First version: February 25, 2019
Second version: April 22, 2019
Minor revision: May 10, 2019

Pass-through of oil supply shocks to domestic gasoline prices: evidence from daily data[1]

Etsuro Shioji
(Hitotsubashi University)[2,3]

[1] Key Words: Key Words: SVAR-IV, High Frequency Identification, news shock, oil price, pass-through, JEL Codes: E31, F41, Q41, Q43, C32
[2] Research for this paper has been funded by Hitotsubashi University, Hitotsubashi Institute for Advanced Study, the Grant-in-aid for scientific research S-24223003, C-15K03418, and A-17H00985, and Nomura Foundation. I would like to acknowledge excellent research assistance by Shiro Yuasa. I thank Professor Eiji Ogawa and all the other members of the RIETI research group for their valuable insights.
[3] Department of Economics, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo, 186-8601, Japan. E-mail: shioji (at) econ.hit-u.ac.jp.

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 percent of the entire adjustment process is completed within just 18 business days.

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 afterwards, the Nikkei Newspaper in Japan was reporting that gasoline stations were raising their prices. In such an environment, a standard 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 on gasoline prices in Japan. This is obtained from a popular price comparison site specialized in gasoline diesel oil, to which tens of thousands of registered drivers, as well as shop managers, report prices posted at gasoline stations from all over Japan.

In our analysis, we take seriously the claim made in the literature that 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 (2019) 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 (sometimes 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 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, and, if so, magnitudes of the surprises. This is done by computing daily changes in WTI 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 spot crude oil prices (WTI), the exchange rate (between the USD and the JPY) and domestic gasoline prices. Note that our 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.

Another characteristic of this paper is that it takes into account the presence of heavy taxes on gasoline in Japan. As they are mostly levied by the quantity of purchases (not the value), the tax component of gasoline prices is insensitive to oil price changes. To see the consequence of this treatment, I follow Shioji and Uchino (2011) to estimate and remove the tax component from the observed price, and include this adjusted price as an alternative index of gasoline prices in Japan.

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 30 percent if we use gasoline prices that are not adjusted for taxes, but this number goes up to 60 percent when we use the tax-adjusted data. And this pass-through is fast: as already mentioned, over 60 percent of the long run pass-through occurs within the first three weeks after the shock.

The rest of the paper is organized as follows. Section 2 gives a brief overview of the related literature. Section 3 introduces our new dataset on gasoline prices at Japanese gas stations. Section 4 reviews the idea behind the SVAR-IV approach. Section 5 explains the construction of our External Instruments. Section 6 presents the estimation results, and Section 7 concludes.

2. Background

Importance of disentangling supply and demand

Economists have long been interested in the effects of oil price changes on the domestic economy, at least since the First Oil Crisis of the early 1970s. The most important progress in the literature in recent years has come from 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, which will be reviewed below.

SVAR approach by Kilian (2009)

The new methodology of Kilian (2009) and Kilian and Park (2009) 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, Hirakata and Sudo (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. This approach has been applied to G7 countries, including Japan, by Kilian (2008a).

Relationship of this paper with Kilian (2009)

The approach proposed in this paper differs from the popular methodology of Kilian (2009) in some ways. Most notably, this paper tries to identify only one out of many kinds of driving forces behind world oil prices: we do not seek to completely decompose the entire movements in oil prices into a few specific types of structural shocks. And the type of shocks that this paper is trying to identify, namely shocks to the expectations for future world oil supply, 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 of Kilian (2009)). And, to identify such type of shock, we employ an event-study like approach which has been popularized mainly in the empirical literature on monetary policy.

SVAR-IV approach by Känzig (2019)

Among existing studies, Känzig (2019) is the closest to this paper both in terms of the research objective and the methodology[4]. Like this paper, his study aims to identify shocks to expectations about future oil supply. His paper tries to achieve the identification with the SVAR-IV methodology, using the response of the WTI futures to news about future oil supplies as the external instrument, as is done in this paper.

The two papers differ in how the news dates are picked. His paper lists every date within the sample period on which there was a formal announcement after the OPEC meeting. This has an advantage of avoiding any subjective judgement in choosing the dates. On the other hand, this paper argues that the news has to be "important enough" to reduce the 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 also takes the view that the real news might not happen on the day of an official announcement. For example, the media might report the likely outcome of an OPEC meeting before the last day of the meeting. Also, important news about OPEC might 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. 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.

Another difference is the types of the domestic (meaning oil importer country's) variables that we investigate. Känzig (2019) studies responses of US aggregate variables, such as GDP and CPI inflation, to the oil news shock. This requires him to convert the frequency of the instrumental variable from daily to quarterly. Here, as already stated, the objective is to uncover how day-to-day gasoline prices in Japan are affected by those shocks. As a consequence, there is no need for frequency conversion, which avoids the resulting loss of information.

