

Does Drawing Down the U.S. Strategic Petroleum Reserve Help Stabilize Oil Prices?

Lutz Kilian

Federal Reserve Bank of Dallas
CEPR

Xiaoqing Zhou

Federal Reserve Bank of Dallas

January 28, 2019. This version: December 19, 2019

Abstract: We study the efficacy of releases from the U.S. Strategic Petroleum Reserve (SPR) within the context of fully specified models of the global oil market that explicitly allow for storage demand as well as unanticipated changes in the SPR. Using novel identifying strategies and evaluation methods, we examine seven questions. First, how much have exogenous shocks to the SPR contributed to the variability in the real price of oil? Second, how much would a one-time exogenous reduction in the SPR lower the real price of oil? Third, are exogenous SPR releases partially or fully offset by increases in private sector oil inventories and how does this response affect the transmission of SPR policy shocks? Fourth, how effective were actual SPR policy interventions, consisting of sequences of exogenous changes in the SPR, at lowering the real price of oil? Fifth, are there differences in the effectiveness of SPR emergency drawdowns and SPR exchanges? Sixth, how much did the creation and expansion of the SPR contribute to higher real oil prices? Finally, how much would selling half of the oil in the SPR, as recently proposed by the White House, lower the global price of oil (and hence the U.S. price of motor gasoline) and how much fiscal revenue would it generate?

JEL code: Q38, Q43, E62.

Key words: SPR, crude oil, oil inventories, storage, expectations, policy intervention, fiscal policy

Acknowledgments: The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of Federal Reserve Bank of Dallas or the Federal Reserve System. We thank Ana Maria Herrera, Martin Stuermer, the editor, and three anonymous referees for helpful comments. The authors have no conflict of interest to disclose.

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

Xiaoqing Zhou, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. Email: xqzhou3@gmail.com.

1. Introduction

Since October 1977, the United States has maintained a Strategic Petroleum Reserve (SPR) that includes substantial stocks of crude oil to be used to stabilize the oil market in the event of a disruption of U.S. crude oil imports, of unexpected oil supply shortfalls in the U.S. refining market, or of geopolitical disruptions in global oil markets. There is a widespread belief among policymakers that releases of crude oil from the SPR help calm the oil market and stabilize the real price of oil.

This paper proposes an econometric methodology to evaluate the empirical support for this view. Quantifying the effect of exogenous SPR policy interventions is essential for policymakers because it helps pin down the sequence of interventions required to achieve the targeted change in the real price of oil. In fact, without a quantitative structural model, a policymaker is unable to determine to what extent the changes in the real price of oil observed after a policy intervention can be attributed to the policy intervention and to what extent they would have occurred even in the absence of the intervention.

In the absence of such counterfactuals, policymakers' perceptions of the success of past policy interventions have been based on casual empiricism and strong prior views more than empirical evidence. For example, the U.S. Department of Energy concluded that the partial drawdown in 1990/91 helped restore stability to world oil markets during the Persian Gulf War, but many critics have contended that the SPR releases came too late to matter and that the steep drop in the price of oil in 1991 was caused by the success of Allied military operations and by the ability of OPEC to ramp up oil production instead (e.g., Bordoff, Halff and Losz 2018). In recognition of this uncertainty, even policymakers' own assessments of the success of more recent U.S. SPR releases have tended to be vague and guarded. For example, after the coordinated International Energy Agency (IEA) crude oil release in response to the Libyan crisis in 2011, the IEA took the position that its actions "played at least a partial role in helping avoid a damaging price spike" during the summer of 2011.¹

¹ Testimony of Richard H. Jones, deputy executive director of the International Energy Agency, US Senate Committee on Energy and Natural Resources, January 31, 2012, https://www.energy.senate.gov/public/index.cfm/files/serve?File_id=B31E1AD9-73EB-4BB8-ACD3-354B5E3ED9B7.

In this paper, we are concerned with constructing explicit counterfactuals that help guide policymakers. We study the efficacy of SPR policy interventions within the context of a fully specified structural model of the global oil market that explicitly allows for storage demand as well as unanticipated changes in the SPR. This framework also allows us to examine the extent to which the effects of SPR policy shocks depend on alternative sets of identifying assumptions.²

Our study breaks new ground in several dimensions. First, unlike earlier studies, we explicitly allow for a response of private sector oil inventories to exogenous SPR interventions, which dampens the effects of SPR interventions on the price of oil. Although we find evidence that increases in private oil inventories tend to offset exogenous reductions in the SPR in the short run, we show that this response is not important for the substance of our conclusions about the effectiveness of SPR releases. Second, we show that SPR policymakers systematically respond to the state of the global oil market much like monetary policymakers respond to the state of the domestic economy. On average, only 55% of the variation in changes in the SPR is accounted for by exogenous SPR policy shocks. Third, although exogenous SPR releases are associated with additional drawdowns in non-SPR oil stocks at longer horizons, the response of the real price of oil to exogenous SPR releases is only weakly negative.

Fourth, although exogenous SPR shocks are not an important determinant of the variability of the real price of oil on average, we show that historically SPR policy interventions, defined as sequences of exogenous SPR shocks during selected periods, have helped stabilize the price of oil. Their effect on the price of oil, however, has been modest. Unlike previous studies that focused on the effect of one-time SPR policy shocks, we represent the major policy interventions since 1990 as sequences of exogenous SPR shocks within the structural moving average representation of a VAR model. This allows us to quantify their cumulative effect on the real price of oil based on the structural moving average representation of the VAR model.

² Our approach contrasts with recent work by Stevens (2014) and Newell and Prest (2017), whose models are partially identified in that they impose a recursive ordering, employ scheduled SPR purchases or releases as extraneous instruments, or rely on high-frequency futures returns around SPR release announcements to identify exogenous variation in the SPR.

We contrast our posterior median estimates of these cumulative effects with perceptions among policymakers about the success of these policy interventions. We find, for example, that the cumulative effect of the SPR releases after the invasion of Kuwait in 1990 was a reduction of only \$2/barrel in the real price of oil after 7 months. Likewise, the SPR release after Hurricane Katrina in 2005 cumulatively lowered the price of oil by only \$3 after one quarter, whereas the SPR intervention prompted by the Libyan crisis in 2011 succeeded in lowering the real price of oil by as much as \$12 cumulatively after 4 months. The modest estimates for 1990 and 2005 are noteworthy, given the folk wisdom that these policy interventions were highly successful and prevented much larger oil price increases. We also highlight differences between SPR exchanges and emergency drawdowns, showing that SPR exchanges ultimately raise the real price of oil, whereas emergency drawdowns lower the real price of oil. In addition, we quantify the costs of building and expanding the SPR, as measured by its cumulative effect on the real price of oil. These costs are shown to be modest. We estimate, for example, that the expansion of the SPR between 2001 and 2005 cumulatively raised the price of oil by only \$4/barrel after 43 months.

Fifth, we show how our approach may also be used to quantify the effects of hypothetical policy interventions. We provide a detailed analysis of the benefits of the 2018 White House proposal to sell off half of the SPR within the next decade. Our concern is not whether this proposal is economically sensible, but whether it would reach the objectives of raising fiscal revenue and of lowering oil and gasoline prices. We estimate the cumulative effect of this proposal by means of a novel counterfactual policy simulation, which accounts for estimation uncertainty. We show that the expected fiscal benefits of this plan are somewhat higher than the revenue of \$16.6 billion dollars projected by the White House. The median cumulative revenue is \$17 billion dollars (with a 68% credible region of between \$14 and \$19 billion dollars). The implied median cumulative change in the U.S. real price of motor gasoline is -13c/gallon (with a 68% credible region of between -50c and 9c). Given that the U.S. has a nearly balanced trade account in crude oil and petroleum products, however, one would not expect the White House plan to have important additional stimulating effects on the U.S. economy through its effect on consumer spending.

Finally, our paper also breaks new ground methodologically. First, this paper is one of only two studies in the structural VAR literature that combine sign restrictions, zero restrictions and narrative restrictions to achieve identification.³ The presence of zero and narrative restrictions invalidates the conventional approach to estimating sign-identified VAR models as discussed in Rubio-Ramirez, Waggoner and Zha (2010). We deal with this problem by combining the importance samplers recently developed by Arias, Rubio-Ramirez and Waggoner (2018) and Antolin-Diaz and Rubio Ramirez (2018). There have been few applications of these importance samplers in the literature and none, to our knowledge, that have used them in conjunction.

Second, we introduce a computationally more efficient algorithm for drawing from the space of rotation matrices that exploits the block recursive structure of some of our models.⁴

Third, we not only propose a new approach to evaluating policy interventions that involve sequences of policy shocks rather than one-time interventions, but we propose a novel approach to evaluating the posterior distribution of the cumulative effects of actual and hypothetical policy interventions that can be applied more generally in other set-identified structural VAR models and hence is of broader interest to applied users.⁵ Conventional summary statistics for the set of admissible structural models based on cumulative sums of pointwise posterior medians (or of products of pointwise posterior medians) tend to conflate estimates from different structural models, calling into question their interpretation (Kilian and Lütkepohl 2017). In contrast, we first cumulate the simulated effects before computing their posterior median, when assessing the cumulative effect on the real price of oil. Likewise, we cumulate the product of prices and quantities first, before computing the posterior median of the fiscal revenue of policy interventions. Our approach is designed to preserve both the dynamics of the responses and the conditional co-movement of the responses over time.

