

A Decade of Learning: The Role of Beliefs in Oil Futures Markets During the 2000s

Sylvain Leduc

Kevin Moran

Robert Vigfusson*[†]

FRBSF

Université Laval

Federal Reserve Board

May 5, 2015

Abstract

We examine the role of learning in accounting for the movements in oil price futures during the 2000s, a period during which the oil market experienced important changes. We show that a simple unobserved component model in which investors must form beliefs about whether the source of oil price movements is transitory or permanent accounts remarkably well for the fluctuations in oil price futures. Our simple framework notably accounts for the relatively slow increase in futures prices at the beginning of the past decade and their unprecedented run-up between 2004 and 2008. Even during the first half of 2008, a period during which oil prices reached historic highs, the model predicts a level of futures prices that is broadly in line with the data. Our estimates suggest that, through learning, investors revised up the contribution of permanent shocks to the variance of oil prices throughout this period. Using a DSGE model in which oil is storable and used in production, we then show that this learning process may have significantly muted the effects of oil shocks on the economy during that period.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, any other person associated with the Federal Reserve System.

[†]Comments and suggestions can be directed to robert.j.vigfusson@frb.gov.

1 Introduction

The large fluctuations in the price of oil over the past 15 years have renewed interest in the usefulness of futures markets in anticipating these movements. Overall, the futures market failed to predict that events during the past decade, from the growing importance of China in the world economy to onset of the financial crisis, would radically change the outlook for oil prices. These developments led many to call into question the usefulness of oil-price futures as predictors of future oil price movements.¹ Whereas many commentators have attributed the associated forecasts errors to speculation or market inefficiency, this paper provides an explanation of the movements in oil-price futures based on learning and show through a DSGE model that this learning process can significantly alter the quantitative effects of oil shocks on the economy.

In particular, our analysis indicates that the developments in the oil futures markets during the 2000s are consistent with investors learning about whether underlying shocks are transitory or permanent. We first provide evidence using the Kalman filter to infer the permanent and transitory components of shocks to spot oil prices, and show that this simple form of learning can explain remarkably well the observed behavior of futures prices. Our simple framework accounts for the relatively slow increase in oil-price futures at the beginning of the past decade and their unprecedented run-up between 2004 and 2008. Even during the first half of 2008, a period during which oil prices reached historic highs, the model predicts a level of futures prices that is broadly in line with the data. Our estimates suggest a significant and steady increase in the variance of the permanent component of shocks to spot oil prices from 2002 to 2008, accompanied by a similarly steady decline in the variance of the temporary component. In turn, this translates into a growing contribution of permanent shocks to the variance of oil prices over this period.

Although the Kalman filter has the advantage of being a straightforward approach to model learning, it is also somewhat restrictive in that it assumes that the parameters of the model are constant for a given data sample. As a result, we also assess our benchmark

¹See, for instance, Alquist and Kilian (2010) and Alquist, Kilian, and Vigfusson (2012).

results by allowing for a more general, nonlinear, learning process with time-varying parameters. In particular, we consider a variant of the unobserved components model with stochastic volatility of Stock and Watson (2007), with the unobserved components estimated via the particle filter. An additional advantage of using this richer framework is that we can also examine the presence of a time-varying risk premium as an additional source of fluctuations in futures prices (see, among others, the recent analyses of Hamilton and Wu (2012), Singleton (2014), and Baumeister and Kilian (2015)).

We find that our baseline results derived using the Kalman filter are little changed under the unobserved components model with stochastic volatility and continue to indicate a growing importance of permanent shocks to oil prices from the early 2000s to the onset of the financial crisis. As with the Kalman filter, this more general framework fits the movements in the futures price of oil very well. Moreover, we find that variations in the risk premium appear to play little role in accounting for the movements in oil-price futures, a conclusion in line with the findings of Hamilton and Wu (2012).

Finally, we contrast our baseline results to those obtained under a constant-gain learning process under which recent developments are weighted more heavily than earlier data. We show that this specification significantly worsens the fit of the model, predicting much larger upward movements in futures prices during the 2003-2005 period than those actually observed. These results indicate that participants in the oil markets continued to place substantial weight on developments in the 1980s and early 1990s when forming their expectations of future oil prices during the past decade.

Our findings have important implications regarding the impact of oil shocks on the macroeconomy. Using a DSGE model in which oil is storable and used in production, we show that agents' learning of the form suggested by our empirical results can substantially cushion the recessionary effects of oil shocks. Consistent with our empirical results, we calibrate two scenarios that capture market participants' perceptions regarding the persistence of oil prices in 2003, when changes in oil prices were largely thought to be transitory, and 2007, when oil price changes were expected to be much more persistent.

Compared to a framework with full information, we show that the recessionary effects of oil-price increases are roughly halved in the year following the rise in the price of oil under the 2003 scenario. A similar dampening effect occurs under our 2007 scenario, though it tends to be more short lived. We show that part of this dampening effect is due to an interaction between learning and storage. For instance, a misperception that oil shocks are transitory leads agents to draw down inventories, which mutes the rise in oil prices and the associated fall in economic activity.

The more muted effects of oil shocks under learning and storage may partly explain the relatively weak impact that the runup in oil prices in the mid-2000s had on growth. As such, our results complements Kilian's (2008) emphasis on the importance of the source of oil shocks in understanding their effects on GDP. In particular, Kilian finds that the run-up in oil prices starting in 2003 was largely driven by an increase in world aggregate demand and thus had a mitigated effect on U.S. GDP. Our results would add that the effects on growth were also muted because market participants initially failed to correctly assess the persistence of the increase in oil prices. Similarly, the output effect of the increase in oil prices in the Spring of 2008 may have been heightened because market participants anticipated the rise to be nearly permanent.²

1.1 Related Literature

A growing literature interested in understanding the large movements in the spot price of oil and its changing relationship with the futures market has recently emerged. In particular, our paper relates to some of the recent work on the role of financial markets in driving oil prices.³ For instance, Hamilton and Wu (2012) argue that increased participation by index-fund investing has reduced oil futures premia since 2005, accounting for the smaller gap between spot and futures prices observed in the data between 2005 and 2008. Similarly, Buyuksahin et al. (2008) argue that increased market activity

²Hamilton (2009) analyses the contribution of the oil shock of 2007-2008 to the Great Recession.

³This literature is large and growing. Some of the many papers discussing this issue include Singleton (2012), Irwin and Scott (2012) and Fattouh, Kilian and Mahadeva (2012), Kilian and Murphy, Alquist and Gervais (2012).

by commodity swap dealers, hedge funds, and other financial traders, has helped link crude oil futures prices at different maturities. Acharya et al. (2013) emphasize limits to arbitrage and their effects on spot and futures prices in commodity markets. In their environment, speculators face capital constraints in commodity markets, which limits commodity producers' ability to hedge risk and is reflected in commodity prices. As such, these papers attribute developments in the futures markets to the increased financialization of commodity markets or to speculators' risk appetite, while we show that part of these movements can be attributed to learning.

Our paper also complements the work of Alquist and Kilian (2010) on the role of a convenience yield associated with oil inventories in accounting for the large and persistent fluctuations in the oil futures spread. Using a theoretical model, they argue that, under greater uncertainty about future oil supplies, the presence of a convenient yield may explain the excess variability of oil futures prices relative to that of the spot price, which underlies the poor predictive performance of oil-price futures. We add to this literature by highlighting the interaction between learning and oil storage.

Our approach also shares with Milani (2009) the emphasis on learning. In a DSGE model in which oil is used in production, Milani (2009) studies the effect of learning on the changing relationship between oil prices and the macroeconomy since the 1970s, assuming that agents in the model learn about the parameters of the model and the underlying shocks through constant-gain learning. He shows that learning is important to account for the changing effects of oil shocks on output and inflation over time. In comparison, we show that a simple model of learning about the persistence of underlying shocks to spot prices accounts for the evolution of oil futures prices since the late 1990s strikingly well, and that this type of learning can also result in more muted effects on the macroeconomy. Lastly, our work relates to Singleton (2012) who emphasizes informational frictions and learning about economic fundamentals as important factors behind the boom and bust cycle of oil prices in the previous decade.

The rest of the paper is organized as follows. After describing in more detail develop-

ments in oil prices, we layout our empirical framework, estimating a time series model of permanent and temporary shocks to oil prices. Given our assumptions regarding learning and using only spot price data, we estimate that the predicted permanent component of spot prices matches very well the observed developments in futures markets. We then develop a theoretical framework to examine the impact of learning about the persistence of oil market developments on the economy, calibrating the learning process to match our empirical findings in 2003 and 2007. After describing our theoretical findings, we conclude in the last section.

