The Strategic Petroleum Reserve and Crude Oil Prices

Reid Stevens*

University of California, Berkeley

November 15, 2014

Abstract

Have U.S. oil market policy interventions succeeded in lowering the price of crude oil? This paper uses a structural vector autoregression (VAR) model of the U.S. oil market to estimate the effect of the Strategic Petroleum Reserve (SPR) on crude oil prices. In contrast to existing literature, I show that unanticipated oil releases from the SPR have no measurable effect on oil prices. However, unanticipated oil purchases for the SPR raise oil prices 1.5 percent. I confirm these estimates with two alternate VAR identification strategies that use external information, rather than the traditional recursiveness assumption. First, the SPR purchase schedule is used as an instrument in the VAR model to identify the effect of SPR purchase shocks on oil prices. I verify the statistically significant purchase effect and show that standard identifying assumptions have a downward bias. Second, I identify SPR release shocks with high-frequency futures market data. Using several futures contracts to estimate unanticipated SPR policy shocks, I find that SPR releases do not have an effect on oil prices. To explain this puzzle, I estimate an interacted VAR with time-varying policy effects. This model shows that SPR purchases only increase oil prices when uncertainty is high, but SPR releases do not lower oil prices at any level of uncertainty.

Keywords: Energy Prices, Oil and the Macroeconomy, Government Policy, Oil Shocks

JEL: Q35, Q38, Q41, Q43, Q48

---

*Department of Agriculture and Resource Economics, University of California, Berkeley.
E-mail: reidstevens@berkeley.edu

1 Introduction

The Strategic Petroleum Reserve (SPR) was created in response to the 1975 Arab Oil Embargo, and was intended to, "Store petroleum to reduce the adverse economic impact of a major petroleum supply disruption to the United States" (U.S. Congress, 1975). Since the SPR was constructed in 1977, over 150 million barrels of crude oil have been released from the reserve. During these releases, crude oil from the SPR made up 5 to 12 percent of weekly domestic oil supply. The government has purchased over 850 million barrels of crude oil for the reserve, which accounted for between 2 to 4 percent of U.S. crude oil consumption during purchase weeks.

Despite this long history of purchases and releases, little is known about the effect of the SPR on crude oil prices. Much of the existing literature on the SPR focuses on optimal purchase and release strategies, and does not directly estimate the effect of the SPR on crude prices (see Tolley and Wilman, 1977; Nichols and Zeckhauser, 1977; Teisberg, 1981; Chao and Manne, 1983; Zhang et al., 2009). These papers generally assume a SPR crude oil release during a supply disruption will lower oil prices, but SPR purchases when the oil market is calm will not affect prices. Unsurprisingly, these models show that under optimal management, the benefits of the SPR far outweigh the costs of construction, oil acquisition, and storage. The estimated benefits of the SPR are even higher when the negative macroeconomic effects of oil price spikes are taken into account (Weimer, 1982; Bohi and Toman, 1996; Leiby and Bowman, 2000). When Congress recently approved plans to expand the reserve from its current level of 727 million barrels to 1 billion barrels, SPR proponents praised its ability to lower oil prices and stabilize the economy during oil supply disruptions (Government Accountability Office, 2006; Congressional Research Service, 2006).

The few empirical estimates of the SPR price effect vary widely: SPR releases are estimated to lower the price of crude oil by 3 to 32 percent, while SPR purchases are estimated to increase the price of crude oil by 0.4 to 32 percent (U.S. Senate, 2003; Verleger, 2003; Considine, 2006). The lack of consistency among these estimates is evidence that identifying the effect of SPR policy on oil prices is difficult. Some SPR crude oil releases occurred in response to severe oil market supply disruptions (e.g., 1991 Desert Storm sale) while others occurred because oil prices were relatively low and stable (e.g., 1996 sales to reduce the deficit). Similarly, the largest SPR purchases were a response to high oil prices caused by the Arab Oil Embargo. Purchases for the SPR were also made in the late 1980s, while there was excess capacity in oil markets and prices were relatively low.

Isolating the effect of SPR policy is challenging because the policy depends, in part, on the state of the oil market. Changes in the price of oil following an SPR policy action reflects both the policy and market conditions to which the policy is responding. The endogeniety of crude oil supply, crude oil demand, and SPR policy has likely led to the widely varying estimates of the SPR's effect on the price of oil.

Given these difficulties, I use several structural vector autoregression (VAR) models of the U.S. oil market to estimate the effect of SPR policy. Since Kilian (2009), structural VAR models have been used to model both the global and U.S. oil markets (Apergis and Miller, 2009; Lippi and Nobili, 2012; Baumeister and Peersman, 2013b; Jo, 2014). This paper contributes to this literature by adding oil market policy variables to the VAR model, following the voluminous monetary policy literature (Bernanke and Blinder, 1992; Sims, 1992; Bernanke and Mihov, 1998; Christiano et al., 1999). As a starting point, I identify SPR policy effects by restricting the structural shocks through the recursiveness assumption. For example, the exclusion restrictions in this model imply that oil supply and oil demand respond only with a lag to SPR policy shocks. To justify this assumption, I estimate the model with weekly data, rather than the monthly or quarterly data typically used in the oil market literature. In this benchmark model, an unanticipated SPR purchase raises the price of the price of oil 1.5 percent over 20 weeks following purchase, but an unanticipated SPR release does not have a statistically significant effect on the price of oil.

Though VAR models in both the oil market and monetary policy literature are identified by assuming recursiveness (see Kilian, 2009; Christiano et al., 1996), this identifying assumption is controversial (Rudebusch, 1998). Papers on oil markets have increasingly relied on sign restrictions to identify structural shocks (Baumeister and Peersman, 2013b; Kilian and Murphy, 2012). I take a different approach and identify policy shocks using information that is external to the VAR model, but correlated with unanticipated policy actions (Stock and Watson, 2008; Olea et al., 2012). Following work by Romer and Romer (1989), who use the minutes from Federal Reserve meetings to identify monetary policy actions that were exogenous with respect to market conditions, many papers have constructed external instruments to identify structural shocks (Stock and Watson, 2012). I construct two separate instruments to partially identify the effects of SPR purchases and releases in a VAR framework.

First, I use the SPR purchase schedule as an instrument to identify unanticipated SPR purchases. From 1999 to 2010, the Department of Energy (DOE) hired private contractors to purchase oil for the SPR. These contractors were required to deliver a specified amount of oil over a fixed time period (e.g., a contractor would agree deliver 9.4 million barrels of oil from April to July 2005). The purchase schedules were not immediately announced and exemptions from the schedule were rarely granted, which makes the purchase schedule correlated with unanticipated SPR purchases. The schedule was also set well ahead of the delivery window, so the purchase schedule is uncorrelated with other structural shocks at the time of purchase. Using the purchase schedule as an instrument to partially identify the VAR model, I estimate that SPR purchase shocks cause crude oil prices to rise 6 percent over 15 weeks following purchase.

Next, I use crude oil futures data to identify the effect of SPR releases on crude oil prices. This method follows work by Bagliano and Favero (1999) and Cochrane and Piazzesi (2002), who estimate monetary policy shocks with Federal Funds futures data. In these models, the change in futures prices immediately following a policy announcement is used to estimate the unanticipated component of the policy. For example, an unanticipated SPR release announcement would be immediately followed by a drop in the crude oil futures price, whereas a wholly anticipated SPR release announcement would elicit no response from the crude oil futures market. Using the benchmark U.S. crude oil futures contract, West Texas Intermediate (WTI), to estimate policy shocks confirms that unanticipated crude oil releases from the SPR do not lower the price of crude oil.

To ensure the change in oil futures prices following an SPR release reflects only the policy effect, and not other shocks to the oil market, I construct a novel financial instrument: the spread between the U.S. (WTI) and Canadian (WCS) crude oil futures. Given the similarities and proximity of the U.S. and Canadian crude oil markets, a drop in the U.S. crude futures price relative to Canadian crude futures immediately following a SPR release announcement should capture only the effect of SPR policy, and not other shocks to the oil market. This instrument also shows that unanticipated SPR releases do not lower the price of crude oil.

Under a variety of empirically plausible identifying assumptions, SPR purchase shocks raise oil prices, but SPR release shocks have no measurable impact on prices. The asymmetric effect of SPR policy is puzzling. Following the growing uncertainty literature (Aastveit et al., 2013; Bekaert et al., 2013), I augment the VAR model with an uncertainty variable and policy-uncertainty interaction terms, which allows SPR policy to have time-varying price effects. I find that an unanticipated SPR purchase has no measurable effect on prices when oil market uncertainty is low, while an unanticipated purchase increases the price of oil 3 percent when oil market uncertainty is high. The price of crude oil is unaffected by SPR release shocks at all levels of oil market uncertainty. These results are robust to using alternate measures of uncertainty based on stock market volatility and policy uncertainty. The asymmetric effect of SPR purchases under uncertainty is consistent with the literature on the value of real options under uncertainty (Dixit and Pindyk, 1994; Bloom, 2009).

The central result of my paper–SPR release shocks do not lower oil prices but SPR purchase shocks raise oil prices–is at odds with the academic literature and has a number of implications for SPR management. First, SPR purchases should be avoided during periods of high oil market uncertainty. Though policymakers tend to be interested in filling the reserve when oil markets are volatile (Government Accountability Office, 2006), I find those purchases have a large, immediate cost and have no measurable effect on prices when the oil is eventually released. Second, crude oil releases from the SPR should not be relied on to lower oil prices. A policymaker facing a spike in oil prices should explore other policy options to lower the price of crude oil.

The remainder of the paper is organized as follows. Section 2 provides an overview of the related literature. Section 3.1 presents a benchmark VAR model of the U.S. oil market with SPR policy variables. Sections 3.2 and 3.3 show that the results of the benchmark model hold under alternative identifying assumptions that incorporate external information. Section 4 presents an uncertainty model with time-varying policy effects. Concluding comments are contained in Section 5.

2 Literature Review

The increasing frequency of oil supply disruptions from the 1950 to the 1970s led economists to propose a government controlled crude oil reserve to protect the economy from insecure oil suppliers (Cabinet Task Force on Oil Import Control, 1970; National Petroleum Council, 1973; Nordhaus, 1974). These calls for the U.S. government to create a publicly managed crude oil reserve intensified during the Arab Oil Embargo (Tolley and Wilman, 1977). The theory behind these policy proposals is simple: the government could offset supply disruptions by increasing oil supply with the SPR rather than by restricting oil demand with unpopular quotas, tariffs, or taxes. Since oil price spikes had consistently preceded economic downturns in the postwar economy (Hamilton, 1983), a policy mechanism that controlled oil prices had broad political appeal (Weimer, 1982).

Much of the early work on the SPR focused on optimal management strategy (Teisberg, 1981; Balas, 1981; Chao and Manne, 1983; Samouilidis and Berahas, 1982). Most of these papers assumed that crude oil released from the SPR during a supply disruption would lower prices, but crude oil purchases during normal market conditions would not significantly raise prices. These assumptions could not be tested with data available at the time. Few additional papers were written on the topic until China began building its strategic oil reserve in 2004, and economists estimated optimal management strategies for China's strategic reserve using similar, untested assumptions (see Wu et al., 2008; Zhang et al., 2009; Bai et al., 2012).

Only a few papers have directly estimated the effect of strategic reserves on crude oil prices. Considine (2006) used a simultaneous equation model of the SPR to estimate its effect on oil prices between 1992 and 2005. Consistent with the assumptions in the literature, Considine found that filling the reserve increased prices imperceptibly (0.4 percent), while releasing oil during a supply shock can lower prices by 3.5 percent. Other estimates, which use models of the forward price curve, find that SPR purchases and releases have a symmetric effect on prices and can change the price of oil by as much as 32 percent (Verleger, 2003).

The key difficulty in the empirical literature is to isolate the effect of SPR policy from the state of the oil market to which the policy responds. I contribute to the literature by using a structural VAR model of the U.S. oil market to disentangle these effects. Beginning with Killian's (2009) structural VAR model of the global oil markets, VAR models have been used to study country-level oil markets (Peersman and Van Robays, 2009, 2012; Baumeister and Peersman, 2013a; Kilian and Murphy, 2014), volatility in oil markets (Van Robays, 2012; Baumeister and Peersman, 2013b; Jo, 2014), and storage in oil markets (Kilian and Murphy, 2014). I extend these models by explicitly including oil market policy variables in a VAR model of the U.S. oil market.

To estimate the effect of SPR policy in a VAR framework, I rely on the extensive work on VAR models of monetary policy (see Christiano et al., 1999). In this literature, the effects of monetary policy have been estimated using many different identification strategies and estimation methods (Sims, 1986; Bernanke and Blinder, 1992; Leeper et al., 1996; Bernanke and Mihov, 1998). Applying the time series models developed in this field to oil markets allows me to answer the question, "How do oil prices respond to unanticipated SPR purchases and releases?" with several sets of plausible identifying assumptions.

This paper also applies identification methods recently developed to incorporate external information in VAR models (Stock and Watson, 2008; Olea et al., 2012; Stock and Watson, 2012). Beginning with Romer and Romer (1989), who used minutes from Federal Reserve meetings to identify unanticipated monetary policy actions, many papers have constructed exogenous instruments to identify structural shocks (Ramey and Shapiro, 1999; Barro and Redlick, 2011; Stock and Watson, 2012; Mertens and Ravn, 2013). However, oil market VAR models are typically identified using recursive timing or sign restrictions. These models rely on restrictions on the coefficients or residuals in the model–what Stock and Watson (2008) call "internal restrictions." I construct two separate external instruments to partially identify the effects of SPR purchases and releases in a VAR framework, without internal restrictions.

A final contribution of this paper is an analysis of SPR policymaking under uncertainty. The uncertainty literature has shown the effects of economic policy vary with uncertainty (Bloom, 2009; Bekaert et al., 2013; Baker et al., 2013; Aastveit et al., 2013). Specifically, the stimulative effect of monetary policy can be dampened by high economic uncertainty when producers face non-convex adjustment costs. I extend VAR models of oil market uncertainty (Jo, 2014; Baumeister and Peersman, 2013a) by including SPR policy variables and policy-uncertainty interaction terms. This model allows me to show that unlike other economic policies, the effect of SPR purchases is amplified by uncertainty.