Pass-through literature

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

[4] 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.

Japan. On the other hand, in Shioji and Uchino (2009), we study pass-through from world oil prices to domestic prices, using monthly data.

3. Overview of the gasoline data

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. For this analysis, between January 2013 (which is the beginning of the sample) and mid-July 2018, I used the nationwide average series for “Regular Fuel (meaning not high-octane), Cash-only (i.e., non-member prices)”. Between mid-July to the end of year 2018, I switch to the average of non-member and member prices[5]. For January 2019 (which is the last part of the current sample), due to a technical problem, I switched back to non-member prices. Those different series are connected using growth rates.

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

[5] During this switching process, we lost a few days’ observation by error. Those needed to be estimated through interpolation.

series will be denoted as **gasNON**, while the second one will be called **gasADJ**.
In each of the two panels in Figure 2, I plot each of those series, together with the spot price of crude oil (WTI in this case), denoted WTIspot, converted into the units of the Japanese Yen using the market exchange rate. 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 panels suggests the extent to which the Japanese tax system helps dampen volatilities in gasoline prices that consumers face.

4. 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 $y_{2,t}$:

$$Y_t = AY_{t-1} + v_t \qquad (1)$$

where

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

In the above, v_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:

$$v_t = B\varepsilon_t \text{, where } \varepsilon_t \equiv \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}. \qquad (2)$$

Then, assuming invertibility, we can write:

$$Y_t = C(L)B\varepsilon_t \quad \text{where } C(L) = (I - AL)^{-1}. \qquad (3)$$

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:

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

and

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

Then we obtain the following:

$$Ev_tZ_t = \begin{bmatrix} b_{11}\alpha \\ b_{21}\alpha \end{bmatrix}. \tag{6}$$

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:

$$y_{2,t} = b_{21}y_{1,t} + d_1y_{1,t-1} + d_2y_{2,t-1} + b_{22}\varepsilon_{2,t}. \tag{7}$$

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

$$\hat{C}(L) = (I - \hat{A}L)^{-1}. \tag{8}$$

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

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

5. Construction of instruments

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, when 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 (2019) 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 (2019).

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 (2019). 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 some of 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 (2019): as my sample period (from 2013 to early 2019) is much shorter and more recent than his (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 Obama administration of the US 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.

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 (2019), as explained earlier, uses *all* the dates on which there was a formal announcement 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" about it 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.

Figure 2 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.

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 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 (2019) 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 was 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 1 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 (2019), 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[6].

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] 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".

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 final step, suppose, for example, that OPEC has just made an announcement that it would cut the member countries' production by 10 percent. 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 percent rather than 10 percent. 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 the market response to a news by the log change in the closing price of WTI futures on the news date from the previous business day's closing price. This procedure is basically the same as the one in Känzig (2019). The only exception (as also noted in Table 1) is April 18, 2016, when an important news appeared over a weekend: for this case only, I decided to take log difference between Friday's closing and the following Monday's opening.

Resulting instruments

The above procedure produces two indicators, which will be used as external instruments in the SVAR-IV estimation. The first one measures the market reaction to OPEC-related news; this will be called IV1 (OPEC) in what follows. The second one has to do with Iran-sanction-related news, and will be called IV2 (Iran). Figure 2 plots them together with WTIspot, which is the spot price of WTI, 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.

6. Estimation details and results

6.1 Estimation details

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

(1) World crude oil price = **WTIspot**: Cushing, OK WTI Spot Price FOB (Dollars per Barrel); taken from the web site of the US Energy Information Administration.
(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 New York).
(3) Gasoline price in Japan = **gasNON** or **gasADJ**, as explained in Section 3. Note that they are measured in the Japanese Yen.

We take log first differences of all the series. As for the External Instrumental Variable, I use either one of the following three:

[1] **IV1** (OPEC) = Equals to the rate of change (log difference) of WTI futures prices on the days when the OPEC-related news arrived. Otherwise, it is equal to zero.
[2] **IV2** (Iran) = Similar index for the news related to Iranian sanctions.
[3] **IV** = IV1 + IV2.

In total, the model is estimated in six different specifications (as we have two alternative gasoline variables and three alternative instruments).

The dataset is daily but includes only the US business days when the market was open (the Japanese gasoline data contains observations on all days, including weekends and Japanese holidays), and starts from the beginning of January 2013 and ends on 28 January, 2019. The number of lags is set at 20 (or about four weeks).
To take into account the time differences between Japan and the US (Tokyo is 13-14 hours ahead of New York), we lead the gasoline variable by one day; that is, the gasoline price included as a contemporaneous variable in the system is actually that of the following day (in the Japanese calendar).
The estimation is carried out with Ambrogio Cesa-Bianchi's VAR Toolbox for Matlab, made public through his web site (https://sites.google.com/site/ambropo/MatlabCodes).
The estimated impulse responses to an identified oil supply shock are presented in Figures 4-7. Now I will examine them in turn.

