Informing SPR Policy Through Oil Futures and Inventory Dynamics

Richard G. Newell Brian C. Prest *

May 15, 2017

Abstract

This paper examines how information on the time pattern of expected future prices for crude oil, based on the term structure of futures contracts, can be used in determining whether to draw down, or contribute to the Strategic Petroleum Reserve (SPR). Such price information provides insight on expected changes in the supply-demand balance in the market and can also facilitate cost-effective transitions for SPR holdings. Backwardation in futures curves suggests that market participants expect shocks to be transitory, creating a stronger case for SPR releases. We use vector autoregression to analyze the relationship between the term structure of futures contracts, the management of private oil inventories, and other variables of interest. This relationship is used to estimate the magnitude of the impacts of SPR releases into the much larger global inventories system. Impulse response functions estimate that a strategic release of 10 million barrels will reduce spot prices by up to 4% and mitigate backwardation by approximately 1.5 percentage points. Historical simulations suggest that past releases reduced spot prices by 20% to 30% and prevented steep backwardation in futures curves, relative to a no-release counterfactual. This research can help policymakers determine when to release SPR reserves based on economic principles informed by market prices. It also provides an econometric model that can help inform the amount of SPR releases needed to achieve given policy goals, such as reductions in prices or spreads.

*Newell: Resources for the Future, 1616 P St NW, Washington, DC 20036, Duke University, Box 90328, Durham, NC 27708, and NBER (email: newell@rff.org). Prest: Duke University, Box 90328, Durham, NC 27708 (email: brian.prest@duke.edu). We acknowledge funding from the Department of Energy and NBER under grant 36269.00.04.00. We thank Christiane Baumeister and Lutz Kilian for helpful comments and data.

1 Introduction and Motivation

Understanding the relationship between crude oil prices, inventories, and market expectations is crucial for understanding not only private market dynamics, but also the value and effectiveness of public inventories, particularly strategic reserves. If a significant portion of oil supply is disrupted, how would releasing oil from the U.S. Strategic Petroleum Reserve (SPR) and other international oil stock holdings affect oil prices? This depends on a number of factors, including whether the disruption is expected to intensify or attenuate, the elasticity of oil demand and supply, and whether members of the Organization of the Petroleum Exporting Countries (OPEC) are likely to react to such a release. While these questions are important to guiding SPR drawdown policy, economic research on this topic is highly limited. This leaves policymakers with little economically-driven guidance. As a result, SPR release policy is typically informed by a combination of back-of-the-envelope estimates, qualitative market assessment, and simply "counting barrels", whereby policymakers release an amount of crude oil equal to the estimated size of the disruption.

This paper argues that the term structure of futures prices for crude oil can provide an indicator to policymakers on the severity of oil disruptions, and its relationship to inventories can help guide the magnitude of any response using strategic reserves. In particular, the spread between long- and short- dated crude futures prices reflect market expectations about the persistence of supply and demand shocks.[1] In particular, SPR releases are more effective and appropriate in response to temporary supply shocks, and less so in the face of persistent shocks. The degree of contango (a positive spread between future and current prices) or backwardation (a negative spread) in crude oil markets provides market-based information about how persistent supply shocks are likely to be, with a higher degree of backwardation suggesting transitory shocks. This differs from the common focus on the *levels* of spot prices, but economic theory clearly implies that price *spreads* are the more important factor for stockpiling decisions (e.g., Fama and French 1987, 1988; Pindyck 2001).[2] This suggests that policymakers should look to oil price spreads when considering SPR releases. It also suggests that better understanding of the dynamic relationships between spreads and private inventories can offer insights for when to release strategic reserves.

Aside from the question of *whether* to release oil reserves, policymakers also face the question of *how much* to release to achieve given policy goals. To consider this question, this paper also uses a vector auto-regression

[1] Throughout this paper, the spread is defined as the difference between the 12-month and 1-month real crude oil futures prices: $p_{12} - p_1$, where prices are expressed in log terms to represent percentage differences. We also note the important caveat discussed in Baumeister and Kilian (2016b) that futures prices also reflect a risk premium. We discuss this caveat more below, and present a sensitivity analysis in the appendix.

[2] We use the term "spot" price informally as a shorthand to refer to the prompt-month oil future price, distinguishing it from longer-term futures prices.

(VAR) model to empirically analyze the dynamic relationships between oil spot prices, price spreads, and privately-held commercial inventories. Once estimated, we use this model to estimate the impacts of a shock to private inventories (say, from an emergency infusion from public inventories) on oil prices, spreads, and other variables of interest. These estimated responses show strong relationships between inventories and spreads, with a 1% crude oil inventory shock—equivalent to approximately 10 million barrels of crude oil—temporarily increasing spreads by 1 to 2 percentage points (i.e., a reduction in backwardation, or increase in contango). This inventory shock reduces both spot and futures prices of crude oil. Spot prices temporarily fall by about 3% to 4%, while 12-month futures prices fall by a smaller amount, explaining the increasing spread. Prices fall despite some potential offset by increased OPEC supply, which is imprecisely estimated in our model.

We also use the model to simulate counterfactual release scenarios during relevant historical periods. These periods include times when oil was released from the SPR and times when such releases were considered but not executed. Simulations for 2005 (Hurricane Katrina) and 2011 (Arab Spring) suggest that the SPR releases in those years did indeed prevent spreads from entering steep backwardation in the short run as well as prevented temporary steep rises in oil prices. A simulation for 2003, when oil futures entered steep backwardation but no SPR release occurred,[3] suggests that a release during this period could have muted the strong backwardation and undone some of the price increases. These simulations also provide guidance on the size of SPR releases needed to achieve given policy goals, such as dollar-reductions in oil prices or changes in price spreads.

2 Background and Literature

Governments have long held strategic petroleum reserves. In the wake of the 1973 OPEC oil embargo, the International Energy Agency (IEA) was established with a key purpose being the coordination of strategic reserves to respond to oil supply disruptions. The U.S. Strategic Petroleum Reserve was established shortly thereafter in 1975 for the same reason. Emergency reserves have been released on three major occasions since then: during the 1991 Gulf War, during Hurricane Katrina in 2005, and during the 2011 Arab Spring.[4] There were also other times when SPR releases were considered but ultimately not used, such as in early 2003 when the Venezuelan oil strike and the Iraq War disrupted oil production.[5]

[3]This was amid the Venezuelan oil strike and run-up to the Iraq War, which together disrupted millions of barrels per day. Some experts argued that this case clearly called for SPR releases, but no releases were ultimately made.

[4]Reserves have been released on other occasions, but not on an emergency basis. These mostly include test sales and loans.

[5]See Dobbs, Michael. "Oil Reserve Is 'First Line of Defense' for U.S.", *New York Times*, 18 Feb. 2003.

Despite the importance of SPR release decisions, the economics literature provides little guidance for optimal policy. Instead, policymakers typically rely on expert judgment and back-of-the-envelope estimates of the impacts of releases on oil prices. A 2006 GAO report asked a group of experts about SPR policy, who generally suggested that strategic releases should be used to offset supply disruptions that the market cannot make up on its own (e.g., through private inventory drawdowns or supply responses in other regions).[6] In other words, a policy of "counting barrels" that are disrupted on net, and releasing that amount of reserves. Indeed, in its recent strategic assessment and benefit analysis for the SPR, the Department of Energy assumed exactly this decision rule.[7]

This policy of counting barrels is inadequate because it does not incorporate market expectations about the expected persistence of supply disruptions. Further, it embodies the outdated notion that the SPR is for addressing physical shortages. In a globally integrated market, the SPR primarily addresses the economic impacts of price shocks—both its macroeconomic and distributional effects. If the goal of the SPR then is to dampen or eliminate price shocks associated with supply disruptions, how do policy-makers decide whether a particular situation warrants release and how large it should be, given that simply counting barrels is insufficient? There is no literature on this issue that we are aware of.

The initial literature on SPR management dates back to the late 1970s and early 1980s–the years following the oil shocks of the Arab oil embargo, the Iranian Revolution, and the establishment of the reserve. These studies include Teisberg (1981) and Chao and Manne (1983), which use dynamic programming to model optimal SPR filling and release policies. Those studies illustrate that releasing reserves is more desirable when the disruption is likely to be temporary.