³ The other is Kilian and Zhou (2019).

⁴ For a similar approach see Mumtaz and Surico (2009), whose model has a structure that resembles one of our block recursive models, but not the other. Their econometric evaluation of that model, however, was shown to be invalid in Arias et al. (2018).

⁵ A closely related approach has been used in a different context by Herrera and Rangaraju (2019) for characterizing the cumulative response of real GDP to oil demand and oil supply shocks.

The remainder of the paper is organized as follows. Section 2 provides institutional background on the U.S. SPR. In Section 3, we propose three alternative joint econometric models of the global oil market and of changes in the SPR and discuss the estimation of these models. In section 4, we provide empirical evidence that changes in the SPR are endogenous with respect to the oil market and quantify the impact of exogenous SPR policy shocks on the real price of oil and other oil market variables. In section 5, we examine the question of how successful the major SPR policy releases since 1990 have been in lowering the real price of oil. We also examine the cumulative effect of building and expanding the SPR. In section 6, we quantify how much selling half of the oil in the SPR, as recently proposed by the White House, would lower the global price of oil (and hence the U.S. price of motor gasoline) and how much fiscal revenue it would generate. The concluding remarks are in section 8.

2. The U.S. Strategic Petroleum Reserve

The U.S. Strategic Petroleum Reserve was created by the Energy Policy and Conservation Act signed into law by President Ford in late 1975, following the Arab oil embargo of 1973/74. It is a complex of deep underground caverns along the Texas and Louisiana Gulf Coast that is used to stockpile crude oil for national emergencies. The U.S. SPR can hold up to 727 million barrels of crude oil. In November 2018, it contained about 650 million barrels of mostly light sweet crude, which corresponds to 7 days of global oil production and 82 days of U.S. crude oil imports. Although many other countries maintain strategic oil reserves, the U.S. strategic oil reserve is the largest of its kind, has a long history, and has been utilized extensively, making it a natural setting to study the effect of SPR releases.⁶

The size of the U.S. SPR is tailored to the risk of disruptions of U.S. oil imports and oil supplies. From 1977 to 1989, the size of the SPR steadily increased. It stabilized at close to 600 million barrels in the 1990s. In November 2001, President George W. Bush ordered the SPR to be increased to approximately 700 million barrels. When this target was reached in 2005, Congress directed the SPR to

⁶ Most members of the International Energy Agency maintain privately and publicly owned strategic oil reserves covering near 90 days of net oil imports. Other countries including China and India recently have begun building their own strategic oil reserves. There are no official data on the size of the Chinese government-controlled strategic oil reserves, but Chinese government-controlled oil stocks are likely to rival those of the United States and to surpass those of all other IEA members by 2020.

take actions to further expand the SPR to one billion barrels. The latter plan was terminated in 2011.

Although releases from the SPR receive disproportionate attention, the month-to-month change in the SPR is close to zero on average with a standard deviation of 3.4 million barrels. The distribution is fat-tailed, but not skewed. Figure 1 plots the change in the oil stocks held in the U.S. SPR from November 1977 to November 2018. Periods of releases of oil from the SPR alternate with periods of stock building. SPR oil releases have been a regular occurrence since 1990. Oil in the SPR may be drawn down in four distinct ways: (1) An emergency drawdown may be authorized by the President of the United States, as occurred during the Persian Gulf War in 1991, after Hurricane Katrina in 2005, and in response to the Libyan unrest in 2011. (2) The Department of Energy may conduct test sales. Test sales occurred in 1985, 1990 and 2014. (3) The Department of Energy may also implement exchange agreements with private companies, providing for a loan in the form of oil to be repaid with “interest” in kind within a certain date. This approach is a common response to temporary oil supply shortfalls caused by hurricanes, pipeline blockages, and ship channel closures. (4) Finally, lawmakers may direct the Department of Energy to sell oil from the SPR even in the absence of an emergency, typically in an effort to reduce the federal budget deficit.

3. Structural VAR Models of the Effect of Changes in the SPR

The starting point of our analysis is the global oil market model of Kilian and Murphy (2014), which has become the workhorse model for assessing the relative importance of oil demand and oil supply shocks for the evolution of the real price of oil (e.g., Kilian and Lee 2014; Kilian 2017; Cross 2019; Herrera and Rangaraju 2019; Zhou 2019).

Our VAR model specification includes the percent change in the global production of crude oil (Δq_t), as reported by the U.S. Energy Information Administration; a measure of cyclical variation in global real economic activity (rea_t) originally proposed by Kilian (2009); a proxy for the change in global crude oil inventories (including U.S. SPR oil stocks), as discussed in Kilian and Murphy (2014) and Kilian and Lee (2014) (Δinv_t); and the log real price of oil (p_t) obtained by deflating the

U.S. refiners' acquisition cost for imported crude oil by the U.S. CPI for all urban consumers.⁷ We also include the change in the U.S. stock of crude oil contained in the U.S. Strategic Petroleum Reserve (Δspr_t), as reported by the U.S. Energy Information Administration. Δspr_t and Δinv_t are expressed in barrels of crude oil. All data are monthly and have been seasonally adjusted. The sample extends from 1977.10 to 2018.10, reflecting the availability of the SPR data.⁸

Let $y_t = (\Delta q_t, \text{rea}_t, p_t, \Delta inv_t, \Delta spr_t)'$ be generated by the covariance stationary structural VAR(24) process of the form $B_0 y_t = B_1 y_{t-1} + \dots + B_{24} y_{t-24} + w_t$, where the stochastic error w_t is mutually uncorrelated white noise and the deterministic terms have been suppressed for expository purposes. Setting the lag order to 24 allows the model to capture long cycles in the real price of oil. The reduced-form errors may be written as $u_t = B_0^{-1} w_t$, where B_0^{-1} denotes the structural impact multiplier matrix, $u_t = y_t - A_1 y_{t-1} - \dots - A_{24} y_{t-24}$, and $A_l = B_0^{-1} B_l$, $l = 1, \dots, 24$. The $\{ij\}$ th element of B_0^{-1} , denoted b_{ij}^0 , represents the impact response of variable i to structural shock j , where $i \in \{1, \dots, 5\}$ and $j \in \{1, \dots, 5\}$. Given the reduced-form estimates, knowledge of B_0^{-1} suffices to recover estimates of the structural impulse responses, variance decompositions and historical decompositions from the reduced-form estimates, as discussed in Kilian and Lütkepohl (2017).

Let $w_t = (w_t^{\text{flow supply}}, w_t^{\text{flow demand}}, w_t^{\text{storage demand}}, w_t^{\text{other oil demand}}, w_t^{\text{spr}})'$, where $w_t^{\text{flow supply}}$ denotes a shock to the flow supply of oil, $w_t^{\text{flow demand}}$ denotes a shock to the flow demand for oil, $w_t^{\text{storage demand}}$ denotes a shock to storage demand (or, equivalently, speculative demand). Finally, w_t^{spr} denotes an exogenous policy shock to the stock of crude oil in the U.S. SPR, and $w_t^{\text{other oil demand}}$ is a conglomerate denoting all other shocks to

⁷ The Kilian index of global real economic activity is based on data for bulk dry cargo ocean shipping freight rates. It is the most widely used indicator of global real economic activity in the oil market literature. As discussed in Kilian and Zhou (2018), this index has several conceptual advantages compared with proxies for global industrial production when it comes to modeling the global market for crude oil. We use the corrected version of the index available from <https://sites.google.com/site/kilian2019/research/data-sets>, as discussed in Kilian (2019a).

⁸ Since $\Delta inv_t \equiv \Delta spr_t + \Delta spr_t^{\text{non}}$, the inclusion of Δspr_t and Δinv_t in the reduced-form VAR model is equivalent to specifying a model including $\Delta spr_t^{\text{non}}$ and Δspr_t . The advantage of our specification is that it facilitates the imposition of identifying restrictions on the change in total oil inventories.

the demand for oil such as shocks to preferences for oil or shocks to the oil inventory technology. All shocks are normalized to represent a shock that raises the real price of oil. Our analysis focuses on the shocks that have an explicit structural interpretation.

3.1. Identifying Restrictions

One of the challenges in assessing the effect of exogenous changes in the SPR is that different identifying assumption may yield different results. For example, much depends on how privately held oil inventories respond to SPR policy interventions. In the existing literature, it has been common to impose strong assumptions about the response of the private sector. One of our innovations is that we examine how alternative assumptions affect the estimates. Much also may depend on the extent to which changes in the SPR evolve endogenously in response to the evolution of the global oil market. We therefore consider three alternative structural models that span the universe of economically plausible specifications. We evaluate each model on its own and assess how much the response estimates are sensitive to the identifying assumptions. All models have in common that they consist of two blocks. One block consists of four variables that describe the global oil market. This block is the same in all three structural models. The other block consists of the change in the U.S. SPR. The structural models differ in how exogenous variation in the U.S. SPR is identified.