2 Oil Prices During the 1990s and 2000s

We start our analysis by presenting some evidence of the oil market's evolving views on the persistence of the underlying shocks hitting the world economy. To do so, consider the movements in the spot and futures prices of oil since the early 1990s depicted in Figures 1 and 2. In each figure, the solid line shows the evolution of spot prices, while the dotted lines depict the futures prices path at a given point in time. During the 1990s (Figure 1), the spot price of oil tended to gyrate around fairly steady oil-price futures, suggesting that market participants viewed economic developments impacting the oil market as mainly temporary. Underlying shocks would tend to move the spot price of oil, at times substantially, but the spot price would rapidly return to roughly \$18 per barrel, about the average price during that period. Clearly, whatever the disturbances affecting the world economy, market participants did not view them as persistent enough to substantially alter their view of oil prices in the future.

However, Figure 2 suggests that the relationship between spot and futures price changed during the early 2000s. Between 2000 and 2008, the spot price of oil rose steadily, from roughly \$27 per barrel to more than \$135 per barrel. In contrast, oil-price futures remained initially low, consistently fluctuating below \$20 per barrel until 2003. Then, oil-price futures started a gradual rise, bringing them to roughly \$50 per barrel by the mid-2000s. Spot and futures prices then tended to move in lockstep between 2005

and 2008.

One possible interpretation of this pattern is that, between 2000-2003, market participants thought that movements in spot prices would likely be temporary, as indeed was the case throughout the 1990s. However, after being consistently surprised by the persistence of the rise in spot prices, market participants reassessed their views, placing more weight on the possibility that the increase in oil prices was persistent rather than transitory.

According to this interpretation, by the time the oil-market reached its peak in the spring of 2008, market participants largely anticipated the movements in spot prices to be highly persistent, remaining about \$135 per barrel over the next 5 years, as indicated in Figure 2. In this paper, we argue that changes in the perceived persistence of oil prices between 2000 and 2008 may help explain the more muted impact of rising oil prices on economic activity in the early part of that decade compared to the effects in 2008, which have been singled out as an important contributing factor to the Great Recession (Hamilton (2009)).

Figure 2 also suggests that the collapse in economic activity in the United States and many other countries during the fall of 2008 led to a reassessment of the long-run equilibrium price of oil. In addition, as in the 1990s, market participants now appear to view most of the fluctuations in the spot price of oil as being largely transitory. That is, the long-run futures price between the end of the Great Recession through 2012 remained fairly stable despite significant movements in the spot price.⁴

3 Empirical framework

In this section, we develop a simple unobserved components model to account for the role of permanent and temporary shocks in determining oil-price futures. We postulate that spot oil prices are the result of movements in a permanent and a transitory components

⁴We stop our analysis in 2012 because, by then, the large increases in North American oil production had resulted in the price of WTI no longer being indicative of global demand.

and that market participants use the Kalman filter to assess the relative importance of these two components over time. In addition, under our baseline model, we allow the model parameters to evolve with the data sample. Below, we also assess the robustness of our baseline model by using of a more general framework that explicitly allows for time-varying parameters.

3.1 A simple model

Consider the following linear process relating the spot price of oil s_t (expressed in logs) to a permanent component, e_t^P , and a stationary one, e_t^τ ,

$$s_t = e_t^P + e_t^\tau. \quad (1)$$

The permanent component is modeled as a random walk with drift

$$e_t^P = \mu + e_{t-1}^P + v_t, \quad (2)$$

where v_t is an independently and identically normally distributed disturbance with mean zero and constant variance σ_p^2 . The temporary component is assumed to follow an AR(1) process

$$e_t^\tau = \phi_\tau e_{t-1}^\tau + \varepsilon_t \quad (3)$$

where ε_t is an homoskedastic, Gaussian error term: $\varepsilon_t \sim N(0, \sigma_\tau^2)$ and with $|\phi_\tau| < 1$.

Assuming full information at time t about the temporary and permanent components underlying oil prices, the k -period ahead futures price at time t , $f_{t,k}$, is given by the following expression

$$f_{t,k} = E_t s_{t+k} = k\mu + e_t^P + \phi_\tau^k e_t^\tau. \quad (4)$$

In contrast, absent full information on the current levels of e_t^P and e_t^τ , the futures price will be based on the best forecasts given past values of s_t

$$f_{t,k} = E_t (s_{t+k} | \{s_{t-i}\}_{i=t}^1) = E_t (k\mu + e_t^P + \phi_\tau^k e_t^\tau | \{s_{t-i}\}_{i=t}^1). \quad (5)$$

In the next section, we propose that market participants use the Kalman filter to assess the importance of permanent and transitory shocks when forming expectations about future oil prices. To determine the relative importance of permanent shocks, a simple statistic can be derived from the expression for the change in the spot price of oil

$$\Delta s_t = \mu + v_t + (\phi_\tau - 1) e_{t-1}^\tau + \varepsilon_t,$$

which implies that the variance of Δs_t can be expressed as

$$\text{var}(\Delta s_t) = \sigma_p^2 + \sigma_\tau^2 + \frac{(1 - \phi_\tau)^2}{1 - \phi_\tau} \sigma_\tau^2.$$

Therefore, the fraction of the growth rate's variance that is due to permanent shocks can be calculated as

$$\frac{\sigma_p^2}{\sigma_p^2 + \sigma_\tau^2 + \frac{(1 - \phi_\tau)^2}{1 - \phi_\tau} \sigma_\tau^2}.$$

3.2 Learning

We assume that market participants use the Kalman filter to form expectations of future oil prices. In particular, define ξ_t as the unobserved state vector of the model above, comprising the trend, μ , as well as the permanent and temporary components: $\xi_t = \left(\mu \quad e_t^P \quad e_t^\tau \right)'$. Given values of the model's parameters, $\Gamma = [\sigma_p^2, \sigma_\tau^2, \mu, \phi_\tau]$, the Kalman filtering equation relates how the observed variables y_t (such as spot oil prices) respond to the changes in the unobserved state vector ξ_t .

The equations for the dynamics of the observed variables y_t are given by the following system

$$\begin{aligned} y_t &= H\xi_t \\ \xi_t &= F\xi_{t-1} + \varepsilon_t \end{aligned} \tag{6}$$

where F and H are vectors of parameters. In turn, the unobserved state vector evolves according to the following standard equation

$$\xi_{t|t} = \xi_{t|t-1} + P_{t|t-1}H(H'P_{t|t-1}H)^{-1}(y_t - H'\xi_{t|t-1}) \quad (7)$$

given initial estimates of $\xi_{t|t-1}$ and $P_{t|t-1}$, where the forecast error $(y_t - H'\xi_{t|t-1})$ is used to update the estimates of the size of the permanent and transitory components via the term $P_{t|t-1}H(H'P_{t|t-1}H)^{-1}(y_t - H'\xi_{t|t-1})$. Thus, the value of $P_{t|t-1}$ governs whether a given surprise is assumed to be part of the permanent or the transitory component, which is influenced by the values of σ_p^2 and σ_τ^2 .

Whether movements in the spot price of oil are perceived to be permanent or temporary will have an important impact of futures prices as well. If a one percent increase in the spot price of oil is interpreted as purely transitory, then the futures price k period ahead, $E_t(s_{t+k} | \{s_{t-i}\}_{i=t}^1)$, only increases by ϕ_τ^k , as indicated by equation (5) above. In contrast, if the same increase is interpreted as purely permanent, then the value of $E_t(s_{t+k} | \{s_{t-i}\}_{i=t}^1)$ increases by value of the permanent component, e_t^P . If a shock is actually permanent but mistaken to be temporary, then the value of the futures price will include this error.