But they did not explore how to determine the probabilities of persistence in practice (rather, the authors assumed constant transition probabilities). One reason for this is that literature did not have access to the information provided by oil futures markets, simply because the studies pre-dated the existence of those markets. Hubbard and Weiner (1983, 1985, 1986) use econometric methods to model the impacts of SPR releases, but they also pre-dated the establishment of long-dated futures contracts. Short-run oil futures markets were established in 1983, and liquid markets for long-dated contracts did not arise until the late 1980s. In light of this, Devarajan and Hubbard (1982) was ahead of its time in arguing that the SPR should sell futures contracts to accelerate the impact of strategic releases. The establishment of long-dated contracts

[6] See GAO, "Strategic Petroleum Reserve", 2005.

[7] See Department of Energy, "Long-Term Strategic Review (LTSR) of the U.S. Strategic Petroleum Reserve (SPR) Report to Congress", 2016. In particular, the benefit analysis assumed that net supply disruptions would be offset one-for-one by coordinated IEA releases.

creates more information about expected future spot prices, to which the existing dated literature did not have access.

Literature conceptualizing the relationship between inventories and price spreads dates back to Fama and French (1987, 1988) and was summarized by Pindyck (2001). Intuitively, these studies use a now standard no-arbitrage result from finance that in equilibrium the futures price spread (the benefit from holding inventories) must equal the net cost of storage and lost interest accounting for the risk premium (together representing the cost of holding inventories). While inventory activity upholds this equilibrium, inventories do not appear explicitly in this result or as a function of spreads. We directly estimate this relationship.

There is a recent empirical literature on oil inventories, including Kilian and Lee (2014), Kilian and Murphy (2014), and Baumeister and Kilian (2016a), and we follow those studies in using vector autoregression to model oil market dynamics methods. However, those studies estimate the relationship between inventories and oil price *levels*, not spreads. Those studies are also not focused on informing SPR policy.

This study also contributes to a larger literature that uses time series methods to study oil market dynamics. These include Hamilton (2009); Kilian (2016, 2017a,b) and Herrera (forthcoming), which generally consider dynamics between oil prices and supply and demand factors. Those studies do not, however, focus on the relationship of inventories and price spreads.

One other study that is relevant for this research is Baumeister and Kilian (2016b), which distinguishes between oil price futures and expected future oil prices. While often conflated, these two differ according to the risk premium. Conceptually, this matters when using futures prices as a proxy for expected prices in the future and hence the expected persistence of supply shocks. Econometrically, this matters when computing the price "spread" against the prompt-month oil price. To explore the impact of this distinction on our results, we estimate our models using both methods. We first estimate the model with the spread computed as the difference between long-dated and prompt-month futures prices. We then run sensitivities by estimating the model with spreads computed as the difference between the expected future spot prices (estimated by Baumeister and Kilian 2016b) and prompt-month futures prices.

In summary, there is little to no recent economic literature guiding SPR policy, and also no literature that we are aware of that considers the relationship between oil inventories and price spreads. This paper contributes significantly in these areas, and provides important guidance to future SPR release policy.

3 Methods and Data

In this section, we overview the methods and data we use in our estimation and simulations. First, we provide an overview of the vector autoregression model, the computation of cumulative impulse response functions, and how we run counterfactual simulations of SPR release scenarios. This section also includes a summary of the data used and its sources.

3.1 Methods

3.1.1 VAR Model

We estimate a vector autoregression (VAR) with $K = 5$ variables: real global industrial production, OECD commercial crude inventories, OPEC spare production capacity, West Texas Intermediate (WTI) oil prices, and price spreads. All variables are expressed in log-changes, implying percentage changes after taking first differences (or percentage points in the case of spreads).[8] It is natural to include crude inventories, prices, and spreads based on the discussion above. Industrial production is included as an indicator of economic activity and broader commodity demand factors. OPEC spare capacity (i.e., available production capacity that is currently idle) is included because OPEC supply responses are a potential moderating factor in inventory dynamics and SPR releases.[9]

[8] We take first differences because the series are generally non-stationary in levels, but are stationary in differences.

[9] Our results are qualitatively similar if we were to exclude industrial production and spare capacity. We include them nonetheless because they are important factors affecting oil markets. Further, including them allows us to simulate the dynamics of these variables as well, including OPEC supply responses to inventory shocks.

$$\mathbf{y}_t = \mathbf{A}_0 + \sum_{\ell=1}^{24} \mathbf{A}_\ell \mathbf{y}_{t-\ell} + \boldsymbol{\varepsilon}_t, \tag{1}$$

where

$$\mathbf{y}_t = \begin{pmatrix} \Delta\text{Global Industrial Production}_t \\ \Delta\text{OECD Commercial Crude Inventories}_t \\ \Delta\text{OPEC Spare Capacity}_t \\ \Delta\text{WTI Prompt-Month Price}_t \\ \Delta\text{WTI 12-Month Spread}_t \end{pmatrix}, \tag{2}$$

and variables are measured in logs. Spreads represent the log difference between the 12-month and prompt-month real crude futures price, so it can be interpreted as a percentage point. In interpreting our results, we implicitly treat the 12-month futures price as a market-based expectation of future spot prices. However, as documented by Baumeister and Kilian (2016b), futures prices differ from expected spot prices due to a time-varying risk premium. Because we take first differences, this affects our results to the extent that the risk premium varies over time. As a sensitivity, we also estimated the VAR using a time series of 12-month-ahead expected spot prices generously provided by Baumeister and Kilian (2016b) to construct our spread variable. The estimates are very similar and can be found in the appendix.[10] We nonetheless present the results using actual futures prices because that is the data that would be available to a policymaker during an emergency oil supply disruption.

We also include four seasonal dummy variables in the estimation, but do not include them in the equations below for notational convenience. The sample period is February 1991 through June 2016, monthly.[11] We use 24 lags (2 years), following work by Hamilton and Herrera (2004); Kilian (2009); Kilian and Lee (2014) and Kilian and Murphy (2014), which stress the need for including many lags to account for slow-moving dynamics in oil markets.

[10]The impulse response functions were qualitatively unchanged when using the Baumeister and Kilian (2016b) data to construct the spreads. The only significant impact was that the responses of spreads to all impulses are somewhat larger when using the Baumeister and Kilian (2016b) data. However, caution should be taken in comparing these, since they represent conceptually different variables in that one measure includes the risk premium and the other does not.

[11]We are constrained by the industrial production series, which begins in January 1991. As a result, our first observed monthly change is February 1991.

3.1.2 Cumulative Impulse Response Functions

In what follows, we overview impulse response functions for those less familiar with them. Readers comfortable with IRFs and OIRFs may wish to skip to the next subsection, 3.1.3. We present the results through cumulative orthogonalized impulse response functions (IRFs). These represent how each of the variables change over time $(t+s)$ in response to a one-unit orthogonalized shock in to one of the variables (j) at time t. The formula for IRFs can be derived analytically through the moving average representation of the VAR:

$$\mathbf{y}_t = \Phi_0 \boldsymbol{\varepsilon}_t + \Phi_1 \boldsymbol{\varepsilon}_{t-1} + \Phi_2 \boldsymbol{\varepsilon}_{t-2} + \ldots,$$

where $\Phi_0 = I_K$, $\Phi_1 = \mathbf{A}_1$ and the subsequent Φ_s matrices are defined recursively as

$$\Phi_s = \sum_{\ell=1}^{s} \Phi_{s-\ell} \mathbf{A}_\ell.$$

The (i,j) element of Φ_s gives the expected response of variable i, s periods after a unit shock to only variable j. Of course, the shocks to the five variables are likely to be correlated, meaning a shock to variable j is likely to be accompanied by shocks to other variables. For example, expectedly large inventory shocks are likely to put downward pressure on oil prices immediately, and indeed we find a negative correlation between these residuals. The orthogonalization process accounts for the correlation in these shocks, de-correlating the $\boldsymbol{\varepsilon}_t$ vector by multiplying it by the Cholesky decomposition of the residual covariance matrix, L (so $LL' = \Sigma_{\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}'} = cov(\boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_t')$). This results in uncorrelated shocks, denoted $\boldsymbol{u}_t = L\boldsymbol{\varepsilon}_t$. The orthogonalized IRF (OIRF) is

$$\Psi_s = \Phi_s L,$$

representing the expected responses of the variables to these orthogonalized shocks.[12] Intuitively, this accounts for the immediate relationships between the variables through the correlation of the residuals.