All three models employ sign restrictions on the oil market block of the structural impact multiplier matrix. We also impose bounds on the one-month price elasticities of oil demand and oil supply, which may be expressed as inequality restrictions on functions of selected impact responses. Finally, we impose dynamic sign restrictions on selected structural impulse response functions and narrative sign restrictions on the historical decompositions. These inequality restrictions render the models set-identified. The identifying restrictions on the responses within the oil market block are discussed at length in Kilian and Murphy (2014). We make only two minor changes.

First, we follow the more recent literature in imposing the sign restriction that the response of the real price of oil to the first three shocks is positive not only on impact, but for the first 12 months (e.g., Inoue and Kilian 2013; Kilian 2017). The same type of dynamic sign restriction is also imposed on the

response of the real price of oil to an exogenous increase in the SPR. In addition, we impose the dynamic sign restriction that a positive price response to a storage demand shock must be associated with a positive oil inventory response for the first year.

Second, we follow Zhou (2019) in imposing a more conservative upper bound of 0.04 on the impact price elasticity of oil supply than Kilian and Murphy (2012, 2014) did. This bound is consistent not only with economic theory, which implies that the one-month price elasticity of oil supply should be close to zero, given the high costs of shutting down and reopening conventional oil wells (see Anderson, Kellogg, and Salant 2018); it is also consistent with extraneous microeconomic evidence. Most of the oil sold in global markets is conventional crude oil. Our upper bound of 0.04 on the one-month aggregate global oil supply elasticity is about four standard errors larger than the estimate of the one-quarter price elasticity of the supply of conventional crude oil in the United States, as reported by Newell and Prest (2019) based on a comprehensive analysis of oil producers in Texas, North Dakota, California, Oklahoma and Colorado.⁹ Newell and Prest's corresponding supply elasticity estimate for shale oil is even lower, which makes economic sense given the time it takes to complete shale oil wells, allowing us to bound the overall elasticity by 0.04 (see Golding 2019).

3.1.1. Model 1

Model 1 treats innovations to the SPR as contemporaneously exogenous with respect to the global oil market, but allows for unrestricted feedback within the month to the oil market variables. This model allows us to empirically assess the response of non-SPR oil inventories (including, notably, privately held oil inventories), conditional on the premise that SPR shocks are predetermined with respect to the oil market.

The latter assumption is not unreasonable to the extent that SPR releases are driven primarily by domestic political considerations such as upcoming elections that are plausibly exogenous with respect to the oil market. Moreover, SPR releases in response to oil market events are not automatic, but

⁹ Similar results hold for oil producers outside the United States. For example, Kilian (2019b) shows that Saudi Arabia's price elasticity of oil supply is close to zero.

discretionary. For example, there was no SPR response to the 2003 Iraq War, the outbreak of the 1980 Iran-Iraq War, or the 2019 drone attack on Saudi Arabia, suggesting that SPR releases tend to be unpredictable. Even when SPR releases occurred in response to events in the oil market, there were often considerable delays, motivating the imposition of delay restrictions. For example, there was no direct SPR response to the invasion of Kuwait in August of 1990 until the oil release of January 1991.¹⁰

The absence of contemporaneous feedback from the oil market variables to the SPR in model 1 implies a block recursive structural impact multiplier matrix. The sign and exclusion restrictions on the elements of B_0^{-1} are described in expression (1):

$$\begin{pmatrix} u_t^{\Delta q} \\ u_t^{\text{rea}} \\ u_t^{\text{p}} \\ u_t^{\Delta \text{inv}} \\ u_t^{\Delta \text{spr}} \end{pmatrix} = \begin{bmatrix} - & + & + & b_{14}^0 & b_{15}^0 \\ - & + & - & b_{24}^0 & b_{25}^0 \\ + & + & + & b_{34}^0 & b_{35}^0 \\ - & - & + & b_{44}^0 & b_{45}^0 \\ 0 & 0 & 0 & 0 & b_{55}^0 \end{bmatrix} \begin{pmatrix} w_t^{\text{flow supply}} \\ w_t^{\text{flow demand}} \\ w_t^{\text{storage demand}} \\ w_t^{\text{other oil demand}} \\ w_t^{\text{exogenous spr}} \end{pmatrix}. \quad (1)$$

The parameter b_{45}^0 pins down the response of total oil inventories to an unexpected increase in the SPR stock of crude oil. If the overall inventory increase in response to a positive exogenous SPR shock is smaller than the increase in SPR stocks, indicating a partial offset of exogenous SPR shocks by changes in private oil inventories, we have that $b_{45}^0 < b_{55}^0$, yielding an additional cross-equation restriction. This situation arises if the private sector views public and private oil inventories as imperfect substitutes, as discussed in Wright and Williams (1982). In the absence of a response of private inventories (or zero substitutability), as postulated by Newell and Prest (2017), we have that $b_{45}^0 = b_{55}^0$.¹¹

3.1.2. Model 2

¹⁰ Although the U.S. Department of Energy ordered a small test sale in October 1990 in response to the invasion of Kuwait, even that test sale, which predated the Desert Storm emergency release, occurred a full two months after the invasion of Kuwait, consistent with the delay restriction in model 1.

¹¹ To be precise the difference between SPR oil stocks and total oil stocks includes both private inventories and strategic oil reserves held by other countries. It is theoretically possible for total inventories to decline on impact, for example, if there is no response of private inventories to an exogenous SPR release, but other governments systematically react to the U.S. release by releasing oil from their own strategic reserves. We abstract from this possibility because most U.S. SPR releases have occurred in response to domestic oil shortages and were not coordinated with other countries. The notable exceptions were the emergency drawdowns of 1991, 2005, and 2011. In the latter cases, however, the releases of other IEA member countries often occurred in the form of fuels rather than crude oil.

Of course, we cannot rule out a priori the possibility of contemporaneous feedback from shocks in the oil market to the SPR. Model 2 allows for such feedback. It is identified in the same manner as model 1, except that the block-recursive ordering is reversed, reflecting an alternative view of the economic structure of the global oil market:

$$\begin{pmatrix} u_t^{\Delta q} \\ u_t^{\text{rea}} \\ u_t^{\text{p}} \\ u_t^{\Delta \text{inv}} \\ u_t^{\Delta \text{spr}} \end{pmatrix} = \begin{bmatrix} - & + & + & b_{14}^0 & 0 \\ - & + & - & b_{24}^0 & 0 \\ + & + & + & b_{34}^0 & 0 \\ - & - & + & b_{44}^0 & 0 \\ b_{51}^0 & b_{52}^0 & b_{53}^0 & b_{54}^0 & b_{55}^0 \end{bmatrix} \begin{pmatrix} w_t^{\text{flow supply}} \\ w_t^{\text{flow demand}} \\ w_t^{\text{storage demand}} \\ w_t^{\text{other oil demand}} \\ w_t^{\text{exogenous spr}} \end{pmatrix}. \quad (2)$$

Model 2 allows SPR innovations to respond within the current month to all shocks in the global oil market, but it imposes that exogenous changes in the SPR affect the oil market variables only with delay of one month. The latter restriction amounts to imposing the assumption that non-SPR oil inventories (and, in particular, privately held oil inventories) change to exactly offset any SPR policy intervention in the impact period. Since, under this assumption, aggregate oil inventory holdings do not change in response to an exogenous SPR shock, the effect on the real price of oil and on other oil market variables is zero on impact.

The need to account for the response of private sector oil inventories is well established in the literature (Wright and Williams 1982). Clearly, if the private sector expects oil stocks from the SPR to be available in the event of an unexpected shortfall, it has no reason to maintain the same level of private inventories, as in the absence of SPR oil stocks. Thus, public and private oil stocks are substitutes to some extent. The assumption in model 2 that oil market participants view public and private oil stocks as perfect substitutes is best viewed as a limiting case. As shown in Wright and Williams (1982), this assumption makes sense if public and private storage involve the same cost function. In that case, in the absence of distributional concerns and market distortions, either the government or competitive profit-maximizing private firms can achieve the optimal level of storage. Although this situation is unlikely to hold in real life, we include model 2 for expository purposes and compare its impulse responses to alternative models.

3.1.3. Model 3

A more plausible identifying restriction is that changes in private inventories only partially offset exogenous changes in the SPR. The latter case seems particularly likely in light of the theoretical arguments in Wright and Williams (1982). Model 3 relaxes the assumption of a complete offset of exogenous changes in the SPR in model 2 without requiring the predeterminedness of SPR shocks imposed in model 1. The identification in model 3 is based on sign restrictions only:

$$\begin{pmatrix} u_t^{\Delta q} \\ u_t^{\text{rea}} \\ u_t^{\text{p}} \\ u_t^{\Delta \text{inv}} \\ u_t^{\Delta \text{spr}} \end{pmatrix} = \begin{bmatrix} - & + & + & b_{14}^0 & + \\ - & + & - & b_{24}^0 & - \\ + & + & + & b_{34}^0 & + \\ - & - & + & b_{44}^0 & + \\ b_{51}^0 & b_{52}^0 & - & b_{54}^0 & + \end{bmatrix} \begin{pmatrix} w_t^{\text{flow supply}} \\ w_t^{\text{flow demand}} \\ w_t^{\text{storage demand}} \\ w_t^{\text{other oil demand}} \\ w_t^{\text{exogenous spr}} \end{pmatrix}. \quad (3)$$

The exogenous SPR shock is defined as a shock that increases both SPR oil stocks and overall oil inventory holdings, which means that the change in the SPR is not completely offset by private oil inventory changes. If it were, then we would revert back to model 2. Since there is an overall increase in storage demand, the impact responses of global real activity and of the real price of oil must be positive, while that of global oil production is negative, exactly as in the case of a positive storage demand shock.