6.2 Estimation results with IV1 (OPEC)

Figure 4 reports the results for the case in which IV1, the OPEC-related index, is used as the external instrument. This figure consists of two panels. Panel A corresponds to the case in which gasNON was used as the domestic gasoline price variable, while Panel B uses gasADJ instead. 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 boxes, the solid line represents the median response based on 10,000 bootstrap draws, while the dotted lines are the 95 percent confidence bands. All the responses are normalized so that the initial response of WTIspot is equal to 1. As a consequence, the underlying shock should be interpreted as a negative shock to oil supply. Responses of up to 120 periods (i.e., business days), which corresponds to around five to six months, are shown.
Starting from Panel A, in the first box, the identified oil supply shock has an immediate and permanent effect on WTIspot, the spot price of crude oil. In fact, the response immediately reaches the peak and stays flat thereafter. It 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.

Finally, in the third box, the response of gasNON is practically zero at the beginning, but turns significantly positive in just eight business days. The long run response is around 0.30, based on the median. In other words, the long-run pass through rate to gasNON is about 30 percent. What is most notable is how fast this pass-through occurs over time. Out of the entire long run response of 0.30, over 30 percent of it would be completed in just ten business days after the shock. This number goes up to over 70 percent on the 18^{th} business day, and crosses the 90 percent mark on the 36^{th} business day. By the 65^{th} business day after the shock, 99 percent 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.

Shifting our attention to Panel B, we see that the responses of WTIspot and USDJPY, reported in the first two boxes, are quite similar to the ones we saw in Panel A. In the third box, the response of gasADJ, the index of gasoline price after removing the effects of taxes, is presented. Its shape is similar to the one for gasNON that we saw in Panel A, but the size is much larger.

In Figure 5, to facilitate the comparison, I put the median response of gasNON (from Panel A of Figure 4) and that of gasADJ (from Panel B of the same figure) in a single box. The long run pass-through rate for gasADJ is 0.58, which is nearly double the estimated value of 0.30 we obtained by using 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 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 67th day after the shock that over 90 percent of the long run pass-through is realized, as opposed to the 65th day.

6.3 Estimation results with IV2 (Iran)

Figure 6 presents results when IV2, based on the Iran-related news dates, is used as the external instrument. Although they are broadly similar to those in Figure 4, two differences are notable. First, the response of USDJPY is still positive but smaller, and is no longer significant. Second, the confidence bands around the response of gasADJ are wider.

Those two observations might indicate that IV2 is slightly 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 increase the demand for the Japanese Yen, which is widely considered as a safe haven currency, which could in turn lead to its appreciation, rather than depreciation.

6.4 Estimation results with IV (=IV1+IV2)

Figure 7 demonstrates what happens when we merge the two instruments by adding them up. The result is closer to the one in Figure 6 (i.e., the case with IV2) in that the response of USDJPY is insignificant. However, the widths of the confidence bands around the responses of the domestic gasoline prices are more similar to those in Figure 4 (the case with IV1).

7. Conclusions

This paper has utilized a new dataset on gasoline prices in Japan to estimate the degree of pass-through from the world oil prices to the domestic prices, based on daily data. 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 30 percent.
(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 60 percent.
(4) The pass-through is fast. About 70 percent of the entire process occurs within 18 business days, and 90 percent of the price adjustment is completed in just 65-67 business days.
(5) The above observation suggests that there is a high value in utilizing daily data in this kind of study. A study that relies on monthly data might not be able to capture the entire dynamics depicted in this paper.
(6) There are reasons to believe that the external instrument related to sanctions on Iran might be 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 future work, I intend to explore possibilities of utilizing other types of news in the market for oil, to see if they would help strengthen the identification strategy developed in this paper. Also, it would be interesting to see if making use of even higher frequency information would add value to the current analysis. For example, we could think of looking into changes in cure oil futures' prices between hours before certain announcement (by, for example, OPEC) and hours afterwards.

References

Fukunaga, Ichiro, Naohisa Hirakata, and Nao Sudo (2011) "The Effects of Oil Price Changes on the Industry-Level Production and Prices in the United States and Japan". in Takatoshi Ito and Andrew K. Rose A. (eds), *Commodity Prices and Markets, East Asia Seminar on Economics, Vol.*20, University of Chicago Press, Chicago, IL, 195–231.