As is well known, the ordering of the variables matters in the orthogonalization because the Cholesky decomposition is lower triangular. Variables placed earlier in the ordering are effectively assumed to be "slower moving" than variables placed later.[13] For example, industrial production, being placed first in equation 2, is not allowed to respond within-period to any shocks in other variables, only responding with a lag through

[12] While the multiplication by L itself scales the shocks to one standard deviation shocks (instead of one-unit), we rescale these responses by the inverse of the standard deviations to once again represent unit changes.

[13] The alternative would be to use non-orthogonalized shocks. This would implicitly make the stronger assumption that shocks are contemporaneously unrelated, contrary to the observed correlations.

the lag structure in equation 1. OPEC spare capacity, being placed in the middle, can immediately respond to shocks to industrial production and crude inventories, but it only responds to shocks to WTI price and spreads through the lag structure in equation 1. WTI prices and spreads are placed at the end, allowing them to response to shocks in all preceding variables.[14] We argue that this ordering is reasonable. Prices and spreads can and do change instantly in response to news, justifying their placement at the end. OPEC spare capacity (placed third) is clearly slower moving than prices, and OPEC's historic status as swing producer suggests that it can adjust more quickly than commercial inventories (placed second), which requires changes to flow supply or demand and the movement of physical barrels. Lastly, industrial production is a very slow moving series, justifying its placement first in the VAR.

Because all of our variables are in monthly changes, Ψ_s represents the impacts of a shock on monthly changes in variables after s periods. Since we are interested primarily in the levels of variables of interest (e.g., oil prices) and not their monthly changes, we present cumulative impulse response functions ($\sum_{\ell=1}^{s} \Psi_\ell$). This accumulates the effects on monthly changes in the variables to show the impact on their (log) levels over time. We construct 95% confidence intervals through residual-resampling bootstrap.

3.1.3 Simulations

We conduct counterfactual historical simulations aimed at estimating how the SPR releases would affect other variables over time, including crude oil prices and spreads. We model a strategic release as an unexpected increase in commercial crude inventories.

For each simulation, we first choose a month (τ) and size (r_τ) for incremental SPR release we are simulating.[15] We compute new "counterfactual" shock that month, $\boldsymbol{\varepsilon}_\tau^{cf}$, by adjusting the observed residual that month $\boldsymbol{\varepsilon}_\tau$, according to an orthogonalized inventory shock of r_τ:

[14] The relative ordering of spreads and prices prevents prompt-month prices from immediately responding to spread shocks, implicitly treating spread shocks as shocks to futures prices alone. Interchanging the order of these two variables has no effect on the key OIRF of interest, which is the response to inventory shocks.

[15] In simulations where we are simulating a counterfactual of no SPR release, r_τ is set to equal to the negative of the actual observed historical release. This effectively "subtracts" the observed historical release.

$$\boldsymbol{\varepsilon}_\tau^{cf} \equiv \boldsymbol{\varepsilon}_\tau + \frac{1}{L_{(2,2)}} L \begin{pmatrix} 0 \\ r_\tau \\ 0 \\ 0 \\ 0 \end{pmatrix} \tag{3}$$

The division by $L_{(2,2)}$ (which is approximately the standard deviation of the inventory residual) is required because the pre-multiplication by L scales the magnitude of shock into one of r_τ standard deviations. This conversion ensures that the inventory residual, $\boldsymbol{\varepsilon}_{2,\tau}$, changes by precisely the assumed SPR release size, r_τ. The orthogonalization captures the immediate impacts that inventory releases may have on oil prices and spreads, based on the correlation of the VAR residuals.[16]

All other residuals are left unchanged: $\boldsymbol{\varepsilon}_t^{cf} \equiv \boldsymbol{\varepsilon}_t$ for $t \neq \tau$.

With this one residual changed, we recursively simulate subsequent counterfactual values of $\mathbf{y}_t^{cf}$ as

$$\mathbf{y}_t^{cf} \equiv \mathbf{A}_0 + \sum_{\ell=1}^{24} \mathbf{A}_\ell \mathbf{y}_{t-\ell}^{cf} + \boldsymbol{\varepsilon}_t^{cf}, \tag{4}$$

where $\mathbf{y}_t^{cf} = \mathbf{y}_t$ for $t < \tau$ and the seasonal dummies are suppressed for notational convenience. While $\mathbf{y}_t^{cf}$ are in log-differences, we aggregate up to levels when presenting the simulation results (e.g., barrels for inventories or dollars per barrel for oil prices) to make them easier for the reader to interpret.

3.2 Data

Our data come from four sources. Oil prices are obtained from Bloomberg L.P.[17] Daily prices are averaged to the monthly level and converted to real 2014 dollars using the CPI All Urban Consumer (All Items) index. Spreads are computed as the log of real 12-month futures prices minus the log of real prompt-month prices. This log difference thus approximately represents spreads in percentages. It also creates an exact linear relationship between log prices and spreads, which aids in interpreting the results.

[16] For example, a 1 percent shock in inventories is estimated to be associated with an immediate reduction in spot prices of approximately 1 percent and an increase in spreads of approximately 0.6 percent. That is, with $r_\tau = 0.01$, the last term in equation (3) is approximately $(0, 0.01, 0.008, -0.010, 0.006)'$ (with the final two elements representing percentage changes in oil prices and spreads, respectively).

[17] We use WTI prices, which use the CL1 Comdty and CL12 Comdty tickers for the prompt-month and 12-month contracts, respectively.

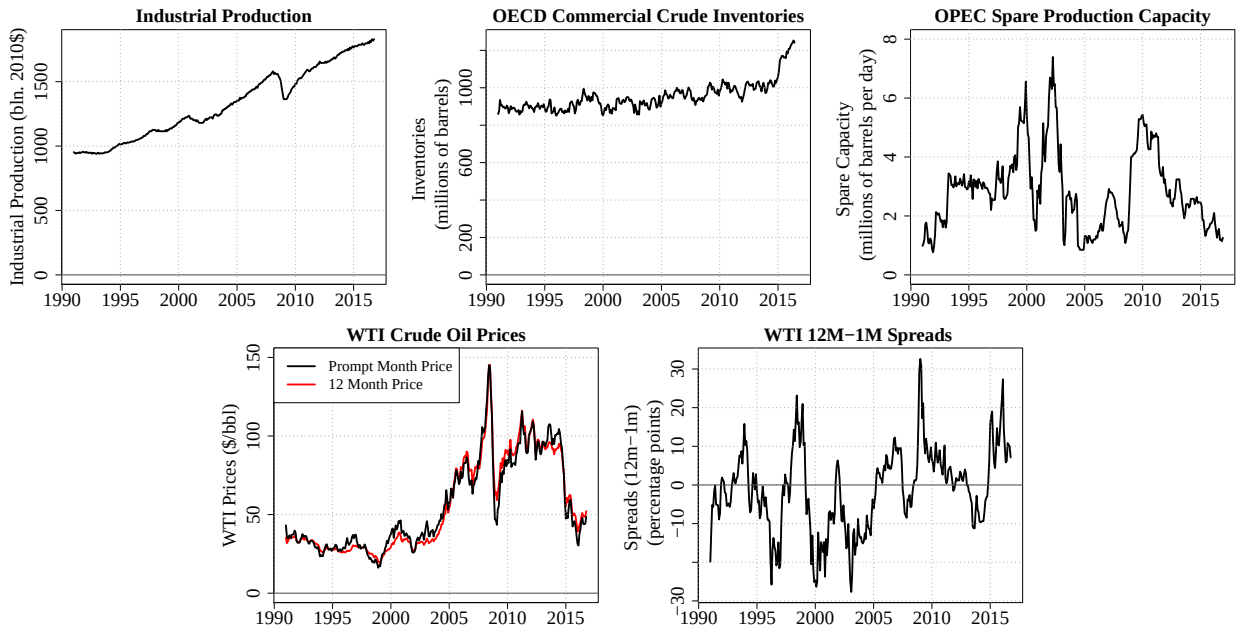


Figure 1: Time Series of Industrial Production, OECD Commercial Crude Inventories, OPEC Spare Capacity, WTI Oil Prices, and WTI Spreads

The OECD commercial crude oil inventory series is from the International Energy Agency (IEA). It corresponds to the crude stocks reported in their monthly Oil Market Reports, typically shown in their table titled "OECD Industry Stocks and Quarterly Stock Changes".[18]

The real global industrial production series is from the World Bank (2017) Global Economic Monitor and represents constant 2010 US$, seasonally adjusted.