So far our discussion assumed that the parameter values were known with certainty, but in practice we will need to estimate the model parameters $\Gamma = [\sigma_p^2, \sigma_\tau^2, \mu, \phi_\tau]$, to derive forecasts of future oil prices. To do so, we follow two approaches. First, we assume that market participants use all available information and estimate the model's parameters using the standard likelihood function

$$LL_T = - \sum_{t=1}^T \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_t\| + (s_t - Es_t) V_t^{-1} (s_t - Es_t) \right) \quad (8)$$

where $V_t = H'P_{t|t-1}H$ is the variance of the prediction errors. Alternatively, we assume that market participants put more weight on recent observations than on more distant ones, possibly because of concerns about structural breaks and time-variations in the model parameters (in the learning literature, this approach is referred to as constant gain learning). In the spirit of the recursive least squares algorithm in Cho, Williams

and Sargent (2002), we therefore modify the likelihood function as follows

$$LL_T = (1 - \chi_T) LL_{T-1} - \chi_T \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|F_t\| + (s_t - Es_t) V_t^{-1} (s_t - Es_t) \right) \quad (9)$$

If $\chi_t = \frac{1}{t}$ then all observations have the same weight, equivalent to the standard likelihood function described above. In contrast, if χ_t is a constant, then recent observations are more important than lagged observations. In particular, for a dataset of T observations, the first observation contributes $\prod_{t=1}^T (1 - \chi_t) \chi_1$, whereas the most recent observation (observed at time T) has a much greater weight of χ_T .

4 Baseline Results

In this section, we present our model's predictions for futures prices assuming that market participants form expectations about the permanent and temporary components of oil prices through the Kalman filter. We estimate our model assuming that market participants form their beliefs using univariate methods, i.e., using only data on the spot price of oil. Although extracting information about the components of oil prices only looking at previous spot prices may be suboptimal, our emphasis on univariate methods has the benefit of simplicity and share similarities with the learning algorithm used in the monetary policy literature (see, for instance, Orphanides and Van Norden (2001) and Primiceri (2006)). Using the average monthly price of West Texas Intermediate oil at the end of each quarter from 1980Q1 to 2009Q4, we construct the model-implied estimates of the one-year and two-year ahead futures prices and compare them to the actual futures price data.⁵

We are particularly interested in the behavior of futures prices from the late 1990s until 2010. As such, we first estimate the model from 1980Q1 to 1998Q4 and start calculating futures prices from this period hence, using an expanding window of data. Thus, the futures prices at the beginning of 2000Q1 are calculated using the model estimated over

⁵As described earlier, we stop the estimation sample in 2012 given that WTI is no longer the global benchmark. We begin the sample in 1980 when U.S. oil production was deregulated. .

the period from 1980Q1 to 1999Q4. Similarly, the sample 1980Q1-2004Q3 would be used to estimate the model and compute forecast of futures prices in 2004Q4.

To compute the k -period ahead futures prices we apply the following three-step procedure. First, we use spot oil prices observed up to time $t - 1$ and estimate the model parameters $\Gamma = [\sigma_p^2, \sigma_\tau^2, \mu, \phi_\tau^k]$ using the standard (or modified) likelihood function. Second, we apply the Kalman filter using the estimated model parameters and observed prices through time t to get estimates of the unobserved permanent and temporary components e_t^P and e_t^τ . In the third step, we use the estimated e_t^P and e_t^τ and Γ , to construct $f_{t,k}$.

We first present the behavior of the model's parameter estimates over the expanding estimation window in Figure 3. The solid black lines report the estimated coefficients, while the grey intervals are two-standard-deviation confidence intervals. The results are broadly in line with the narratives of Figures 1 and 2. The left panel of Figure 3 reports the estimated value of μ . Using just the pre-2000 part of the sample, the point estimate of μ is slightly below zero, implying a negative trend for the nominal oil price. However, as the estimation sample includes more of the twenty-first century, the estimated trend first turns positive and then begins to increase, though the uncertainty around the estimated value is large. The maximum value of the trend occurs for the sample ending in the second quarter of 2008, before starting to decline and stabilizing around 1.3 percent per quarter or at an annual rate of just over 5 percent. In contrast to μ , the value of ϕ_τ , the autoregressive coefficient of the temporary component, is more stable, varying only slightly around a value of 0.7.

Figure 4 reports the estimation results for the variance of the permanent and temporary components, σ_p and σ_τ , respectively. These estimated coefficients do vary considerably as the estimation sample period expands and are also more precisely estimated. In particular, the estimated value of σ_p is notably zero for the initial sample ending in 1998Q4, implying that market participants perceived oil prices to be solely driven by temporary factors. As more data from the 2000s are included in the estimation sample,

the perceived contribution of the permanent component steadily increases, peaking in the second quarter of 2008, before the financial meltdown and ensuing global recession. In contrast, the estimated value of σ_τ broadly follows the opposite pattern.

The evolution of these model parameters can also be viewed by considering the role of permanent shocks in affecting the variance of Δs_t , which is reported in Figure 5. The figure shows that the estimated contribution of permanent shocks only slowly increases over time. In the early part of the sample, because the standard deviation of innovations to the permanent component is extremely small, the permanent component's contribution to the variance of Δs_t is negligible, so that the temporary component is the main driver of changes in the spot price. These results are very much in line with Figures 1 and 2, where the futures curves show transitory increases that ultimately return to a long-term trend during the 1990s and early 2000s. However, the estimated contribution of permanent shocks rises roughly steadily between 2002 and the first half of 2008, where it accounted for more than 60 percent of the variance of Δs_t . Thereafter, the sharp fall in oil prices in the last half of 2008 resulted in lower estimates of the role of permanent shocks.

Given the estimates of the model's parameters, we now construct forecasts of the one- and two-year ahead futures prices, using the Kalman filtering formula (7) for each quarter from 1999Q1 onward. Figure 6 illustrates the evolution of the estimated and actual 2-year-ahead futures prices, as well as the spot price of oil.⁶ Comparing the actual futures prices with the estimated value of the actual futures price, the figure shows that our simple framework model does remarkably well.⁷ As the actual futures prices,

⁶For the actual futures prices, we use the two-year ahead futures price observed on the business day closest to the beginning of the quarter, so that our futures price forecast is consistent with the information available to market participants when the actual futures price was determined.

⁷The grey bands indicate the confidence interval which is defined as the following set

$$\left\{ E_t \left(s_{t+8} \mid \{s_{t-i}\}_{i=t}^1, \tilde{\Gamma}_{t-1} \right) \mid \left(\Gamma_{t-1} - \tilde{\Gamma}_{t-1} \right)' W \left(\Gamma_{t-1} \right)^{-1} \left(\Gamma_{t-1} - \tilde{\Gamma}_{t-1} \right) \leq 4.5 \right\}$$

where Γ_{t-1} is the maximum likelihood estimate estimate and $W \left(\Gamma_{t-1} \right)$ is the corresponding estimated variance covariance matrix. The critical value of 4.5 is chosen as the 66 percentile of the chi-squared distribution with 4 degrees of freedom.

our estimated futures prices are particularly well below the observed spot prices in the early 2000s, suggesting again that market participants viewed the underlying factors pushing spot prices up to be mostly transitory. By the mid-2000s, our estimated futures price move closely together with the spot price. In line with the rising estimate of the contribution of the permanent component to the variance in oil price changes in Figure 5, changes in the spot price of oil by the mid 2000s are perceived as being mostly permanent and are thus being reflected rapidly in futures prices. As the financial crisis intensified in mid 2008, the spot price of oil fell rapidly, but this decline was much more pronounced than the fall in the actual futures price, which is well captured by our estimated value. From the perspective of our model, these movements partly reflect the fact that investors perceived the effects of the crisis on oil prices to be somewhat temporary.

Overall, our results indicate that learning about the persistence of shocks to the spot price of oil influenced market participants in forming expectations about future oil prices. Although our model is simple and only use information from past movements in the spot price of oil, it accounts remarkably well for the fluctuations in oil-price futures over the past decade or so. In particular, it captures the more muted movements in futures prices in the early 2000s and during the Great Recession compared to the fluctuations in the spot price of oil during those periods. Our results suggest that informational frictions and learning of the sort emphasized here may play an important role behind the perceived shortcomings of the oil futures market in generating accurate price forecast.

4.1 Constant Gain Learning

Our baseline results highlight the importance of time-variations in the model’s parameter estimates. This finding suggests that investors may be concerned with structural break and may choose to weight recent observations relatively more than distant ones. As a result, we now consider the sensitivity of our estimated futures prices assuming that market participants’ use the modified likelihood function (9). In conducting this exercise, we consider different constant values of χ_T in the weighted maximum likelihood estimation.

As a starting point, we use information from the literature on learning and monetary policy to parameterize the gain. We first set χ_T to 2 percent based on the value reported by Orphanides and Williams (2005) who estimate the (constant) gain that best fits the inflation forecasts from the Survey of Professional Forecasters. This value implies that an observation 8 years in the past gets only half as much weight in the likelihood as the current observation. Figure 7 compares our baseline results to the estimate of the two-year ahead futures price when $\chi_T = 0.02$.⁸ The figure shows that discounting past observations at this rate generally leads to worse forecasts. Relative to our baseline model, it significantly over predicts oil-price futures from 2004 onward, predicting a peak of \$180 per barrel in the second quarter of 2008, well above the actual peak value.