3 Vector Autoregression Models

3.1 Benchmark Model

In this section, I use a system of equations to model the U.S. oil market that include SPR policy variables. Endogeneity between SPR policy and other oil market variables complicates estimation. The exogenous components of SPR policy, SPR policy shocks, are identified using restrictions on the contemporaneous effects of oil market variables. These exogenous policy shocks are then used to estimate the effect of SPR policy on crude oil prices.

The SPR policy function is modeled as:

$$\text{SPR}_t = P(\Omega_t) + \varepsilon_t^{\text{SPR}} \tag{1}$$

where SPR_t represents the government's oil market policy instruments: crude oil purchases and crude oil releases. Following Christiano et al. (1999), SPR policy at time t is related to the policymaker's time t information set, Ω_t, by some linear function P. The portion of SPR policy that reflects the policymaker's systematic reaction to changes in the oil market is given by $P(\Omega_t)$. The residual component of SPR policy that is not accounted for by changes in the oil market is the SPR policy shock, $\varepsilon_t^{\text{SPR}}$. This unanticipated component of policy is assumed to be exogenous to the information set, Ω_t.

In the remainder of the paper, I rely on this distinction between SPR policy *shocks* and SPR policy *actions* (Bagliano and Favero, 1999). A SPR policy action is determined by the state of the oil market through the policy function. A SPR policy shock is the component of observed SPR policy that is not predicted by the policy function. Since SPR policy shocks are deviations from the policy rule that are exogenous to the state of the oil market, they can be used in causal analysis. Observed SPR policy is the sum of a SPR policy action and a SPR policy shock.

There are several economic interpretations of these SPR policy shocks (Christiano et al., 1999). The shocks could reflect an exogenous change in policymaker preferences. For example, a new presidential administration could change the weight given to oil supply shocks relative to oil demand shocks in their policy function. An exogenous policy shock could also reflect the policymaker's response to market expectations of SPR policy. If the President announced an SPR release due to private sector expectations, rather than oil market conditions, that would reflect an exogenous SPR shock in this model. Finally, policymakers may rely on preliminary oil market data that is measured with error. The exogenous variation in the policymaker's reaction function induced by measurement error would be an exogenous policy shock. Regardless of interpretation, these shocks are the exogenous component of SPR policy that will be used to identify the effects on the oil market.

To estimate SPR policy shocks, the vector of oil market variables, Y_t, is partitioned as follows,

$$Y_t = \begin{bmatrix} Y_t^{\text{slow}} \\ \text{SPR}_t \\ Y_t^{\text{fast}} \end{bmatrix} \tag{2}$$

These vectors are contained in the policymaker's information set, Ω_t. I assume the variables in the Y_t^{slow} vector do not respond contemporaneously to SPR policy shocks, and SPR policy does not respond immediately to the variables in the Y_t^{fast} vector. But, the response of the variables in the Y_t^{fast} vector to policy shocks is unconstrained. These recursive timing restrictions are the key assumptions to identify policy shocks in this model, but will be relaxed in Sections 3.2 and 3.3.

The U.S. oil market is modeled with variables that capture on crude oil supply, demand, and prices (Kilian, 2009). The Y_t^{slow} vector contains the oil supply and demand variables for the U.S. market. Oil supply is the sum of domestic oil production and imports, expressed as a log difference. This measure of oil supply captures domestic and foreign supply disruptions. Oil demand is measured by the Aruba, Diebold, and Scotti (ADS) business conditions index published by the Philadelphia Federal Reserve (Aruoba et al., 2009). This index combines high- and low-frequency economic activity measures into a daily estimate, and is regularly used to estimate economic activity at high-frequencies (Berge and Jordà, 2011; Andreou et al., 2013; Monteforte and Moretti, 2013).

The SPR_t vector, plotted in Figure 1, contains separate policy variables for SPR purchases and releases. Since the SPR opened in 1977, more than 850 million barrels have been purchased over 769 weeks, which covers 46 percent of the sample period. There have been 19 separate SPR releases that have taken place over 243 weeks, about 15 percent of the sample period. The variables in the SPR_t vector are the log differences of the SPR inventory during purchases and sales, respectively. Finally, the Y_t^{fast} vector contains the oil price variable. The price of crude oil is measured by the log spot price of West Texas Intermediate (WTI) crude, deflated by the Consumer Price Index.

The vector of oil market variables is modeled as a finite order structural VAR:

$$A_0 Y_t = \alpha + \sum_{i=1}^{40} A_i Y_{t-i} + \varepsilon_t \tag{3}$$

where ε_t are mean zero and serially uncorrelated structural shocks,

$$E(\varepsilon_t | Y_{t-1}, Y_{t-2}, \dots) = 0 \tag{4}$$

$$E(\varepsilon_t \varepsilon_t') = \Sigma_\varepsilon = \begin{bmatrix} \sigma_1^2 & & 0 \\ & \ddots & \\ 0 & & \sigma_5^2 \end{bmatrix} \tag{5}$$

This benchmark model is estimated with 40 lags using weekly data from August 1, 1983 through October 5, 2014.[1] The 9 month lag length is consistent with Hamilton and Herrera (2004), who show the importance of long lag length in models of oil markets. Other VAR models of oil markets use lag lengths ranging from 4 months (Baumeister and Peersman, 2013b; Jo, 2014) to 24 months (Kilian, 2009; Kilian and Murphy, 2014).

The reduced form of the model,

$$Y_t = \underbrace{A_0^{-1}\alpha}_{\beta} + \sum_{i=1}^{40} \underbrace{A_0^{-1}A_i}_{B_i} Y_{t-i} + \underbrace{A_0^{-1}\varepsilon_t}_{e_t} \tag{6}$$

can be estimated with standard least-squares methods. However, the reduced form coefficients (β's) and residuals (e_t) cannot be used to estimate the causal effect on SPR policy on oil prices because each of the reduced form residuals is a weighted average of all structural shocks. Parameter values from the structural model are required to estimate causal effects.

The coefficients of the structural model (A_i's) cannot be recovered, but the structural shocks (ε_t) can be estimated using the reduced form residuals. I assume A_0, and by extension A_0^{-1}, is lower triangular, so the reduced form errors can be decomposed using $e_t = A_0^{-1}\varepsilon_t$,

$$e_t \equiv \begin{bmatrix} e^{\text{oil supply}} \\ e^{\text{econ activity}} \\ e^{\text{SPR purchase}} \\ e^{\text{SPR release}} \\ e^{\text{oil price}} \end{bmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 & 0 \\ a_{41} & a_{42} & a_{34} & a_{44} & 0 \\ a_{51} & a_{52} & a_{35} & a_{45} & a_{55} \end{bmatrix} \begin{bmatrix} \varepsilon^{\text{supply shock}} \\ \varepsilon^{\text{demand shock}} \\ \varepsilon^{\text{purchase shock}} \\ \varepsilon^{\text{release shock}} \\ \varepsilon^{\text{price shock}} \end{bmatrix} \tag{7}$$

[1]Though the SPR purchases began in 1977, weekly data on the SPR are only available beginning in 1983. These results are robust to including weekly SPR data interpolated from the monthly 1977–1982 data series.

The restrictions on A_0^{-1} are called the *recursiveness* assumption (Christiano et al., 1999). This assumption identifies the model by limiting the simultaneous responses among variables. For example, oil supply and demand are assumed to take at least one period to respond to SPR policy shocks.

The recursiveness assumption used to identify the elements of A_0^{-1} cannot be tested directly in this VAR model; rather, its validity comes from economic theory. As shown in equation 2, these exclusion restrictions divide the oil market variables into three categories: slow-moving variables, policy variables, and fast-moving variables. As in Killian (2009), oil supply and oil demand are slow-moving with respect to other variables in the model. U.S. oil supply (the sum of domestic production and imports) is fixed at the weekly level with respect to aggregate demand, SPR policy, and oil price shocks. This assumption is consistent with the high costs of short-term changes to oil production and import schedules. Similarly, U.S. oil demand (domestic economic activity) can respond immediately to oil supply shocks, but does not respond within a week to SPR policy or oil price shocks. These assumptions are reflected in the zeroes in the first two rows of A_0^{-1}.

SPR policymakers can respond immediately to unanticipated shocks to oil supply and economic activity, but are assumed to respond to oil price innovations with a one week lag. Since SPR policy decisions are the products of meeting-filled political processes, this assumption is reasonable. Purchases are usually scheduled well ahead of actual delivery and releases typically take more than one week to approve. The fastest SPR response to a specific oil supply disruption occurred on September 2, 2005 in response to Hurricane Katrina, seven days after the Hurricane entered the Gulf of Mexico and oil prices spiked. The zeroes in the rows three and four of A_0^{-1} reflect these assumptions. Finally, the fast-moving variable of the model, the price of oil, can immediately respond to all shocks, hence there are no zeros in the last row of A_0^{-1}.

Table 1 summarizes the SPR policy shocks over time. During the initial SPR fill (1983–1989), the average purchase shock was somewhat negative (-0.02) and the average release shock was negligible. During this period, crude oil deliveries to the SPR were slowed because of construction delays, which made the market overestimate the amount of oil to be purchased. And the only crude oil release during this period was a small test sale in 1985. Negative purchase shocks continued into the 1990s, though the SPR releases during the first Gulf War (1990–1993) were somewhat larger than markets anticipated. In the 2000s, purchase shocks were nonnegative. During that period, SPR purchases were consistently larger than market expectations and SPR release shocks were mixed.

The impulse response functions in Figures 2 and 3 display the causal effect of an unanticipated SPR policy shock over time, holding all other structural shocks constant at zero.

The 95 percent confidence interval is included in the figures, and is constructed with the semiparametric fixed-design wild bootstrap, accounting for conditional heteroskedasticity of an unknown form (Goncalves and Killian 2004). The paths plotted in these figures are robust to reordering the variables within Y_t^{slow} and SPR_t vectors.

Figure 2 shows the impulse response function for a one-standard deviation unanticipated SPR purchase. SPR purchase shocks have a sustained, statistically significant effect on the real price of oil which peaks at 1.5 percent 20 weeks after purchase before declining. However, Figure 3 shows that an unanticipated SPR release has a small (less than 0.25 percent) negative effect on the real price of oil that is not statistically significant. These estimates can be scaled to give the oil price response in terms of dollars per barrel released or purchased for the SPR. The peak effect of a one-million barrel unanticipated SPR purchase raises the price of a $100 barrel of oil by $1.76, while an unanticipated SPR release of the same magnitude lowers the price of a $100 barrel of oil by $0.28.

3.1.1 The SPR and Private Crude Oil Stocks

Early work on the SPR pointed out that increases in public oil storage could influence behavior of private oil storers (Wright and Williams, 1982). If private crude oil stocks declined in response to SPR acquisitions, or rose in response to SPR releases, the effect of SPR policy could be dampened. The gradual decline in private petroleum storage since the SPR opened in 1977 suggests that the SPR growth has been offset by private storage declines. But, there has been no empirical work on the relationship between the SPR and private storage. I use the policy shocks from the VAR model to estimate the causal effect of SPR purchases and releases on private stocks.

Firms purchasing oil from, or for, the SPR pay the transportation cost to, or from, the delivery sites in the Gulf Coast. The SPR is directly connected to the oil pipeline network that runs throughout the Midwest (PADD 2) and the Gulf Coast (PADD 3).[2] Transporting oil to, or from, the SPR through these pipelines is fast and relatively inexpensive. However, other PADDs can only move oil to, or from, the SPR by ship, which is more expensive and takes 8 to 24 days. In the most recent SPR releases, over 90 percent of SPR sales were delivered to oil companies in the Gulf Coast or Midwest. Likewise, oil purchased for the SPR tends to come from oil companies in PADD 2 and PADD 3. For this reason, I examine the effects of SPR purchases and releases on PADD 2 and PADD 3 stocks.

PADD 2 and PADD 3 stock data are only available at the monthly frequency over the

[2]Data published by the Energy Information Administration (EIA) divides the U.S. into five Petroleum Administration for Defense Districts (PADD): East Coast (PADD 1), Midwest (PADD 2), Gulf Coast (PADD 3), Rocky Mountain (PADD 4), and West Coast (PADD 5).

entire 1983 to 2014 time period. Following Kilian (2009), I convert the weekly policy shock series estimated in the VAR model into a monthly series, $\hat{\eta}_m$, by averaging the weekly shocks by month,

$$\hat{\eta}_m^{\text{Purchase}} = \frac{1}{4}\sum_{j=1}^{4} \hat{\varepsilon}_{j,m}^{\text{Purchase}} \tag{8}$$

$$\hat{\eta}_m^{\text{Release}} = \frac{1}{4}\sum_{j=1}^{4} \hat{\varepsilon}_{j,m}^{\text{Release}} \tag{9}$$

where the week j structural shock in month m is given by $\hat{\varepsilon}_{j,m}$. These monthly policy shocks can be used to estimate the causal effect of SPR policy on private stocks if they are exogenous with respect to private stocks. The process used to construct the weekly shocks in Section 3.1 guarantees endogeniety with respect to all variables in Y_t. Since private stocks are not in the information set, additional assumptions are required to use the monthly shocks to estimate the causal effect on stocks.

Identifying the causal effect of SPR policy on private stocks requires the assumption that the monthly policy shocks are exogenous with respect to private stocks. In other words, there can be no feedback between policy shocks and private stocks within a month. Though this cannot be directly tested, evidence can be provided by comparing unanticipated policy changes to innovations in private stocks. Again, following Kilian (2009), I model private stocks as an autoregressive process, and find the residuals are uncorrelated with SPR policy shocks. Since this correlation is low, I assume SPR policy shocks are fixed at the monthly level with respect to private stocks.

The response of private stocks, s_t, to policy shocks, $\hat{\eta}_m$, is modeled as,

$$s_t = \gamma + \sum_{m=1}^{9} \phi_m \hat{\eta}_m + \theta_m \tag{10}$$

where ϕ_m is the month m stock response and θ_m is the residual. The coefficients in these regressions give the response of private crude oil stocks in PADD 2 and PADD 3 to SPR purchase and release shocks. The impulse response functions based on these coefficients are plotted in Figures 4 and 5, along with the 95 percent confidence interval.