Iwaisako, Tokuo and Hayato Nakata (2015). "Oil Price, Exchange Rate Shock, and the Japanese Economy." RIETI Discussion Paper Series 15-E-028.

Känzig, Diego R. (2019) "The Macroeconomic Effects of Oil Supply Shocks: New Evidence from OPEC Announcements", (Version January 22, 2019, Available at SSRN: https://ssrn.com/abstract=3185839), (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).

Kilian, Lutz. (2008a). "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries." *Journal of the European Economic Association*, 6(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?" *Review of Economics and Statistics,* 90(2): 216–40.

Kilian, Lutz. (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99 (3): 1053-69.

Kilian, Lutz and Cheolbeom Park (2009) "The Impact of Oil Price Shocks on the U.S. Stock Market." *International Economic Review*, 50(4), November 2009, 1267-1287

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

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), pp. 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 E., and Uchino T. (2011) Pass-through of oil prices to Japanese domestic prices, in Takatoshi Ito and Andrew K. Rose A. (eds), *Commodity Prices and Markets, East Asia Seminar on Economics, Vol.*20, University of Chicago Press, Chicago, IL, 155–89.

Stock, J.H. and M.W. Watson, 2012. Disentangling the Channels of the 2007-09 Recession. *Brookings Papers on Economic Activity*, no. 1: 81-135.

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

Table 1: List of news dates

(Rows in white = OPEC-related news, Rows in 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.

Figure 1A: Time series plot of WTIspot (in JPY) vs gasNON

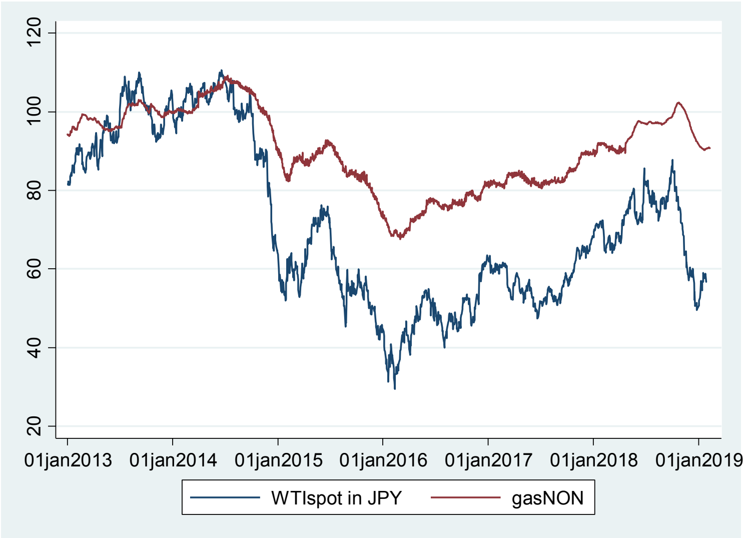

Figure 1B: Time series plot of WTIspot (in JPY) vs gasADJ

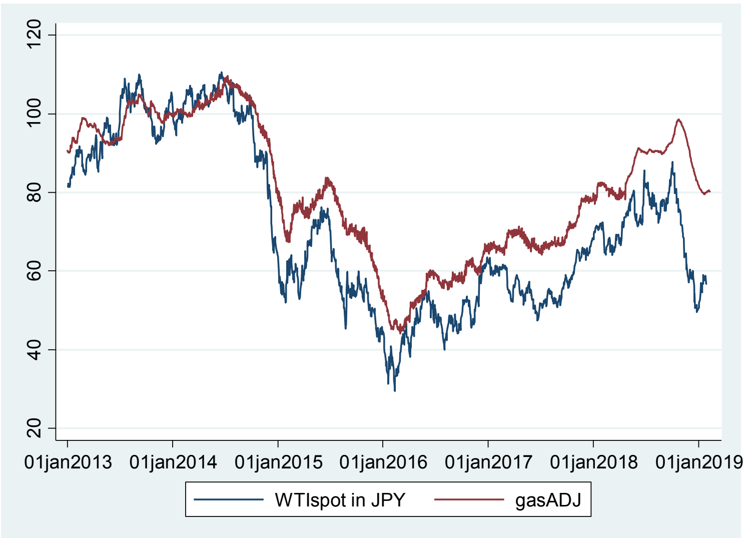

Note: "WTIspot in JPY" = WTI spot price (WTIspot), converted into the JPY units; "gasNON" = gasoline price in Japan not adjusted for taxes (this corresponds to prices paid at gas stations). "gasADJ" = same but adjusted for taxes (this is the before-tax price series). All are normalized to equal 100 at the beginning of January 2014.