The OPEC spare production capacity series comes from the US Energy Information Administration (EIA). This series is publicly available beginning in 1997 from the Short Term Energy Outlook Custom Table Builder.[19] Earlier data was obtained by request from EIA.[20]

These time series are shown in Figure 1. That figure shows the series in levels, but as previously mentioned we use their log-changes in the estimation.

[18]See `https://www.iea.org/oilmarketreport/reports/` for an archive of these reports. Our data does not precisely match the data in the archived reports because the series is continually updated by IEA. We obtained the series from IEA in August 2016, and it is complete through June 2016.

[19]This is available at `https://www.eia.gov/outlooks/steo/query/`.

[20]This earlier series represents world spare capacity, not just OPEC. However, during that time period, OPEC accounts for virtually all spare capacity, so the distinction is minimal. Further, we compared the two series during the period when they overlap (1997-2002) to ensure the series were consistent. The correlation between the two series' monthly log-changes is very strong, at 0.9. We construct the full series by concatenating the monthly log-changes in the two series: world spare capacity for January 1990 to December 2002, and OPEC spare capacity from January 2003 to December 2016.

4 Results

4.1 Cumulative Impulse Response Functions

While there are many impulse-response relationships implied by our VAR estimation, the primary relationship of interest is the effect of a shock to inventories. This is because we wish to model the effect of SPR releases, which amounts to making additional strategic inventories available to the market. Figure 2 shows the responses of each variable to a 1% impulse to commercial crude inventories.[21] When the responses are statistically significant, their direction and size generally comport with economic theory and intuition. As seen in the top middle chart (the effect of the inventory shock on inventories), the inventories remain elevated after the shock, although some of the surplus is gradually drawn down.[22] The impacts on WTI prices and spreads are shown in the bottom two charts. We see a small immediate reduction in spot prices which grows in magnitude to -3% to -4% over several months, before reverting towards zero after a year. At the same time, spreads temporarily increase by approximately 1.5 percentage points as this spot price reduction is accompanied by a smaller reduction in futures prices. Price spreads revert back toward baseline conditions with about 18 months. [23] Under the assumption that releases of strategic reserves can be modeled as surprise inventory builds, this suggests that SPR releases can reduce prices and mitigate backwardation, relative to a counterfactual of no release.

The response of OPEC spare capacity is temporarily positive, implying OPEC reducing output in response to inventory builds, although it reverts to zero and the confidence intervals include zero for much of the response window. In particular, the OPEC response is statistically insignificant in the first few months following an inventory shock, which is the most relevant period for short-term price effects in response to strategic releases. The effect on industrial production is also small and generally statistically insignificant.[24]

We explored the other impulse response functions, which estimate the effects of impulses in each of the other variables (e.g., industrial production, spare capacity, and spreads) on each other. These OIRFs can be found in the appendix. When the responses are statistically significant, their direction and size

[21] A 1% inventory shock is on the order of 10 million barrels of crude oil, since private inventories are on the order of 1 billion barrels (see Figure 1). Because of the linearity of the VAR model, all effects scale linearly with the inventory shock. For example, an inventory shock of 2% would be estimated to produce responses that are exactly twice as large as those from a 1% shock.

[22] This is consistent with Hubbard and Weiner (1986), which finds that private inventories do not fully absorb strategic releases.

[23] While the 12-month futures price does not explicitly appear in the VAR, it appears implicitly through its additive relationship between spot prices and spreads. Namely, since spreads are computed as the log-difference between future and spot prices ($Spread = \ln(Future) - \ln(Spot)$), the effect on futures prices is the sum of the effect on spot prices and spreads ($\ln(Future) = \ln(Spot) + Spread$, which implies $\frac{\partial \ln(Future)}{\partial u_{j,t}} = \frac{\partial \ln(Spot)}{\partial u_{j,t}} + \frac{\partial Spread}{\partial u_{j,t}}$). So, for example, after 12 months, a -4.2% change in spot prices accompanied by a 1.3% change in spreads implies a -2.9% change in futures prices ($= -4.2\% + 1.3\%$).

[24] If we exclude industrial production and/or spare capacity from the model completely, the responses of the remaining variables (oil prices, spreads, and inventories) are qualitatively similar but slightly smaller.

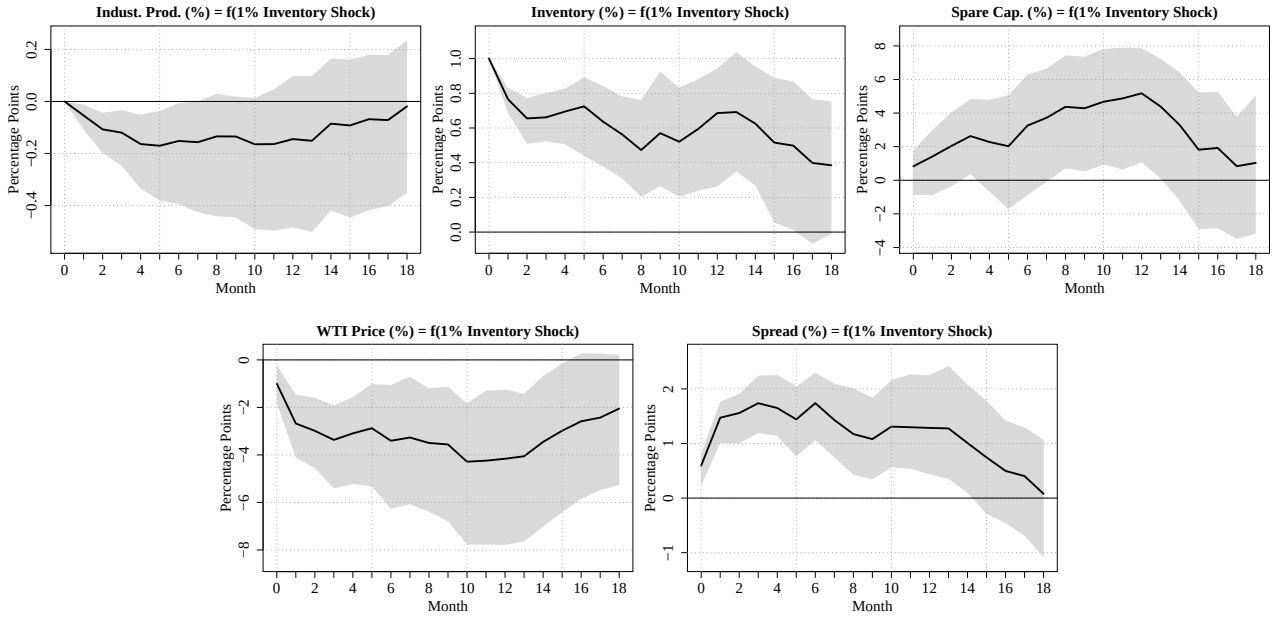


Figure 2: Cumulative Impulse Responses from a 1 Percentage Point Shock to Inventories

Notes: Shaded areas represent 95% confidence intervals computed through 100 bootstrap iterations.