These results suggest that investors on average placed more weight on more distant observations than the constant-gain value of 2 percent indicates. Therefore, as Figure 8 shows, investors perceived oil-price movements during this period as being less driven by permanent factors compared to the case under constant-gain learning. The figure shows that permanent shocks drive all of the variance of Δs_t under a 2 percent constant gain, more than twice as much as under our baseline model.

Is there a value of χ_T such that the weighted maximum likelihood estimation results in a better model fit? The value of χ_T that best matches the two-year ahead futures path is 1.5 percent. Remarkably, this is also the value used by Primeceri (2005) in a model of U.S. inflation in which policymakers learn about the natural rate of unemployment using a constant gain algorithm. However, as shown in Figure 7, even with this optimized value, the constant-gain estimation tends to over predict the run-up in futures prices between 2004 and 2008 relative to our baseline model, partly reflecting the large weight ascribed to permanent shocks as source of oil price fluctuations during this period (Figure 8). A similar pattern emerges since the end of the Great Recession.

⁸We concentrate on the two-year-ahead futures prices, the results for the one-year-ahead futures prices being very similar.

4.2 Particle Filter and the Risk Premium

The three-step procedures that we used to derive our baseline results has the benefit of simplicity, but it also presents some potentially important limitations. It assumes that for a given sample period the model's parameters, Γ , are constant. As a result, while the procedure allows investors to learn about the importance of temporary and permanent shocks to the spot price of oil, it restricts the evolution of the parameter values by requiring the parameter values to fit the entire sample rather than just recent observations. Second, our procedure abstracts from the presence of a time-varying risk premium, which could be an alternative source of movements in oil futures prices, as recently emphasized, for instance, by Hamilton and Wu (2012), Singleton (2014), and Baumeister and Kilian (2015)

In this section, we address these two limitations by assessing the robustness of our baseline results by extending our analysis to a more general learning process. In particular, we consider a variant of Stock and Watson's (2007) unobserved component model with stochastic volatility, in which we introduce an additional temporary shock to the level of oil prices. Although the model is similar to our baseline framework, it differs by allowing for time variation in Γ . Therefore, as before, the model for the spot price of oil is

$$s_t = e_t^P + e_t^T, \quad (10)$$

where we continue to assume that the permanent component follows a random walk with drift

$$e_t^P = \mu_t + e_{t-1}^P + v_t. \quad (11)$$

However, in contrast to our baseline framework, we allow for time-variation in μ_t and assume that the disturbance v_t is Gaussian with time-varying variance $\sigma_{p,t}^2$: $v_t \sim N(0, \sigma_{p,t}^2)$. Moreover, we postulate that the drift parameter, μ_t , follows a random walk

$$\mu_t = \mu_{t-1} + \xi_{\mu,t} \quad (12)$$

and that $\sigma_{p,t}^2$ evolve according to

$$\ln \sigma_{p,t}^2 = \ln \sigma_{p,t-1}^2 + \xi_{P,t} \quad (13)$$

where $\xi_{\mu,t}$ and $\xi_{P,t}$ are Gaussian disturbances with zero mean and constant variance $\sigma_{\xi_{\mu}}^2$ and $\sigma_{\xi_P}^2$, respectively.

As before, the temporary component follows an AR(1) process

$$e_t^\tau = \phi_{\tau,t} e_{t-1}^\tau + \varepsilon_t^\tau \quad (14)$$

where $\phi_{\tau,t}$ is allowed to vary through time and $\varepsilon_t^\tau \sim N(0, \sigma_{\varepsilon_t^\tau}^2)$. Again, we assume that

$$\ln \sigma_{\tau,t}^2 = \ln \sigma_{\tau,t-1}^2 + \xi_{\tau,t}, \quad (15)$$

where $\xi_{\tau,t}$ is an homoskedastic, Gaussian error term with zero mean and constant variance $\sigma_{\xi_\tau}^2$.

In addition, we depart from our baseline model by letting the k -period ahead futures prices to also be influenced by movements in a time-varying risk premium, Φ_t ,

$$f_t^k = E p_{t+k} + \Phi_t. \quad (16)$$

As for the temporary component, we assume that the risk premium follows an AR(1) process

$$\Phi_t = \rho \Phi_{t-1} + \xi_{\Phi,t}, \quad (17)$$

with the random disturbance $\xi_{\Phi,t}$ is assumed to be independently and identically normally distributed with zero mean and time-varying variance $\sigma_{\xi_{\Phi}}^2$, which evolves according to

$$\ln \sigma_{\xi_{\Phi},t}^2 = \ln \sigma_{\xi_{\Phi},t-1}^2 + \xi_{\Phi,t}, \quad (18)$$

where $\xi_{\Phi,t} \sim N(0, \sigma_{\xi_{\Phi}}^2)$

To bring our model to the data, we use the following set of equations consisting of the growth rate of the spot price of oil

$$\Delta s_t = [\mu_t + \varepsilon_t^P + (\phi_\tau - 1) e_{t-1}^\tau] + \varepsilon_t^\tau, \quad (19)$$

the expression for the spread between the k -period ahead futures price and the spot price

$$f_t^k - s_t = [k\mu_t + \rho\Phi_{t-1} + (\phi^k - 1)\phi e_{t-1}^\tau] + \phi^k \varepsilon_t^\tau + \xi_{\Phi,t}, \quad (20)$$

as well as (11), (14), (12), (13), (15), and (18). Given the presence of a risk premium shock, we complement the use of the price of West Texas Intermediate oil used for estimating our baseline model with the one-year-ahead futures price. We then use the model's estimates to forecast two-year-ahead futures prices. As before our estimation period covers 1980Q1 to 2009Q4. The estimation is done using the particle filter as described in Creal (2012).⁹

Figure 9 compares the expected futures prices using the Kalman filter model to those from the particle filter. The figure shows that, overall, both models have very similar predictions, especially in the 2002-05 period. In 2006 and 2007, the particle-filter-implied futures estimates are slightly higher, but the differences are relatively slight.

Figure 10 reports the estimated risk premium Φ . In order to provide a quantitative scale, we also report the future-spot spread $f_t^k - s_t$, the black line. The estimated risk-premium explains relatively little of the variation of the spread.

4.3 2010 and Beyond

We have shown that unobserved component models capture well the fluctuations in oil-price futures during the 2000s, a decade with important changes in the oil market that witnessed an unprecedented run-up in oil prices and their collapse during the Great Recession. We now examine the performance of our models since 2010. In Figure 11, we extend our sample to include the most current available data. In particular, the sample now also includes the rapid decline in oil prices that began in the summer of 2014.

The figure shows that our models continue to track the movements in futures prices very well through 2012. However, our baseline model misses the nearly flat trajectory of futures prices during 2013 and the first half of 2014 that is combined with higher spot prices, predicting a steady rise in oil-price futures instead.

⁹See Appendix A for details.

The pattern of rising spot prices coupled with relatively constant futures prices observed between 2013 and mid-2014 is reminiscent of the early 2000s. Our model could account for this pattern during the earlier episode since movements in the spot price of oil were the perceived to be largely transitory. Our model's misses during the more recent period reflects the fact that, by 2012, our framework ascribes a relatively important weight to permanent shocks as drivers of the spot price of oil. Therefore, our estimates of futures prices rose through that period, pushed by the perception that higher spot prices were relatively permanent. Our estimated futures prices then rapidly declined since then, in line with the plunge in the spot price of oil, bringing our estimates more in line with actual futures prices.

In contrast, our estimates derived using the particle filter continue to track the 2-year-ahead futures prices well post 2010, but this largely reflects the use of information embedded in the one-year-ahead futures price.

5 A DSGE Model with Learning

We showed that investors' perception of the persistence of oil prices clearly changed during the past decade and that these changes are captured well by a learning process about the role of transitory and persistent factors in the economy. We now assess the importance of this learning process for the impact of oil shocks on economic activity, using a DSGE model in which agents learn about the persistence of oil-price movements via the Kalman filter.¹⁰ One novel aspect of our approach is the use of a storage model to examine the effects of learning about the persistence of oil shocks. This framework is a priori appealing for this purpose, since oil prices are directly linked to expectations of future oil prices. The model consists of households that supply labor and rent capital to firms and save over time by holding one-period, pure discount bonds and by accumulating capital. We follow Arseneau and Leduc (2013) and assume that households directly hold

¹⁰We use the Kalman filter both for simplicity and because our empirical results with the particle filter were broadly consistent with those using the Kalman filter.

oil inventories. As is typical in the rational expectations storage literature, speculation in inventory holdings allows the household to smooth volatility in the oil market. The production side of the model is composed of firms producing a consumption good using labor, capital, and oil.