Figure 2A: Google Trends search results for "OPEC"

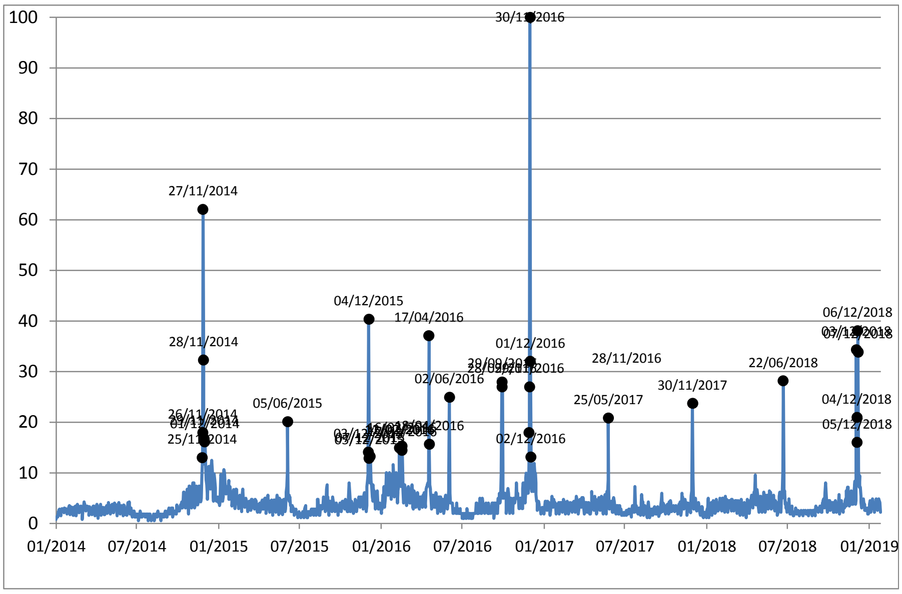

Figure 2B: Google Trends search results for "Iran Sanction"

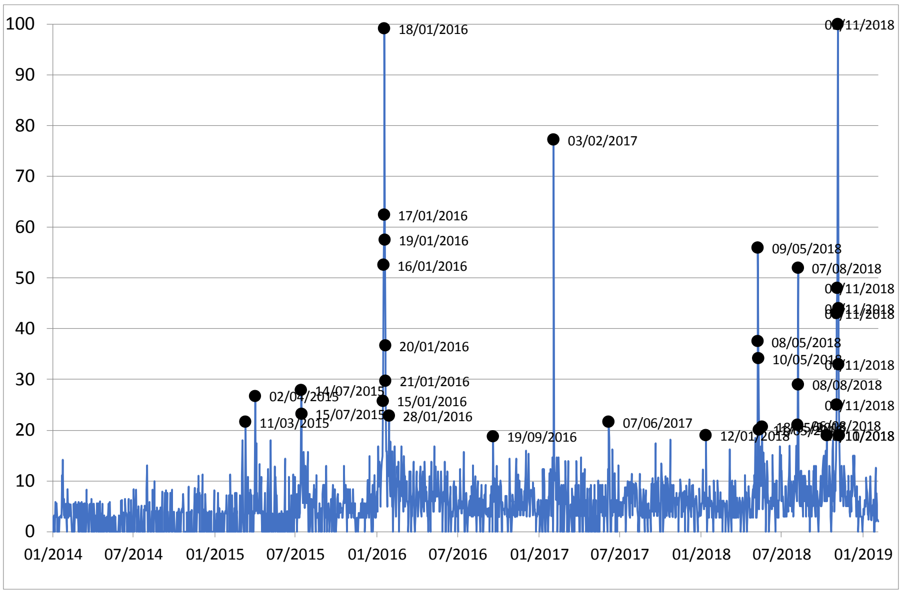

Note: Based on the number of internet searches from the entire world according to Google Trends. Both series are normalized to equal 100 at the maximum. Black dots with dates are observations exceeding two standard deviations above the mean.

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.

Figure 3: Time series plot for External Instruments (with WTIspot)

