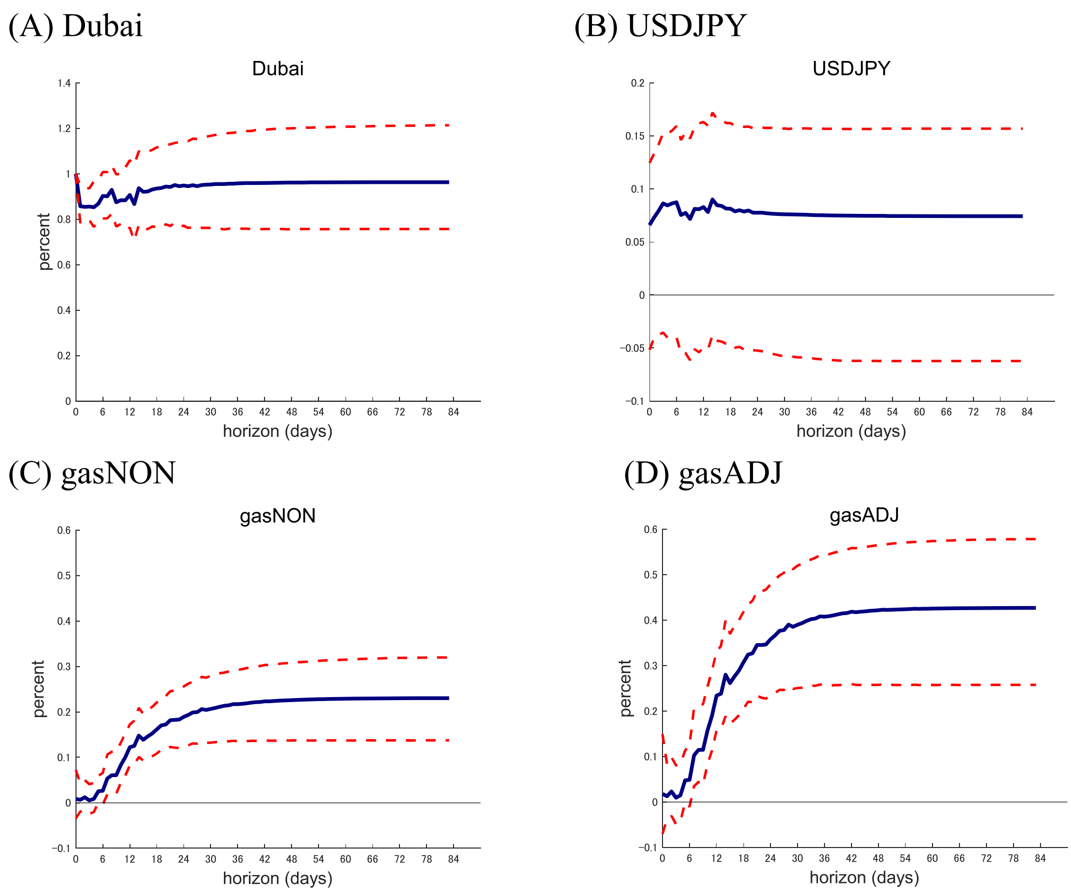

Fig. 6. Impulse responses to an oil supply news shock, case with IV1 (OPEC). Note: All the endogenous variables in the VAR are in log first differences, and their cumulative responses are shown here. The initial response of Dubai is normalized to equal 1. Solid lines are the point estimates, and dashed lines are the 95% confidence bands based on 5000 bootstrap draws.

seven days. The long run response is around 0.22. What is most notable is how fast this pass-through occurs over time. Out of the entire long run response of 0.22, over 30% of it would be completed in just ten days after the shock. This number goes up to over 70% on the 19th day, and crosses the 90% mark on the 31st day. By the 56th day after the shock, 99% of the price adjustment would be completed.

The fact that pass-through is so rapid suggests that there is a high value in the usage of daily data such as the one utilized in this paper. A conventional analysis which typically relies on monthly data would not be able to capture this entire dynamics, much of which is completed within the month after a shock hits the global market.

Panel D shows the response of gasADJ, when this variable is used in place of gasNON as the third variable in the VAR. The shape is similar to the one for gasNON that we saw in Panel C, but the size is much larger. To facilitate the comparison, I use the same scaling for the two panels. The long run pass-through rate for gasADJ is 0.41, which is nearly double the estimated value for gasNON. The difference is explained by the feature of the Japanese system of taxes on gasoline. As discussed in Section 2, most of those taxes are proportional to the volume of gasoline that one purchases, not the value. As a result, the tax component of the price paid by the consumer is insensitive to changes in the cost of producing gasoline. The above comparison reveals the extent to which

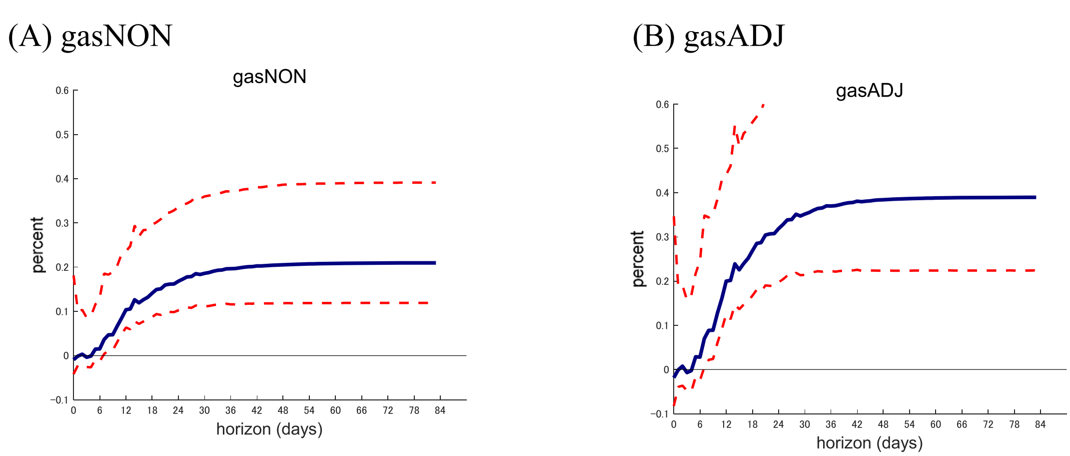

Fig. 7. Impulse responses, case with IV2 (Iran). Note: refer to the note for Fig. 6.

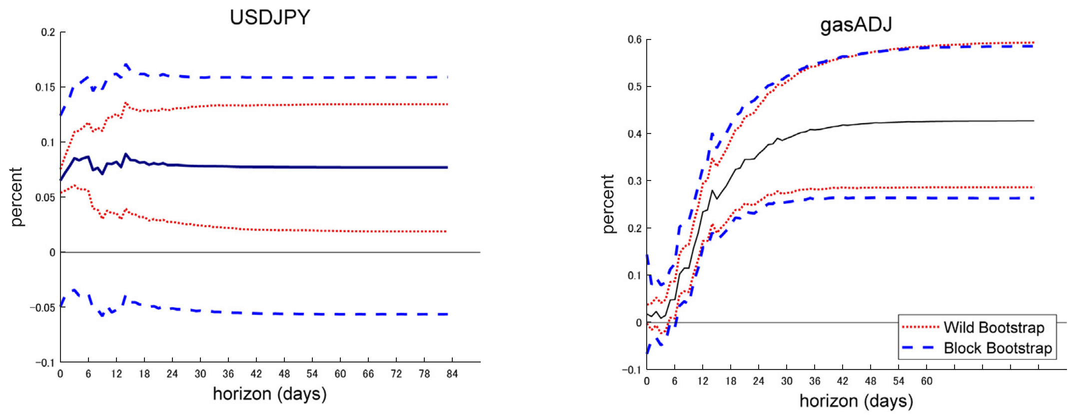

Fig. 8. Comparing confidence bands derived from different bootstrap methods. Note: refer to note for Fig. 6.

this dampens the fluctuation of the prices that the Japanese consumers are paying at gas stations.