Unanticipated SPR purchases have no measurable effect on private stocks in PADD 2 and PADD 3. Given the transportation costs discussed above, unanticipated SPR purchases are likely filled by domestically produced or imported oil and not private oil stocks in PADD 2 and PADD 3. However, private storers respond to unanticipated SPR releases by raising stocks nearly 2 percent, and this effect is statistically significant at the 90 percent level. This model shows that private storers dampen the immediate effect of SPR releases by increasing their reserves.

3.2 Identification Using External Instruments

Though the estimated SPR policy effects are robust to alternate specifications that include private stocks and reordering variables within Y_t^{slow} and SPR_t, this approach is still subject to Rudebusch's (1998) critique of VAR model of policy shocks. Namely, the policy shocks estimated from linear, time-invariant functions with a small information set may not be related to actual unanticipated policy changes. Alternate identifying assumptions used to model the oil market, like sign restrictions, are subject to the same critique. In Sections 3.2 and 3.3, I depart from the oil market literature by using external information, which is related to unanticipated policy changes, to identify SPR policy shocks (Stock and Watson, 2008). These instruments confirm the key insights from this benchmark model: unanticipated SPR releases do not have a measurable effect on oil prices but SPR purchase shocks have a statistically significant, positive effect on oil prices.

The recursiveness assumption identifies unanticipated policy changes by restricting the transmission of structural shocks in the oil market. If, for example, oil production responds immediately to policy shocks, structural shocks could not be identified. Even if the recursiveness assumption holds, the SPR policy function in equation 1 could be misspecified by omitting variables. If policymakers used unmodeled variables in their "true" policy function, the policy shocks estimated in Section 3.1 would not be exogenous. In either case, the policy shocks estimated in Section 3.1 would confound the effects of SPR policy changes and the oil market conditions that cause the policy changes, and could not be used to estimate the causal effect of policy shocks (Rudebusch, 1998).

To address these issues, I partially identify the VAR presented in equation 6 with an external instrument for SPR purchases. This identification method follows the narrative approach developed by Romer and Romer (1989, 2004, 2010) to identify monetary policy shocks, which has been used regularly in subsequent research (see Ramey and Shapiro, 1999; Barro and Redlick, 2011; Olea et al., 2012). The omitted variables and simultaneity critiques do not apply to the policy shocks identified in this model if the instrument is correlated with unanticipated purchases and uncorrelated with other shocks. This approach only identifies the effects of SPR purchase shocks; in the next section I use a separate instrument to identify SPR release shocks.

Identifying the causal effect of SPR purchases on oil prices requires an instrument with the following properties: the instrument is not included in Y_t, the instrument is correlated with unanticipated purchases, and the instrument is uncorrelated with contemporaneous oil market conditions (excluding policy effects). The SPR purchase instrument exploits the procedures used by the Department of Energy (DOE) to acquire crude oil for the SPR. From 1999 to 2010, DOE did not directly purchase crude oil from the SPR, rather they negotiated contracts with private companies to purchase and deliver crude oil to the SPR (U.S. Senate, 2003; Congressional Research Service, 2006).

The DOE's crude oil acquisition procedure gives the purchase schedule two useful properties: the schedule is correlated with unanticipated SPR purchases but uncorrelated with other structural shocks. The DOE deliberately kept oil market participants in the dark about the timing and amount of SPR purchases. If information about the purchase schedule were available to market participants, speculators could front-run the contractors and drive up prices. Since it was in the interest of the DOE and its contractors to purchase crude cheaply, it is reasonable to assume that the purchase schedule is correlated with unanticipated purchases.

The purchase schedule is assumed to be exogenous with respect to oil market conditions at the time of purchase because the schedule was set well in advance of actual purchases. The DOE rarely granted contractors exemptions from the schedule when the oil market tightened, so contractors were forced to provide oil for the SPR despite their requests to postpone purchases (Government Accountability Office, 2006).

These data are collected and coded from several official sources. To start, I obtained data on DOE purchase contracts from 1999 to 2010 through a Freedom of Information Act request. I confirm these data using the documents submitted pursuant to the 2003 Senate hearing on the SPR (U.S. Senate, 2003) and information on the DOE website. In these documents, a purchase schedule specifies the total amount of crude oil to be purchased and the time period of delivery. For example, in early 2005, DOE contracted for 9.4 million barrels of oil to be delivered from April to July 2005. These schedules are converted into a weekly series by assuming purchases are divided evenly over the delivery window.[3]

Between 1999 and 2010, the SPR grew by 154 million barrels and I have contracts for 126 million of those barrels.[4] Figure 6 plots SPR purchases along with the constructed purchase schedule, which shows that the purchase schedule has reasonable predictive power for SPR

[3]The delivery window gives contractors some flexibility in purchasing crude oil. However, the physical infrastructure of the SPR limits the maximum amount of oil that can be delivered in a single day. Typical crude oil deliveries take place over several months, even when filling the SPR at the maximum daily rate.

[4]The DOE also acquires crude oil for the SPR via short-term oil loans that are repaid with interest. The repayment of loans are not subject to the same procedures as SPR purchases and are not included in this instrument.

purchases. A regression of SPR purchases on the purchase schedule gives an R-squared of 0.35.

To identify the effect of SPR purchases on oil prices using this instrument, I estimate the model in equation 6 with standard least-squares methods. Then, following the approach outlined by Stock and Watson (2008), I partition the data into the into SPR purchases (S_t) and the other variables (X_t),

$$Y_t = \begin{bmatrix} X_t \\ S_t \end{bmatrix} \tag{11}$$

where X_t is a (4×1) vector and S_t is a scaler. The order of variables within X_t is not relevant because the recursive timing assumptions (used in Section 3.1) are not used in this model. The structural errors and reduced form residuals are similarly partitioned,

$$e_t = \begin{bmatrix} e_t^X \\ e_t^S \end{bmatrix}, \varepsilon_t = \begin{bmatrix} \varepsilon_t^X \\ \varepsilon_t^S \end{bmatrix} \tag{12}$$

Analogous to the A_0 matrix in the benchmark model, I assume there exists some matrix Q, such that $Qe_t = \varepsilon_t$, where Q can be partitioned as

$$Q = \begin{bmatrix} Q_{XX} & Q_{XS} \\ Q_{SX} & Q_{SS} \end{bmatrix} \tag{13}$$

Solving $Qe_t = \varepsilon_t$ for the reduced form errors yields,

$$e_t^X = -Q_{XX}^{-1} Q_{XS} e_t^S + Q_{XX}^{-1} \varepsilon_t^X \tag{14}$$

$$e_t^S = -Q_{SS}^{-1} Q_{SX} e_t^X + Q_{SS}^{-1} \varepsilon_t^S \tag{15}$$

There are four steps to recover the structural purchase shock, ε_t^S in equation 15, from this system of equations. First, the purchase schedule (Z_t) is used as an instrument for e_S^t in equation 14. This gives an unbiased estimate of $-Q_{XX}^{-1} Q_{XS}$ because the instrument is, by assumption, uncorrelated with all other structural shocks ($E(Z_t \varepsilon_t^X) \neq 0$) and correlated with unanticipated purchases ($E(Z_t e_t^S) \neq 0$). Next, ε_t^x is estimated up to scale using, $\widehat{\varepsilon_t^X} = e_t^X + \widehat{Q_{XX}^{-1} Q}_{XS} e_t^S$. This estimate of ε_t^X is then used as an instrument for e_t^X in equation 15, which gives an estimate of $\widehat{-Q_{SS}^{-1} Q}_{SX}$. Finally, exogenous purchase shocks are estimated, up to scale, using $\widehat{\varepsilon_t^S} = e_t^S + \widehat{Q_{SS}^{-1} Q}_{SX} e_t^X$.

With these estimated structural shocks, the dynamic causal effect of unanticipated SPR purchases on the price of oil can be estimated. The impulse response function plotted in Figure 7 shows the positive, statistically significant effect of SPR purchases on oil prices identified using the purchase instrument is quite similar to the effect identified with the recursiveness assumption. Both effects are estimated to peak around 20 weeks, though the effect estimated with the instrument is somewhat larger than the effect estimated in the benchmark model–6 percent as opposed to 1.5 percent. In this partially identified model, a one-million barrel unanticipated SPR purchase raises the price of a \$100 barrel of oil by \$7.12.

Though the SPR purchase effect identified with the instrument is larger than the purchase effect identified in Section 3.1, it is important to keep in mind that this instrument is only available between 1999-2010. During that period, oil prices jumped sharply, then collapsed during the financial crisis and recession. Excluding the commodity price boom and financial crisis from the sample, by dropping data from 2006 to 2010, lowers the peak effect on oil prices to about 3.5 percent.

It is also important to note that this method assumes that there is no heterogeneity among SPR purchase shocks. That is, an unanticipated SPR purchase shock during the 2003 Iraq War has the same effect as an unanticipated purchase while oil markets were calm in the late 1990s. If purchase shocks are heterogenous, then this model estimates the *local* average treatment effect of SPR purchase shocks, not the average treatment effect. I address the issue of heterogenous policy shocks in Section 4, where the effect of SPR policy is modeled as a function of oil market uncertainty.

3.3 Identification Using Daily Futures Market Data

As a robustness check of the SPR release effect estimated in Section 3.1, I use crude oil futures market data to identify SPR release shocks. This section offers a different approach to the simultaneity and omitted variables problems discussed in Section 3.2. Recall that the identification restrictions used in the benchmark model do not allow SPR policy to respond within a week to oil price shocks. If policymakers respond immediately to oil price changes, the release effect estimated in Section 3.1 would be biased. The benchmark VAR model also assumes a simple SPR policy function that only takes into account oil supply, demand, and price. All other factors that determine SPR policy (e.g., political opinion polling or gasoline prices) are omitted. Though monetary policy functions are similarly modeled (Christiano et al., 1999), there are concerns that these policy functions omit key variables used by policymakers (Rudebusch, 1998).

Following Bagliano and Favero (1999) and Cochrane and Piazzesi (2002), I use daily oil futures data to overcome these difficulties. Crude oil futures market expectations incorporate far more information about oil markets than could be modeled in a VAR. This identification strategy uses changes in crude oil futures prices immediately after SPR release announcements to estimate whether the policy announcement was unanticipated. This structural shock estimate is then used to construct an impulse response function.

This model for estimating unanticipated SPR releases has two key assumptions: intraday changes in the crude oil futures price measure SPR policy shocks and the SPR release shock is the only oil market shock at the time of announcement. The first assumption is based on the idea that futures prices aggregate information relevant to oil prices and futures prices only respond to unanticipated information. If the market anticipated some portion of the SPR release prior to announcement, I assume that would be included in the futures price. The second assumption is based on the idea that at a small enough time interval, the SPR release announcement is the only shock to the oil market. This assumption may not hold if the SPR release announcement reveals information about the government's view of the oil market. This method places no other restrictions on the determinants of oil price expectations or oil market dynamics.

Figure 8 illustrates this method with two examples from the WTI futures market around two SPR release announcements: the 1996 deficit reduction sales and the 2011 Libya sale. On April 25, 1996, one day prior to the deficit reduction sale announcement, the one-month WTI futures contract closed at $14.20 (Figure 8b). The futures price did not respond to the April 26 release announcement, closing the day at $14.28, which suggests the futures market anticipated this release. Since this SPR sale was ordered by Congress to reduce the deficit and the release amount and approximate timing had been discussed publicly for months, it is not surprising that the release was already priced into futures contracts.

On June 22, the day before the 2011 SPR release announcement, the WTI futures market closed at $42.27. Following the 9:30 am release announcement, the futures price dropped sharply, and the market closed at $40.32 on June 23 (Figure 8a). This release came at the end of the Arab Spring after Libyan oil production dropped sharply. Though there was speculation about the release, the size and timing were not clear prior to the announcement. The nearly $2 drop in the crude oil futures price suggests this policy announcement was unanticipated by the market.

The validity of the assumptions relies on the choice of futures contract. Two financial instruments are used to estimate the structural SPR release shocks. The first is the one-month WTI futures price. An unanticipated SPR release occurs when the WTI futures price falls following the announcement. Since the movements in the WTI futures price could reflect any oil market supply or demand shock, I use another financial instrument to isolate the SPR effect: the one-month US-Canadian crude futures (WTI-WCS) spread. The U.S. and Canadian oil markets are similar and subject to similar shocks. However, only the U.S. has a strategic oil reserve, and Canadian oil companies cannot purchase oil from the SPR. Differencing U.S. and Canadian oil should remove the effects of non-SPR supply and demand shocks to the North American oil market, which allows me to directly estimate the magnitude of the SPR release shock.

The US-Canadian crude futures spread identifies unanticipated releases that were not followed by a decline in the WTI futures price. For example, the WTI futures price edged up \$0.40 (1.6 percent) following the Hurricane Ivan release (Figure 9). This SPR announcement would not be considered an SPR release shock using the WTI futures price. However, the WCS futures price rose \$2.05 (5.3 percent) following the release announcement. Since the futures trends pre- and post-announcement are nearly identical, the change in the futures spread indicates the WTI futures price rose less than expected. Since WTI became less expensive, relative to WCS, following an SPR release announcement, the release was unanticipated according to the WTI-WCS spread.

There are two steps to estimate the causal effect of the SPR releases in this model. First, SPR release shocks are estimated for the entire sample of 19 SPR releases and the shock variable is set to zero on weeks with no announcement.[5] I construct the structural shock variable by taking the difference between the futures price at market close the day prior the announcement and the futures price at market close the day of the announcement.[6] Using the WTI futures contract, seven SPR releases were unanticipated, while the WTI-WCS futures spread estimates that eight of the releases were unanticipated. Second, I regress the change in the spot price of crude oil, $P_{t+j} - P_t$, on the series of estimated structural shocks, $\varepsilon_t^{release}$. The coefficients of this regression give consistent estimates of the causal effects of unanticipated SPR releases on the price of oil, which I use in the impulse response functions.