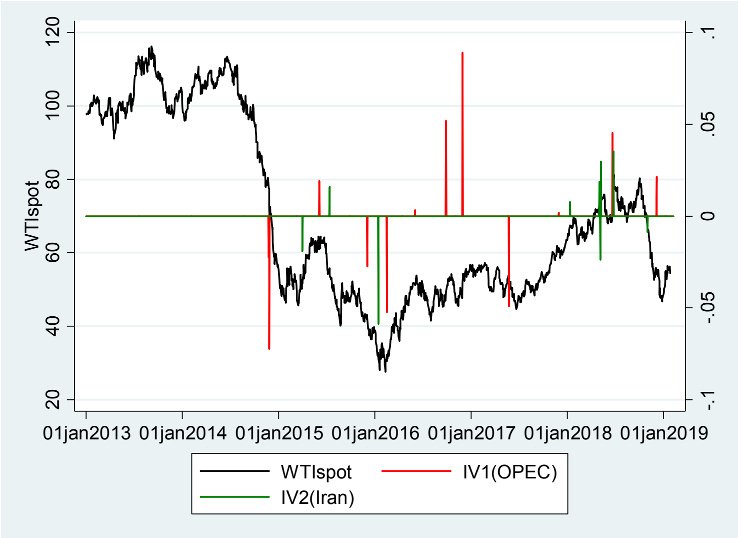

Note: "WTIspot" = WTI spot price (in US Dollars), "IV1(OPEC)" = External Instrument 1 (related to news about OPEC), "IV2(Iran)" = External Instrument 2 (related to news about the US-led coalitions' sanction on oil exports from Iran). WTIspot is normalized to equal 100 at the beginning of January 2014.

Figure 4: Impulse responses, case with IV1 (OPEC)

(A) Gasoline price = gasNON (actual prices that consumers pay)

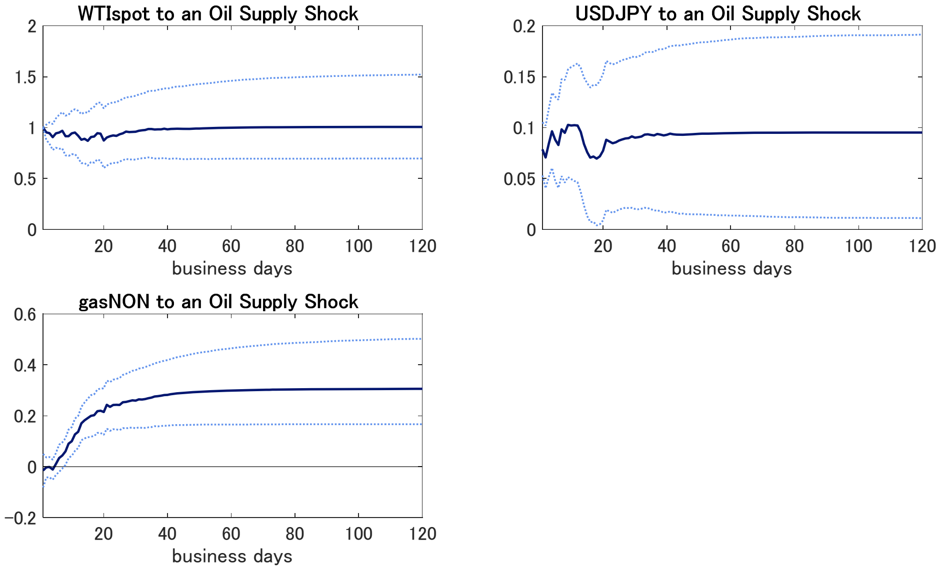

(B) Gasoline price = gasADJ (after removing the effects of taxes)

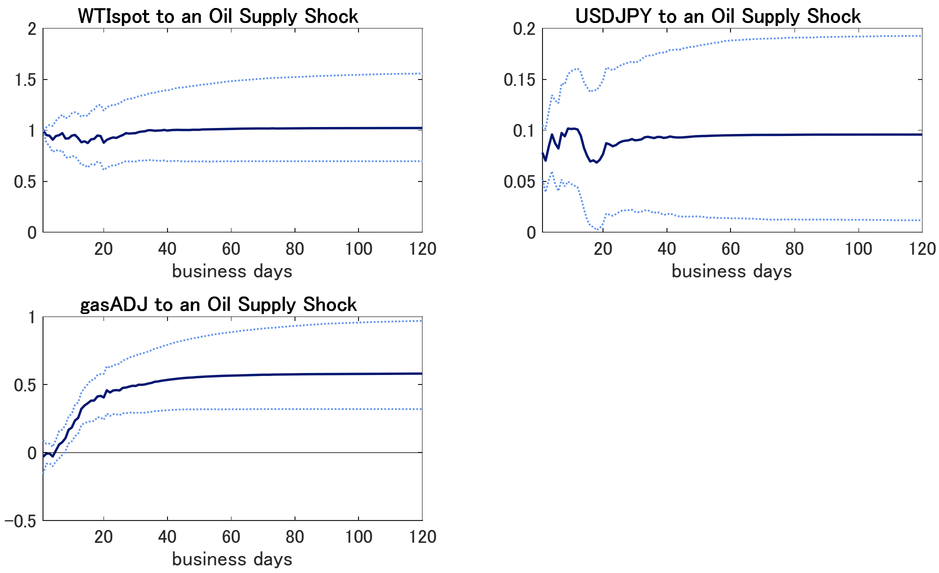

Note: All the endogenous variables in the VAR are in log first differences and the plots are their cumulative responses. The initial response of WTIspot is normalized to equal 1. Solid lines are the medians, and the dotted lines are the 95 percent confidence bands based on 10,000 bootstrap draws.