On the other hand, even with this tax-adjusted series, it remains true that the pass-through is fast: almost all the above comments on this issue apply here, with a very minor difference that it is now the 28th day after the shock that over 90% of the long run pass-through is realized, as opposed to the 31st day.

7.3. Estimation results with IV2 (Iran) and IV1 + IV2

Fig. 7 presents the estimated responses of gasNON and gasADJ when IV2, which is based on the Iran-related news dates, is used as the external instrument. Although they are broadly similar to those in Fig. 6, the confidence bands around the responses of both gasNON and gasADJ are slightly wider and a little more asymmetric. This might indicate that IV2 is less effective as an external instrument compared to IV1. As sanctions usually come in packages, they are rarely purely about oil. In addition, those events tend to have broader implications about the world's political as well as military stability. For example, if imposing a new sanction is perceived to have increased the degree of uncertainty in the global market, it might reduce the worldwide demand for crude oil.

The case with IV1 + IV2 as the external instrument turns out to be quite similar to the one with IV1. I omit the impulse response figures to save space.

7.4. Effects of different methods to draw confidence bands

As mentioned earlier, this paper uses the moving block bootstrap method to draw the confidence bands. It has been proposed by Jentsch and Lunsford (2019) as an alternative to the wild bootstrap method of Mertens and Ravn (2013). Fig. 8 explores how the results look different between the two approaches. Starting from Panel (B), the results for gasADJ appear to be broadly similar across the methods. Bands from wild bootstrap are slightly wider, which seems consistent with the view that this method leads to underestimation of the true degree of uncertainty.

A much more important difference emerges in panel (A), which reports the response of USDJPY. The bands are much wider with the block bootstrap method. In fact, with the wild bootstrap, one would conclude that the response is significantly positive. The panel thus seems to suggest the merit of taking this methodological debate seriously.

8. Evidence from weekly and monthly data

In this section, the preceding results are compared to the cases in which the data is aggregated up to either weekly or monthly frequencies.

8.1. Estimation with weekly data

Weekly data on Dubai, USDJPY, gasNON and gasADJ is constructed by taking averages of their daily values. I also use gasGOV, the official statistic on weekly gasoline prices that was introduced in section 4 (see also Fig. 2(A)). I also take weekly averages of the daily IVs. A "week" here is defined to end on Monday, to be consistent with the way gasGOV is measured.

I redo the VAR analysis of the previous section with this weekly data set. The number of lags is set equal to 2 based on AIC. The first-stage F-statistics, with gasADJ as the gasoline variable and IV1 as the instrument, and with the Newey-West lag of 5 (which was chosen following the standard formula), was 17.96. This is much lower than the daily data case, but is still higher than the conventional threshold of 10.

Fig. 9 shows the impulse responses, estimated with IV1 as the instrument. The results are similar with the other instruments. Panel (A) is the response of Dubai, which is slightly larger in medium to long runs compared to the daily data case. In Panel (B), the response of the exchange rate remains insignificant. Panel (C) is the response of gasNON. It is qualitatively similar to the case of the daily data. The size of the response is larger, probably reflecting the larger response of Dubai. The bands are now wider. This seems to indicate that identification becomes less sharp with weekly data. In Panel (D), the third endogenous variable in the VAR is replaced by gasADJ. Again, the results are similar, qualitatively, but the bands are wider. In Panel (E), I use gasGOV in place of my gasoline price indicators. The estimated impulse response is quite similar to the one in Panel (C).

8.2. Estimation with monthly data

As already discussed, at the monthly frequency, we can make use of the two additional official statistics, which we called gasCPI and gasCGPI. I build a monthly data set on Dubai, USDJPY, gasNON, gasADJ, and the IVs by aggregating their daily observations up to the monthly level. I also take monthly averages of the weekly variable gasGOV.

Fig. 10 reports the impulse responses from the monthly data analysis. The number of lags is set at 1 based on AIC. Now the error bands are really wide, especially in medium to long runs, and they are noticeably asymmetric. Although the responses are "significant", technically speaking, it seems difficult to put much faith in the results. This seems to point to a failure of identification.

The reason behind these disappointing results seems to be a weak instrument problem. We have found that the F statistics for the first stage regression analyses goes down considerably when we move to monthly data. For example, with gasADJ as the gasoline price variable and IV1 as the instrument, and with the Newey-West lag of 3 (again following the customary formula), the above-mentioned statistics was is 7, which is below the conventional threshold of 10.