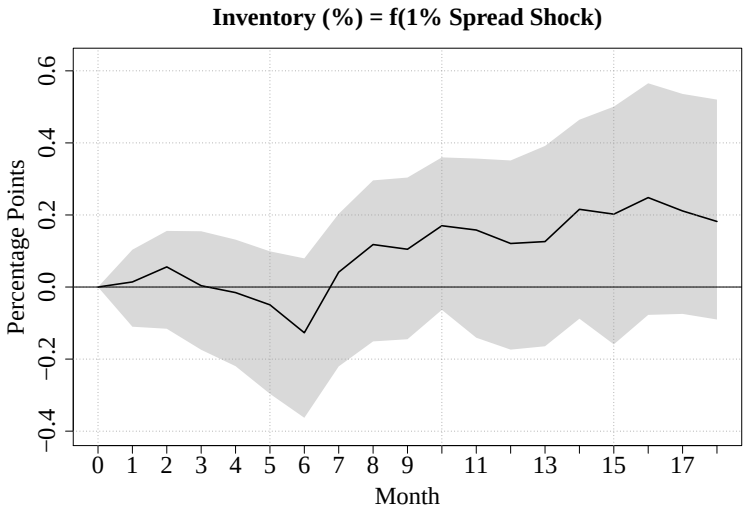


Figure 3: Cumulative Inventory Response from a 1 Percentage Point Shock to Spreads

Notes: Shaded areas represent 95% confidence intervals computed through 100 bootstrap iterations.

generally comport with economic theory and intuition. Of interest, we find that the effect of spread shocks on inventories is small and statistically insignificant, as shown in Figure 3. This suggests that the causal relationship runs from inventories to spreads, not vice versa.[25] This is intuitive, since market prices can and do change instantly in response to news about inventories, whereas inventory builds require adjustments to flow supply and/or demand and the movement of physical barrels.

4.2 Simulations

We use our estimated VAR to conduct three counterfactual simulations to illustrate the impacts of strategic releases during three historical periods: the Iraq War in 2003, Hurricane Katrina in 2005, and the Arab Spring in 2011. In the first scenario, we simulate the impact of an SPR release at a time when it was considered but not executed. The latter two scenarios involve times when strategic reserves were released, so we simulate a no-release counterfactual.

4.2.1 Simulation 1: 2003 Iraq War

The first simulation is in spring 2003, when there were two significant supply disruptions in a short period of time: the Venezuelan oil strike and the Iraq War. The Venezuelan oil strike began in late 2002, disrupting over 2 million barrels per day in supply (Nov-2002 to Jan-2003). This disruption was in the process of being resolved in early 2003, when the invasion of Iraq disrupted 1.3 million barrels per day (Mar-03 to Apr-03). Some of this disruption was offset by a 0.5 million barrel per day increase in production from other OPEC nations, so the net disruption from OPEC members was 0.8 million barrels per day. Spot crude prices rose sharply, sending the futures market into deep backwardation in excess of 20% (during Jan-03 to Mar-03). We model a strategic release of 48 million barrels, equal to two months of disrupted supply (0.8 million barrels per day for 60 days). Other recent strategic releases were of similar order of magnitude (typically 60 million barrels), so this figure is a reasonable approximation of what an actual release might have been. We simulate the release occurring in February 2003, which was the month of steepest backwardation.

Figure 4 shows the results of simulation, with actual values shown in black and simulated values shown in green. The simulated release of crude oil temporarily reduced spot prices by about $6, a decrease of 15-20%. Simulated spreads show significantly reduced backwardation, falling to about -5 to -10% instead of -15%. After one-to-two years, the effect of the release diminishes somewhat, suggesting a temporary effect.[26]

[25] This finding is robust to the ordering of the variables in the VAR, implying it is not simply the result of the orthogonalization.

[26] This temporary nature of the effect is not an inevitable consequence of the VAR model or the assumed number of lags (24). First, as is clear from the moving average form of the VAR, the model allows to persist for much longer. Second, these responses

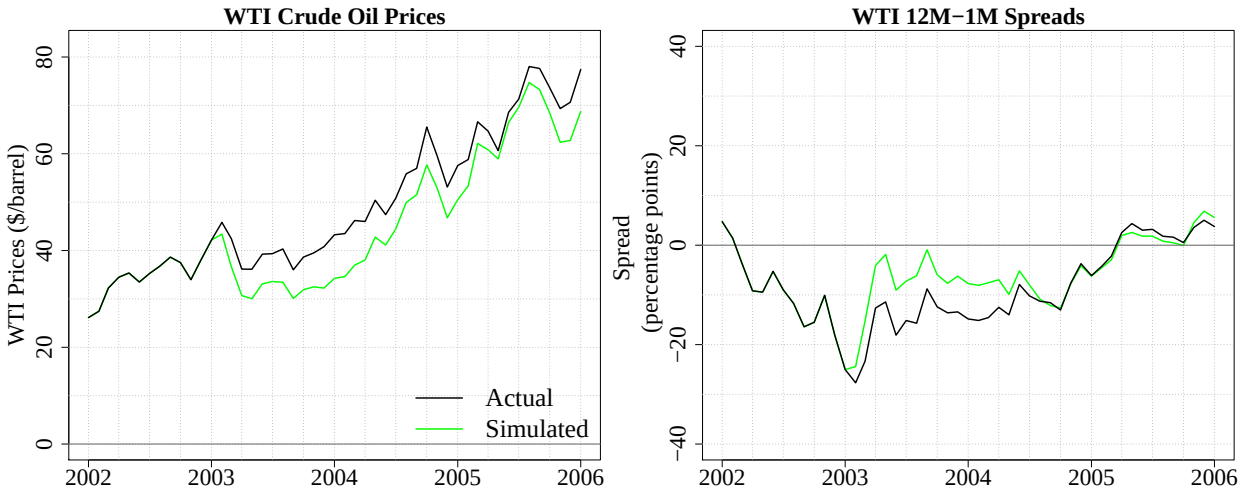


Figure 4: Simulation **with** Strategic Release in February 2003

Notes: The assumed inventory shock is +48 million barrels of crude oil ($\approx 5.5\%$ of inventories then), equal to the net disruption in production of 0.8 million barrels per day for a period of 2 months.

Actual spreads eventually recover to the simulated levels, and the counterfactual price rises somewhat towards approaches actual prices.

In summary, this simulation suggests that strategic releases could have mitigated the short-term price increases and backwardation brought on by the Iraq War.

4.2.2 Simulation 2: 2005 Hurricane Katrina

The second simulation considers the strategic release that occurred in the aftermath of Hurricane Katrina in 2005. The hurricane shut in offshore production in the Gulf of Mexico, cutting off about one quarter of U.S. oil production. It also caused severe damage to oil production infrastructure and refineries. Oil prices rose, but gasoline prices spiked even more due to the temporary inability to refine and transport refined products. While crude oil futures markets did not enter backwardation, gasoline futures did (not shown here) due to refinery closures. In response, in September 2005 the United States and the IEA conducted a coordinated emergency release, making available 60 million barrels of both crude oil and refined products.[27]

We simulate a counterfactual case with no strategic release during this period. We implement this by subtracting out a shock of 60 million barrels in September 2005, which is equal to the total amount made available by all IEA member countries. This ignores the distinction between crude oil and refined products,

represent accumulated effects on monthly changes. While the effect on monthly changes asymptotes to zero in the long run, the accumulation of these changes need not do so.

[27]Not all of the SPR crude oil that was made available was ultimately delivered because of a lower number of successful bids. However, during this period, refined product supply (e.g., gasoline) was a larger constraint than that of crude oil. In addition, the administration conducted short-term in-kind loans of crude oil to allow refineries with disrupted supply lines to keep operating.

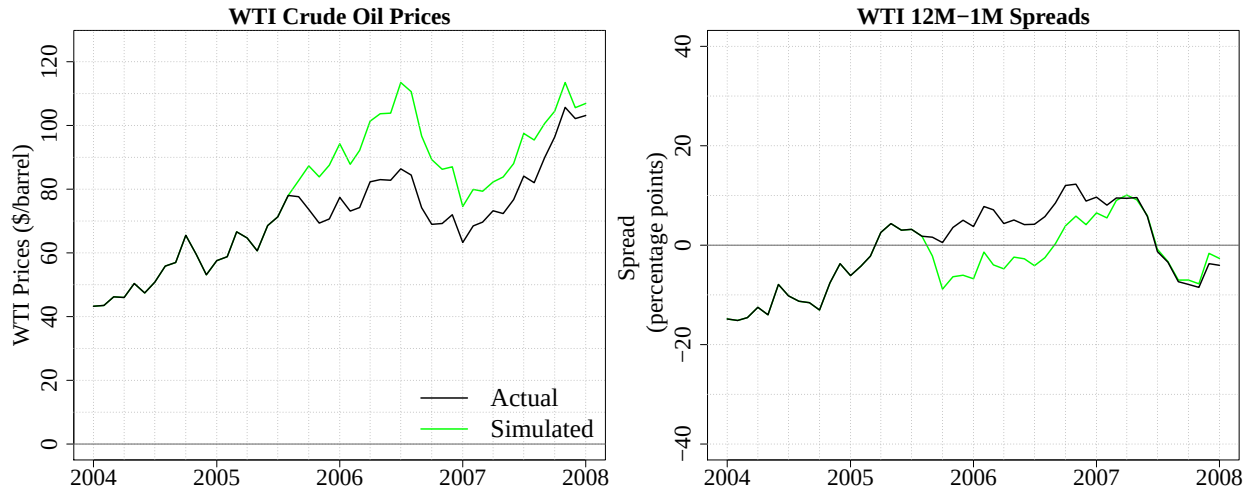


Figure 5: Simulation **without** Strategic Release in September 2005

Notes: The assumed inventory shock is -60 million barrels ($\approx -5.9\%$ of inventories then), equal in magnitude to the total IEA release that year.

as well as the fact that not all reserves offered were ultimately delivered. This only changes the magnitude of the results, not their direction, but we believe it would tend to overstate the magnitude of the impact of the actual releases.