Figure 5: Comparison of the Median impulse responses, gasNON vs gasADJ (Case with IV1 (OPEC))

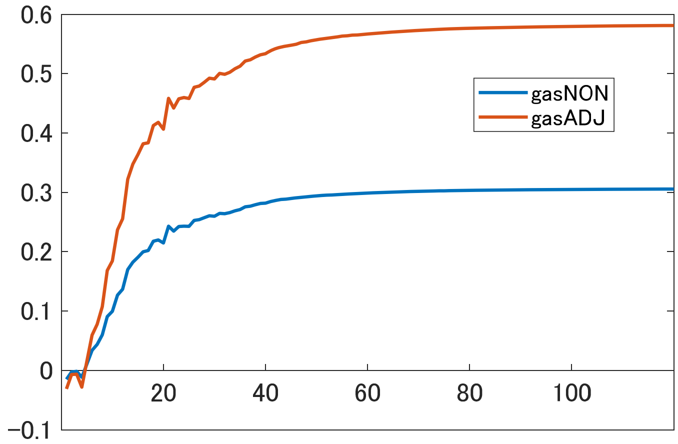

Figure 6: Impulse responses, case with IV2 (Iran)

(A) Gasoline price = gasNON (actual prices that consumers pay)

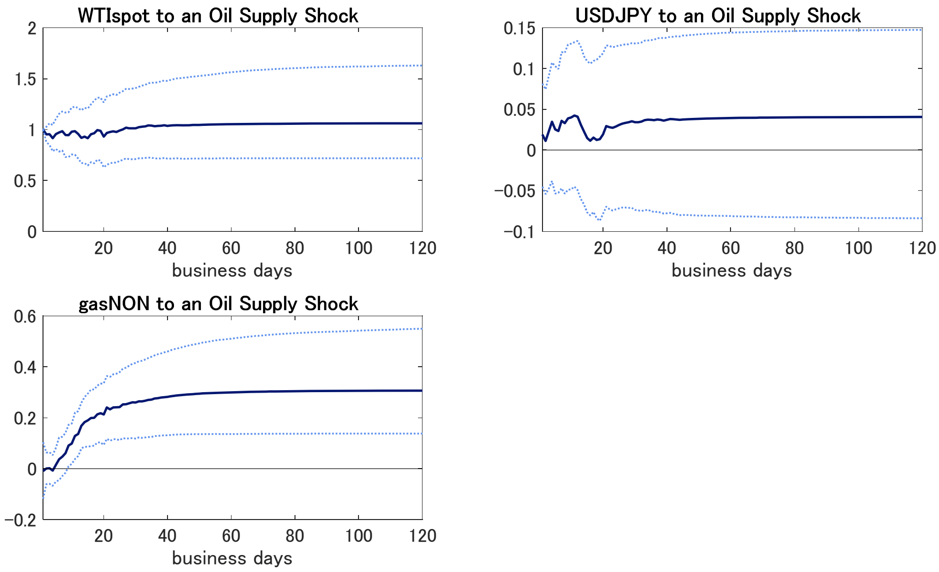

(B) Gasoline price = gasADJ (after removing the effects of taxes)

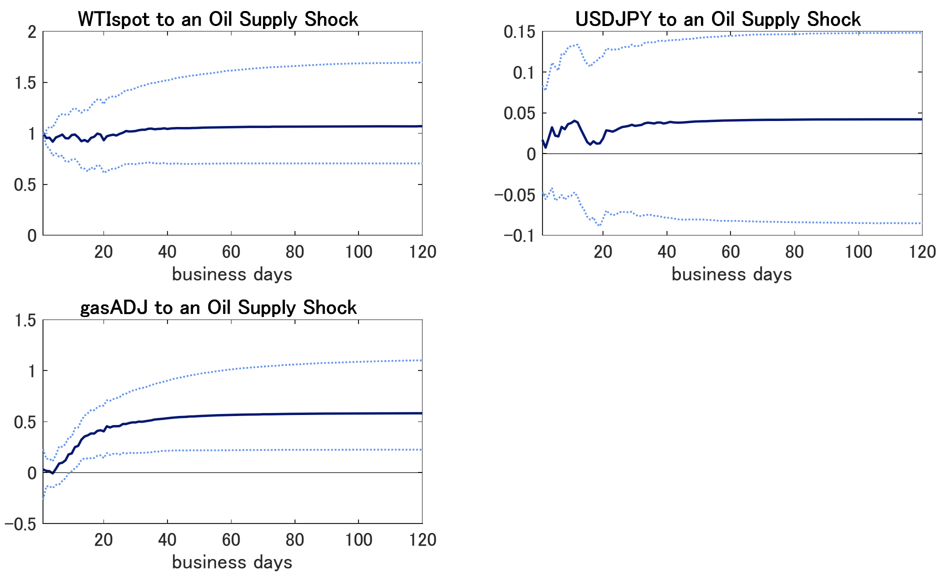

Note: All the endogenous variables in the VAR are in log first differences and the plots are their cumulative responses. The initial response of WTIspot is normalized to equal 1. Solid lines are the medians, and the dotted lines are the 95 percent confidence bands based on 10,000 bootstrap draws.