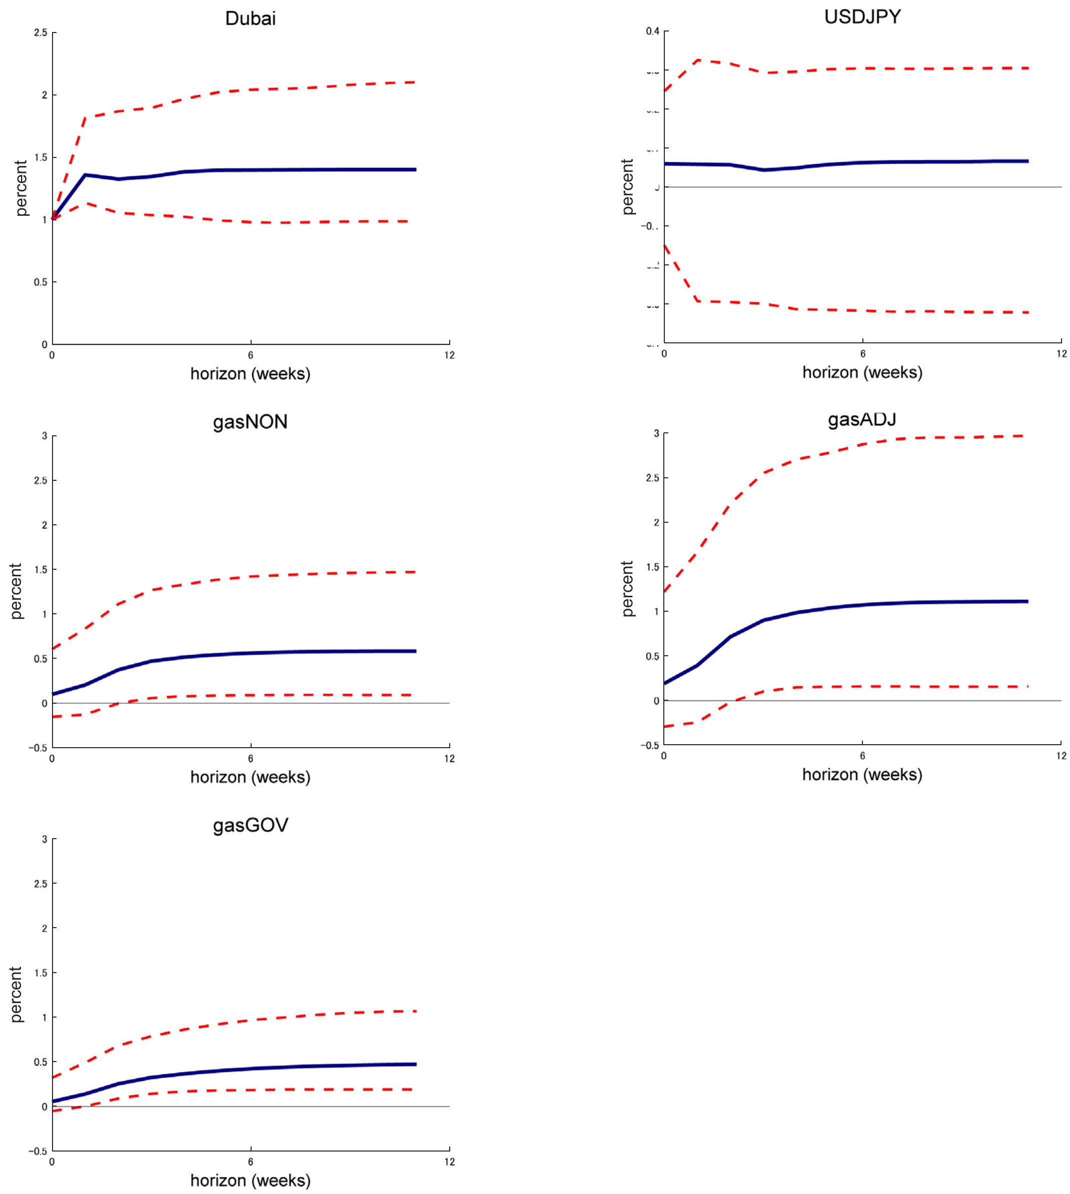

Fig. 9. Impulse responses estimated with weekly data, with IV1 (OPEC). Note: refer to note for Fig. 6.

8.3. Discussion

The above result suggests that there is a potentially sizable advantage in the use of data at relatively high frequency (to the extent that they are available). Even if the IVs are based on high frequency information and are not contaminated by other types of shocks, many other types of shocks hit the economy (in this case, the world oil market) within a month. Hence, the correlation between the IVs and innovations to the endogenous variables in the VAR model can be weakened considerably. This turns out to be the case here. Our study thus shows that utilizing high frequency data, when available, can be quite beneficial in macroeconomic research.

9. Evidence from micro level data

9.1. On the data set

As stated, the gasoline price data that we have used so far comes from a price comparison site on the web. The site also provides information on prices at individual gasoline stations in Japan. I will now explore what can be learned from this vast amount of information.[6]

As I started collecting this data only recently, the data set is much shorter along the time dimension. Concretely, the shop-level data spans the period of 535 days, between August 6, 2018 and January 7, 2020. Number of shops which had at least one observation during this period is 21,932 (in comparison, according to the Agency for Natural Resources and Energy of Japan, as of March 2019, there were 30,070 gasoline stations in Japan, and around 29,000 of them were registered with the web site). The total number of observations in the data set is 657,534. This means that the average number of observations per shop (with at least one observation) is 30. This implies that there are many missing values in the sample. I have thus given up the idea of

[6] This micro level data set was originally constructed for an ongoing project with Shiro Yuasa (Shioji and Yuasa (2020)). I thank him for allowing me to use this data for this paper.