To assess the role of agents' perceptions about the oil market and learning, we use two parameterizations of the model based on our empirical results. First, we assume that agents in our economy perceive shocks to oil prices as being mostly transitory, as was the case in the early 2000s. We refer to this case as the 2003 scenario. Second, we calibrate the economy such that agents' perceptions is that oil-price movements are mostly permanent, in line with investors' actual perceptions on the eve of the Great Recession. We label this case the 2007 scenario. We then simulate the impact of an oil demand shock on economic activity under both calibrations and contrast the results to capture the effect of agents' perceptions of oil-price persistence. We focus on an oil demand shock, since it was highlighted as an important source of the run-up in oil prices during the past decade (see, Kilian (2008)). However, the gist of our results does not depend on the source of oil price fluctuations and generalizes, for instance, to a model in which oil prices are exogenous and subject to random fluctuations as in Kim and Loungani (1992) or Leduc and Sill (2004).

5.1 Households

The representative household's utility function is defined over consumption of a final good, c_t , and hours worked, n_t

$$U_t = \sum_{t=0}^{\infty} \beta^t \frac{[c_t^\theta (1 - n_t)^\theta]^{1-\sigma}}{1 - \sigma} \quad (21)$$

where $\beta \in (0, 1)$ denotes the subjective discount factor, the parameter θ controls the relative weight of consumption and leisure, $(1 - n_t)$, in the utility function, and σ represents the coefficient of relative risk aversion.

We assume that households can purchase a one-period real discount bond, denoted

b_{t+1} , and can also purchase s_t units of the commodity to hold as storage until the next period. Because households cannot borrow oil from the future, inventories must be non-negative. However, to simplify the numerical analysis below, we assume that the economy fluctuates around a steady state characterized by inventory holdings sufficiently large that the non-negativity constraint is never binding (see, Williams and Wright (1991) and Arseneau and Leduc (2013) for partial and general equilibrium analyses directly tackling the nonnegativity constraints on inventories). We verify that this condition is met in our simulations below. Holding inventories entails a per-unit cost $\phi(s_t)$, with $\phi'(s_t) > 0$ in terms of oil. Households also received a fixed endowment of oil each period, \bar{e} , which they can sell to firms on the spot market for oil.

The household also accumulates capital, k_t , which evolves according to the following law of motion

$$k_{t+1} = i_t + (1 - \delta)k_t,$$

where i_t denotes investment and δ represents the depreciation rate of capital. The household supplies labor and capital to the firm and derives labor income, $w_t n_t$, and capital income, $r_t k_t$, with the real wage and rental rate denoted by w_t and r_t , respectively.

The household chooses sequences of c_t , n_t , k_{t+1} , s_{t+1} , and b_{t+1} to maximize (21) subject to an infinite sequence of flow budget constraints given by:

$$c_t + p_{b,t+1} b_{t+1} + k_{t+1} + p_t s_{t+1} + p_t \phi(s_{t+1}) s_{t+1} = w_t n_t + r_t k_t + b_t + p_t s_t + (1 - \delta)k_t + p_t \bar{e}, \quad (22)$$

where $p_{b,t+1}$ is the price of the real discount bond; κ is the cost of storage valued in units of the aggregate consumption good; p_t is the relative price of oil.

The optimal demand for oil inventories is characterized by the following efficiency condition

$$p_t(1 + \phi(s_{t+1}) + \phi'(s_{t+1})s_{t+1}) = \beta E_t \left[\frac{\lambda_{t+1}}{\lambda_t} p_{t+1} \right], \quad (23)$$

where λ_t denotes the marginal utility of wealth.¹¹ This expression states that the house-

¹¹The complete set of equilibrium conditions are derived in Appendix A.

hold will accumulate oil inventories up until the marginal cost of holding one additional unit, inclusive of the cost of storage, is equal to the expected gain from holding the commodity for one period and then selling it at tomorrow's expected future spot price. If the spot price of oil today is higher than the discounted futures price net of the cost of storage, it is optimal to sell oil stocks today. Alternatively, if the spot price is lower than the discounted futures price net of the cost of storage, the household will accumulate inventories with the goal of selling oil in the future at the higher price. The household accumulates/decumulates inventories up until the point at which the intertemporal arbitrage condition holds.

Finally, we define the rational futures price of the commodity as $f_{t+1|t} \equiv \beta E_t \left[\frac{\lambda_{t+1}}{\lambda_t} p_{t+1} \right]$, or the expected future spot price of the commodity discounted back to today.

5.2 Production

The final good is produced using three inputs into production: labor, capital, and oil according to the following production function

$$y_t = zn_t^\gamma ((1 - \omega_t) k_t^\alpha + \omega_t q_t^\alpha)^{\frac{1-\gamma}{\alpha}} \quad (24)$$

where z denotes total factor productivity, q_t is the firm's usage of oil, and ω_t is a shock to the demand for oil, similar to that in Bodenstein and Guerrieri (2010). We assume that the oil demand shocks evolves according to the following AR(1) process

$$\begin{aligned} \omega_t &= \bar{\omega} \varepsilon_{\omega,t} \\ \varepsilon_{\omega,t} &= \varepsilon_{\omega,t}^P + \varepsilon_{\omega,t}^\tau \end{aligned}$$

where $\varepsilon_{\omega,t}$ denotes an innovation composed of a permanent component, ε_t^P , and a transitory one, ε_t^τ . For simplicity, the permanent component is modeled as a near-random walk

$$\varepsilon_{\omega,t}^P = \rho_\omega^P \varepsilon_{\omega,t-1}^P + u_{\omega,t}^P, \quad (25)$$

where we set ρ_P arbitrarily close to, but slightly below, 1 and where $u_{\omega,t}^P$ is a Gaussian error term with mean zero and standard deviation σ_{ω}^P . Similarly, the temporary component is assumed to follow an AR(1) process

$$\varepsilon_{\omega,t}^{\tau} = \rho_{\omega}^{\tau} \varepsilon_{\omega,t-1}^{\tau} + u_{\omega,t}^{\tau},$$

where ρ_{τ} denotes the persistence of the process and u_t^{τ} is a Gaussian disturbance with mean zero and standard deviation σ_{ω}^{τ} . Importantly, agents do not directly observe the permanent and transitory shocks to oil prices. Consistent with empirical results above, we assume that agents use the Kalman filter to learn about the source of oil-price movements.

The production function dictates that labor hours and capital services have a unit elasticity given by the Cobb-Douglas function, but that the capital stock and oil used are combined according to a CES function. The flexibility of the CES function will allow us to calibrate the elasticity of substitution between capital and oil to be in line with the empirical evidence of low substitutability between these two inputs.

The firm chooses hours and the amount of capital and oil to maximize per-period profits given by

$$\Pi_t^F = (y_t - w_t n_t - r_t k_t - p_t q_t). \quad (26)$$

5.3 Equilibrium

Taking as given the exogenous shocks ω_t , the equilibrium of the model is a sequence of $\{y_t, c_t, n_t, k_t, q_t, s_t, w_t, p_t, r_t\}$ that satisfy: the household optimality conditions; the optimality conditions of the final goods producing firm; the bond market clearing condition, the oil market clearing conditions $(1 + \phi(s_{t+1}))s_{t+1} - s_t + q_t = \bar{e}$; and the resource constraint $c_t + i_t = y_t$.

5.4 Calibration

We calibrate the structural parameters to match several steady-state observations. For those structural parameters that do not affect the model's steady state, we calibrate their

values to be consistent with other empirical studies in the literature. The structural parameters to be calibrated include β , the subjective discount factor; θ , the weight on consumption in the utility function; δ , the depreciation rate of capital; γ , the elasticity of substitution between labor and capital services in the production of the final good; α , the parameter governing the elasticity of substitution between capital and oil; $\bar{\omega}$ the parameter governing the share of oil in output; and $\phi(s)$, the storage cost function. In addition, we need to calibrate the parameters of the shock process, ρ_ω^τ , σ_ω^P and σ_ω^τ (ρ_ω^P is set arbitrarily close to 1, by assumption).