The only difference from a positive storage demand shock is that a positive SPR shock implies an increase in SPR oil stocks, whereas a positive storage demand shock implies a decline in SPR oil stocks on impact. In short, SPR oil stocks are drawn down on impact in response to an unanticipated increase in storage demand, but not enough to overturn the increase in overall oil inventories. The restriction $b_{33}^0 < 0$ is necessary to differentiate between storage demand shocks and SPR shocks. We can rule out the alternative assumption of a systematically positive impact response of the SPR stocks as economically implausible on a priori grounds, and the case of no response in the SPR stock is already covered by model 1.

Finally, if the overall inventory increase in response to an exogenous SPR shock is smaller than the increase in SPR stocks, indicating a partial offset of exogenous SPR shocks by changes in private oil inventories, we have that $b_{45}^0 < b_{55}^0$, yielding an additional cross-equation restriction. This situation would arise if the private sector views public and private oil inventories as imperfect substitutes, complementing

the case of perfect substitutes covered by model 2. In the absence of a response of private inventories (or zero substitutability), we have that $b_{45}^0 = b_{55}^0$.

3.1.4. The Role of Anticipation

The typical procedure for an SPR emergency release is a presidential declaration of an emergency, followed by a notice of sale issued by the Department of Energy about its willingness to accept offers for a certain amount of oil by a pre-specified date. An obvious concern with our approach to defining monthly SPR shocks is that pre-announced SPR releases may be anticipated by private agents before the SPR release takes place (Stevens 2014). This anticipation is not an issue for our models if there is little delay between the notice of sale and the actual SPR interventions. In particular, if both the advance notice and the SPR release fall into the same month, our identification of the effect of SPR shocks works even in the presence of anticipation. For example, following Hurricane Katrina in 2005, President Bush declared an energy supply disruption on September 2, the U.S. Department of Energy issued a notice of sales on September 6 and, by September 14, the results of the SPR sale were announced. In contrast, to the extent that the notice of sales occurs in a month prior to the actual release, as happened during the Libyan Crisis of 2011, the announcement effect may be partially absorbed into the storage demand shock. In that case, we would underestimate the overall effect of SPR releases on the real price of oil and our estimates may be viewed as lower bounds.

3.1.5. Narrative sign restrictions

For all three structural VAR models, the inequality and exclusion restrictions are complemented by additional narrative restrictions. Narrative restrictions refer to restrictions on the signs or relative magnitudes of structural shocks or of historical decompositions. They were first employed by Kilian and Murphy (2014) for selecting the most economically plausible models among the set of admissible structural models. This idea was subsequently generalized and formalized by Antolin-Diaz and Rubio-Ramirez (2018) and applied in Zhou (2019).

Motivated by the reasoning in Kilian and Murphy (2014, p.460, 469) and Kilian and Lee (2014, p.74), we postulate (1) that storage demand shocks cumulatively raised the log real price of oil by

at least 0.2 (or approximately 20%) between May and December 1979, consistent with evidence of a surge in inventory building in the oil market during that time, (2) that storage demand cumulatively lowered the log real price of oil by at least 0.15 between December 1985 and December 1986, after OPEC collapsed, and (3) that storage demand shocks raised the log real price of oil by at least 0.1 cumulatively between June 1990 and October 1990, reflecting market expectations that Iraq would invade its neighbors. The motivation for the latter restriction is that there must have been an increase in storage demand in the third quarter of 1990 because an oil supply disruption alone would have been associated with a reduction in oil inventories during this period, which is not what we see in the data (see Kilian and Murphy 2014).¹² Flow supply shocks are assumed to have raised the log real price of oil cumulatively by at least 0.1 between July and October of 1990, reflecting the invasion of Kuwait and the cessation of Iraqi and Kuwaiti oil production in early August. Finally, the cumulative effect of flow demand shocks on the log real price of oil between June and October of 1990 is bounded by 0.1, given that the oil price spike of 1990 was not associated with the global business cycle. These thresholds are chosen to be conservative. For example, few observers would attribute only a 10% increase in the real price of oil in 1990 to storage demand shocks. Likewise, the impact of oil supply shocks on the real price of oil in 1990 is likely to be much higher than 10%. Nevertheless, these weak narrative sign restrictions in conjunction are helpful in narrowing the range of admissible structural models. Our results are robust to reasonable variation in these bounds.¹³

3.2. Generating Draws from the Posterior Distribution

Although the oil market block in our three structural models resembles the model of Kilian and Murphy (2014), the econometric methods we employ differ from those originally employed by Kilian and Murphy. For expository purposes, we start with model 3, which is estimated by state-of-the-art Bayesian

¹² We include the month leading up to this war, given evidence in Kilian and Murphy (2014) that rising political tensions in the Middle East increased storage demand even before the war broke out.

¹³ In constructing the posterior distribution we bound the dominant root of the VAR process at 0.990655. This restriction implies that the effect of a one percent shock at the beginning of the sample on the model data is reduced to at most 1% at the end of a sample. This bound ensures that posterior draws from the historical decomposition closely resemble the actual historical data for the real price of oil. Without this bound, no meaningful analysis of the cumulative effects of the structural shocks on the real price of oil is possible. Nor would it make sense to impose narrative sign restrictions.

methods for sign-identified models subject to additional narrative restrictions, as discussed in Antolin-Diaz and Rubio-Ramirez (2018). We postulate a diffuse Gaussian-inverse Wishart prior for the reduced-form VAR parameters and a uniform prior for the rotation matrix Q . Let $A = [A_1, \dots, A_p]$ denote the autoregressive slope parameters and Σ_u the residual variance-covariance matrix. For a given realization of A and of the lower triangular matrix $P = chol(\Sigma_u)$ with positive diagonal elements, we draw realizations of the matrix Q from the space of $K \times K$ orthogonal matrices by generating at random many $K \times K$ matrices W consisting of $NID(0,1)$ draws, where K is the number of model variables. For each W , we apply the QR decomposition $W = QR$ with the diagonal of the upper triangular matrix R normalized to be positive. Then a candidate solution for B_0^{-1} is PQ , since $QQ' = I_K$. We use each of these candidate solutions in conjunction with A to construct the candidate structural models and their structural impulse responses.

Given the diffuse Gaussian-inverse Wishart prior distribution for the reduced-form parameters, this procedure may be repeated for a large number of reduced-form posterior draws for (A, Σ_u) to account for parameter estimation uncertainty. We retain the set of models that satisfy the static and dynamic restrictions on the impulse responses. The imposition of additional narrative restrictions requires us to reweight these draws with weights obtained from the importance sampler described in Antolin-Diaz & Rubio-Ramirez (2018). This reweighting is necessary because bounds on the historical decomposition restrict the space of the structural errors and hence the likelihood.

For model 2, this algorithm is modified further to account for the presence of additional zero restrictions on B_0^{-1} . For a given realization of A and P , we draw realizations of the matrix Q from the space of $K \times K$ orthogonal matrices by generating at random many $(K-1) \times (K-1)$ matrices W consisting of $NID(0,1)$ draws. For each W , we apply the QR decomposition $W = \bar{Q}R$ with the diagonal of the upper triangular matrix R normalized to be positive, and let

$$Q = \begin{bmatrix} \bar{Q} & 0 \\ 0 & 1 \end{bmatrix}.$$

Then a candidate solution for B_0^{-1} is PQ , as before. This algorithm ensures that every draw for B_0^{-1} is block recursive, as postulated in model 2. We construct the implied structural impulse responses and retain the structural models whose responses satisfy the static and dynamic sign restrictions. These candidate solutions must be reweighted in two steps. First, the imposition of zero restrictions on Q renders draws from the space of all possible Q matrices non-uniform and requires the use of an importance sampler based on numerical derivatives (see Arias, Rubio-Ramirez & Waggoner 2018). Second, as in the case of model 3, the imposition of narrative sign restrictions requires the use of another importance sampler.

Finally, for model 1, without loss of generality, we reorder the model variables such that Δspr_t is ordered first. For a given realization of A and P , we draw realizations of the matrix Q from the space of $K \times K$ orthogonal matrices by generating at random many $(K-1) \times (K-1)$ matrices W consisting of $NID(0,1)$ draws. For each W , we apply the QR decomposition $W = \bar{Q}R$ with the diagonal of the upper triangular matrix R normalized to be positive, and let

$$Q = \begin{bmatrix} 1 & 0 \\ 0 & \bar{Q} \end{bmatrix}.$$

Then the candidate solution PQ satisfies the exclusion restrictions imposed on B_0^{-1} in model 1. We then proceed exactly as for model 2 by implementing the two importance samplers and reweighting the draws. It should be noted that our approach to drawing at random from the space of Q matrices for models 1 and 2 is computationally more efficient than the procedure discussed in Arias et al. (2018) because it exploits the block recursive structure of the model.