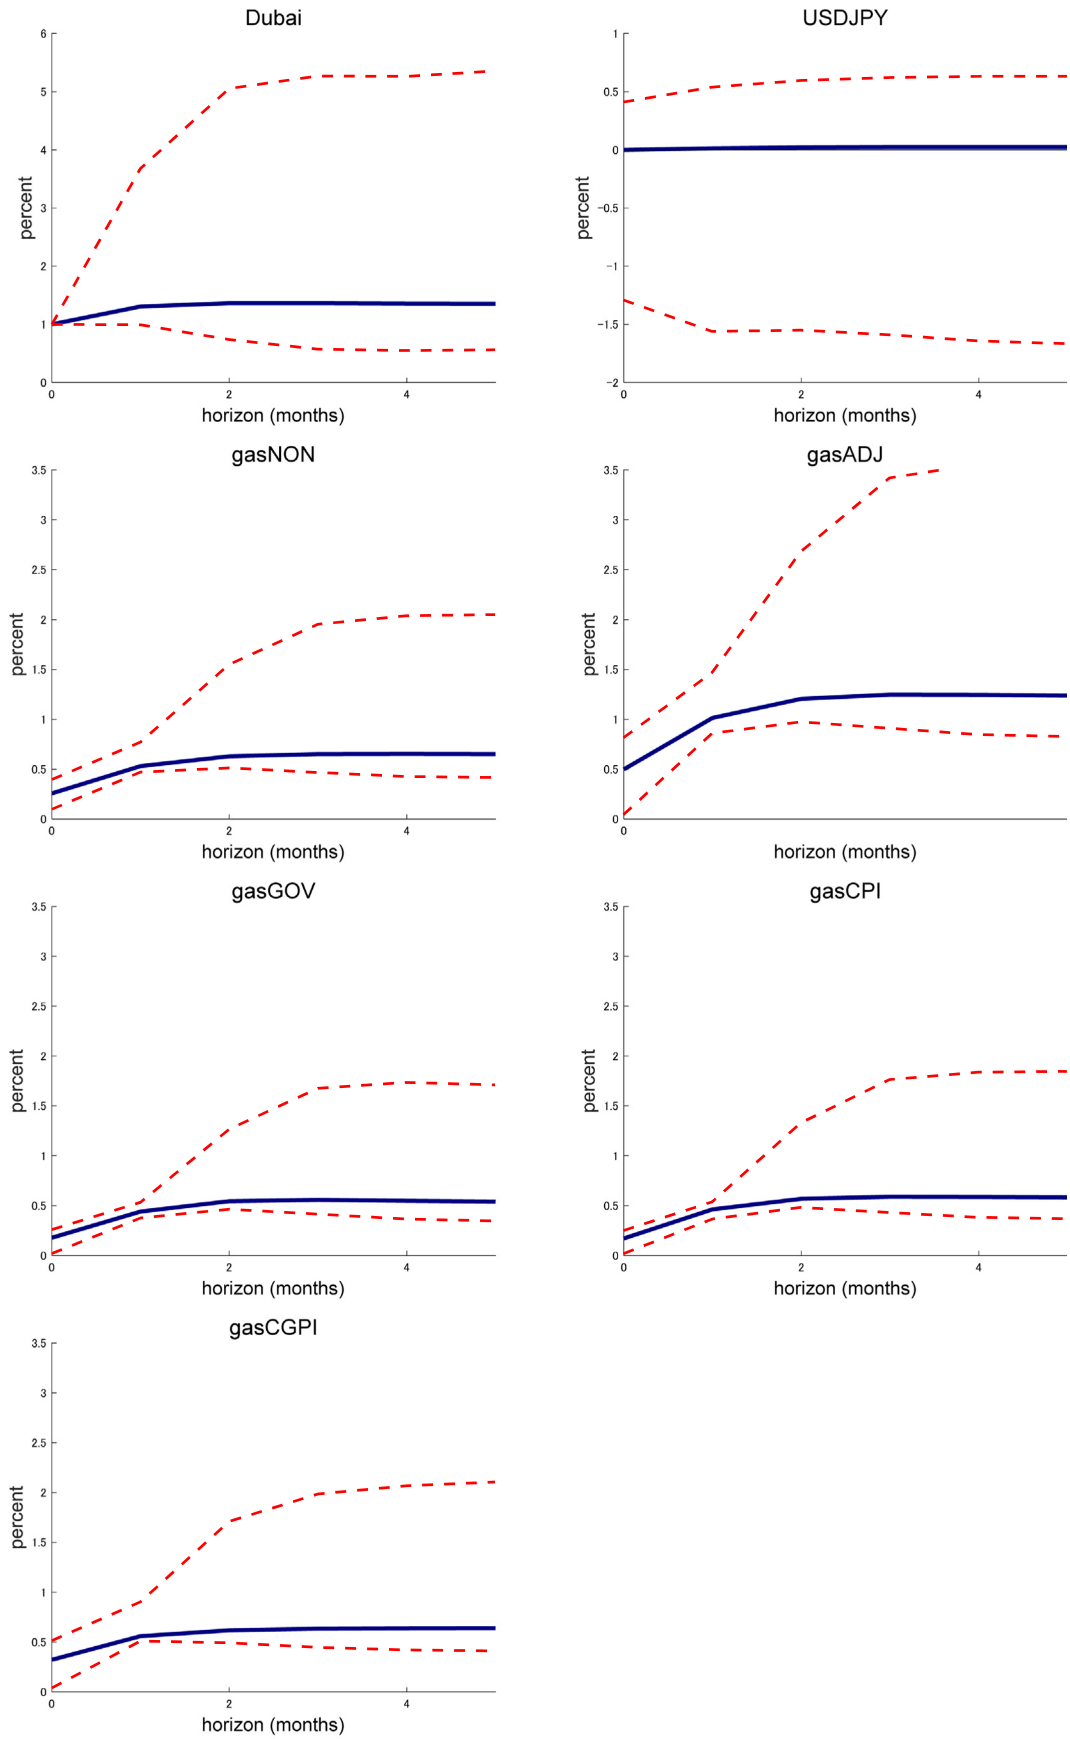

Fig. 10. Impulse responses estimated with monthly data, with IV1 (OPEC). Note: refer to note for Fig. 6.

applying a panel data technique to this data, and decided to treat it as a repeated cross section data.

It is beyond the scope of this paper to explain all the detailed features of this data set. Interested readers are invited to refer to Shioji and Yuasa (2020). Here, I will report some results that are more relevant to the main analysis of this paper.

It is of interest to see if one of the main results of this paper, namely the finding of a fast pass-through, is driven by a certain set of gasoline stations with similar characteristics. In this section, I investigate this issue along three dimensions. First. I ask if the tendency of a fast pass-through is uniform across regions. For example, I ask if stations in urban areas tend to respond faster to an oil price shock compared to non-urban shops. Second, I group the stations based on their affiliation with corporate groups. I ask if those shops that belong to big company groups tend to be quicker in their responses to an oil price increase. Third, I will try to see if there is any difference based on other characteristics of the shops.

9.2. Pass-through estimation method

9.2.1. First stage: pooled OLS

The estimation proceeds in two steps. In the first stage, I run a pooled OLS regression in which the dependent variable is the station and the explanatory variables are shop characteristics, such as regions which the shop is located in, corporate group affiliations, if the shop is located on highways or not, and dummies for various types of services provided at the shop, including if it is a self-service station. After purging the effects of those shop specific features, I extract the residuals from this regression, and move on to the second stage. However, whether I include this first stage analysis or not seems to have virtually no impact on the second stage results. I also note that, in this process, I limit the sample to those shops with at least 15 observations. As a consequence, the overall number of sample for the pooled regression was down to 599,467 with 10,863 shops included.

9.2.2. Second stage: computing group averages

In the next step, I compute group-specific means for each of the days within the sample period. For example, I compute the means of the residuals from the first stage estimation across shops that are located in central Tokyo. This produces a time series of the mean gasoline price for that region (exclusive of the effects of shop specific characteristics). I apply the same procedure to compute the national average.

9.2.3. Third stage: local projection estimation

The final step is a time series estimation. As the data is too short along the time dimension, I could not apply the news-based identification strategy that I conducted in the previous sections: there were too few pieces of oil supply related news during this sample period. I thus adopt a simpler approach, in which the shock variable is the crude oil price itself. The estimation is based on the Local Projection method of Jordà (2005). For that, I obtain and modify the matlab code provided by Jordà on his web site. The left hand side variable is either the national average or one of the group means from the second stage. On the right hand side, I include the daily Dubai futures, to estimate the response of the left hand side variable to this variable. I also include USDJPY. All the variables enter in their log differences. I considered the possibility of including lagged dependent variables, as well as lagged values of Dubai and USDJPY, on the right hand side. I considered 0, 7, 14, 21, and 28 as candidates for the number of lags. AIC preferred 7 and BIC suggested 0. Considering the shortness of the sample period, I decided to go for 0, meaning not including any lagged variables. But the results are almost identical if I include the lags, including very long ones.

9.3. Regional characteristics and pass-through

I omit the results from the first stage pooled OLS, and go directly to the third stage results. In each of the panels in Fig. 11, solid line is the cumulative response of the national average, and the grey area is the 95% confidence bands. Hence, they are identical across the panels. The response is insignificant initially, but quickly turns positive and significant. It eventually becomes flat. Those features are qualitatively similar to the findings of the previous sections, though the point estimates are slightly larger, the bands are wider, and it takes a little longer for the response to flatten.

In panels (A)-(C) of the figure, the dashed line is the response of the mean of gasoline prices across shops that belong to a certain regional category. In panel (A), I show that for the shops in the urban areas in Japan. They consist of the entire prefectures of Tokyo and its three neighbors (i.e., Saitama, Chiba and Kanagawa), Aichi (which includes the big industrial city of Nagoya), Osaka and its two neighbors (i.e., Kyoto and Hyogo), and all the "designated cities" outside those prefectures. The response turns out to be almost identical to that of the national average. We can see that there is practically no difference between the response of urban shops and their non-urban counterparts.

That is not to say that there is no cross-regional heterogeneity. In panel (B), I show the response of the mean of the shops located in central Tokyo (its 23 "wards"), which is the most densely populated area in the country. It is slightly faster than the national average, though the point estimate falls within the confidence bands for the most part. Panel (C) presents the response of Hokkaido, which is the most sparsely populated area. Again, the response is slightly faster than the national average.

Hence, I conclude that there are some signs of cross-regional heterogeneity in the speed of price adjustment. However, it is not simply a difference between urban areas and non-urban places. I intend to investigate the determinants of those different regional characteristics in future research.

9.4. Corporate group affiliations and pass-through

I now shift my attention to corporate affiliation of each shop. The Japanese retail gasoline market is characterized by a relatively high degree of vertical integration. That is, much of gasoline is distributed to individual gasoline stations directly from one of large refinery groups. This is in contrast to the US market in which "jobbers" who specialize in distribution play important roles. Due to repeated mergers between those refinery groups, the market is increasingly being concentrated: practically speaking, there are only five at the moment. According to *Nikkei Value Search*, as of March 2020, over 70% of all the gasoline stations in Japan were affiliated with one of the three major petroleum refinery-cum-distributors. Top two groups had a combined share of above 80% in sales. On the other hand, "private brand" gasoline stations consist of those that are owned by either a large General Trading Company or the Japan Agricultural Cooperatives, and others which are mostly shops that have no corporate affiliation (I shall call them independents).

It is of interest to see if shops with different types of affiliation exhibit different behaviors. For example, Deltas (2008) argues that a gasoline station that faces a less severe local market competition is slower to adjust its price. As shops that belong to large groups are more likely to be insulated from competitive pressures, their prices might be more sluggish. On the other hand, as they are more likely to be able to obtain high-quality information about the prospects for future cost movements from the big corporations, they might be faster to respond.