We set $\beta = 0.99$, so that the model implies a steady-state real interest rate of 4 percent per year and calibrate θ such that the household spends one third of his time working in steady state. We let capital depreciate 2.5 percent per quarter, while we set γ , which determines the labor share of income, to 64 percent. We choose $\bar{\omega}$ such that the share of oil in output is 5 percent. We set α such that the elasticity of substitution between capital and oil is roughly 0.2.

We assume that the storage cost function has the following form: $\phi(s) = \frac{\phi}{2}s$. We follow Unalmis et al (2012) and set ϕ such that oil stocks contribute 10 percent of oil absorption in steady state.

The relative importance of permanent and transitory shocks to oil demand, which is determined by parameters of the shock processes, is key for determining agents' inference about the persistence of oil-price movements. We consider two calibrations, which we label the 2003 and 2007 scenarios, based on our empirical findings. In a nutshell, we calibrate σ_ω^P , σ_ω^τ , and ρ_ω^τ so that the agents perceive oil-price movements as being largely transitory under the 2003 scenario and largely permanent under the 2007 scenario.

More precisely, we use our empirical estimates for the standard deviations and persistence of the permanent and transitory components to oil prices in 2003 and 2007 to calibrate the parameters of our oil-demand shocks. In particular, we set σ_ω^P , σ_ω^τ , and ρ_ω^τ so that, in response to oil demand shocks, our model with complete information can reproduce the magnitude and persistence that permanent and transitory innovations have

on oil prices in the data.

For instance, for the 2003 scenario, Figure 4 shows that our 2003 estimates of σ_p and σ_τ is 0.06 and 0.14, respectively. Similarly, the 2003 estimate of ϕ , the persistence of the transitory shock in equation (3), is 0.75, as shown in Figure 3. These coefficient estimates imply that oil prices increase by 0.09 following a one standard-deviation transitory shock to the demand for oil, with a half life of 1.6 months.¹² Using our model with complete information, we set σ_ω^τ , and ρ_ω^τ to be in line with this evidence. Similarly, we use our model to set σ_ω^P (again with ρ_ω^P being set arbitrarily close to 1) such that oil prices increase by 0.06 following a one-standard deviation permanent shock to oil demand. We follow a similar strategy to calibrate our 2007 scenario. Since our empirical estimates suggest that by 2007 σ_ω^P had significantly increased while σ_ω^τ had decreased, the 2007 scenario will naturally reflect an environment where permanent shocks are perceived to be a much more important source of fluctuations.

6 Macro effects of learning

We now examine the response of the economy to a near-permanent increase in oil prices, comparing the effect of the shocks between the 2003 and 2007 episodes. We linearize the model's equilibrium conditions around the economy's steady state and examine the effects of a one percent increase in the price of oil. Agents observe the increase in the price of oil, but must infer its persistence based on their perception of the relative importance of permanent and transitory shocks. Over time, as the economy evolves, agents reassess their views regarding the persistence of the oil-price increase using the Kalman filter.¹³

To better understand the model mechanism, we first contrast the transmission of a permanent increase in the price of oil to a transitory one, assuming that agents have complete information about the source of the shock underlying the rise in the price of oil. Figure 12 compares the responses, with the solid line denoting the economy's response

¹²The half-life of a transitory shock is k periods, where $k = -\ln(2 * \phi_\tau)$.

¹³We solve the model with learning via the Kalman filter using the methodology proposed in Andolfatto et al. (2008).

following a permanent shock and the dotted line representing the response following a transitory oil-price increase. Intuitively, the figure shows that the transitory shocks has a substantially more muted impact on output than a permanent shock, partly reflecting the decrease in oil inventories. When the price of oil is temporarily higher, inventory demand declines sharply, helping mitigate the reduction in oil supply and the associated reduction in output. In contrast, inventories decline much less following a permanent increase in oil prices, so that the rise in the price of oil and the associated decline in oil supply leads to a larger decline in output, due to the larger contraction in oil usage.

We now examine the effect of the perceptions of oil shock persistence and their interaction with storage. In the following exercises, we look at the response of the economy to a permanent oil demand shock under our two alternative scenarios. We first investigate the effect of the permanent demand shock when the shock is misperceived to be temporary, as parameterized under our 2003 scenario, and then contrast the result to the 2007 scenario in which agents expect the shock to be nearly permanent. In both cases, we compare the economy with learning to one with complete information.

Figure 13 reports the results under the 2003 scenario. It shows that when agents misperceived the shock as being largely transitory the decline in GDP is roughly one third of that when agents have complete information, partly reflecting the greater decline in inventories and thus the associated smaller rise in the price of oil. With incomplete information, the price of oil rises persistently above the one-year-ahead futures price and is associated with persistent realized expectation errors.

Even when agents expects the shock to be fairly persistent, as under the 2007 scenario, the near-term decline in activity continues to be substantially smaller relative to an economy with complete information, as shown in Figure 14. However, agents learn more rapidly about the persistence of the shock, so that the impact on the economy more rapidly resemble that of an economy with complete information. Under the 2007 scenario, spot prices do not rise as much above futures prices and the magnitude and persistence of the realized expectational errors are less acute.

Overall, the empirical learning process in the futures market that we document during the 2000s has a meaningful effect on the transmission of oil shocks on the economy. In particular, our empirical evidence imply that the negative effect of oil-demand shocks on output is only about one-third of its effect compared to an economy with full information. The fact that agents expected movements in oil prices to be somewhat transitory may account for the more muted effects of oil shocks on the macroeconomy during the 2000s than during previous decades.

7 Conclusions

The futures market failed to predict that the developments at the start of the 21st century would radically change the outlook for oil prices. Whereas many commentators have attributed the associated forecasts errors to speculation or market inefficiency, this paper provides an empirical and a theoretical explanation of the movements in oil prices based on learning. We showed that a simple unobserved component model in which investors must form beliefs about whether the persistence of changes in oil prices accounts remarkably well for the fluctuations in oil-price futures. Our simple framework captures the relatively slow increase in futures prices at the beginning of the past decade and their unprecedented run-up between 2004 and 2008. Even during the first half of 2008, a period during which oil prices reached historic highs, the model predicts a level of futures prices that is broadly in line with the data. Our estimates suggest that through learning investors revised up the contribution of permanent shocks to the variance of oil prices over this period.

We then quantified the effect of this learning process for the response of the economy to oil shocks. Using a DSGE model in which oil is storable and used in production, we show that agents' learning of the form suggested by our empirical results substantially cushions the recessionary effects of oil shocks. Consistent with our empirical results, we calibrate two scenarios that capture market participants' perceptions regarding the persistence of oil prices in 2003, when changes in oil prices were largely thought to be

transitory, and 2007, when oil price changes were expected to be much more persistent. Compared to a framework with full information, we show that the recessionary effects of oil-price increases are roughly halved in the year following the rise in the price of oil under learning.

The more muted effects of oil shocks under learning may partly explain the relatively weak impact that the run-up in oil prices in the mid-2000s had on economic activity, complementing the role of demand factors, changes in monetary policy, and a smaller dependence on oil compared to previous decades.

References

Acharya, Viral V., Lochstoer, Lars A., and Tarun Ramadorai, 2013, “Limits to Arbitrage and Hedging: Evidence from Commodity Markets, *Journal of Financial Economics* 109(2), pp. 441-465.

Alquist, Ron and Lutz Kilian, 2010, “What Do We Learn from the Price of Crude Oil Futures?” *Journal of Applied Econometrics* 25(4).

Alquist, Ron, and Olivier Gervais 2011, “The Role of Financial Speculation in Driving the Price of Crude Oil” Working paper, Bank of Canada.

Alquist Ron, Lutz Kilian and Robert J. Vigfusson, 2011, “Forecasting the Price of Oil,” International Finance Discussion Paper No. 1022, Board of Governors of the Federal Reserve System.

Andolfatto, David and Scott Hendry and Kevin Moran 2008, “Are inflation expectations rational?” *Journal of Monetary Economics* 55(2): 406–422.

Arseneau, David M. and Sylvain Leduc, 2013, “Commodity Price Movements in a general Equilibrium Model of Storage,” *IMF Economic Review* 61(1), pp. 199-224.

Baumeister, Christiane and Lutz Kilian, 2015, “A General Approach to Recovering Market Expectations from Futures Prices With and Application to Crude Oil,” manuscript.

M Bodenstein and L Guerrieri, 2011. *Oil efficiency, demand, and prices: a tale of ups and downs*,

Burnside, Craig, Martin Eichenbaum and Sergio Rebelo. 2009 “Understanding the Forward Premium Puzzle: A Microstructure Approach,” *American Economic Journal: Macroeconomics* 1(2), 127-54.