Figures 10 and 11 plot the impulse response functions using the SPR release shocks estimated using WTI futures and the WTI-WCS futures spread, respectively, along with the 95 percent confidence interval. Consistent with the results from the benchmark model in Section 3.1, unanticipated SPR releases do not have a statistically significant, negative effect on the price of oil using either futures contract. In fact, the effect of unanticipated SPR releases increases the price of oil nearly 2 percent using shocks identified with the WTI-WCS futures spread. This is similar to the results from Cochrane and Piazzesi (2002), who show

[5]To ensure that the release announcement is the only effect on crude futures, I drop any SPR release announcements that occur during OPEC meetings or on days when the Energy Information Administration releases oil market data.

[6]To allow for information leaks preceding the announcement, the two-day change was also used. The results are qualitatively identical.

that unanticipated monetary policy shocks, as measured by futures prices, that are intended to drive down interest rates can have the opposite effect.

4 Uncertainty Model

In this section I explore the mechanism driving the SPR purchase and release effects. I begin by testing whether the effect of SPR policy varies with uncertainty. The growing empirical literature on uncertainty has shown that uncertainty shocks dampen the effects of monetary and fiscal policy (Bloom, 2009; Bekaert et al., 2013; Baker et al., 2013; Aastveit et al., 2013). These results are built on real options theory which has shown that uncertainty lowers the value of options when investments are not reversible (Dixit and Pindyk, 1994). This section presents a framework for SPR policy effects to vary with oil market uncertainty by augmenting the model VAR in Section 3.1 with an oil market uncertainty variable and policy-uncertainty interaction terms.

I cannot directly measure oil market uncertainty, so I use several different variables to capture uncertainty regarding oil prices, economic policy, and economic activity. The 90-day average volatility of the WTI spot price is my primary measure of oil market uncertainty. This common index of oil market uncertainty is the 90-day moving average of the standard deviation of intraday spot oil price changes. The volatility of oil prices is an imperfect measure of oil market uncertainty, so I also use other indices of policy and overall economic uncertainty that have been proposed in the literature. Baker et al. (2013) construct an index of policy uncertainty using a variety of news-based uncertainty measures. This variable is designed to capture uncertainty regarding economic policy. I also use an index of stock market uncertainty, VXO, which estimates market volatility based on S&P 500 options prices, to capture general economic uncertainty. These uncertainty indices are plotted in Figure 12.

It is important to note that oil market uncertainty is not always associated with high prices. For example, oil price volatility spiked following the 2007 financial crisis, but oil prices were at their lowest level in five years. SPR releases have also not been restricted to times of high-uncertainty. Crude oil has been released through several different programs, including emergency drawdowns, infrastructure test sales, sales to reduce the deficit, and short-term leases. Some sales have occurred during severe global oil market supply disruptions with high uncertainty (e.g., 1991 Desert Storm sale), while others have occurred during localized supply disruptions with low uncertainty (e.g., 2004 Hurricane Ivan release). SPR releases have also been conducted while oil prices were relatively low and stable in 1985 to test the

SPR infrastructure and in 1996 to reduce the Federal deficit. Likewise, crude oil purchases have occurred during periods of high and low oil market uncertainty. This variation in SPR policy and oil market volatility is used to estimate whether the effect of SPR policy on oil price fluctuates with uncertainty.

I formalize this approach by including oil market uncertainty and SPR policy-uncertainty interaction terms in the benchmark VAR model from Section 3.1. These uncertainty terms add flexibility to the model allowing the effect of oil market interventions to vary over time with uncertainty, which can shed light on puzzling asymmetric response of oil prices to SPR purchases and releases. Unlike other studies of oil markets with time varying coefficients (Baumeister and Peersman, 2013b), the time varying effects of SPR policy shocks are driven by a specific exogenous variable, uncertainty.

Following Aastveit et al.'s (2013) work on monetary policy under uncertainty, I include an exogenous uncertainty variable, V_t, and uncertainty interaction terms for SPR purchases and releases in the oil price equation:

$$A_0 Y_t = \alpha + \sum_{i=1}^{40} A_i Y_{t-i} + C_i V_t \, Y_{t-i} + D V_t + \varepsilon_t \tag{16}$$

where V_t is the uncertainty measure and $V_t \, Y_t$ is the policy-uncertainty interaction term. To study the effect of uncertainty on the SPR policy effect on prices, the coefficients on both the volatility and volatility interaction variables are constrained as follows,

$$C = \begin{bmatrix} 0_{4\times 2} & 0_{4\times 1} & 0_{4\times 1} & 0_{4\times 1} \\ 0_{1\times 2} & c_1 & c_2 & 0 \end{bmatrix}, D = \begin{bmatrix} 0_{4\times 1} \\ c \end{bmatrix} \tag{17}$$

The C matrix limits the uncertainty interaction terms to the SPR policy variables in the oil price equation. The D matrix limits the uncertainty terms to the oil price equation. Given the sample size, these constraints are necessary to restrict the number of variables in the model.

The reduced form of this model,

$$Y_t = \underbrace{A_0^{-1}\alpha}_{\beta} + \sum_{i=1}^{40} \underbrace{A_0^{-1}A_i}_{\beta_i} Y_{t-i} + \underbrace{A_0^{-1}C_i}_{G_i} V_t \, Y_{t-i} + \underbrace{A_0^{-1}D}_{H} V_t + \underbrace{A_0^{-1}\varepsilon_t}_{e_t} \tag{18}$$

is estimated with standard least-squares methods. As with the VAR model in Section 3.1, the coefficients in the structural model (A_i, C_i, and D) cannot be recovered, but the reduced form residuals (e_t) can be converted into structural shocks (ε_t) with recursive timing restrictions. These estimated shocks are used to construct impulse response functions.

The SPR policy effects shown in the impulse response functions in Figures 13 and 14 use oil market volatility to capture uncertainty and are similar to those estimated in the benchmark model without uncertainty.[7] In both models, SPR purchases increase the price of oil 1 to 2 percent and SPR releases have a small negative, but not statistically significant, effect on oil prices.

The reduced form coefficients can be used to estimate the effect of SPR purchases and releases during periods of high and low uncertainty. Having estimated equation 18, I use the coefficients for V^{high} and V^{low} to simplify the model as follows,

$$Y_t^{high} = \underbrace{\widehat{D_0}^{high}}_{\widehat{B}+\widehat{H}V^{high}} + \sum_{i=1}^{40} \underbrace{\widehat{D_i}^{high}}_{\beta_i+G_i} Y_{t-i} + \widehat{e}_t \tag{19}$$

$$Y_t^{low} = \underbrace{\widehat{D_0}^{low}}_{\widehat{B}+\widehat{H}V^{low}} + \sum_{i=1}^{40} \underbrace{\widehat{D_i}^{low}}_{\beta_i+G_i} Y_{t-i} + \widehat{e}_t \tag{20}$$

These equations allow the impulse response functions to be calculated with fixed uncertainty levels. Specifically, V^{high} is defined as the 90^{th} percentile volatility and V^{low} is defined as the 10^{th} percentile volatility.

SPR releases have no measurable effect on oil prices at any either high or low uncertainty, regardless of the uncertainty measure. However, Figures 15 through 20 show the effect of SPR purchase shocks on oil prices increases with uncertainty, across all uncertainty measures. When volatility is low, unanticipated SPR purchases do not have a statistically significant effect on prices (Figures 15, 16, and 17). This confirms claims by the Department of Energy that crude oil purchases during periods of low oil market volatility do not affect oil prices (U.S. Senate, 2003). However, when uncertainty is high, SPR purchase shocks increase the price of oil from 2 percent, using political uncertainty, to 4 percent, using oil market or economic uncertainty (Figures 18, 19, and 20). Unfortunately, SPR purchases tend to occur during periods of high oil market volatility.

5 Conclusion

This paper analyzes SPR policy in a time series model of the U.S. oil market. I begin with the observation that much of the SPR literature has focused on optimal SPR management and does not directly estimate the effect of SPR policy on oil prices. Rather, these papers

[7]The impulse response functions estimated using the stock market volatility and political uncertainty indices are nearly identical and omitted.

assume that SPR crude oil releases lower the price of oil during supply disruptions while purchases only raise the price slightly, if at all. Contrary to these assumptions, I show in a structural VAR model with recursive timing restrictions that an unanticipated SPR release has no measurable effect on oil prices, while an unanticipated SPR purchase increases the price of oil 1.5 percent over 20 weeks following purchase.

I explore alternate identification methods for SPR policy using information that is not contained in the VAR model. Oil market VAR models are typically identified by restricting the timing or signs of structural shocks (Kilian, 2009; Baumeister and Peersman, 2013a,b; Jo, 2014; Kilian and Murphy, 2014). I show that a VAR model of the oil market can be partially identified under different, less restrictive assumptions. I begin by constructing an instrument for SPR purchases from the purchase schedules set by the Department of Energy. These schedules are set months ahead of actual purchases, making them exogenous with respect to other oil market shocks at the time of purchase. The purchase schedules were not immediately announced publicly and exemptions were rarely granted, so the schedules are correlated with unanticipated SPR purchases. This instrument identifies SPR purchase shocks in the VAR framework without other restrictions on the structural shocks. An unanticipated SPR crude oil purchase, identified with the purchase schedule, increases the price of crude 6 percent over 15 weeks following purchase.

I use a different model to identify the effect of SPR releases with high-frequency crude oil futures data. The change in crude oil futures prices immediately following release announcements is used to measure whether SPR releases are unanticipated by the market. Estimating policy shocks with this method avoids problems with endogeneity and omitted variables that have been raised concerning policy shocks estimated in traditional VAR models (Rudebusch, 1998). Several different crude oil futures contracts are used to isolate the unanticipated policy shocks. This model confirms that unanticipated SPR releases do not have a statistically significant effect on oil prices.

Finally, I explore a mechanism driving the SPR purchase and release effects by testing whether the effect of SPR policy varies with uncertainty. I augment the structural VAR with an oil market uncertainty measure and policy-uncertainty interaction terms. This allows the effects of SPR policy to vary over time with uncertainly. I find that SPR policy shocks only impact oil prices when purchases are made during periods of high uncertainty. SPR purchase shocks increase oil prices 2 to 4 percent when uncertainty is high, depending on the uncertainty measure. Unanticipated purchases when uncertainty is low, and releases at any time, do not have an effect on prices.

The main contribution of this paper is to develop a framework for estimating the effect of U.S. oil market interventions on oil prices. I show that under a variety of identifying assumptions, SPR releases have not lowered the price of oil, while SPR purchases have increased its price. These results have strong implications beyond the U.S. and global oil markets. Over half of global crude oil stocks are held in government-controlled reserves in more than 40 countries. Many of these countries are following the U.S. by expanding their government-controlled crude oil reserves (International Energy Agency, 2014). In light of these results, strategic reserve acquisitions should be limited to periods of low oil market uncertainty to minimize the price impact, and unanticipated releases from these reserves should not be expected to lower oil prices.

References

Aastveit, Knut Are, Gisle James Natvik, and Sergio Sola (2013), "Economic Uncertainty and the Effectiveness of Monetary Policy." Technical report, Norges Bank.

Andreou, Elena, Eric Ghysels, and Andros Kourtellos (2013), "Should Macroeconomic Forecasters use Daily Financial Data and How?" *Journal of Business & Economic Statistics*, 31, 240–251.

Apergis, Nicholas and Stephen M Miller (2009), "Do Structural Oil-Market Shocks Affect Stock Prices?" *Energy Economics*, 31, 569–575.

Aruoba, S Boragan, Francis X Diebold, and Chiara Scotti (2009), "Real-time measurement of business conditions." *Journal of Business & Economic Statistics*, 27.

Bagliano, Fabio C and Carlo A Favero (1999), "Information from Financial Markets and VAR Measures of Monetary Policy." *European Economic Review*, 43, 825–837.

Bai, Y, DQ Zhou, P Zhou, and LB Zhang (2012), "Optimal Path for China's Strategic Petroleum Reserve: A Dynamic Programming Analysis." *Energy Economics*, 34, 1058–1063.

Baker, Scott R, Nicholas Bloom, and Steven J Davis (2013), "Measuring Economic Policy Uncertainty." Unpublished Working Paper.

Balas, Egon (1981), *The Strategic Petroleum Reserve: How Large Should it be?* in Energy Policy Planning, Bayraktar et al. (Eds.), Plenum Press, New York.

Barro, Robert J. and Charles J. Redlick (2011), "Macroeconomic Effects From Government Purchases and Taxes." *The Quarterly Journal of Economics*, 126, 51–102.

Baumeister, Christiane and Gert Peersman (2013a), "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market." *Journal of Applied Econometrics*, 28, 1087–1109.

Baumeister, Christiane and Gert Peersman (2013b), "Time-Varying Effects of Oil Supply Shocks on the US Economy." *American Economic Journal: Macroeconomics*, 5, 1–28.

Bekaert, Geert, Marie Hoerova, and Marco Lo Duca (2013), "Risk, Uncertainty and Monetary Policy." *Journal of Monetary Economics*, 60, 771–788.

Berge, Travis J and Òscar Jordà (2011), "Evaluating the Classification of Economic Activity into Recessions and Expansions." *American Economic Journal: Macroeconomics*, 3, 246–277.

Bernanke, Ben S and Alan S Blinder (1992), "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review*, 82, 901–21.

Bernanke, Ben S. and Ilian Mihov (1998), "Measuring Monetary Policy." *The Quarterly Journal of Economics*, 113, 869–902.

Bloom, Nicholas (2009), "The Impact of Uncertainty Shocks." *Econometrica*, 77, 623–685.

Bohi, Douglas and Michael Toman (1996), *The Economics of Energy Security*. Norwell, MA: Kluwer Academic Publishers.

Cabinet Task Force on Oil Import Control (1970), *The Oil Import Question*. Washington D.C.: U.S. Government Printing Office.