In Panel (D) of Fig. 11, I report the response of the average price across shops that belong to the largest three corporate groups. It is undistinguishable from that of the national average. Panel (E) of the same figure shows the response of the mean price across the

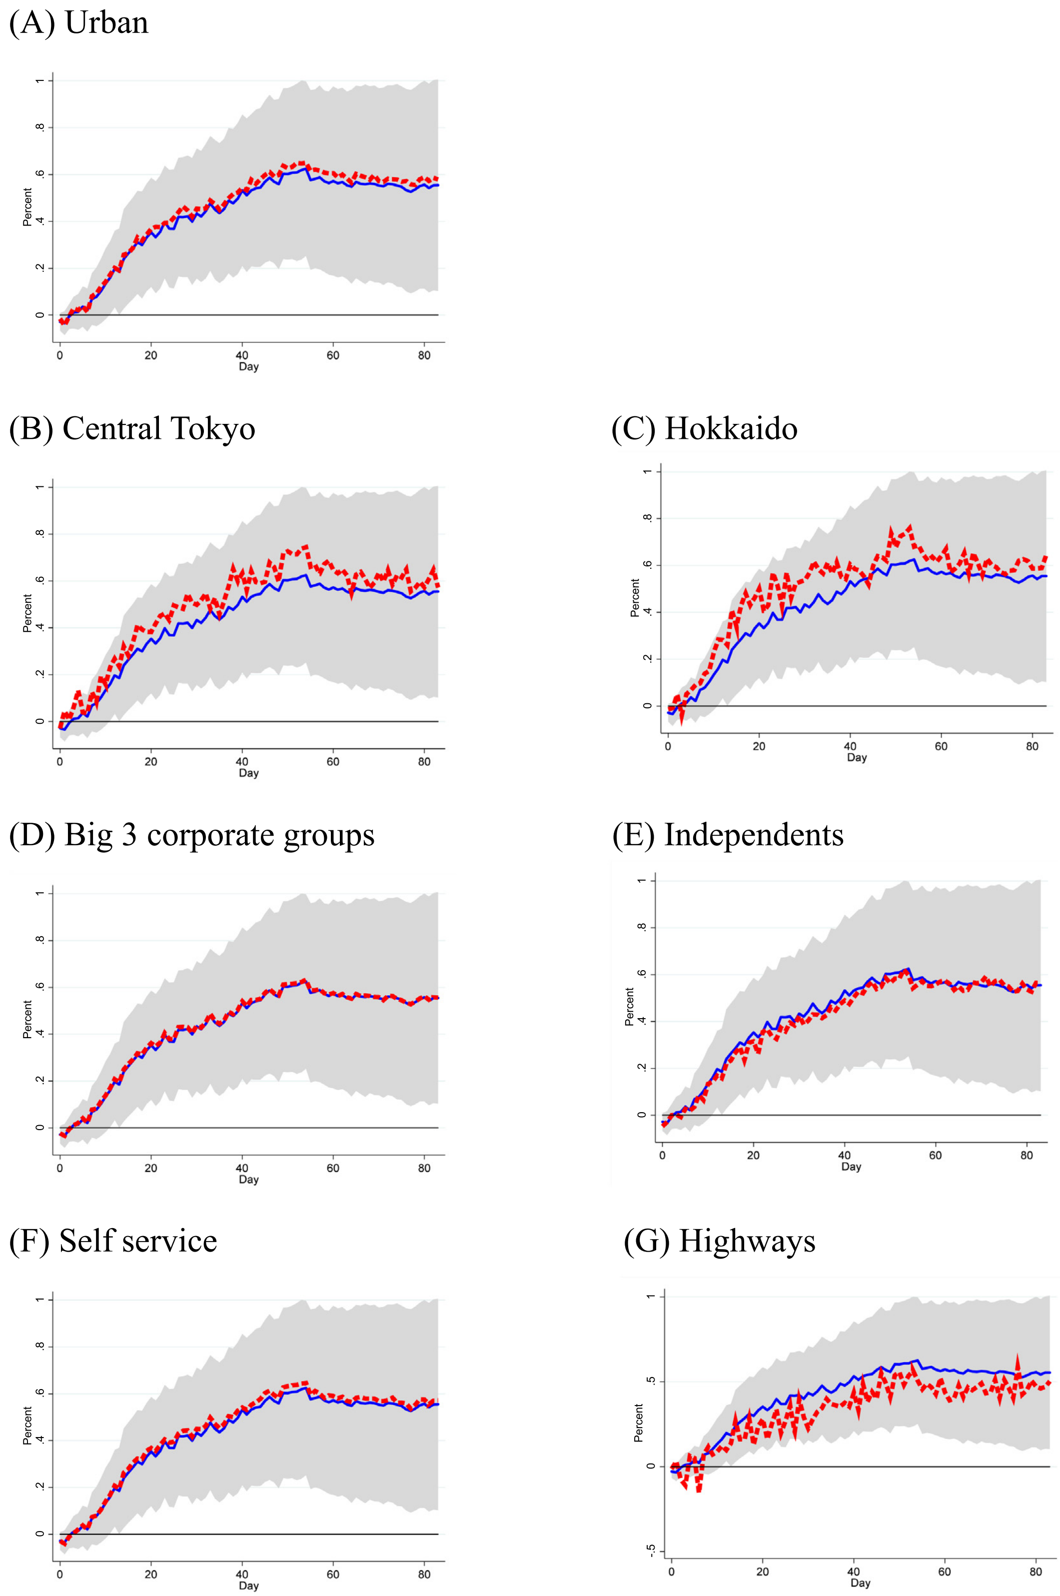

Fig. 11. Responses of group average gasoline prices to an oil price shock. Note: Cumulative responses to a 1% increase in Dubai. Solid lines are for the national average. Dashed lines are for Group averages. Grey areas are 95% bands for the national average. Gasoline prices exclude taxes. Estimated by the Local Projection Method.

independents. Again, it is virtually the same as the national average. It seems that whether a shop belongs to a big corporate group or not does not matter to the price adjustment speed. I also tried similar estimations for each corporate group, including smaller ones. Although I did see some differences across the groups, it was not possible to detect any consistent pattern.