4. Understanding the Effects of Exogenous Shocks to the SPR

Sign-identified VAR models generate no point estimates from which one could construct impulse response functions. Some users report so-called posterior median response functions instead, which are

obtained by connecting across horizons the pointwise posterior medians for individual impulse responses. Several studies have observed that this practice confounds estimates from different structural models and tends to distort the dynamics implied by the estimated models (see, e.g., Fry and Pagan 2011; Kilian and Murphy 2012; Inoue and Kilian 2013; Kilian and Lütkepohl 2017). Moreover, the associated pointwise impulse response error bands understate the true uncertainty about the model estimates. There are readily available econometric solutions to this problem in sign-identified VAR models, as discussed in Inoue and Kilian (2013), but not for models including additional exclusion restrictions. In this section, we therefore begin by reporting the full set of impulse response functions for all admissible structural models.

4.1. Understanding the Role of Exogenous SPR Policy Interventions

The responses to oil demand and oil supply shocks within the oil market for our models look similar to those reported in Kilian and Murphy (2014) or Zhou (2019), for example. We do not discuss these responses beyond noting that they conform with conventional views. The full set of admissible impulse response estimates for each model can be found in the not-for-publication appendix. Here we focus on what sets our models apart from the earlier literature, which is the inclusion of the SPR.

Figure 2 shows the response to a one-time exogenous release of oil from the SPR, normalized to one standard deviation of the SPR shock (or about 1.8 million barrels of crude oil). Although our model is based on data for SPR oil stocks and total oil stocks only, in Figure 2 we also show the implied responses for non-SPR oil stocks. Consider model 1, in which SPR shocks are presumed predetermined. Exogenous SPR releases cause a gradually building persistent decline in the stock of crude oil in the SPR. There is evidence of at least a partial offset by non-SPR inventories in the short run, but over time, non-SPR oil stocks tend to decline, causing an even larger decline in total oil stocks. Although the range of responses in Figure 2 includes models for which the response of total oil stocks is no larger than that of SPR oil stocks, the extent of the overall inventory decline in at least some admissible models is unexpected. This pattern could reflect in part coordination among government releases of strategic oil reserves, although it should be noted that other IEA members as part of their cooperative efforts often release petroleum products rather than crude oil. It may also simply reflect

estimation error. There is no evidence of much of a response in global oil production to exogenous SPR shocks. Likewise, the response of global real economic activity is not precisely estimated. Although model 1 imposes a decline in the real price of oil in response to an exogenous SPR release at intermediate horizons, we cannot rule out that the price response is effectively zero.

These results are remarkably robust to the alternative identifying assumptions employed in model 2, which endogenizes the SPR residuals, suggesting that the extent of the offset between SPR and private oil stocks on impact is less important than one might have thought a priori. Of course, this conclusion applies only under the maintained assumption that public and private inventories are perfect substitutes, which a priori is not very plausible.

Model 3, in contrast, endogenizes the SPR residuals without requiring this assumption. We found, however, that none of the solutions of model 3 among more than 60 billion posterior draws were admissible. To be precise, while otherwise reasonable, none of the model solutions satisfied the narrative sign restriction that the oil price spike in 1990, following the invasion of Kuwait, reflected an oil supply disruption and higher storage demand rather than higher flow demand. Thus, model 3 is not compatible with the data, which is why we do not include solutions for model 3 in Figure 2. Since only models 1 and 3 are consistent with imperfect substitutability between public and private oil inventories, the lack of support for model 3 strongly suggests that we may think of SPR changes as contemporaneously exogenous, as postulated in model 1. In the remainder of the paper, we therefore focus on model 1.

4.2. The average contribution of SPR shocks to the variability of the oil market variables

We now turn to the question of how much of the variability of the model variables is explained by exogenous SPR shocks on average over the sample. Table 1 reports the posterior means of the percent contribution of the flow supply, flow demand, storage demand and exogenous SPR policy shock to the variance of a given variable. By construction, the contributions of all five structural shocks in the model add to 100%. We focus on the posterior mean across admissible models rather than the posterior median because posterior quantiles violate the adding-up constraint.

As in similar oil market models, much of the variability in the real price of oil is explained by

flow demand shocks (49%), storage demand shocks (25%) and flow supply shocks (15%). Exogenous variation in the SPR explains only 8% of the variation in the real price of oil and that contribution is not much larger than its standard deviation. Thus, SPR policy shocks clearly are not an important determinant of the real price of oil. Likewise, global real activity is largely determined by flow demand shocks (56%), followed by flow supply shocks (15%), and storage demand shocks (11%). Exogenous SPR shocks account for only 6% of the variability in the global business cycle. As in the case of the real price of oil, they are not an important determinant and their contribution is only imprecisely estimated. The variability in the growth in global oil production is primarily explained by flow supply shocks (42%). Exogenous SPR shocks contribute only 8%, but that effect is precisely estimated.

The change in overall global oil inventories is primarily explained by storage demand shocks (30%), flow supply shocks (19%), and flow demand shocks (11%). Exogenous SPR shocks with 10% are also modestly important. In contrast, exogenous SPR shocks explain 55% of the variability in the change in U.S. SPR oil stocks, which is consistent with an important exogenous component in the changes in SPR holdings. The fact that flow demand, flow supply, and storage demand shocks jointly account for 27%, however, suggests that there is an important endogenous component in the SPR stock.

5. Case Studies

Although exogenous variation in the U.S. SPR is not very important in explaining the variability of the real price of oil on average, exogenous changes in the U.S. SPR may have important effects on the real price of oil during specific historical episodes. This question cannot be addressed by impulse response analysis, because policy interventions involve a sequence of exogenous SPR shocks rather than a one-time shock. In fact, different types of SPR interventions are characterized by different SPR shock patterns. An exchange agreement, for example, provides for a loan of oil to be repaid, in kind, within a certain date. It involves an initial SPR drawdown followed by a subsequent rebuilding of the SPR stocks, as the crude oil is being returned (possibly with a premium involving additional barrels serving as “interest”). In contrast, an emergency drawdown in response to a disruption of oil supplies involves one or more SPR releases over time.

We begin by identifying four episodes of major SPR releases and two episodes of major SPR additions. The exact timing of each episode is pinned down with the help of historical accounts provided by the U.S. Energy Information Administration. For each episode, we first isolate the sequence of exogenous SPR residuals associated with a given policy intervention period, and then simulate the implied changes in the real price of oil based on the structural moving average representation of the VAR model, setting all SPR shocks before and after the intervention period to zero.¹⁴ Our approach differs from earlier studies that focused on the effect of a hypothetical one-time policy shock rather than any real-life policy intervention. We report both the posterior median of the cumulative effect of the policy intervention at the end of the intervention period and the posterior median of the largest cumulative effect within this period (which allows the cumulative response to be hump-shaped), along with the 10% and 90% posterior quantiles of these statistics. We also report the number of months it takes to reach the largest cumulative effect.

5.1. How successful were SPR releases during key historical episodes?

Traditionally, oil has been released from the SPR in response to unexpected shortfalls in oil production or deliveries. For example, SPR releases took place in response to wars and civil unrest in OPEC member countries, in response to the closing of oil shipping channels, and in response to hurricanes affecting oil production in the Gulf of Mexico. In this section, we provide estimates of the cumulative effect of these policy interventions for selected episodes, allowing us to examine whether the perceived success of these operations is supported by more formal evidence. We show that SPR interventions may help stabilize the real price of oil, but that their effects tend to be considerably more modest than some observers had conjectured.

Desert Storm

The release of oil from the SPR during Operation Desert Storm in 1991, in particular, was widely credited with having calmed oil markets. The extent to which this release actually contributed to the oil

¹⁴ Our timing conventions are based on the “History of SPR Releases” of the U.S. Energy Information Administration at: <https://www.energy.gov/fe/services/petroleum-reserves/strategic-petroleum-reserve/releasing-oil-spr>.

price decline after November 1990 remains an open question, however. Table 2 shows that during the intervention period of 1990.10-1991.4 this policy indeed lowered the real price of oil by 6% cumulatively. Given a price of oil of \$29.88/barrel in September 1990, the median cumulative effect of less than \$2/barrel must be considered a modest success. The 80% credible region bounds the cumulative reduction by \$2.40/barrel. There is no evidence that this intervention prevented a much larger increase in the real price of oil. The advantage of our methodology is that it makes explicit for the first time the magnitude of the benefits of this SPR release.

The Heating Oil Exchange

The heating oil exchange of 2000.9-2004.6 was a different policy intervention from 1990/91 in that it was not an emergency drawdown, but an exchange agreement with the refining industry. It was understood that the oil released from the SPR was to be returned eventually. This difference is reflected in the estimated cumulative effect over the entire intervention period, which is +\$2 per barrel rather than negative. There is evidence, however, that at least temporarily this program lowered the real price of oil by 14% (which translates to a decline of \$4 per barrel). These results show that there are systematic differences in how SPR exchanges and emergency drawdowns affect the real price of oil.