The results are shown in Figure 5. In contrast to the previous simulation, this simulation involves a "no release" counterfactual in green, whereas the actual values (in black) include the effects of the strategic release. The bottom two panels show that absent the release, oil prices are projected to have continued an upward climb apace, quickly exceeding $100 per barrel. The releases are thus estimated to have reduced prices by more than 20% relative to the counterfactual. Futures would have entered into moderate backwardation during the disruption, rather than remaining in slight contango.

This simulation again suggests some success of the observed SPR release in preventing oil price spikes and backwardation in futures curves. However, this particular simulation comes with the caveat that the market for refined products was relatively more important during this period, a feature that our model does not explicitly incorporate.

4.2.3 Simulation 3: 2011 Arab Spring

Our final simulation considers the strategic release of reserves during the 2011 Arab Spring, which disrupted oil supply in Libya and other countries. In June 2011, IEA member countries announced a release of 60 million barrels of oil, which were delivered by August 2011. As in the previous section, we simulate a scenario

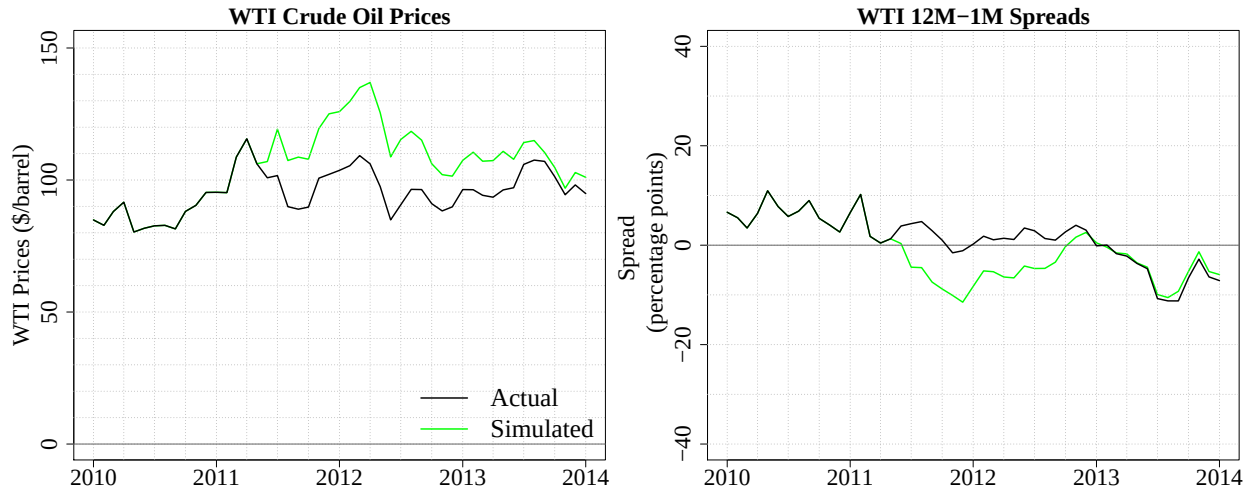


Figure 6: Simulation **without** Strategic Release in June 2011

Notes: The assumed inventory shock is -60 million barrels ($\approx -6.4\%$ of inventories then), equal in magnitude to the total IEA release that year.

without this release by introducing a negative inventory shock of 60 million barrels in June 2011. The results are in Figure 6. The figure suggests that the release prevented a temporary increase in prices and backwardation. With the releases, prices and spreads both remained relatively stable around $100, compared to the counterfactual case where prices rise above $130 per barrel and spreads moved into backwardation of approximately -10%. This suggests that the 2011 strategic releases may have prevented price spikes of 20% to 30%.

As with the previous simulations, this one suggests that strategic releases can prevent oil price increases and swings into backwardation. It should be noted that because the three simulations are based on the same model estimates, they unsurprisingly produce similar results. However, they are not simply repeating exactly the same result because each one incorporates a different series of shocks (residuals) and lagged values. In other words, while it is unsurprising the simulations tell a consistent story, the resulting similarity in the simulations is not simply the inevitable implication of using the same underlying VAR model.

5 Caveats

The above simulations all suggest that strategic releases can effectively moderate price increases and stabilize futures spreads in times of market dislocation. There are, however, important caveats. First, we simulate strategic releases as a surprise shock to private inventories. This implicitly assumes that market actors respond to surprise public releases in the same way that they would respond to surprise increases in private

inventories. If market actors interpret these changes differently (say, because they believe public releases signal future government policy), then they may react differently, although it is unclear which direction this would bias the estimates.

Second, the simulations assume that the strategic releases are unexpected, both in an econometric sense (as deviations from the statistical expectation of the model) and from the perspective of market participants.[28] If strategic releases are expected, then their effects will be incorporated into existing oil prices and spreads. While this suggests imperfections in the timing of the effects found in the simulation, it does not undermine their existence. Indeed, as Devarajan and Hubbard (1982) point out, expectations of future releases help calm markets before they actually occur—so long as policymakers eventually follow through on that expectation. This relates to another issue in the vein of the Lucas critique. If the government commits to a policy of releasing inventories when supply disruptions induce steep backwardation, markets will anticipate this policy, stabilizing oil markets before any particular release is announced. While this expectations channel has desirable anticipatory and stabilizing effects, it also threatens to undermine the value of the price spread itself as an indicator of the severity of supply disruptions.

A third caveat is that the estimated VAR underlying the simulations represents the "average" relationships between the variables over the sample period. The relationships during times of extreme disruption may differ from these average relationships. However, the large variation in the series during these extreme events provides a non-trivial portion of the variation identifying our VAR estimates.

Finally, our model does not directly treat refined product markets. This is mainly a concern in cases where refined markets are disproportionately important, such as during Hurricane Katrina, which severely damaged refining operations. Future work could apply these techniques to more directly explore refined product market disruptions, with relevance to the consideration of the value of refined product reserves.

6 Conclusion

Understanding the dynamic relationships between oil spot prices, price spreads, and inventories is key for both market participants and policymakers considering releases from strategic oil reserves, such as the SPR. First, we make the conceptual point that spreads convey market-based information about expectations regarding the severity and persistence of supply disruptions, with clear implications for when SPR releases are advisable. Namely, steep backwardation suggests disruptions are serious, unlikely to be quickly moderated by private

[28] As pointed out by Kilian and Murphy (2014), with forward-looking investors, these expectations can differ, although VAR models often implicitly equate them.

market forces, and expected to be transitory, indicating a stronger case for releasing reserves. The spread thus provides an indicator guiding policymakers whether to release reserves, or to hold them in case a disruption might deteriorate further.

We use a VAR to estimate the dynamic relationship between oil spot prices, price spreads, and commercial inventories. This relationship illustrates the magnitude of the impacts of SPR releases, with implications for how many barrels to release in times of supply disruption. The estimates suggest that SPR releases—modeled as surprise inventory shocks—can have substantial effects on oil prices and spreads. Impulse response functions suggest that stock releases can reduce spot prices by up to 4% and mitigate backwardation by approximately 1.5 percentage points for each 1% addition to commercial inventories (approximately equal to 10 million barrels). Simulations suggest that past releases may have prevented spikes in oil prices on the order of 20% to 30%. These estimates can guide policymakers in determining how many barrels to release from the SPR in times of supply disruption to achieve specified market impacts.