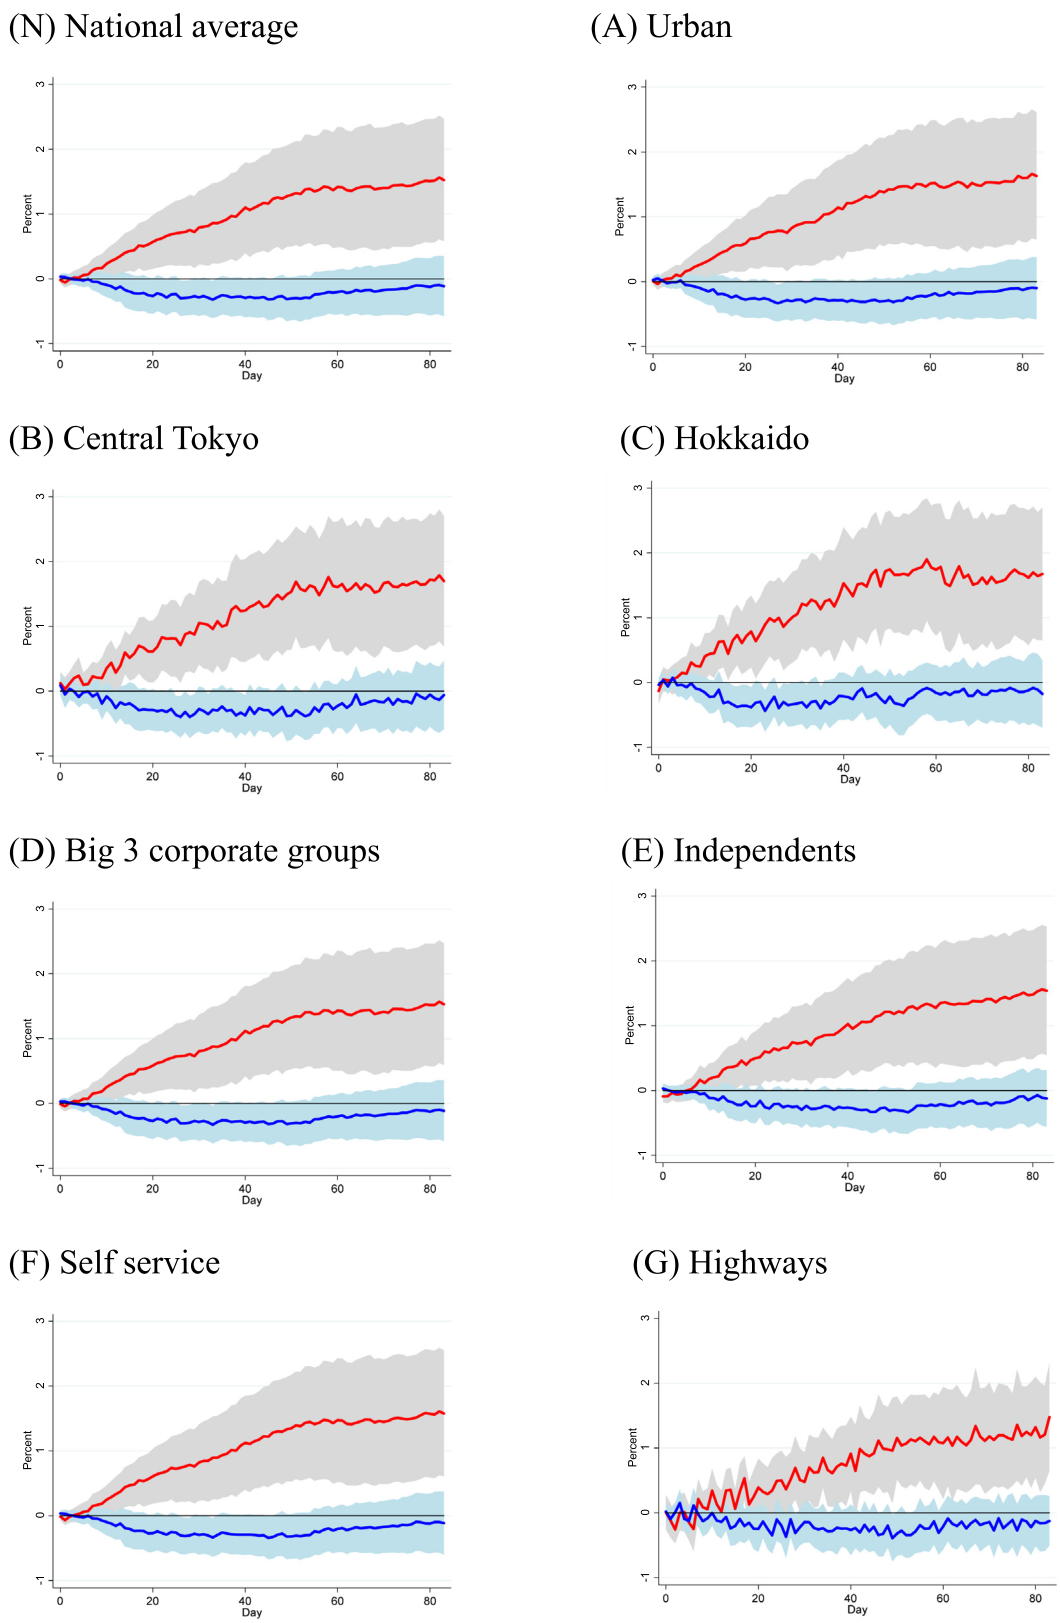

Fig. 12. Allowing for asymmetry in responses. Note: Lines that are (mostly) in the positive areas are for the high oil price period, and the others are for the low oil price period. “High” (“low”) means when the log of oil price is above (below) its sample median.

One possible explanation for this surprising similarity is that even the independents need to rely on one of the large distributors for supplies of gasoline. Hence, they are in more or less the same situations as far as material costs are concerned. Also, as they talk to the same set of distributors irrespective of their formal affiliation, the amount of information they can obtain might not be all that different.

9.5. Roles of other shop characteristics

Panel (F) of Fig. 11 presents the response of the average price at self-service gasoline stations. Once again, it is indistinguishable from that of the national average. Although their gasoline tends to be cheaper on average, percentage responses of their prices are no different from the others. In Panel (G), I show the response of the average price across shops on highways. They are identified by their names: if a shop's name contains words such as SA = "Service Area" or PA = "Parking Area", it is classified as a highway shop. Apparently, those shops are very distinct. The response is noticeably slower compared to the national average. As each of those areas tends to be managed by a public corporation, it is reasonable to expect that those shops are subject to less severe competition. The slow response to a cost shock is probably a result of this unique situation.

9.6. Asymmetry in responses

This sub-section asks if there is any asymmetry in the response of gasoline prices depending on the underlying movements in the cost of oil. An important advantage of our rich data set is that it enables to ask if there is any heterogeneity in the degree of asymmetry in those responses depending on the shop characteristics. For example, we could ask if independents exhibit a higher degree of asymmetry compared to those affiliated with large corporate groups.

The Local Projection estimation of the previous sub-sections is extended to allow for such asymmetry. This is done by replacing the oil price variable on the right hand side by a pair of cross terms between this variable and a dummy variable, each corresponding to one of the two sub-periods that divides the entire sample. We have tried to split the sample in two ways. First, we divided the sample based on the direction of change in crude oil prices. That is, it was divided between the periods when the rate of change in Dubai (its long moving average, to be exact) was positive and when it was negative. Second, we split the sample based on the level of crude oil prices. That is, it is divided between the time when the log of Dubai was above the sample median and when it was below the median. The first approach, which is more in line with the traditional "rockets and feathers" studies (Bacon (1991)), did not result in a clear-cut picture. If anything, responses were larger when oil prices were going down, which is the opposite to the conventional wisdom. In what follows, we focus on the second approach, which is based on the level of oil prices.

In each panel in Fig. 12, the estimated responses when oil prices are high are shown in the same way as in the previous figure. On the other hand, those for the period of low oil prices are shown with the signs flipped, so that the figures would not look too "congested". Hence, in each panel, a line that are mostly in the positive region is for the high oil period, and the other one that is mostly in the negative region is for the time of low oil price, except that the signs are flipped. The first panel denoted (N) is for the national average. The rest is ordered in the same manner as in the previous figure.

There is a very clear sign of asymmetry in all the panels. The response is very strong and significant for the high oil price period. It is much weaker and mostly insignificant when oil prices are low. Quite importantly, we see no difference in the degree of asymmetry across different groups of shops. For example, the response of independents in panel (E) is just as asymmetric as that for the stations affiliated with large corporate groups in panel (D). Hence, we can safely conclude that the asymmetry we observe in the national average (panel (N)) is not an artifact of aggregation.

10. Conclusions

Section 7 of this paper has utilized a new daily dataset on gasoline prices in Japan to estimate the degree of pass-through from the world oil prices to the domestic prices. To disentangle the intertwined relationship between supply and demand which are reflected in observed oil prices, and to identify shocks to the expectations about the future oil supply in the world, this paper has constructed two new indicators. They are based on responses of prices of crude oil futures to arrivals of news related to either (1) OPEC's decisions about future oil supplies, or (2) decisions regarding the US-led efforts to impose sanctions on crude oil exported from Iran. Each one of them has been incorporated into the SVAR-IV model as the external instrument, in turn, to identify shocks to expected future oil supply. The estimation results have indicated the following:

(1) The identified oil news shock has a significantly positive and permanent effect on prices of gasoline in Japan.
(2) The long run pass-through rate to the prices actually paid by the Japanese consumer is around 20%.
(3) But this modest degree of pass-through is, to a large extent, due to a feature of the Japanese gasoline-related taxes, which are mostly linked to the volume of purchases rather than the value. For the non-tax component of the gasoline prices, the long run pass-through rate jumps up to almost 40%.
(4) The pass-through is fast. About 70% of the entire process occurs within 18 business days, and 90% of the price adjustment is completed in just 28 days.
(5) The above observation suggests that there is a potentially large gain from utilizing daily data in this kind of study. This impression is reinforced by our analysis in section 8, where we carried out similar exercises using monthly aggregates; the results apparently suffer from a weak instrument problem.
(6) There are reasons to believe that the external instrument related to sanctions on Iran might be slightly less effective. If so, it could be because those news are correlated with changes in the market perception about issues other than oil supply.