Buyuksahin, Haigh, Harris, Overdahl, and Robe 2008 *Fundamentals, Trader Activity and Derivative Pricing*,

Chernenko Sergey V., Krista B. Schwarz and Jonathan H. Wright 2004, “The Information Content of Forward and Futures Prices: Market Expectations and the Price of Risk,” International Finance Discussion Paper No. 808, Board of Governors of the Federal Reserve System.

Chinn, Menzie D., Olivier Coibion, 2010, “The Predictive Content of Energy Futures,” NBER Working Paper No. 15830 (January).

Cho, I.-K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. *Review of Economic Studies* 69, 1–40.

Creal, Drew 2012 “A Survey of Sequential Monte Carlo Methods for Economics and Finance”, *Econometric Reviews*, Vol. 31(3), pp. 245-296,

Erceg, Christopher J. & Levin, Andrew T., 2003. ”Imperfect credibility and inflation persistence,” *Journal of Monetary Economics*, Elsevier, vol. 50(4), pages 915-944, May.

Engel, Charles, 1996. ”The forward discount anomaly and the risk premium: A survey of recent evidence,” *Journal of Empirical Finance*, Elsevier, vol. 3(2), pages 123-192, June.

Etula, Erkko (2010) Broker-Dealer Risk Appetite and Commodity Returns” Working paper, Federal Reserve Bank of New York.

Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva 2012, “The Role of Speculation in Oil Markets: What Have We Learned So Far?” Working paper, University of Michigan.

J Hamilton 2009 *Causes and Consequences of the Oil Shock of 2007-08*, Brookings Paper

Hamilton, James D. (1994) *Time Series Analysis* Princeton University Press, Princeton, New Jersey.

Hamilton, James D. and Jing Cynthia Wu 2012 “Risk Premia in Crude Oil Futures Prices” *Manuscript*

Irwin, Scott H., and Dwight R. Sanders 2012, “Testing the Masters Hypothesis in Commodity Futures Markets” *Energy Economics* 34, 256-269.

Kilian, Lutz, and Daniel P. Murphy 2011, “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil” Working paper, University of Michigan.

King, Robert and Mark Watson 2002, “System Reduction and Solution Algorithms for Singular Linear Difference Systems under Rational Expectations,” *Computational Economics* 20: 57–86.

Mou, Yiqun 2010, “Limits to Arbitrage and Commodity Index Investment: Front-Running the Goldman Roll” Working paper, Columbia University.

Pierre Perron, Tatsuma Wada, 2009, “Let’s take a break: Trends and cycles in US real GDP”, *Journal of Monetary Economics*, Volume 56, Issue 6, September 2009, Pages 749-765,

Pindyck, Robert S., 2001, “The Dynamics of Commodity Spot and Futures Markets: A Primer,” *Energy Journal* 22(3): 1-29.

Reichsfeld, David A and Shaun K. Roache, 2011 “Do Commodity Futures Help Forecast Spot Prices?” *IMF Working Paper* No. 11/254:

Schwartz, Eduardo S., and James E. Smith (2000) "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science* 46, 893-911.

Singleton, Kenneth J. (2014) "Investor Flows and the 2008 Boom Bust in Oil Prices," *Management Science* 60, 300-318.

Stock, J. H. & Watson, M. 2007, "Why Has U.S. Inflation Become Harder to Forecast?", *Journal of Money, Credit and Banking*, 39(1), 3-33.

Tang,Ke, and Wei Xiong (2011) "Index Investment and the Financialization of Commodities," Working paper, Princeton University.

Unalmis, Deren, Unalmis, Ibrahim, and Derya Filiz Undal, 2012, "On Oil Price Shocks: The Role of Storage," *IMF Economic Review* 60(4), pp. 505-532.

Williams, J.C., and B. D. Wright, 1991, *Storage and Commodity Markets*, Cambridge University Press, Cambridge, U.K.

Appendix I. The Particle Filter

This appendix briefly describes our use of the particle-filter, which draws on the survey by Creal (2012). For a more thorough and advanced treatment, we encourage the readers to consult Creal (2012).

Given the observed data y_t

$$y_t = \left(\Delta \ln p_t \quad \ln f_t - \ln p_t \right)$$

and the model structure, our goal is to infer the distribution of the unobserved state variables

$$x_t = \left(\ln \sigma_{\tau,t}^2 \quad \ln \sigma_{p,t}^2 \quad \mu_t \quad \varepsilon_t^P \quad e_{t-1}^t \quad \Phi_{t-1} \right)$$

To construct estimates of x_t , we will attempt to draw out the distribution. We begin with a set of model parameters

$$\theta = \left(\sigma_{\tau}^2 \quad \sigma_p^2 \quad \sigma_{rp}^2 \quad \sigma_{\mu}^2 \quad \phi \quad \rho \right)$$

(As we describe below, we have chosen values of θ that maximize the log likelihood of the observed data, however for now just assume that we have a set of model parameters.) We start with a set of initial particles x_0^i where $i = 1$ to N . that are drawn randomly from a distribution conditional on θ . For the stationary variables, we draw from the unconditional distribution. For the non-stationary variables, we choose random variables that are reasonable approximations.

$$\mu_0^i = \mu_0 + \varepsilon_t^{\mu} \quad \text{where } \varepsilon_t^{\mu} \sim N(0, \sigma_{\mu}^2)$$

$$\ln \sigma_{\tau,t}^2 = v_t^{\tau} \quad \text{where } v_t^{\tau} \sim N(0, \sigma_{\tau}^2)$$

$$\ln \sigma_{p,t}^2 = v_t^p \quad \text{where } v_t^p \sim N(0, \sigma_p^2)$$

For each i , we then draw x_1^i from the distribution $p(x_1^i | x_0^i; \theta)$

$$x_1^i \sim p(x_1^i | x_0^i; \theta)$$

where the $p(x_1^i | x_0^i; \theta)$ is specified by the model. For example, we draw v_1^i from a normal distribution with variance σ_τ^2 and combine this with $\ln \sigma_{\tau,0}^{2,i}$ to generate $\ln \sigma_{\tau,1}^{2,i}$

$$\ln \sigma_{\tau,t}^{2,i} = \ln \sigma_{\tau,t-1}^{2,i} + v_t^i$$

For each particle we then compute the likelihood of the observed data,

$$P(y_t | x_t^i; \theta)$$

$$y_t = \begin{pmatrix} \Delta p_t \\ f_t^k - p_t \end{pmatrix}$$

where, it is assumed that y_t is randomly distribution, with the conditional mean given above., The variance covariance matrix is also model driven.

$$N\left(\begin{pmatrix} \Delta p_t - [\mu_t + \varepsilon_t^P + (\phi_\tau - 1) e_{t-1}^t] \\ f_t^k - p_t - [k\mu_t + \rho\Phi_{t-1} + (\phi^k - 1) \phi e_{t-1}^\tau] \end{pmatrix}, \begin{pmatrix} \sigma_{\tau,t}^2 & \phi^k \sigma_{\tau,t}^2 \\ \phi^k \sigma_{\tau,t}^2 & \phi^{2k} \sigma_{\tau,t}^2 + \sigma_{rp}^2 \end{pmatrix}\right)$$

We define

$$w_t^i = w_{t-1}^i P(y_t | x_t^i; \theta)$$

note that given the model structure these weights w_t^i are more simple the general expression for w_t^i given in Creal. We normalize the weights

$$\hat{w}_t^i = \frac{w_t^i}{w_t^i}$$

We then report filtered values of variable of interest like μ_t and $\sigma_{\tau,t}^2$, which are calculated as weighted averages of the particles.

$$\begin{aligned} \bar{\mu}_t &= \hat{w}_t^i \mu_t^i \\ \sigma_{\tau,t}^2 &= \hat{w}_t^i \sigma_{\tau,t}^{2,j} \end{aligned}$$

Confidence intervals can be constructed as [insert description]

Each time that we resample, we draw the particles from the distribution with probability \hat{w}_t^i . In line with the discussion in Creal (2012), we follow Liu and Chinn and resample only when the importance weights are unstable. Since excessive resampling should be avoided to minimize Monte Carlo error, we use the effective sample size (ESS) to decide when to resample.

$$ESS = \frac{1}{(\hat{w}_t^i)^2}.$$

If the ESS is less than the critical value $0.5N$, then we resample. For the resampled particles, we then set

$$w_t^i = \frac{1}{N}.$$

We then move on to the next observation.