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

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999), "Monetary Policy Shocks: What Have we Learned and to What End?" *Handbook of Macroeconomics*, 1,

65–148.
Christiano, Lawrence J, Martin Eichenbaum, and Charles Evans (1996), “The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds.” *The Review of Economics and Statistics*, 78, 16–34.
Cochrane, Jone and Monika Piazzesi (2002), “The Fed and Interest Rates - A High-Frequency Identification.” *American Economic Review*, 92, 90–95.
Congressional Research Service (2006), *Energy Policy Act of 2005: Summary and Analysis of Enacted Provisions*. RL33302, Library of Congress.
Considine, Timothy J (2006), “Is the Strategic Petroleum Reserve Our Ace in the Hole?” *The Energy Journal*, 91–112.
Dixit, Avinash and Robert Pindyk (1994), *Investment under uncertainty*. Princeton University Press.
Government Accountability Office (2006), *Available Oil Can Provide Significant Benefits, but Many Factors Should Influence Future Decesions about Fill, Use, and Expansion*. Publication No. GAO-06-872.
Hamilton, James D (1983), “Oil and the Macroeconomy since World War II.” *The Journal of Political Economy*, 228–248.
Hamilton, James D and Ana Maria Herrera (2004), “Comment: Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy.” *Journal of Money, Credit and Banking*, 265–286.
International Energy Agency (2014), *Oil Information, 2014 Edition*. Paris: International Energy Agency.
Jo, Soojin (2014), “The Effects of Oil Price Uncertainty on Global Real Economic Activity.” *Journal of Money, Credit and Banking*, 46, 1113–1135.
Kilian, Lutz (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market.” *American Economic Review*, 99, 1053–69.
Kilian, Lutz and Daniel 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, Lutz and Daniel P. Murphy (2014), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil.” *Journal of Applied Econometrics*, 29, 454–478.
Leeper, Eric M, Christopher A Sims, Tao Zha, Robert E Hall, and Ben S Bernanke (1996), “What Does Monetary Policy Do?” *Brookings papers on economic activity*, 1–78.
Leiby, Paul N and David Bowman (2000), “The Value of Expanding the US Strategic Petroleum Reserve.” *ORNL/TM-2000/179, Oak Ridge National Laboratory, Oak Ridge, Tennessee, November*.
Lippi, Francesco and Andrea Nobili (2012), “Oil and the Macroeconomy: A Quantitative Structural Analysis.” *Journal of the European Economic Association*, 10, 1059–1083.
Mertens, Karel and Morten O Ravn (2013), “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *The American Economic Review*, 103, 1212–1247.
Monteforte, Libero and Gianluca Moretti (2013), “Real-Time Forecasts of Inflation: The Role of Financial Variables.” *Journal of Forecasting*, 32, 51–61.
National Petroleum Council (1973), *Emergency Preparedness for Interruption of Petroleum Imports into the United States*. Department of Interior. Washington D.C.: U.S. Govern-

ment Printing Office.

Nichols, Albert L and Richard J Zeckhauser (1977), "Stockpiling Strategies and Cartel prices." *The Bell Journal of Economics*, 66–96.

Nordhaus, William D (1974), "The 1974 Report of the President's Council of Economic Advisers: Energy in the Economic Report." *The American Economic Review*, 558–565.

Olea, M, J Stock, and M Watson (2012), "Inference in Structural VARs with External Instruments." Technical report.

Peersman, Gert and Ine Van Robays (2009), "Oil and the Euro Area Economy." *Economic Policy*, 24, 603–651.

Peersman, Gert and Ine Van Robays (2012), "Cross-Country Differences in the Effects of Oil Shocks." *Energy Economics*, 34, 1532–1547.

Ramey, Valerie A. and Matthew D. Shapiro (1999), "Costly Capital Reallocation and the Effects of Government Spending." Working Paper 6283, National Bureau of Economic Research.

Romer, Christina D. and David H. Romer (1989), "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." Working Paper 2966, National Bureau of Economic Research.

Romer, Christina D. and David H. Romer (2004), "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94, 1055–1084.

Romer, Christina D. and David H. Romer (2010), "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review*, 100, 763–801.

Rudebusch, Glenn D. (1998), "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review*, 39, 907–931.

Samouilidis, J. and S. Berahas (1982), "A Methodological Approach to Strategic Petroleum Reserves." *Omega*, 10, 565–574.

Sims, Christopher A (1986), "Are Forecasting Models Usable for Policy Analysis?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 10, 2–16.

Sims, Christopher A. (1992), "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy." Cowles Foundation Discussion Papers 1011, Cowles Foundation for Research in Economics, Yale University.

Stock, James H and Mark W Watson (2008), "NBER Summer Institute Minicourse 2008: What's New in Econometrics - Time Series, Lecture 7: Structural VARs."

Stock, James H and Mark W Watson (2012), "Disentangling the Channels of the 2007-2009 Recession." Technical report, National Bureau of Economic Research.

Teisberg, Thomas J (1981), "A Dynamic Programming Model of the US Strategic Petroleum Reserve." *The Bell Journal of Economics*, 526–546.

Tolley, George S and John D Wilman (1977), "The Foreign Dependence Question." *The Journal of Political Economy*, 323–347.

U.S. Congress (1975), "S.622: The Energy Policy and Conservation Act of 1975." Pub. L. No. 94-163. U.S. Government Printing Office.

U.S. Senate (2003), "U.S. Strategic Petroleum Reserve: Recent policy has increased costs to consumers but not overal U.S. energy security (108-18)." Text from: *Committee Reports.* U.S. Government Printing Office.

Van Robays, Ine (2012), "Macroeconomic Uncertainty and the Impact of Oil Shocks." Technical report, CESifo Working Paper: Monetary Policy and International Finance.

Verleger, Phillip (2003), "Measuring the economic impact of an oil release from the Strategic Petroleum Reserve to compensate for the loss of Venezuelan oil production." Included in U.S. Strategic Petroleum Reserve (108-18), Appendix 4.

Weimer, David Leo (1982), *Strategic Petroleum Reserve: planning, implementation, and analysis*. Greenwood Press, Westport, CT.

Wright, Brian D and Jeffrey C Williams (1982), "The Roles of Public and Private Storage in Managing Oil Import Disruptions." *The Bell Journal of Economics*, 341–353.

Wu, Gang, Ying Fan, Lan-Cui Liu, and Yi-Ming Wei (2008), "An Empirical Analysis of the Dynamic Programming Model of Stockpile Acquisition Strategies for China's Strategic Petroleum Reserve." *Energy Policy*, 36, 1470 – 1478.

Zhang, Xiao-Bing, Ying Fan, and Yi-Ming Wei (2009), "A Model Based on Stochastic Dynamic programming for Determining China's Optimal Strategic Petroleum Reserve Policy." *Energy Policy*, 37, 4397–4406.

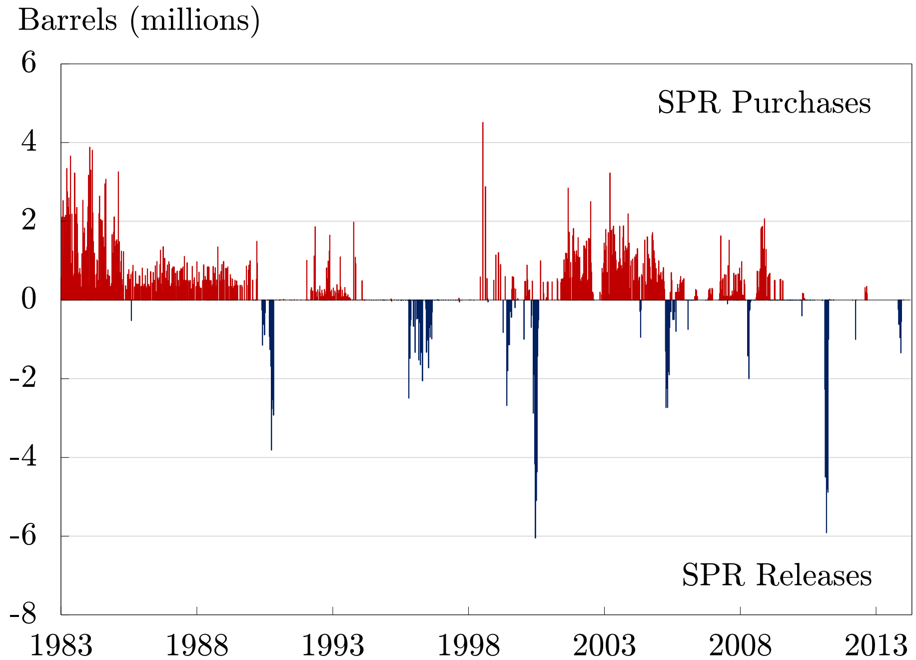

Figure 1: SPR Purchases and Releases (1983–2014)

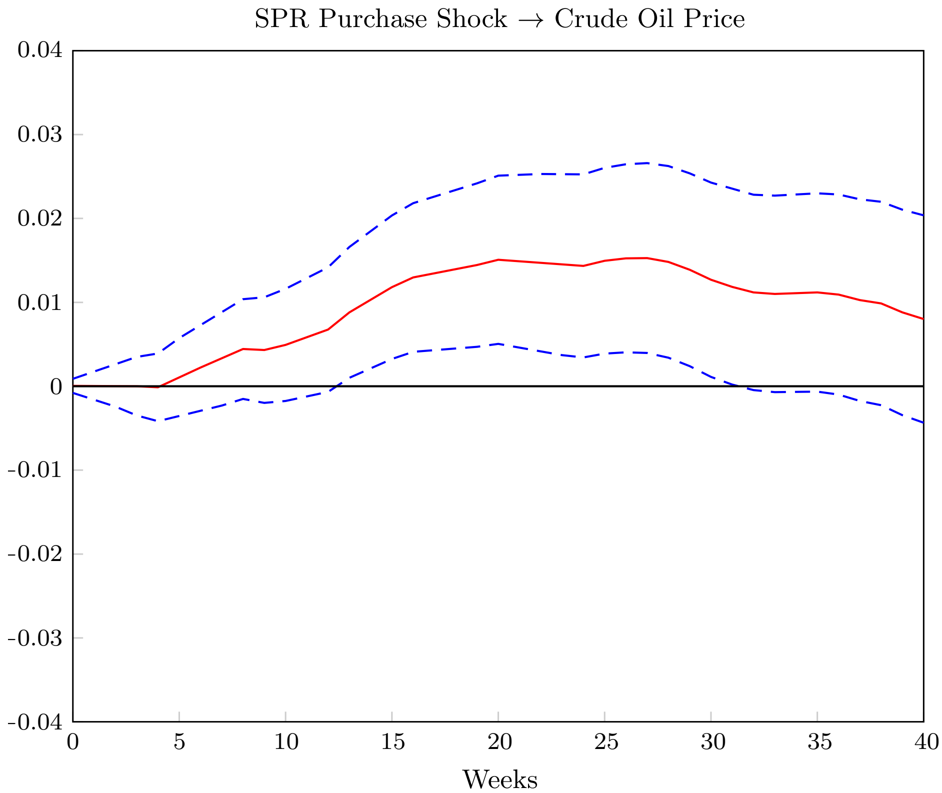

Figure 2: SPR Purchase Impulse Response Function (Benchmark Model)

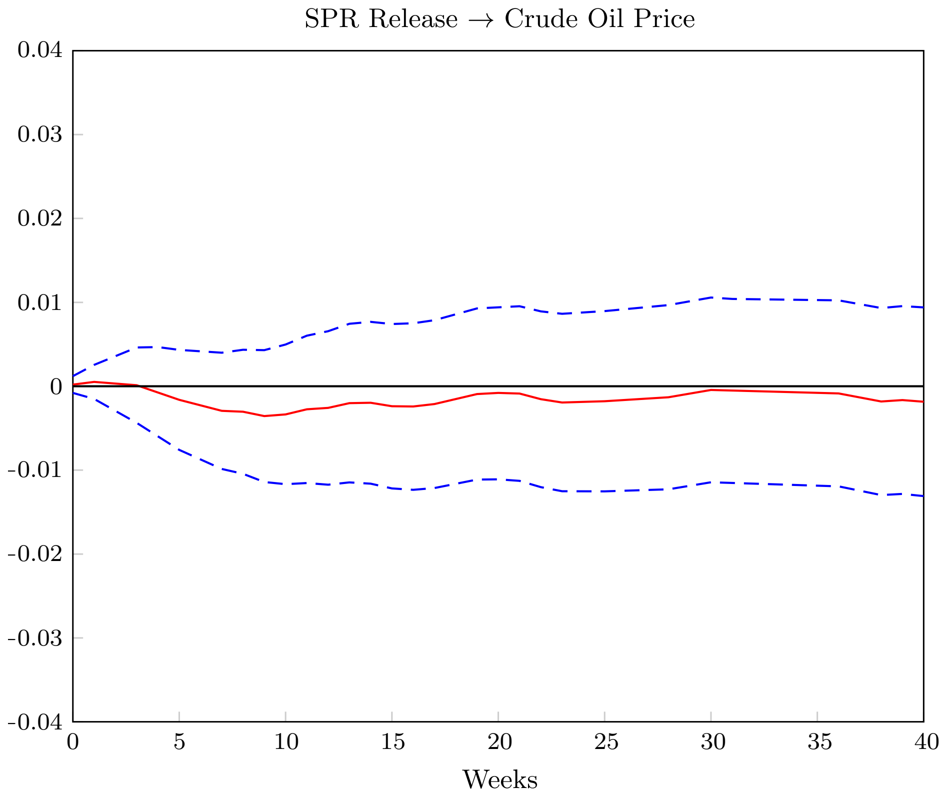

Figure 3: SPR Release Impulse Response Function (Benchmark Model)