In section 9 of the paper, I estimated responses of group-specific means of gasoline prices, constructed from a data set on individual gasoline stations. Although there are signs of heterogeneity across groups, it has not been easy to detect any clear underlying characteristic that would make a shop either more or less responsive to oil prices. Most notably, responses in urban areas are indistinguishable from those of the rest, and shops that are affiliated with large corporate groups are no faster or slower than independent stations in their responses. The only clear finding so far has been that shops on highways tend to exhibit more sluggish price adjustment, probably due to weaker competitive pressures.

In future work, it would be interesting to extend the time series analysis conducted in section 7 by incorporating other types of news in the market for oil, to see if they would help strengthen the identification strategy developed in this paper. Also, I intend to make use of the cross sectional dimension of the micro data studied in section 9 more fully, to characterize the evolution of the distribution of gasoline prices across shops in Japan.

Acknowledgement

Research for this paper has been funded by Hitotsubashi University, Hitotsubashi Institute for Advanced Study, the Grant-in-aid for scientific research C-15K03418, C-18K01605, A-17H00985, and A-20H00073, Japan Center for Economic Research, and Nomura Foundation.

I would like to acknowledge excellent research assistance by Shiro Yuasa. I am grateful to two anonymous referees for their insightful advice and constructive suggestions. I thank Professor Eiji Ogawa and all the other members of the RIETI research group, Tomoo Inoue, as well as the participants at the 2019 Workshop on Energy Economics: Econometric Analysis of Energy Demand and Climate Change, the 15th International Symposium on Econometric Theory and Applications (SETA2019), the Asia Pacific Economic Association Meeting (APEA2019), and the Hitotsubashi Summer Institute (HSI2019), for their valuable insights.

Appendix A. Robustness study: estimation in levels

In the main text, we presented cumulative impulse responses derived from estimation in which all the endogenous variables enter in their first differences. Although we feel this approach to be justified, because of the lack of evidence of cointegrating relationships between the endogenous variables, we present results based on the levels specification. Fig. A-1 reports non-cumulative responses for the case in which IV1 is the external instrument; the other cases are qualitatively similar. We can see that the results with the levels specification are broadly similar to those from the differences specification, in the sense that gasoline prices increase significantly in response to the shock, and that the reaction is quite fast. A difference is that the response of Dubai now appears mean-reverting.

Appendix B. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.eneco.2021.105214.

References

Bachmeier, L.J., Griffin, J.M., 2003. New evidence on asymmetric gasoline price responses. Rev. Econ. Stat. 85 (3), 772–776.

Bacon, Robert W., 1991. Rockets and feathers: the asymmetric speed of adjustment of UK retail gasoline prices to cost changes. Energy Econ. 13 (3), 211–218.

Balke, Nathan S., Brown, Stephen P.A., Yücel, Mine K., 1998. Crude oil and gasoline prices: an asymmetric relationship? Federal Reserve Bank of Dallas. Economic Review 1998 (Q1), 2–11.

Baumeister, Christiane, Hamilton, James D., 2019. Structural interpretation of vector autoregressions with incomplete identification: revisiting the role of oil supply and demand shocks. Am. Econ. Rev. 109 (5), 1873–1910.

Baumeister, Christiane, Kilian, Lutz, 2016. Forty years of oil Price fluctuations: why the Price of oil may still surprise us. J. Econ. Perspect. 30 (1), 139–160. https://doi.org/10.1257/jep.30.1.139.

Blair, Benjamin F., Campbell, Randall C., Mixon, Phillip A., 2017. Price pass-through in US gasoline markets. Energy Econ. 65, 42–49.

Brüggemann, Ralf, Jentsch, Carsten, Trenkler, Carsten, 2016. Inference in VARs with conditional heteroskedasticity of unknown form. J. Econ. 191 (1), 69–85.

Chen, Shiu-Sheng, 2009. Oil price pass-through into inflation. Energy Econ. 31 (1), 126–133.

Chesnes, Matthew, 2016. Asymmetric pass-through in US gasoline prices. Energy J. 37 (1).

Chudik, Alexander, Georgiadis, Georgios, 2019. Estimation of impulse response functions when shocks are observed at a higher frequency than outcome variables. Globalization Institute Working Paper No. 356, Federal Reserve Bank of Dallas.

Deltas, George, 2008. Retail gasoline price dynamics and local market power. Journal of Industrial Economics 56 (3), 613–628 (September).

Deltas, George, Polemis, Michael, 2020. Estimating retail gasoline price dynamics: the effects of sample characteristics and research design. Energy Econ. 92 (2020), 104976.

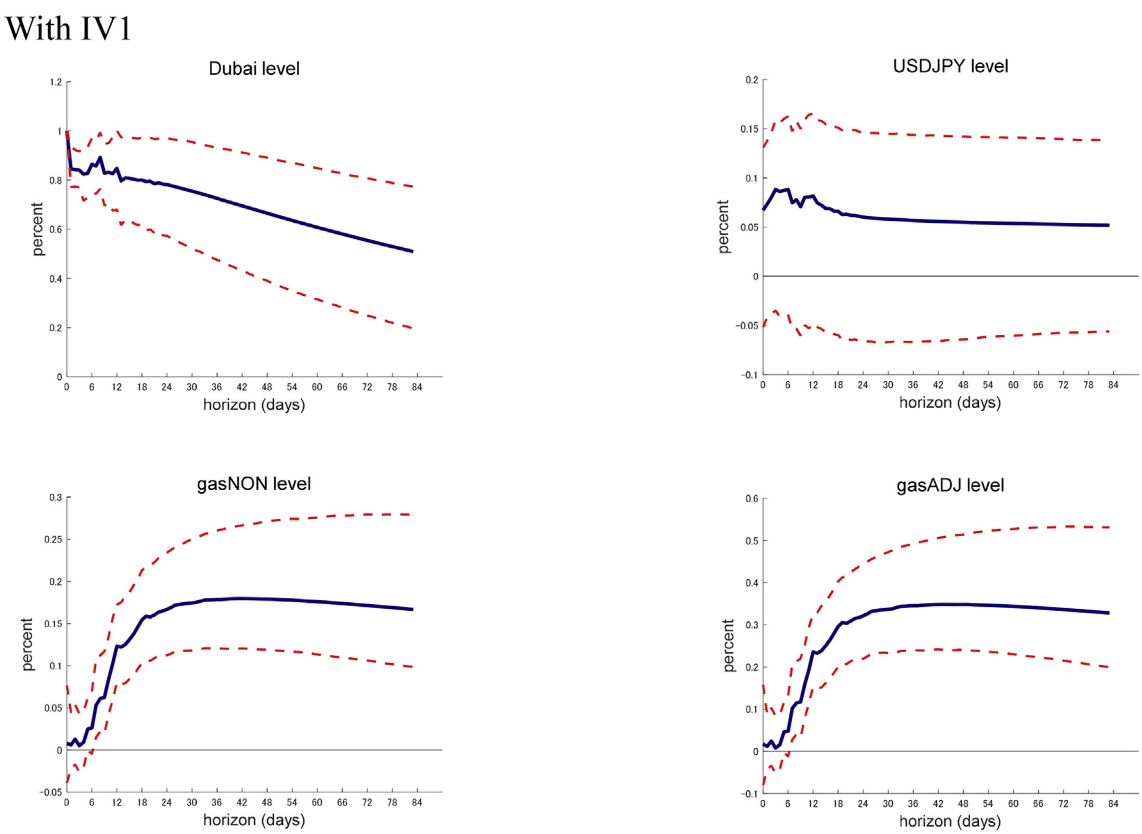

Fig. A-1. Impulse responses based on the levels specification, with the daily data and with IV1 + IV2. Note: All the endogenous variables in the VAR are in their *log levels*. The initial response of Dubai is normalized to equal 1. Solid lines are the point estimates, and dashed lines are the 95% confidence bands based on 5000 bootstrap draws.