References

Baumeister, C., and L. Kilian (2016a): "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us," *The Journal of Economic Perspectives*, 30(1), 139–160.

——— (2016b): "A General Approach to Recovering Market Expectations from Futures Prices with an Application to Crude Oil," *manuscript*.

Chao, H.-P., and A. S. Manne (1983): "Oil Stockpiles and Import Reductions: A Dynamic Programming Approach," *Operations Research*, 31(4), 632–651.

Devarajan, S., and R. G. Hubbard (1982): "Drawing down the Strategic Petroleum Reserve: The Case for Selling Futures Contracts," in *Stockpile Policy*, chap. 11, pp. 187–196. Harvard University, Cambridge, MA.

Fama, E. F., and K. R. French (1987): "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage," *Journal of Business*, pp. 55–73.

——— (1988): "Business Cycles and the Behavior of Metals Prices," *The Journal of Finance*, 43(5), 1075–1093.

Hamilton, J. D. (2009): "Understanding Crude Oil Prices," *Energy Journal*, 30(2), 179–206.

Hamilton, J. D., and A. M. Herrera (2004): "Oil shocks and aggregate macroeconomic behavior: The role of monetary policy: A comment," *Journal of Money, Credit, and Banking*, 36(2), 265–286.

Herrera, A. M. (forthcoming): "Oil Price Shocks, Inventories and Macroeconomic Dynamics," *Macroeconomic Dynamics*.

Hubbard, R. G., and R. Weiner (1983): "The 'sub-trigger' crisis: An economic analysis of flexible stock policies," *Energy Economics*, 5(3), 178–189.

Hubbard, R. G., and R. J. Weiner (1985): "Managing the Strategic Petroleum Reserve: Energy Policy in a Market Setting," *Annual Review of Energy*, 10(1), 515–556.

——— (1986): "Inventory Optimization in the U.S. Petroleum Industry: Empirical Analysis and Implications for Energy Emergency Policy," *Management science*, 32(7), 773–790.

Kilian, L. (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.

——— (2016): "The Impact of the Shale Oil Revolution on U.S. Oil and Gasoline Prices," *Review of Environmental Economics and Policy*, 10(2), 185–205.

——— (2017a): "How the Tight Oil Boom Has Changed Oil and Gasoline Markets," *CEPR Discussion Paper No. DP11876*.

——— (2017b): "The Impact of the Fracking Boom on Arab Oil Producers," *CESifo Working Paper Series No. 5751*.

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 D. P. Murphy (2014): "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29(3), 454–478.

Pindyck, R. S. (2001): "The Dynamics of Commodity Spot and Futures Markets: A Primer," *The Energy Journal*, pp. 1–29.

Teisberg, T. J. (1981): "A Dynamic Programming Model of the U.S. Strategic Petroleum Reserve," *The Bell Journal of Economics*, pp. 526–546.

World Bank (2017): "Global Economic Monitor," http://data.worldbank.org/data-catalog/global-economic-monitor.

A Appendix

A.1 All Impulse Response Functions

In the body we focuses on a selected set of impulse response functions of interest (those relating the inventories and spreads). All impulse response functions are shown in Figure A.1.

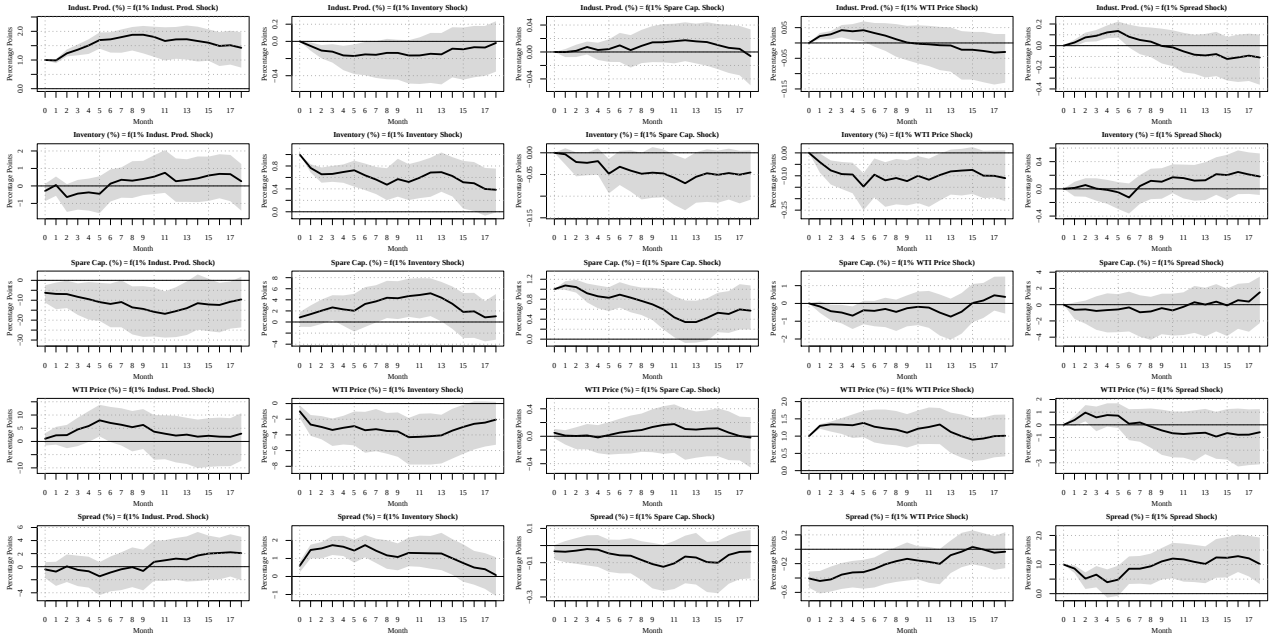


Figure A.1: All Impulse Response Functions using Actual Futures Spreads

Notes: All shocks are 1 percentage point in magnitude. Columns represent impulse variables, rows represent response variables. For example, the 2$^{\text{nd}}$ column shows the response of all variables to an impulse to inventories, corresponding to the graphs shown in Figure 2. And the 4$^{\text{th}}$ row represents how prompt-month oil prices respond to shocks to each of the five variables. Shaded areas represent 95% confidence intervals computed through 100 bootstrap iterations.

A.2 Estimation using Baumeister-Kilian Data

As discussed in section 3.2, we also estimate our model using the expected future (12-month) spot price series from Baumeister and Kilian (2016b) generously provided by those authors. This series is plotted, along with futures prices, in Figure A.2.

We re-estimated our VAR using the spread between the Baumeiseter-Kilian (BK) 12-month expected spot price and the prompt-month future price (i.e., the blue line minus the black line). By contrast, our main specification uses the difference between the actual 12-month future and the prompt month price (the red line minus the black line).