Figure 7: Impulse responses, case with both IV1 (OPEC) and IV2 (Iran)

(A) Gasoline price = gasNON (actual prices that consumers pay)

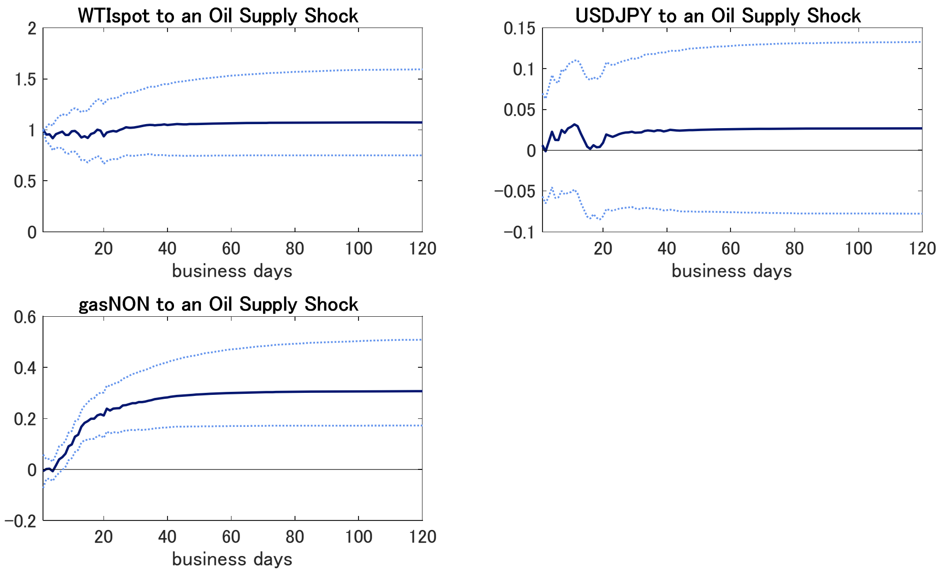

(B) Gasoline price = gasADJ (after removing the effects of taxes)

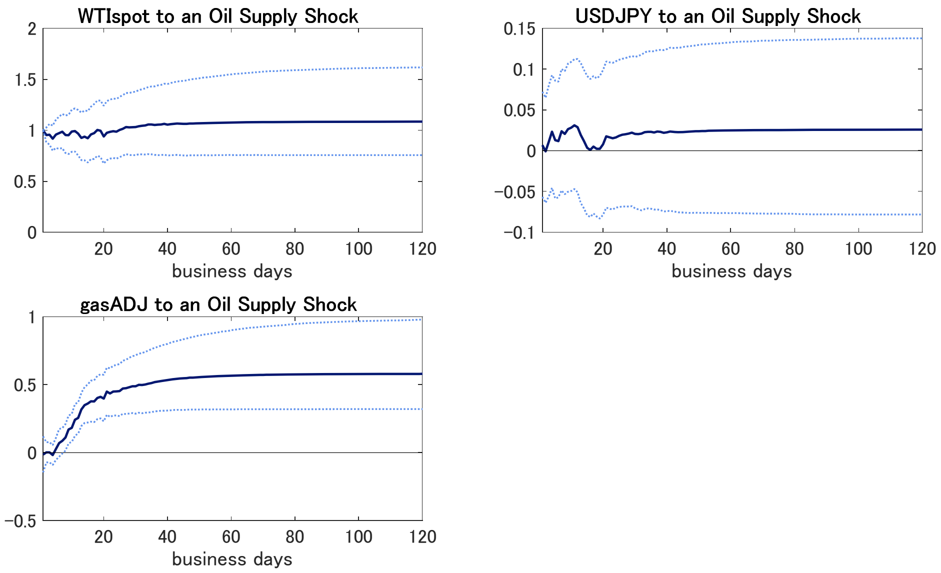

Note: All the endogenous variables in the VAR are in log first differences and the plots are their cumulative responses. The initial response of WTIspot is normalized to equal 1. Solid lines are the medians, and the dotted lines are the 95 percent confidence bands based on 10,000 bootstrap draws.