Hurricane Katrina

Hurricane Katrina in 2005.9-2005.11 prompted the Department of Energy to release oil from the SPR. The motivation for this release was as much psychological as it was real because the hurricane resulted in a shutdown of refiners along the Gulf coast, which reduced the demand for oil from that region by more than the flow of the supply from oil producers in the Gulf region had declined (see Kilian 2010). Thus, the primary effect of the Hurricane Katrina was a gasoline shortage, not a crude oil shortage, making the release of crude oil from the SPR a blunt tool for this crisis. With refiners along the Gulf Coast shut down, injections of crude oil did nothing to stimulate gasoline production by these refiners. The only real effect of the 2005 SPR oil release was to help U.S. refiners outside of the Gulf region to maintain or expand their fuel production.

The Bush administration immediately offered emergency loans of close to 10 million barrels of

crude to refineries whose supply of crude had been disrupted. By early September 2005, the U.S. Department of Energy offered an additional 30 million barrels of crude oil for sale, of which only 11 million barrels were eventually purchased, however, as the market was eager for additional supplies of refined products rather than crude. Table 2 shows a cumulative effect of these SPR oil releases on the real price of oil of -6%. The reduction of \$3/barrel (given the price of oil of \$58.67/barrel in August 2005) was at best modest. Ironically, the 2005 SPR release is often considered among the most successful oil market interventions (see Bordoff et al. 2018). Our evidence suggests that the response of the Department of Energy had quite modest effects on the real price of crude oil.

The Libyan Crisis

In contrast, the SPR intervention during 2011.6-2011.10 prompted by the Libyan oil supply crisis and concerns over the stability of the Middle East during the Arab Spring was much more successful. It resulted in a cumulative reduction of about 11% in the real price of oil. In dollar terms, these reductions amounted to approximately \$12/barrel, relative to a price of 107.98/barrel in May 2011.

5.2. To what extent has the creation of the SPR raised the real price of oil?

Much attention has been given to the reduction in the real price of oil associated with SPR releases. The adverse effects of creating the SPR oil stock in the first place, however, should not be overlooked. In this section we quantify these effects.

The creation of the SPR between 1977 and 1990 and its expansion between 2001 and 2005

The SPR stock was built between 1975.11 and 1990.8. During this period, no oil was released except for a small test sale in 1985. An obvious concern is that building strategic reserves during the oil price surge of the late 1970s may have contributed to rising oil prices. Table 2 confirms that the creation of the SPR cumulatively raised the real price of oil by as much as 29% at its peak, which translates to a cumulative increase of \$4/barrel. The cumulative effect by 1990.8, however, is only 9% (or about \$1/barrel). Likewise, the decision to expand the SPR from 2001.11-2005.7 was associated with a peak cumulative increase in the real price of oil of 21% (or \$4 per barrel). By 2005.7, the cumulative effect had fallen to about 18% (or \$3/barrel), however.

6. The Fiscal Effects of the White House Plan of Selling Half of the SPR within Ten Years

In 2018, the White House proposed selling a significant portion of the SPR to reduce the federal deficit. The Trump administration advocated putting 270 million barrels of SPR crude oil on the market over a ten-year period, stating that a smaller SPR is sufficient to meet international obligations and emergency needs.¹⁵ Similar policies – on a smaller scale – were already implemented under the Clinton and Obama administrations. Recently, the conservative think tank Heritage Foundation even proposed selling off the entire U.S. SPR. Moreover, as noted by Bordoff, Halff and Losz (2018), since 2015, the U.S. Congress has enacted five pieces of legislation calling for the sale of SPR oil. Taken together, these authorized sales will reduce the size of the SPR by 410 million barrels by the end of 2027. In addition to raising fiscal revenue, these policies have the added potential benefit of lowering U.S. retail fuel prices, thereby raising the purchasing power of U.S. consumers and stimulating the U.S. economy (e.g., Edelstein and Kilian 2009).

Table 3 examines the White House plan through the lens of model 1. We address the question of what the cumulative impact of a major reduction in the SPR would be on the real price of oil, on U.S. retail gasoline prices, and on fiscal revenue. For expository purposes, we postulate that the implementation of the proposal starts after the end of our sample in September 2018, when the price of oil was \$66.42 and the U.S. price of motor gasoline was 293 cents/gallon. All dollar figures in the table are in real terms relative to this benchmark. It is assumed that the proposal is implemented by selling the same amount of oil each month for the duration of the policy intervention. The targeted change in the SPR is -2.25 million barrels for each month until September 2028. We start by computing, for each posterior draw, the expected path of the change in SPR inventories in the absence of future shocks based on the structural moving average decomposition of the change in the SPR in the estimated VAR model.

Let

$$y_t \approx \sum_{i=0}^{t-1} \Theta_i w_{t-i},$$

¹⁵ See, e.g., Mufson, S., and C. Mooney, “Trump seeks to sell off half of the Strategic Petroleum Reserve,” Washington Post, May 23, 2017.

where y_t denotes the K -dimensional vector of the data at date $t = 1, \dots, T$, Θ_i is the matrix of impulse response coefficients at horizon i whose elements are denoted by $\theta_{kl,i}$, where $k = 1, \dots, K$, and $l = 1, \dots, K$, and w_t denotes the K -dimensional vector of structural shocks. Then the fitted value is

$$\hat{y}_{kt} = \sum_{j=1}^K \hat{y}_{kt}^{(j)},$$

where $\hat{y}_{kt}^{(j)} = \sum_{i=0}^{t-1} \theta_{kj,i} w_{j,t-i}$, $j = 1, \dots, K$, $K = 5$, and the k th variable is the change in the SPR oil stock.

The expected path in the absence of future shocks is obtained by iterating this expression forward for 120 months beyond the end of the sample, T , based on $w_{T+1} = \dots = w_{T+120} = 0$.

We compare this expected path with the targeted change in the SPR adjusted for the mean change in the SPR, denoted $x_{T+1}, \dots, x_{T+120}$, as shown in Figure 3. We then recursively infer, for each posterior draw, the magnitude of the shock required for the predicted change in the SPR, $\hat{x}_{T+1}, \dots, \hat{x}_{T+120}$, to match this value. Let w_{kj}^c denote the counterfactual exogenous SPR shock at date j , where $w_{kj}^c = w_{kj}$ for $j \leq T$.

Then

$$w_{kj}^c = \frac{(x_j - \hat{x}_{kj})}{\theta_{kk,j}}, \quad j = T + 1, \dots, T + 120,$$

where $\hat{x}_{kt} = \sum_{j=1}^K \hat{x}_{kt}^{(j)}$ and $\hat{x}_{kt}^{(j)} = \sum_{i=0}^{t-1} \theta_{kj,i} w_{j,t-i}^c$ are updated recursively, while $\hat{x}_{kt}^{(j)} = \hat{y}_{kt}^{(j)}$ for $j \neq k$. The resulting counterfactual shock sequences, an example of which is shown in Figure 4, are then imposed in generating future counterfactual realizations of the real price of oil from the structural moving average representation for $T + 1, \dots, T + 120$. This allows us to recover, for each posterior draw, the cumulative percent change in the real price of oil.

We report the posterior quantiles of the cumulative change in the real price of oil as well as the implied change in the U.S. real price of motor gasoline. We also report the posterior quantiles of the implied fiscal revenue. The fiscal revenue, for each posterior draw, is obtained by summing across the 120 months the product of the 2.25 million barrels of oil releases each month and the counterfactual real

price of oil for that month, expressed in 2018.9 dollars. This novel approach is designed to preserve the dynamics as well as the conditional co-movement across variables for each posterior draw. It provides a more accurate representation of the posterior distribution of the object of interest than the practice of computing cumulative changes of pointwise posterior medians (or products of pointwise posterior medians).

Our concern is not whether the White House plan is economically sensible, but whether it would reach its objectives.¹⁶ Table 3 shows that such a policy intervention would lower the real price of oil by 13%. The 68% credible set for the cumulative change in the real price of oil in model 1 ranges from -\$23/barrel to +\$4/barrel with a posterior median reduction of \$6/barrel. Given an oil cost share of about 50% in producing motor gasoline in the United States, the implied reduction in the real price of gasoline is a modest 13 cents/gallon (with the 68% credible set including changes anywhere between -50 cents and +9 cents per gallon).¹⁷

Evidence of reductions in the real price of gasoline are potentially important because selling off the SPR might be expected to stimulate the U.S. economy in addition to the fiscal benefits, as consumers have more pocket money to spend on other goods and services. Such considerations traditionally have played an important role in the policy debate. However, these discretionary income gains are based on the premise that the United States is a net importer of crude oil and gasoline. Given that trade in crude oil and petroleum products is approximately balanced as of late 2019, this effect is likely to be close to zero in the evaluation period. Moreover, to the extent that the real price of oil falls, the aggregate stimulus is likely to be dampened further by cutbacks in the fixed investment of the oil sector.

These real economic effects of reducing the SPR, in any case, are only incidental, as the main objective of the White House plan is to generate federal fiscal revenue. Table 3 shows that the

¹⁶ For a review of the arguments for and against reforming the SPR see, e.g., Bordoff et al. (2018).