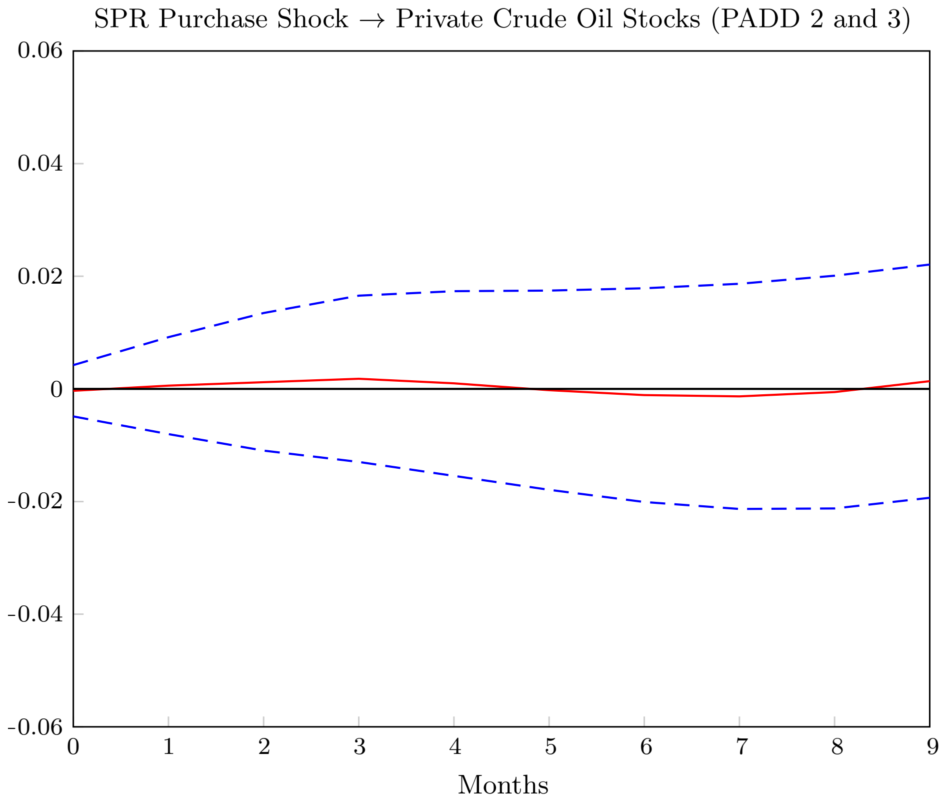

Figure 4: SPR Purchase Impulse Response Function (Private Crude Stocks)

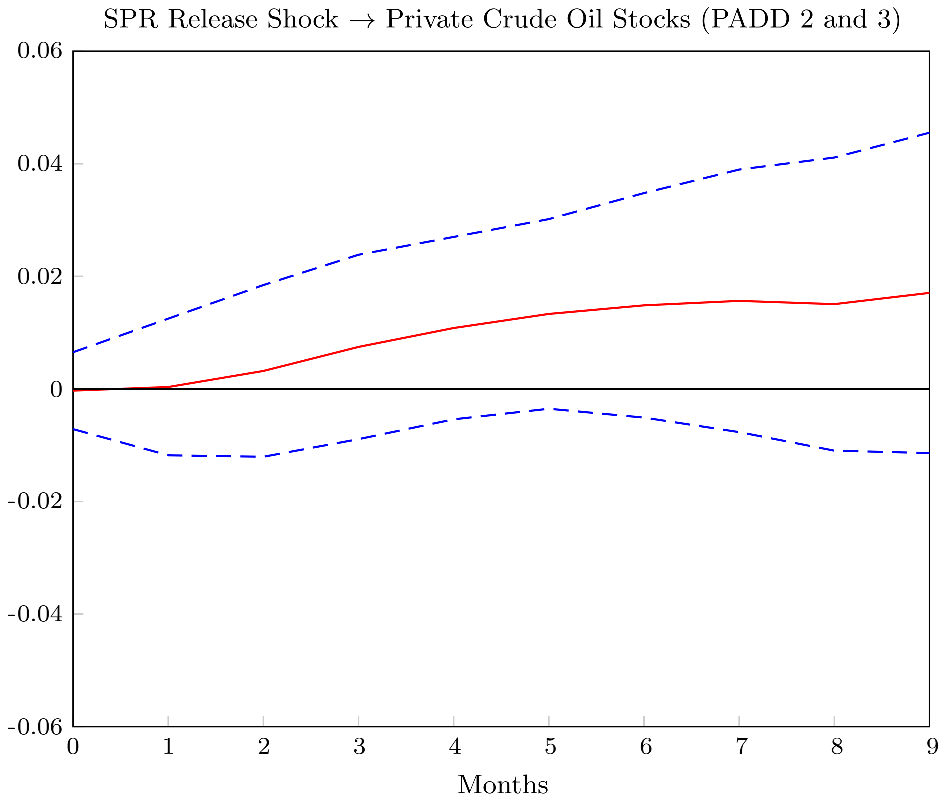

Figure 5: SPR Release Impulse Response Function (Private Crude Stocks)

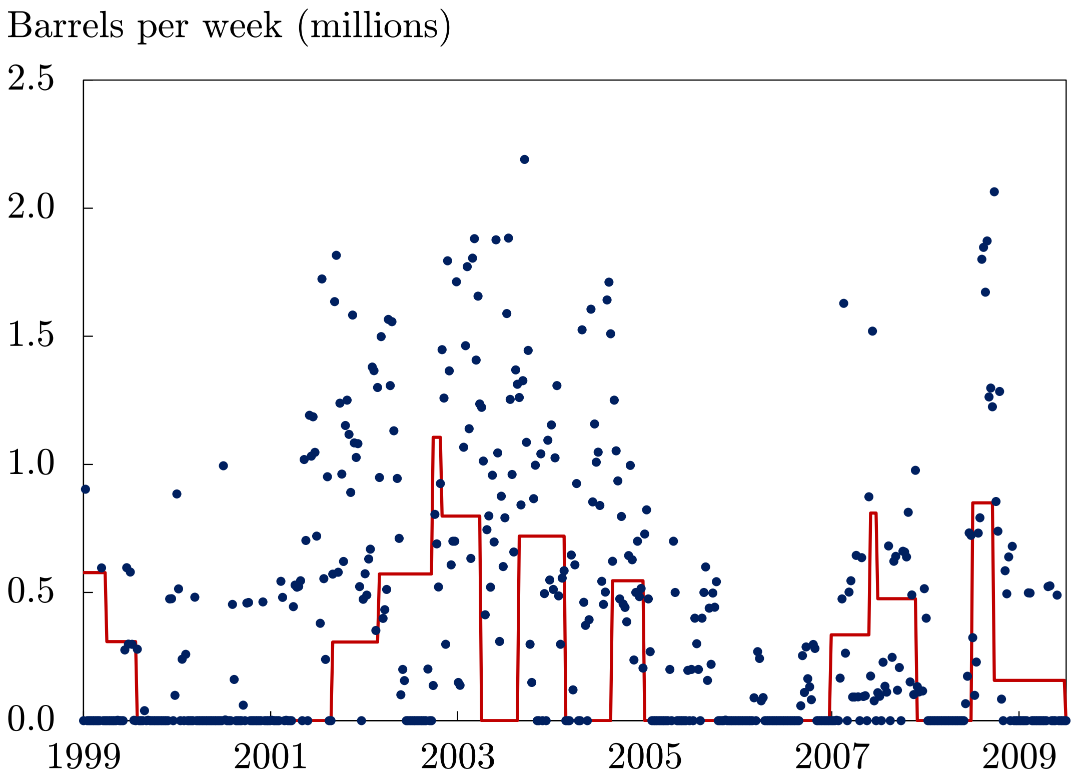

Figure 6: SPR Purchase Schedule (red line) and Actual Purchases (blue dots)

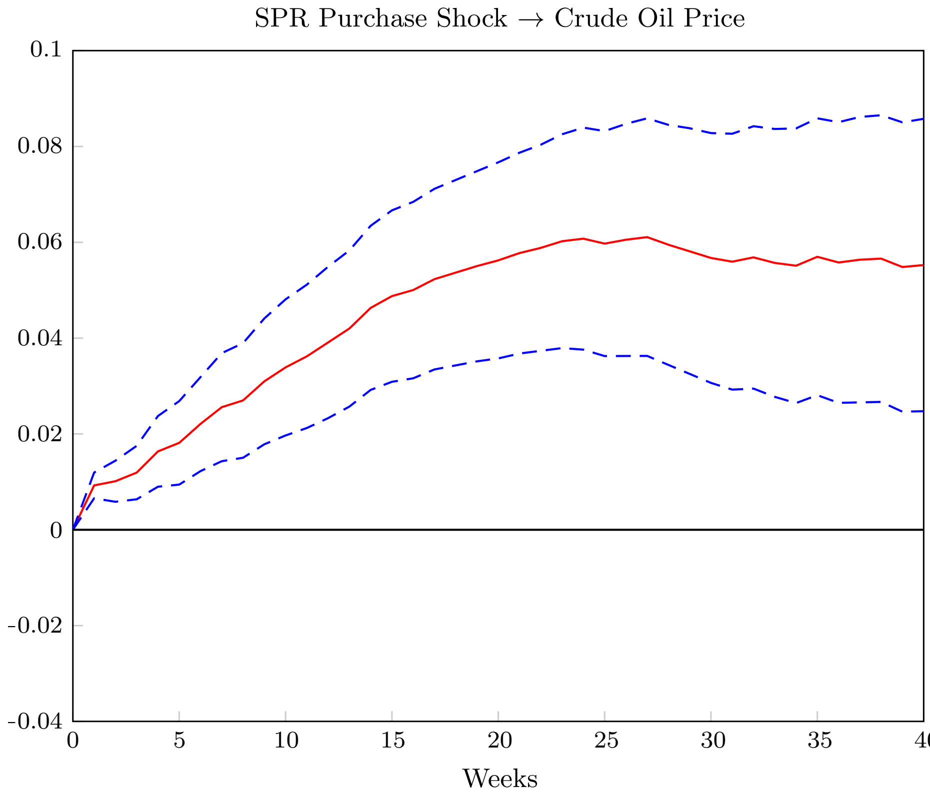

Figure 7: SPR Purchase Impulse Response Function (Identified with External Instrument)