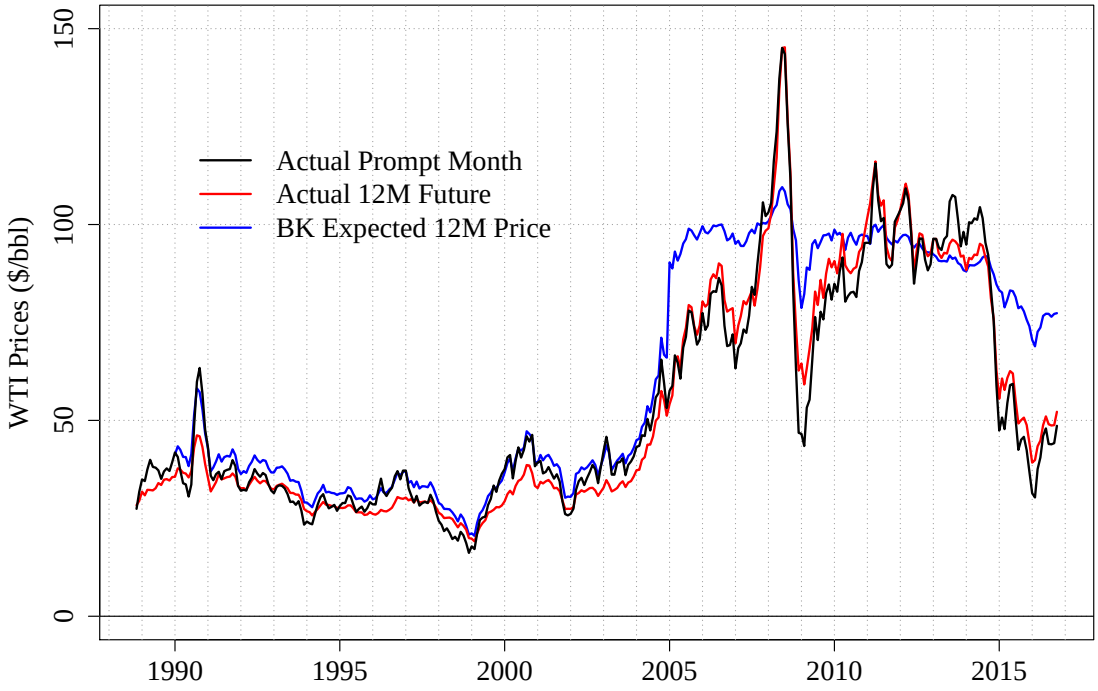


Figure A.2: WTI Oil Futures Prices and Baumeister-Kilian (BK) Expected Spot Prices

Figures A.3 and A.4 show the IRFs using this alternative spread series. These IRFs strongly resemble the corresponding ones in our primary specification, shown in Figures 2 and A.1. Focusing on the IRFs of interest—the effect of an inventory shock shown in Figures A.3—the largest difference is that the spread response is slightly larger when it is computed using the BK series (see last panel, about 2% versus 1.5%, although the confidence intervals mostly overlap). The lack of a substantial difference in these suggests that inventory shocks do not substantially affect the risk premium, since that is the only difference between the two series.

Most of the other IRFs (shown in Figure A.4) are also largely unchanged. The exception is the final column of that figure, showing the effects of a spread shock on the other variables. It is unsurprising that these IRFs show the largest difference because this spread shock represents a different economic concept. Namely, it excludes shocks to the risk premium. However, the standard errors generally overlap zero using either method.

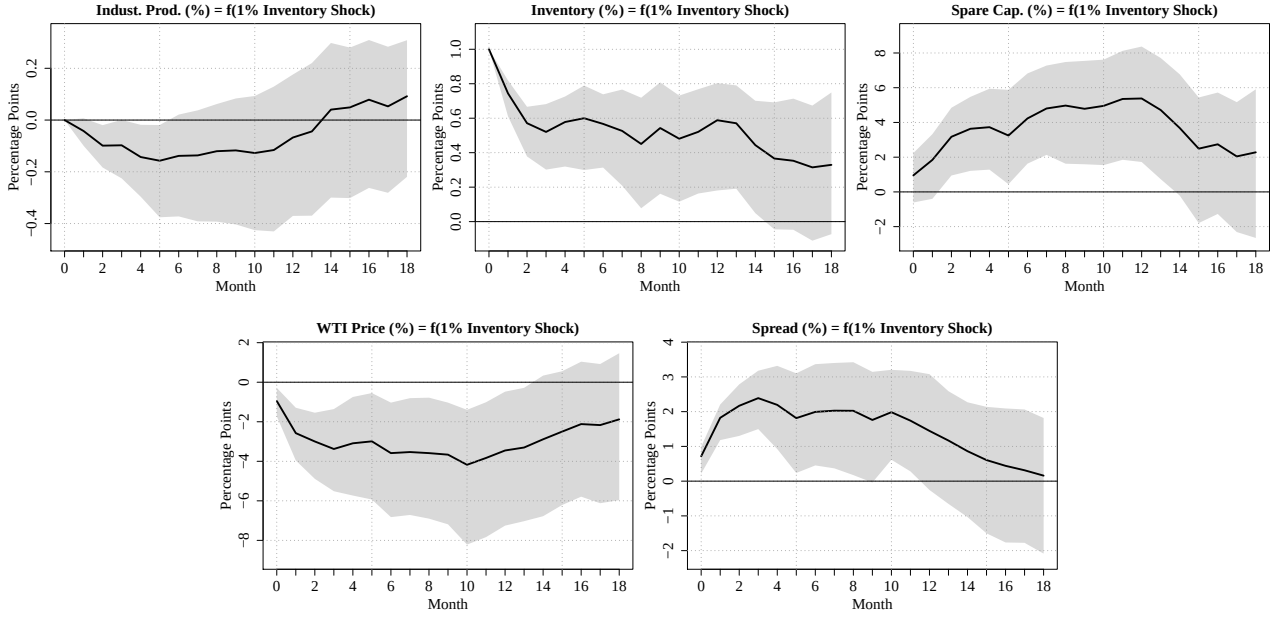


Figure A.3: Cumulative Impulse Responses from a 1 Percentage Point Shock to Inventories, using BK Data

Notes: Compare to IRFs in Figure 2. Shaded areas represent 95% confidence intervals computed through 100 bootstrap iterations.

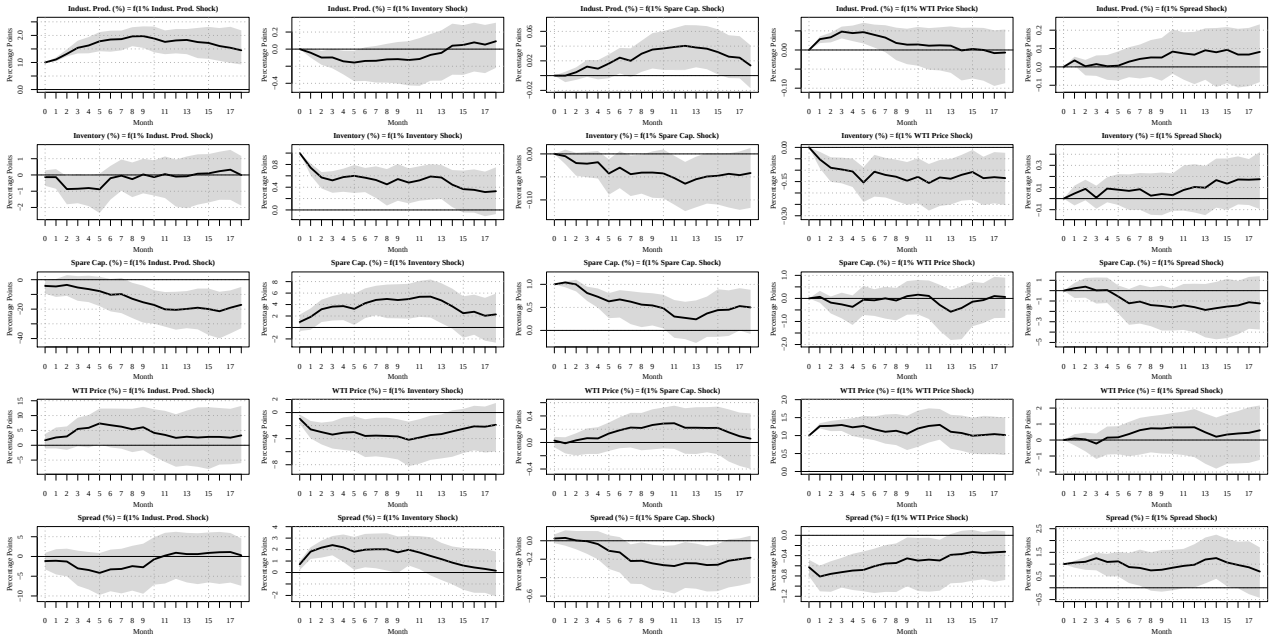


Figure A.4: All Impulse Response Functions, using BK Data

Notes: Compare to IRFs in Figure A.1. All shocks are 1 percentage point in magnitude. Columns represent impulse variables, rows represent response variables. For example, the 2$^{\text{nd}}$ column shows the response of all variables to an impulse to inventories, corresponding to the graphs shown in Figure 2. And the 4$^{\text{th}}$ row represents how prompt-month oil prices respond to shocks to each of the five variables. Shaded areas represent 95% confidence intervals computed through 100 bootstrap iterations.