¹⁷ The cost share of oil in producing gasoline has remained remarkably stable historically, and, hence, may be treated as known. Moreover, historical evidence on the passthrough from crude oil to U.S. gasoline prices suggests that this passthrough is not only rapid, but has not changed much over time, making it reasonable to assume that it will not change over our evaluation period.

posterior median cumulative fiscal revenue over the next decade is \$17 billion in real terms. This estimate only slightly exceeds the White House projection of \$16.6 billion in fiscal revenue. Being able to confirm this estimate illustrates the importance of formal quantitative modeling in assessing the likely effects of SPR interventions. In addition, we are able to assess the precision of these estimates. The 84% quantile of the posterior distribution of the fiscal revenues from the White House plan is \$19.8 billion and the 16% quantile is \$13.6 billion. Thus, even after accounting for the estimation uncertainty about the fiscal benefits of the White House plan, the projected fiscal benefits of this plan of \$16.6 billion appear realistic. Accounting for inflation would raise these estimates only slightly.

Our computations abstract from the cost of implementing the White House proposal, which are negligible compared with the expected benefits. Our analysis also abstracts from any fiscal stimulus that would be caused by spending the fiscal revenue in question. Finally, we do not model the impact of lower gasoline prices on federal gasoline tax revenue. The federal gasoline tax of 18 cents per gallon (which has been constant since 1993) is not ad valorem and hence does not vary with the price of gasoline. Thus, federal gasoline tax revenue would be expected to increase only to the extent that gasoline consumption increases, as the gasoline price falls. Given the expected magnitude of the gasoline price decline in Table 3, this effect is likely to be small.¹⁸

Policy simulations such as the results in Table 3 are subject to important caveats (see Kilian and Lütkepohl 2017). One concern is that the stability of the structural model is in doubt when implementing a policy intervention requires structural shocks that are larger in magnitude than historical structural shocks. Figure 4 shows that this critique is not a concern here because the shocks in question are well within the historical range. The other concern is that long sequences of policy shocks of the same sign may cause agents to change their behavior, causing the structural model coefficients to drift. The latter concern in Figure 4 applies only to the last four years of this policy intervention, when the

¹⁸ In addition, a decline in gasoline prices also implies a reduction in state tax revenue from ad valorem gasoline taxes, which will reduce state revenue in at least some states, offsetting in part the increased federal revenue from selling off the SPR. We abstract from this possibility because the state-level fiscal implications do not seem to play a role in the White House proposal, but note that these offsetting effects deserve additional study in future research.

policy shocks are persistently negative, but also quite small. Thus, one would not expect this particular policy intervention to undermine the stability of the structural model.

7. Conclusion

Quantifying the effect of exogenous SPR policy interventions is essential for policymakers because it helps pin down the sequence of interventions required to achieve the targeted change in the real price of oil. In fact, without a quantitative structural model, a policymaker is unable to determine to what extent the changes in the real price of oil observed after a policy intervention can be attributed to the policy intervention and to what extent they would have occurred even in the absence of the intervention.

We presented an econometric framework that allows the construction of explicit counterfactuals for the estimation of the causal effects of SPR policy interventions. We provided evidence that the size of the SPR endogenously responds to the state of the oil market. We showed that only 55% of the variability in the change in the SPR oil stock is accounted for by exogenous policy shocks. We provided evidence that that exogenous SPR releases are at least partially offset by increases in private sector oil inventories in the short run, which dampens the response of the real price of oil to SPR policy shocks in the short run, but has little effect on the longer-run response. Overall, exogenous shocks to the SPR have not been a major determinant of historical oil price fluctuations.

We estimated that a hypothetical one-time exogenous reduction in the SPR by one standard deviation (or 1.8 million barrels of crude oil) lowers the real price of oil by anywhere between 0% and 3% within one quarter. At longer horizons, the real price of oil may decline anywhere between 0% and 5%. Using a novel econometric approach, we established that some of the SPR oil releases that took place since 1990 lowered the price of oil by as much as \$12/barrel cumulatively. There are important differences between the cumulative effects of emergency drawdowns and SPR exchanges, however. The former interventions ultimately raise the real price of oil, whereas the latter lower the real price of oil. We also showed that the initial creation and subsequent expansion of the SPR caused the price of oil to rise by \$4/barrel.

Finally, we estimated that selling half of the oil in the SPR over ten years, as recently proposed by the White House, would likely raise fiscal revenues by \$17 billion, slightly in excess of the \$16.6 billion projected by the White House. The 68% error band includes values between \$13.6 and \$18.9 billion. In addition, the resulting reduction in the global price of oil is expected to cause a change in the real price of U.S. motor gasoline of -13 cents/gallon. This gasoline price decline is unlikely to generate an additional aggregate stimulus for the U.S. economy, however, given that the U.S. trade balance for crude oil and petroleum products is nearly balanced.

References

- Anderson, S.T., Kellogg, R., and S.W. Salant (2018), “Hotelling under Pressure,” *Journal of Political Economy*, 126, 984-1026.
- Antolín-Díaz, J., and J.F. Rubio-Ramírez (2018), “Narrative Sign Restrictions”, *American Economic Review*, 108, 2802-2829.
- Arias, J.E., Rubio-Ramirez, J.J., and D.F. Waggoner (2018), “Inference Based on Structural Vector Autoregressions Identified by Sign and Zero Restrictions: Theory and Applications,” *Econometrica*, 86, 685-720.
- Bordoff, J., Halff, A., and A. Losz (2018), “New Realities, New Risks: Rethinking the Strategic Petroleum Reserve,” Center on Global Energy Policy, Columbia University.
- Cross, J. (2019), “The Role of Uncertainty in the Market for Crude Oil,” manuscript, BI Norwegian Business School.
- Edelstein, P., and L. Kilian (2009), “How Sensitive Are Consumer Expenditures to Retail Energy Prices?,” *Journal of Monetary Economics*, 56, 766-779.
- Fry, R., and A.R. Pagan (2011), “Sign Restrictions in Structural Vector Autoregressions: A Critical Review,” *Journal of Economic Literature*, 49, 938-960.
- Golding, G. (2019), “Don’t Expect U.S. Shale Producers to Respond Quickly to Geopolitical Supply Disruption,” <https://www.dallasfed.org/research/economics/2019/1003>.
- Herrera, A.M., and S.K. Rangaraju (2019), “The Effect of Oil Supply Shocks on U.S. Economic

- Activity: What Have We Learned?" forthcoming: *Journal of Applied Econometrics*.
- Inoue, A., and L. Kilian (2013), "Inference on Impulse Response Functions in Structural VAR Models," *Journal of Econometrics*, 177, 1-13.
- Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review*, 99, 1053-1069.
- Kilian, L. (2010), "Explaining Fluctuations in U.S. Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market," *Energy Journal*, 31, 87-104.
- Kilian, L. (2017), "The Impact of the Fracking Boom on Arab Oil Producers," *Energy Journal*, 38, 137-160.
- Kilian, L. (2019a), "Measuring Global Real Economic Activity: Do Recent Critiques Hold Up to Scrutiny?," *Economics Letters*, 178, 106-110.
- Kilian, L. (2019b), "Facts and Fiction in Oil Market Modeling," manuscript, Federal Reserve Bank of Dallas.
- Kilian, L., and D.P. Murphy (2012), "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association*, 10, 1166-1188.
- Kilian, L., and D.P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454-478.
- Kilian, L., and T.K. Lee (2014), "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," *Journal of International Money and Finance*, 42, 71-87.
- Kilian, L. and H. Lütkepohl (2017). *Structural Vector Autoregressive Analysis*, Cambridge University Press: Cambridge.
- Kilian, L., and X. Zhou (2018), "Modeling Fluctuations in the Global Demand for Commodities," *Journal of International Money and Finance*, 88, 54-78.
- Kilian, L., and X. Zhou (2019), "Oil Prices, Exchange Rates and Interest Rates," manuscript, Federal Reserve Bank of Dallas.

- Muntaz, H., and P. Surico (2009), “The Transmission of International Shocks: A Factor-Augmented VAR Approach,” *Journal of Money, Credit and Banking*, 41, 71-100.
- Newell, R.G. and B.C. Prest (2017), “Informing SPR Policy through Oil Futures and Inventory Dynamics,” RFF Working Paper 17-19.
- Newell, R.G., and B.C. Prest (2019), “The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata,” *Energy Journal*, 40, 1-30.
- Rubio-Ramirez, J.F., Waggoner, D., and T. Zha (2010), “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference,” *Review of Economic Studies*, 77, 665-696.
- Stevens, R. (2014), “The Strategic Petroleum Reserve and Crude Oil Prices,” manuscript, UC Berkeley.
- Wright, B.D., and J.C. Williams (1982), “The Roles of Public and Private Storage in Managing Oil Import Disruptions,” *Bell Journal of Economics*, 13, 341-353.
- Zhou, X. (2019), “Refining the Workhorse Oil Market Model,” forthcoming: *Journal of Applied Econometrics*.