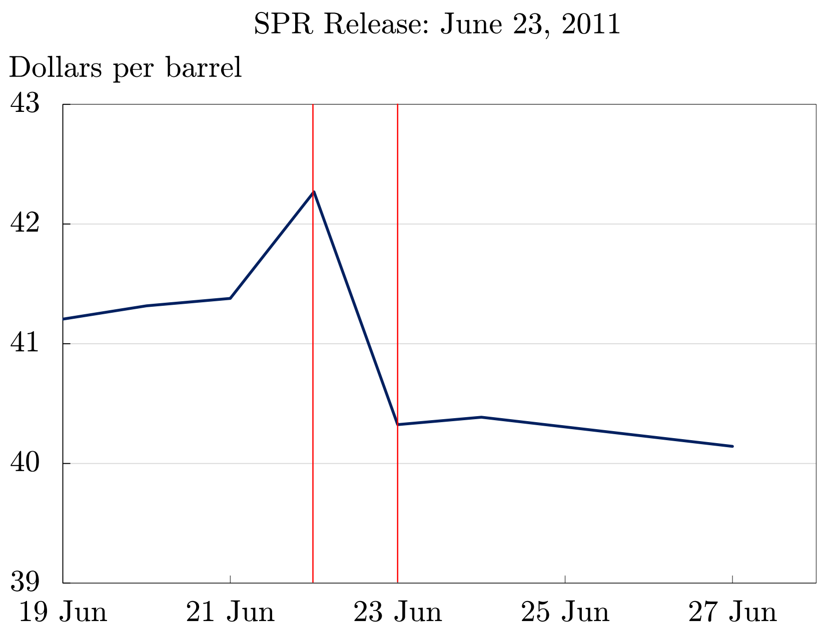

(a) Unanticipated SPR Release

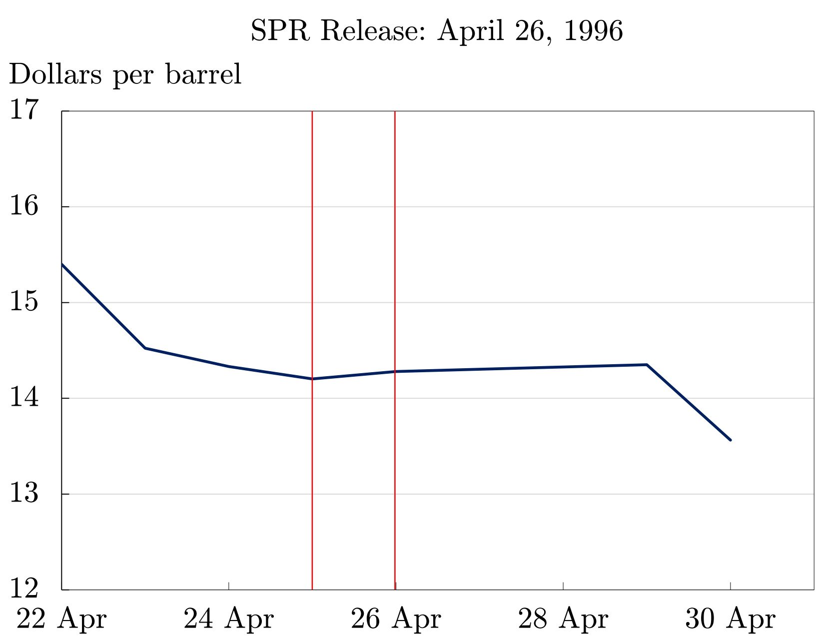

(b) Anticipated SPR Release

Figure 8: Crude Oil Futures Prices Following Release Announcements

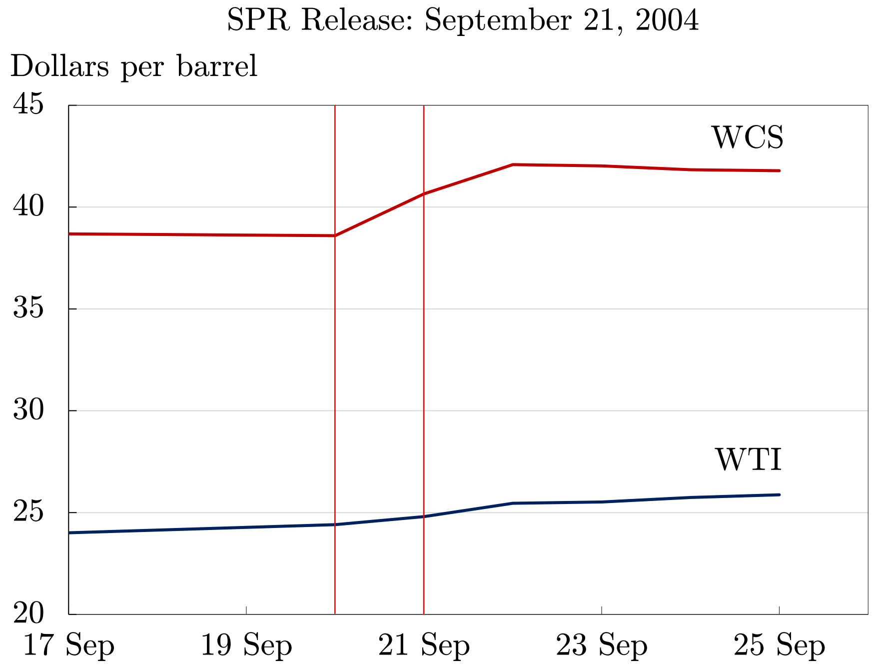

Figure 9: U.S.-Canada Crude Oil Futures Prices Following Release Announcement

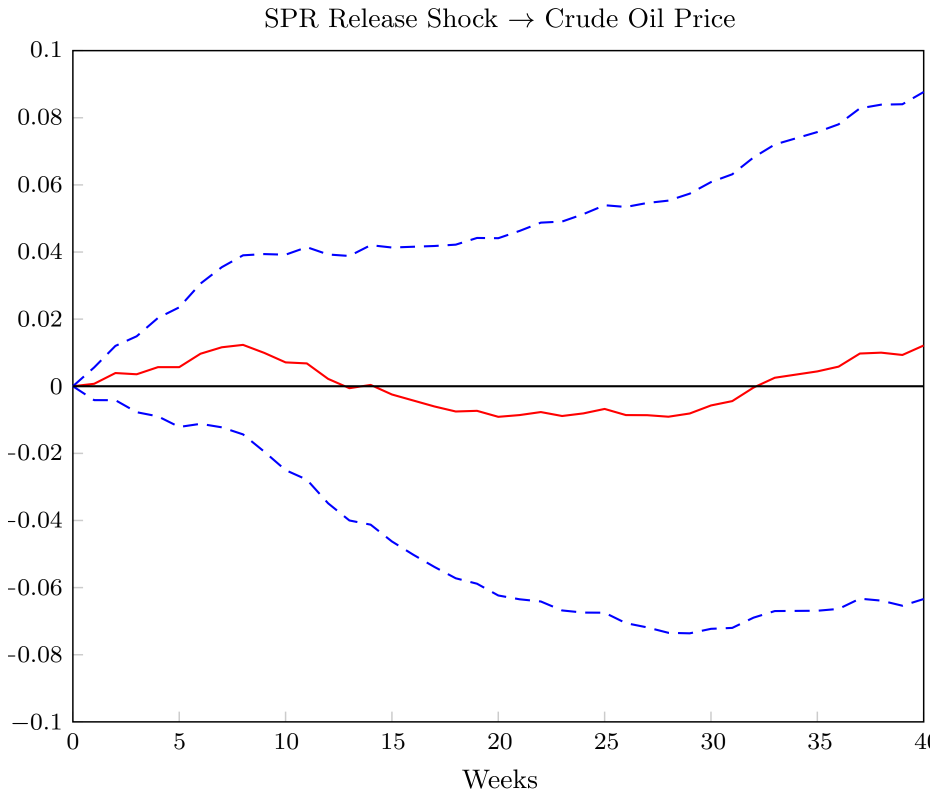

Figure 10: SPR Release Impulse Response Function (identified using WTI futures)

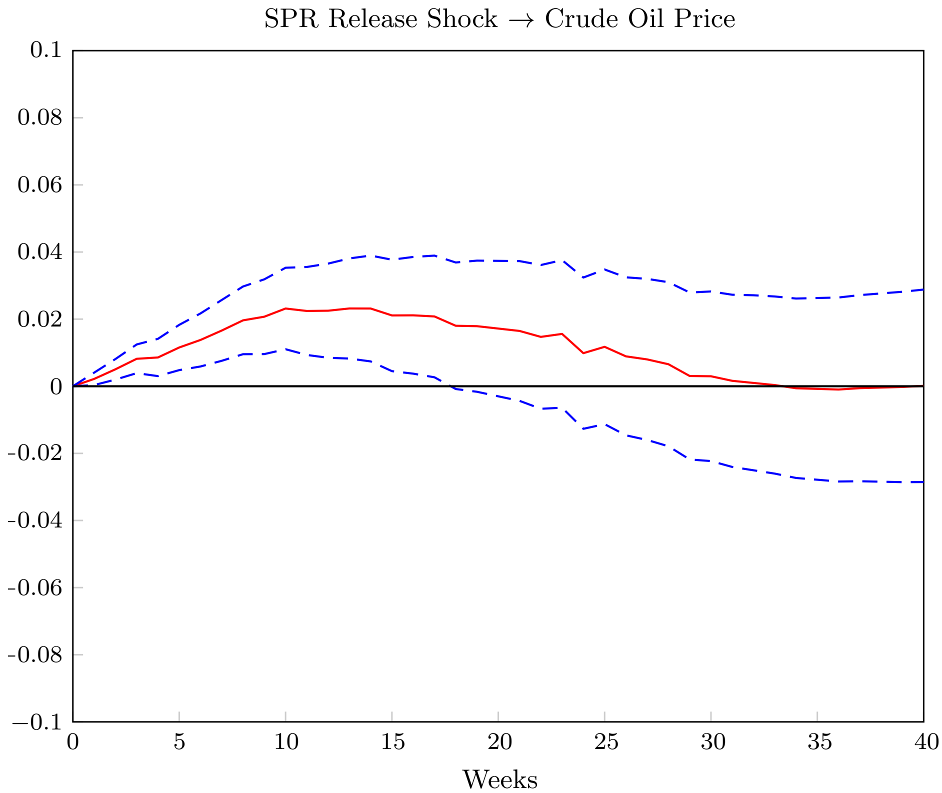

Figure 11: SPR Release Impulse Response Function (identified using the WTI-WCS futures spread)

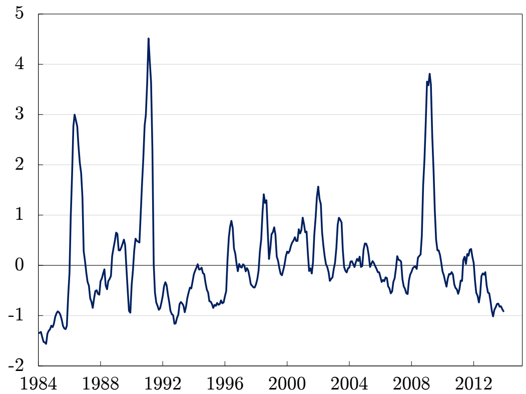

(a) Oil Market Volatility Index

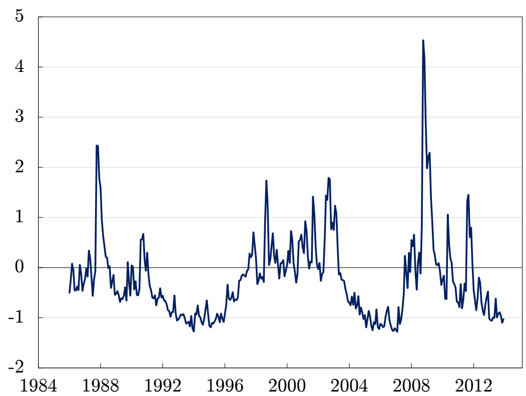

(b) Stock Market Volatility Index

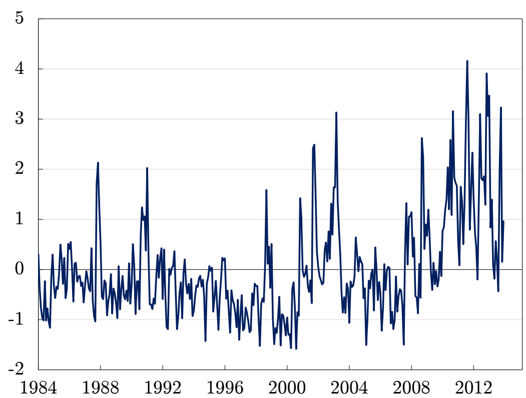

(c) Policy Uncertainty Index

Figure 12: Uncertainty Indices

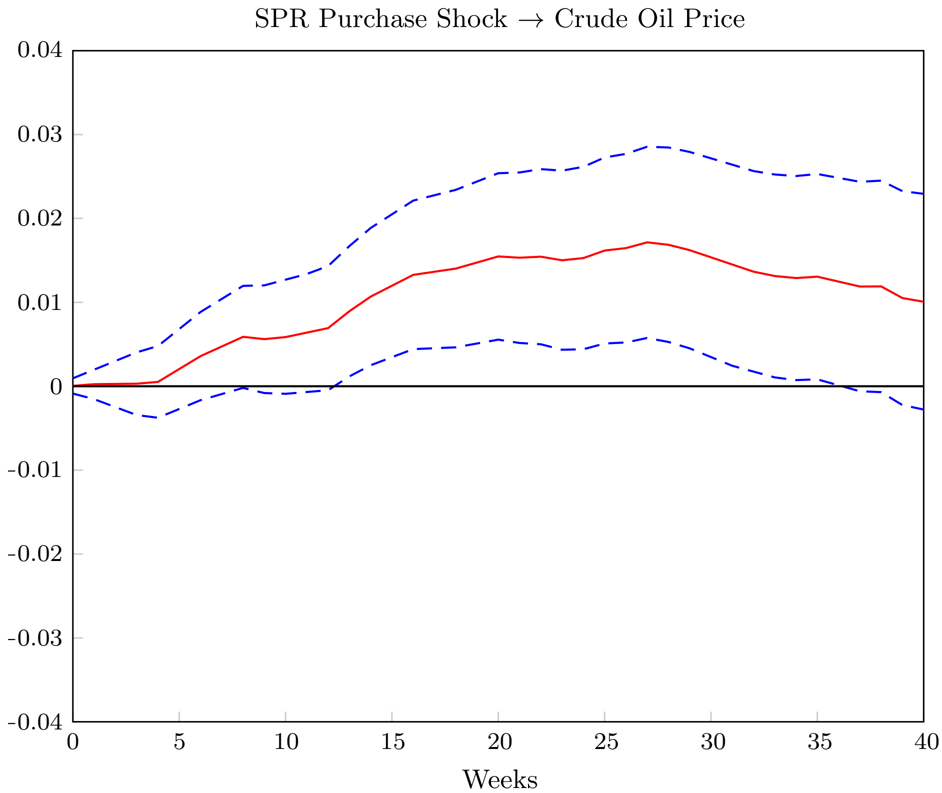

Figure 13: SPR Purchase Impulse Response Function (Mean Oil Market Volatility)

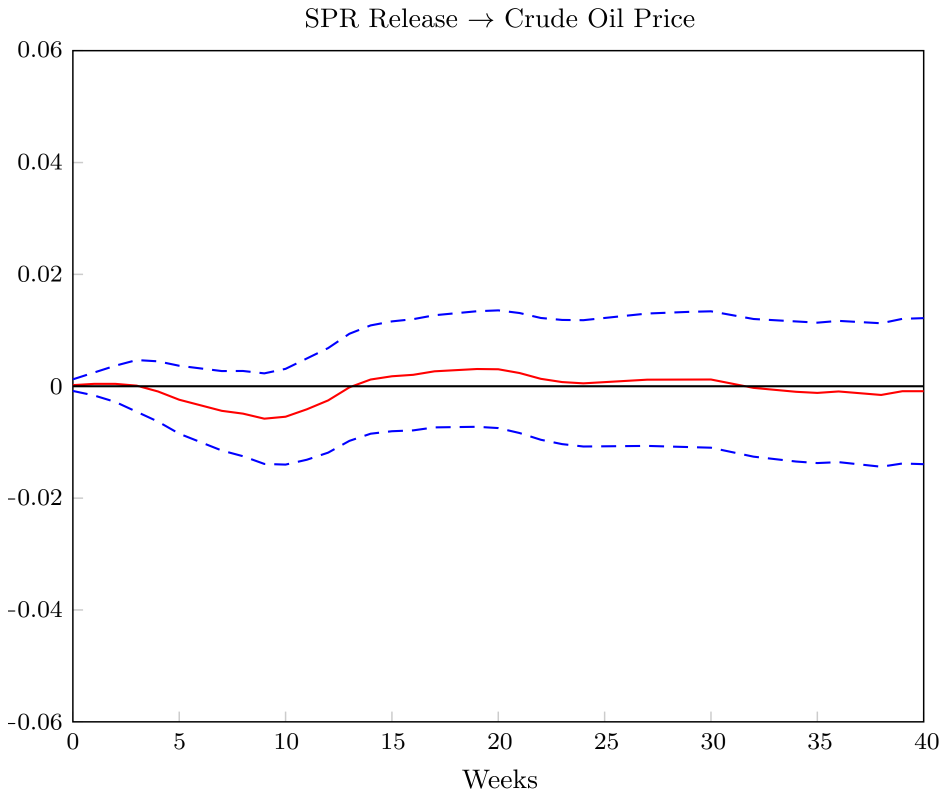

Figure 14: SPR Release Impulse Response Functions (Mean Oil Market Volatility)