Demirer, Rıza, Kutan, Ali M., 2010. The behavior of crude oil spot and futures prices around OPEC and SPR announcements: an event study perspective. Energy Econ. 32 (6), 1467–1476.

Draper, Dennis W., 1984. The behavior of event-related returns on oil futures contracts. J. Futur. Mark. 4 (2), 125–132.

Duffy-Deno, Kevin T., 1996. Retail price asymmetries in local gasoline markets. Energy Econ. 18 (1–2), 81–92.

Fisher, J.D.M., Peters, R., 2010. Using stock returns to identify government spending shocks. Economic Journal 120 (544), 414–436.

Fukunaga, Ichiro, Hirakata, Naohisa, Sudo, Nao, 2011. The effects of oil price changes on the industry-level production and prices in the United States and Japan. In: Ito, Takatoshi, Andrew, K., Rose, A. (Eds.), Commodity Prices and Markets, East Asia Seminar on Economics. vol. 20. University of Chicago Press, Chicago, IL, pp. 195–231.

Gertler, M., Karadi, P., 2015. Monetary policy surprises, credit costs, and economic activity. Am. Econ. J. Macroecon. 7, 44–76.

Godby, Rob, Lintner, Anastasia M., Stengos, Thanasis, Wandschneider, Bo, 2000. Testing for asymmetric pricing in the Canadian retail gasoline market. Energy Econ. 22 (3), 349–368.

Iwaisako, Tokuo, Nakata, Hayato, 2015. Oil price, exchange rate shock, and the Japanese economy. RIETI Discussion Paper Series 15-E-028.

Jentsch, Carsten, Lunsford, Kurt, 2019. The dynamic effects of personal and corporate income tax changes in the United States: comment. Am. Econ. Rev. 109 (7), 2655–2678.

Jordà, Òscar, 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95 (1), 161–182. https://doi.org/10.1257/0002828053828518.

Känzig, Diego R., 2021. The macroeconomic effects of oil supply shocks: new evidence from OPEC announcements. forthcoming in. *American Economic Review*, (The 2019 Dwyer Ramsey award for the best paper presented by a graduate student at the 27th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics).

Karrenbrock, Jeffrey D., 1991. 'The behavior of retail gasoline prices: symmetric or not?' Federal Reserve Bank of St. Louis Review 73 (4), 19–29.

Kilian, Lutz, 2008a. A comparison of the effects of exogenous oil supply shocks on output and inflation in the G7 countries. J. Eur. Econ. Assoc. 6 (1), 78–121.

Kilian, Lutz, 2008b. Exogenous oil supply shocks: how big are they and how much do they matter for the U.S. economy? Rev. Econ. Stat. 90 (2), 216–240.

Kilian, Lutz, 2009. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. Am. Econ. Rev. 99 (3), 1053–1069.

Kilian, Lutz, 2010. Explaining fluctuations in gasoline prices: a joint model of the global crude oil market and the US retail gasoline market. Energy Journal 31 (2).

Kilian, Lutz, 2016. The impact of the shale oil revolution on US oil and gasoline prices. Rev. Environ. Econ. Policy 10 (2), 185–205.

Kilian, Lutz, Park, Cheolbeom, 2009. The impact of oil price shocks on the U.S. stock market. International Economic Review 50 (4), 1267–1287 November 2009.

Kuttner, Kenneth N., 2001. Monetary policy surprises and interest rates: evidence from the fed funds futures market. Journal of Monetary Economics 47 (June), 523–544.

Lin, Sharon Xiaowen, Tamvakis, Michael, 2010. OPEC announcements and their effects on crude oil prices. Energy Policy 38 (2), 1010–1016.

Loderer, Claudio, 1985. A test of the OPEC cartel hypothesis: 1974-1983. J. Financ. 40 (3), 991–1006.

Loutia, Amine, Mellios, Constantin, Andriosopoulos, Kostas, 2016. Do OPEC announcements influence oil prices? Energy Policy 90, 262–272.

Mertens, Karel, Ravn, Morten O., 2013. The dynamic effects of personal and corporate income tax changes in the United States. Am. Econ. Rev. 103, 1212–1247.

Mertens, Karel, Ravn, Morten O., 2019. The dynamic effects of personal and corporate income tax changes in the United States: reply. Am. Econ. Rev. 109 (7), 2679–2691.

Meyler, Aidan, 2009. The pass through of oil prices into euro area consumer liquid fuel prices in an environment of high and volatile oil prices. Energy Econ. 31 (6), 867–881.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Polemis, Michael L., Tsionas, Mike G., 2016. An alternative semiparametric approach to the modelling of asymmetric gasoline price adjustment. Energy Econ. 56, 384–388.

Ramey, Valerie A., 2011. Identifying government spending shocks: It's all in the timing. Q. J. Econ. 126 (1), 1–50.

Ramey A., Valerie, Shapiro, Matthew, 1998. Costly capital reallocation and the effects of government spending. Carnegie-Rochester Conference Series on Public Policy 48, 145–194.

Romer, Christina D., Romer, David H., 1989. Does monetary policy matter? A new test in the spirit of Friedman and Schwartz. NBER Macroecon. Annu. 4 (121), 170.

Schmidbauer, Harald, Rösch, Angi, 2012. OPEC news announcements: effects on oil price expectation and volatility. Energy Econ. 34 (5), 1656–1663.

Shioji, Etsuro, 2012. The evolution of the exchange rate pass-through in Japan: a re-evaluation based on time-varying parameter VARs. Public Policy Review 8 (1), 67–92.

Shioji, Etsuro, 2014. A pass-through revival. Asian Economic Policy Review 9 (1), 120–138.

Shioji, Etsuro, 2015. Time varying pass-through: will the yen depreciation help Japan hit the inflation target? Journal of the Japanese and International Economies 37, 43–57.

Shioji, Etsuro, 2018. Infrastructure investment news and business cycles: evidence from the var with external instruments. 12th International Conference on Computational and Financial Econometrics (CFE 2018).

Shioji, Etsuro, Uchino, Taisuke, 2011. Pass-through of oil prices to Japanese domestic prices. In: Ito, Takatoshi, Andrew, K., Rose, A. (Eds.), Commodity Prices and Markets, East Asia Seminar on Economics. Vol. 20. University of Chicago Press, Chicago, IL, pp. 155–189.

Shioji, Etsuro, Yuasa, Shiro, 2020. Daily dynamics of retail gasoline price dispersion in Japan. 14th International Conference on Computational and Financial Econometrics (CFE2020).

Stock, J.H., Watson, M.W., 2012. Disentangling the channels of the 2007-09 recession. Brookings Papers on Economic Activity. 1, pp. 81–135.

Stock, J.H., Watson, M.W., 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." NBER Working Paper No. 24216.

Yanagisawa, Akira, 2012. Structure for pass-through of oil price to gasoline price in Japan. IEEJ Energy J. 7 (3).

Yilmazkuday, Hakan, 2019. Oil Price Pass-Through into Consumer Prices: Evidence from US Weekly Data (Available at SSRN 3443245).