Table 1: What Drives the Variability of the Model Variables? (Percent)

	Flow Supply	Flow Demand	Storage Demand	SPR policy
Real Oil Price	14.5 (11.2)	48.9 (11.0)	20.0 (9.6)	7.9 (5.6)
Change in Oil Inventories	19.2 (7.1)	11.1 (3.3)	29.9 (7.1)	9.6 (2.5)
Change in SPR Oil Stocks	8.3 (3.8)	10.4 (5.1)	8.2 (3.7)	55.3 (9.4)
Global Real Activity	15.3 (5.9)	56.2 (12.3)	11.2 (5.6)	6.0 (4.8)
Global Oil Production Growth	42.2 (15.3)	7.8 (1.7)	10.7 (2.4)	8.0 (2.7)

NOTES: Posterior means based on model 1 with posterior standard deviations in parentheses. The contribution of the other oil demand shock corresponds to the remainder.

Table 2: Effects of Selected SPR Policy Interventions on the Real Price of Oil (%)

	Intervention	Intervention Period	Cumulative Effect at End of the Intervention Period	Largest Cumulative Effect within Intervention Period	Number of Months until Largest Cumulative Effect
SPR releases	Desert Shield/Storm	1990.10-1991.4	-6.1 (-8.0, -2.2)	-6.2 (-8.0, -2.2)	7
	Heating Oil Exchange	2000.9-2004.6	8.5 (3.5, 24.4)	-13.6 (-18.7, -8.5)	8
	Hurricane Katrina	2005.9-2005.11	-5.7 (-7.1, -2.7)	-5.7 (-7.1, -2.7)	3
	Libyan Crisis	2011.6-2011.10	-11.0 (-17.5, -6.8)	-11.5 (-17.5, -7.3)	4
SPR additions	Creation of SPR	1975.11-1990.8	8.8 (-3.4, 24.3)	29.0 (21.2, 63.3)	86
	Expansion of SPR*	2001.11-2005.7	17.9 (3.0, 34.8)	20.6 (13.7, 36.9)	43


NOTES: The cumulative effects in model 1 are measured by posterior medians with 10% and 90% posterior quantiles in parentheses. * This episode partially overlaps with the heating oil exchange episode. Since the latter ultimately caused an accumulation of SPR stocks, it makes sense to view these episodes in conjunction.

Table 3: The Effect of the White House Proposal of Selling Half of the SPR within Ten Years

	Posterior quantiles		
	16%	50%	84%
Cumulative change in real price of oil	-\$22.7/bb	-\$6.0/bb	\$4.1/bb
Cumulative change in real retail gas price	-50c/gallon	-13c/gallon	9c/gallon
Cumulative fiscal revenue	\$13.6 billion	\$17.0 billion	\$18.9 billion


NOTES: The White House proposal involves selling off 270 million barrels of crude oil from the SPR over 120 months. The implementation of the proposal is assumed to start in September 2018, when the oil price was \$66.42/barrel and the U.S. price of motor gasoline was 293c/gallon. All dollar figures in the table are in real terms relative to this benchmark. It is assumed that the proposal is implemented by selling the same amount of oil each month for the duration of the policy intervention. All computations are based on model 1.

Figure 1: Changes in the Stock of Crude Oil in the U.S. SPR



NOTES: The data source is the U.S. EIA's *Monthly Energy Review*.

Figure 2: Responses to a One-Time Exogenous SPR Release




NOTES: The exogenous SPR release is normalized to one standard deviation (or about 1.8 million barrels of crude oil).

Figure 3: Scenario for the Change in the SPR Over the Next Ten Years



NOTES: The SPR changes have been deseasonalized and demeaned. The vertical line marks the beginning of the implementation of the plan in October 2018.

Figure 4: Representative Counterfactual SPR Shock Sequence for Model 1



NOTES: Sequence of exogenous SPR shocks required to implement the White House plan of reducing the SPR oil stock by 2.25 million barrels each month over ten years. The vertical line marks the beginning of the implementation of the plan in October 2018.

Not-for Publication Appendix

Figure A1: Admissible Impulse Responses in Model 1

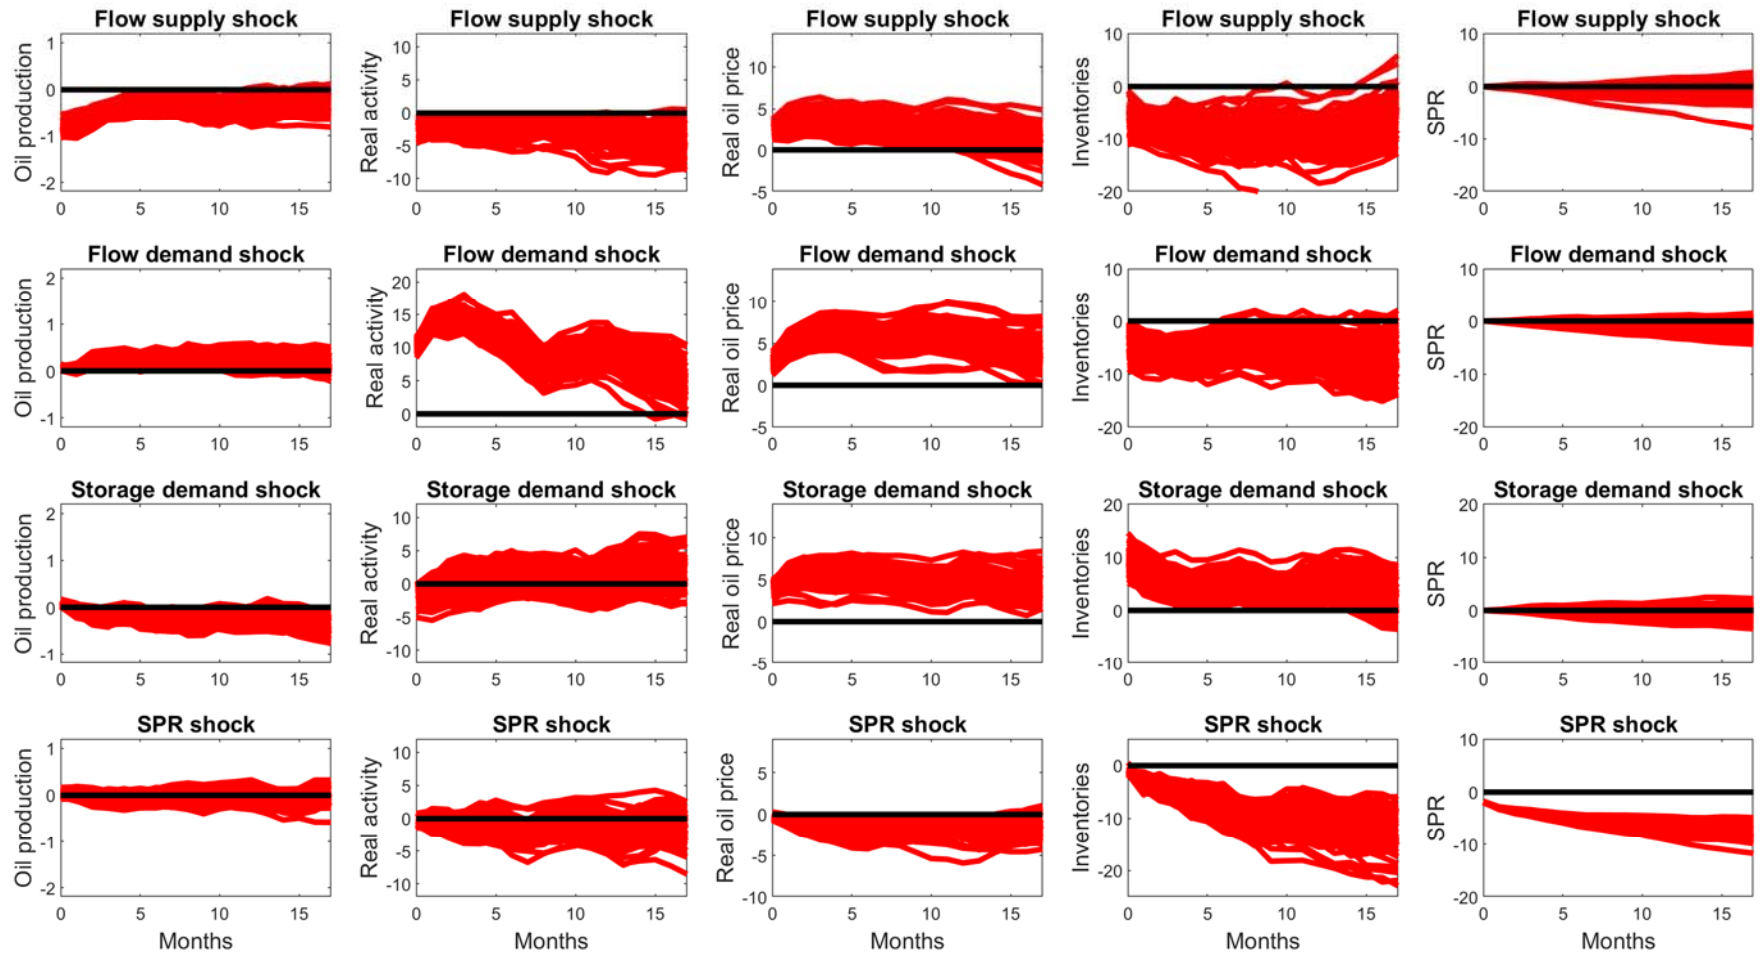


Figure A2: Admissible Impulse Responses in Model 2