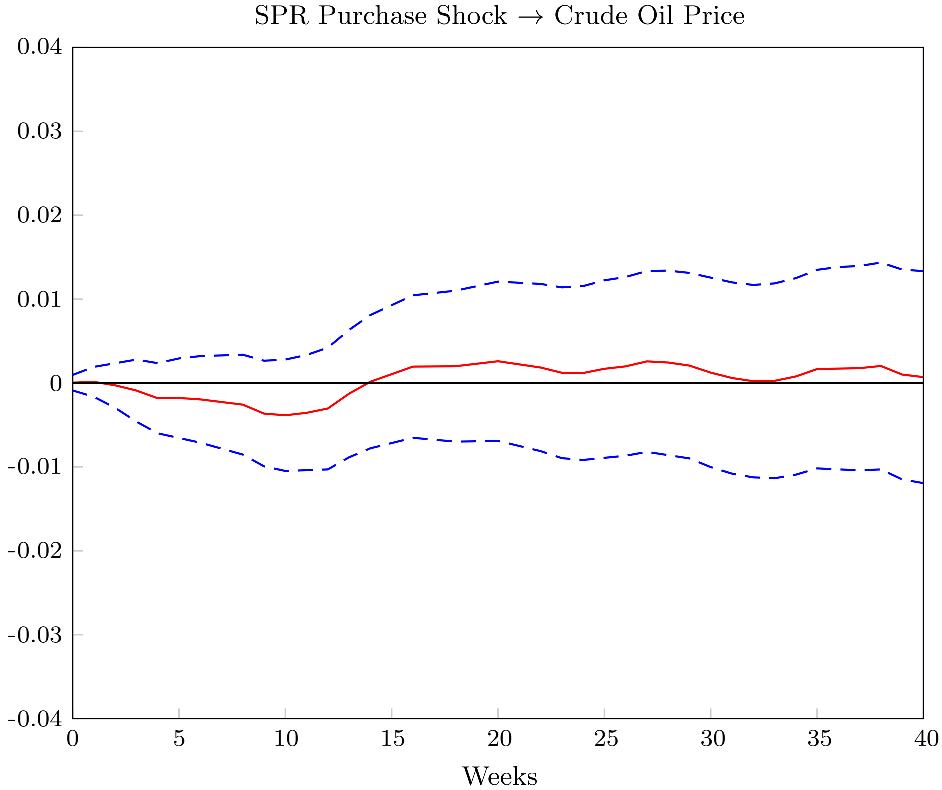

Figure 15: SPR Purchase Impulse Response Function (Low Oil Market Volatility)

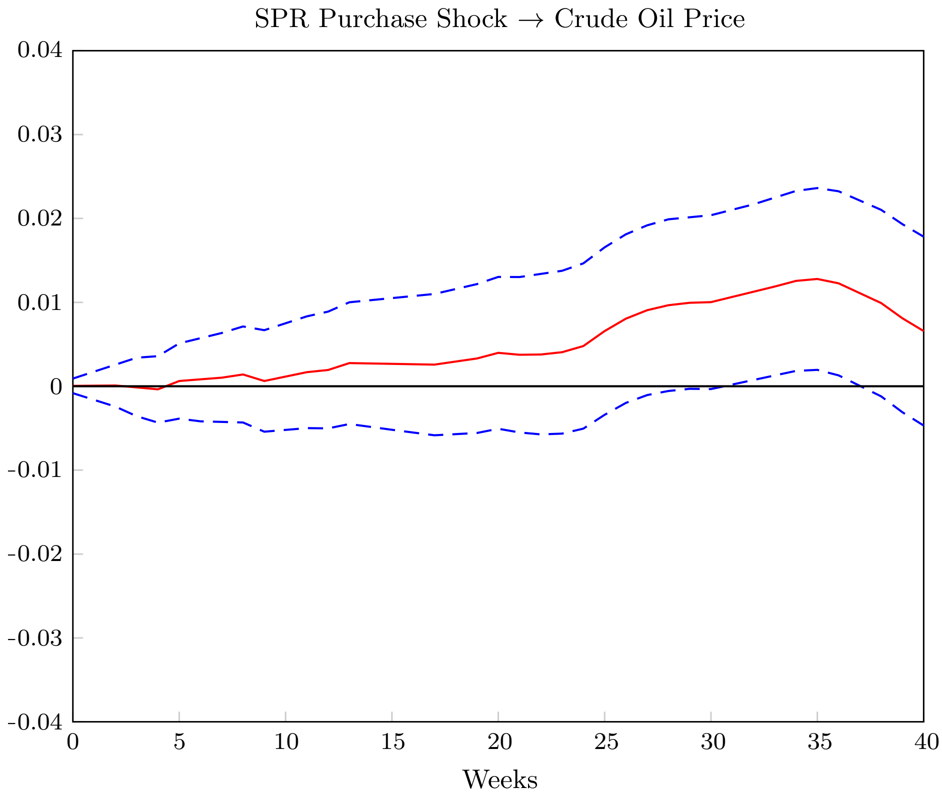

Figure 16: SPR Purchase Impulse Response Function (Low Stock Market Volatility)

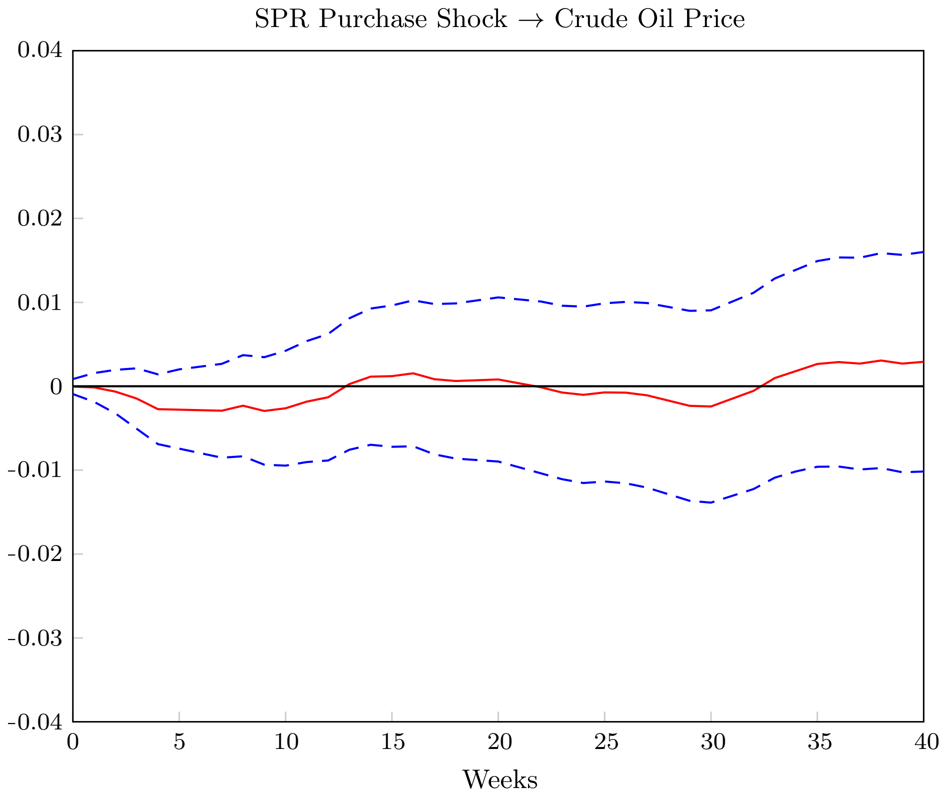

Figure 17: SPR Purchase Impulse Response Function (Low Political Uncertainty)

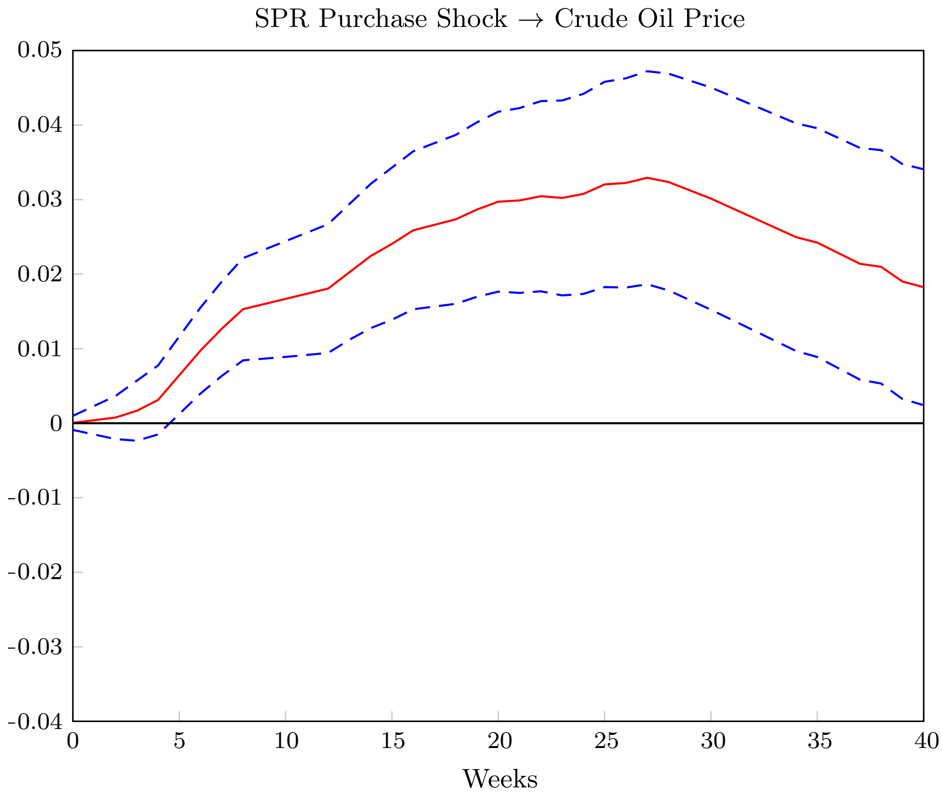

Figure 18: SPR Purchase Impulse Response Function (High Oil Market Volatility)

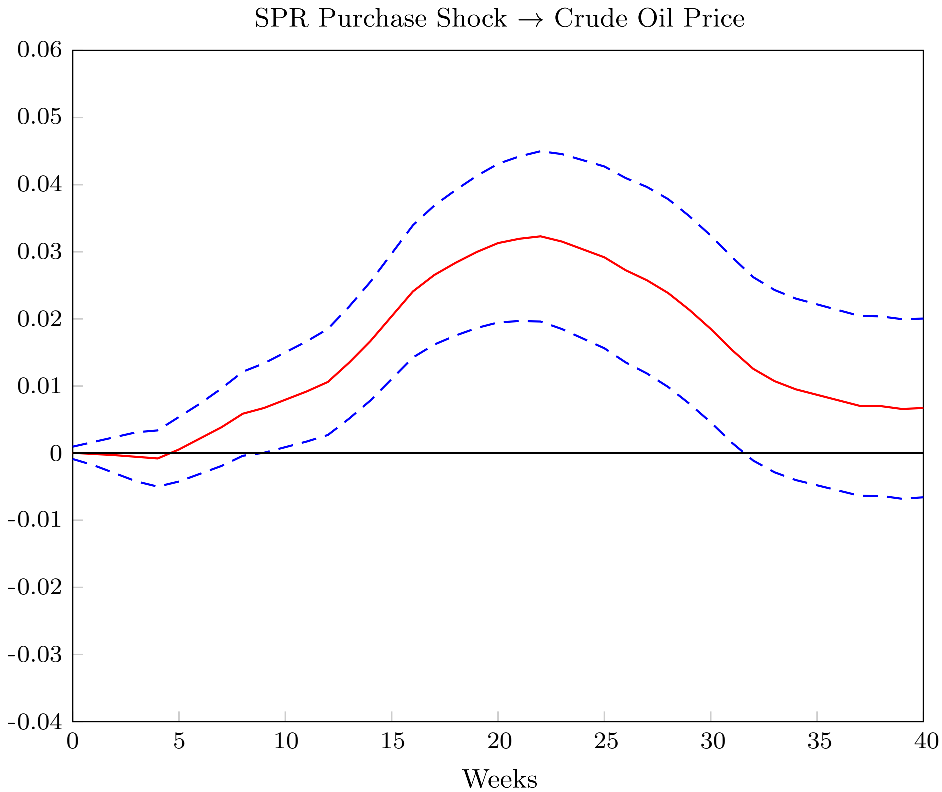

Figure 19: SPR Purchase Impulse Response Function (High Stock Market Volatility)

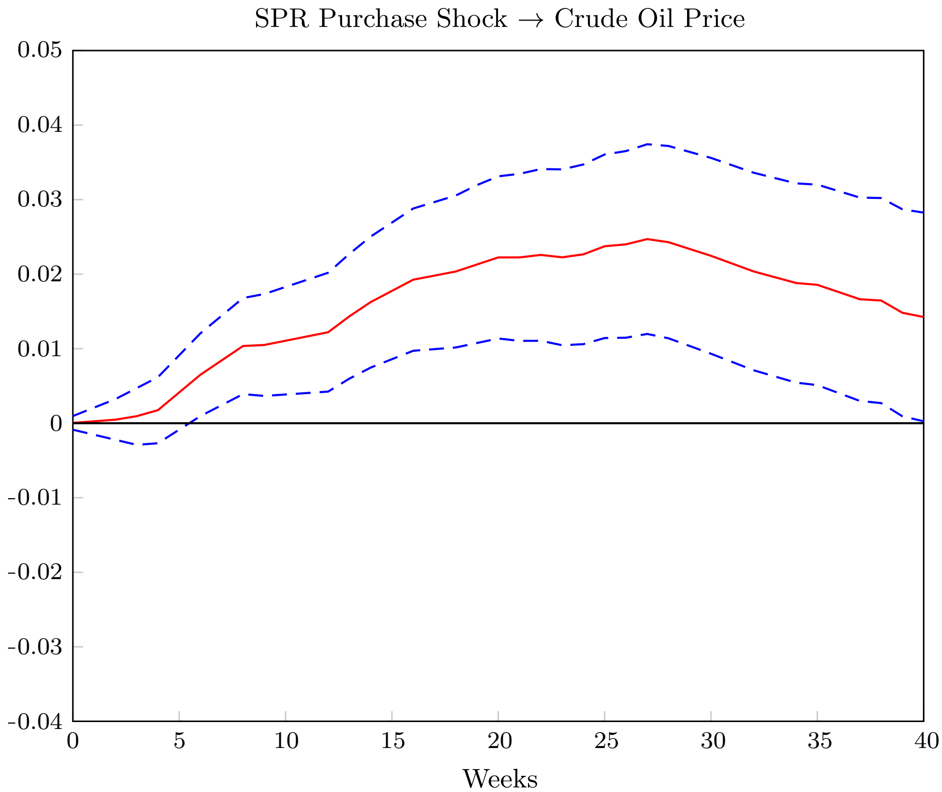

Figure 20: SPR Purchase Impulse Response Function (High Political Uncertainty)

	1983-1989	1990-1994	1995-2001	2002-2008	2009-2014
SPR Purchase	-0.02	-0.07	-0.01	0.06	0.00
SPR Release	0.00	0.02	0.06	-0.05	0.1

Table 1: Average Structural Policy Residuals for Subperiods