Figure 1: Oil Spot and Futures Prices During the 1990s




Figure 2: Oil Spot and Futures Prices Since 2000




Figure 3: Estimates of Oil Price Trend and Persistence of Transitory Shocks




Figure 4: Estimated Standard Deviations of Permanent of Transitory Shocks




Figure 5: The Relative Importance of permanent Shocks

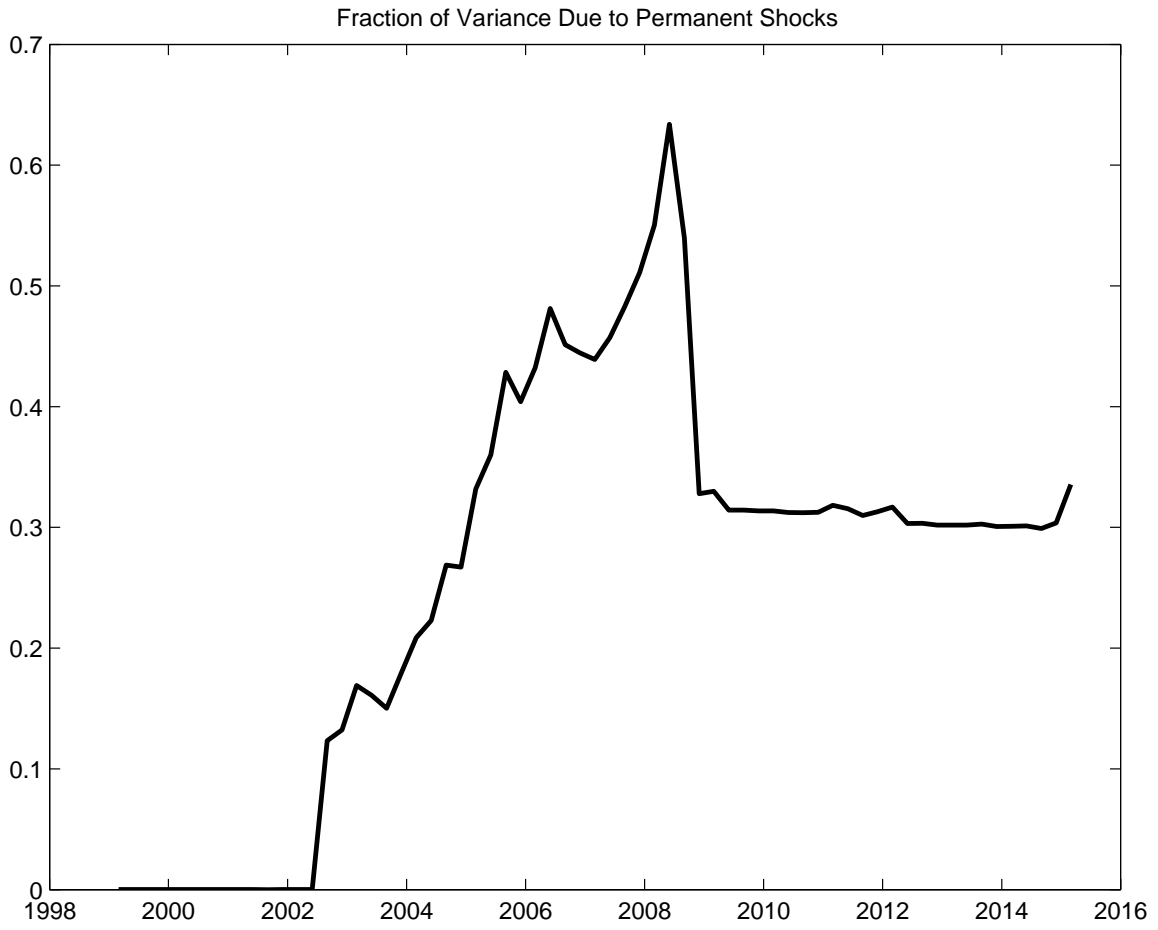


Figure 6: Predicting Futures Prices




Figure 7: Constant Gain Learning and Oil Futures Prices




Figure 8: The Importance of Permanent Shocks Under Constant Gain Learning

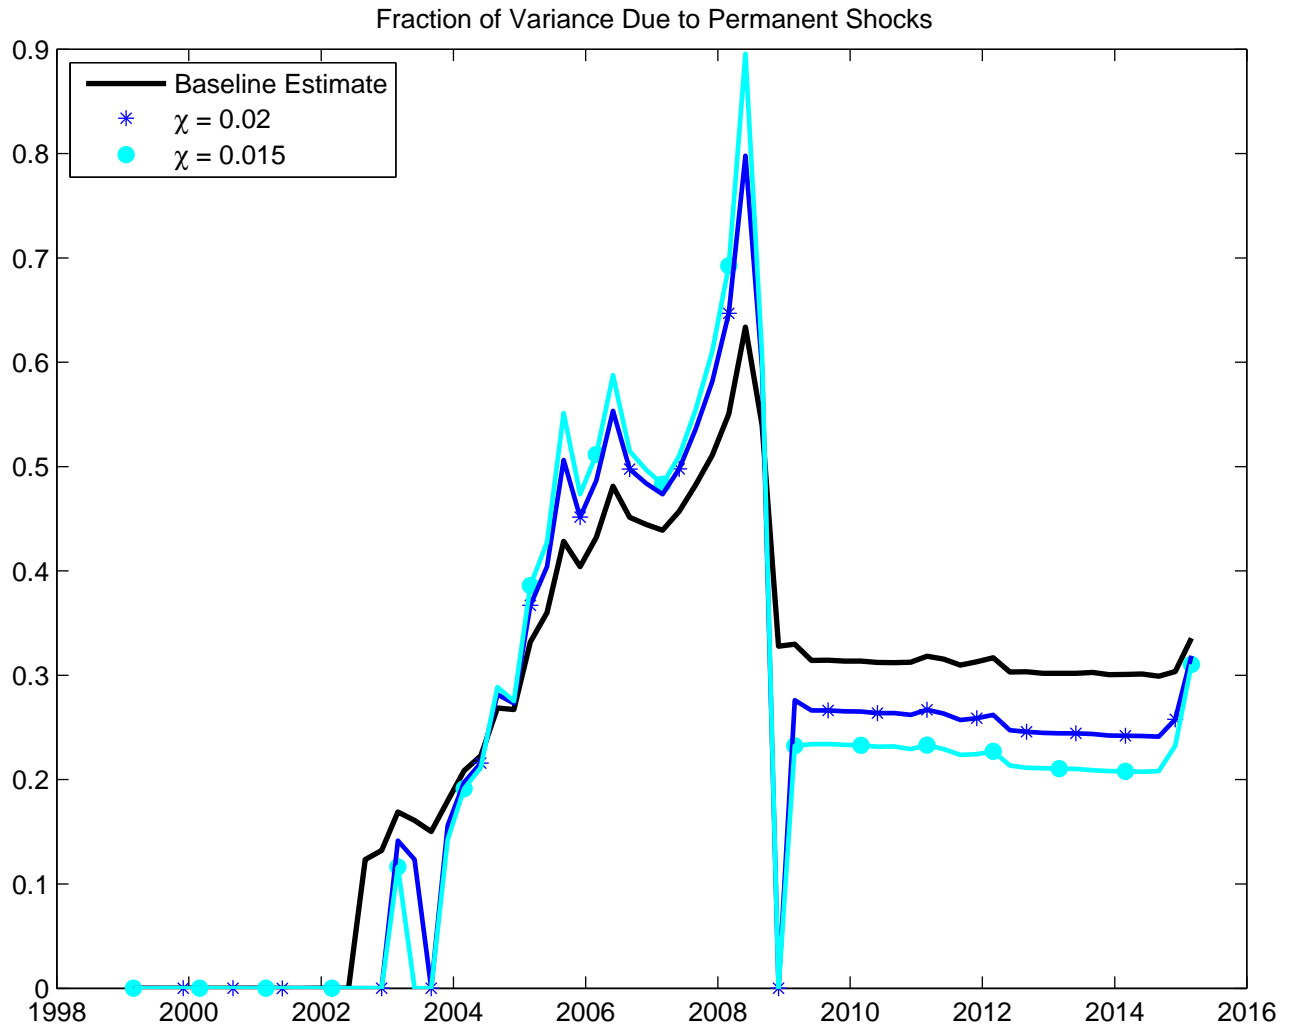


Figure 9: Oil Price Futures Under the Particle Filter

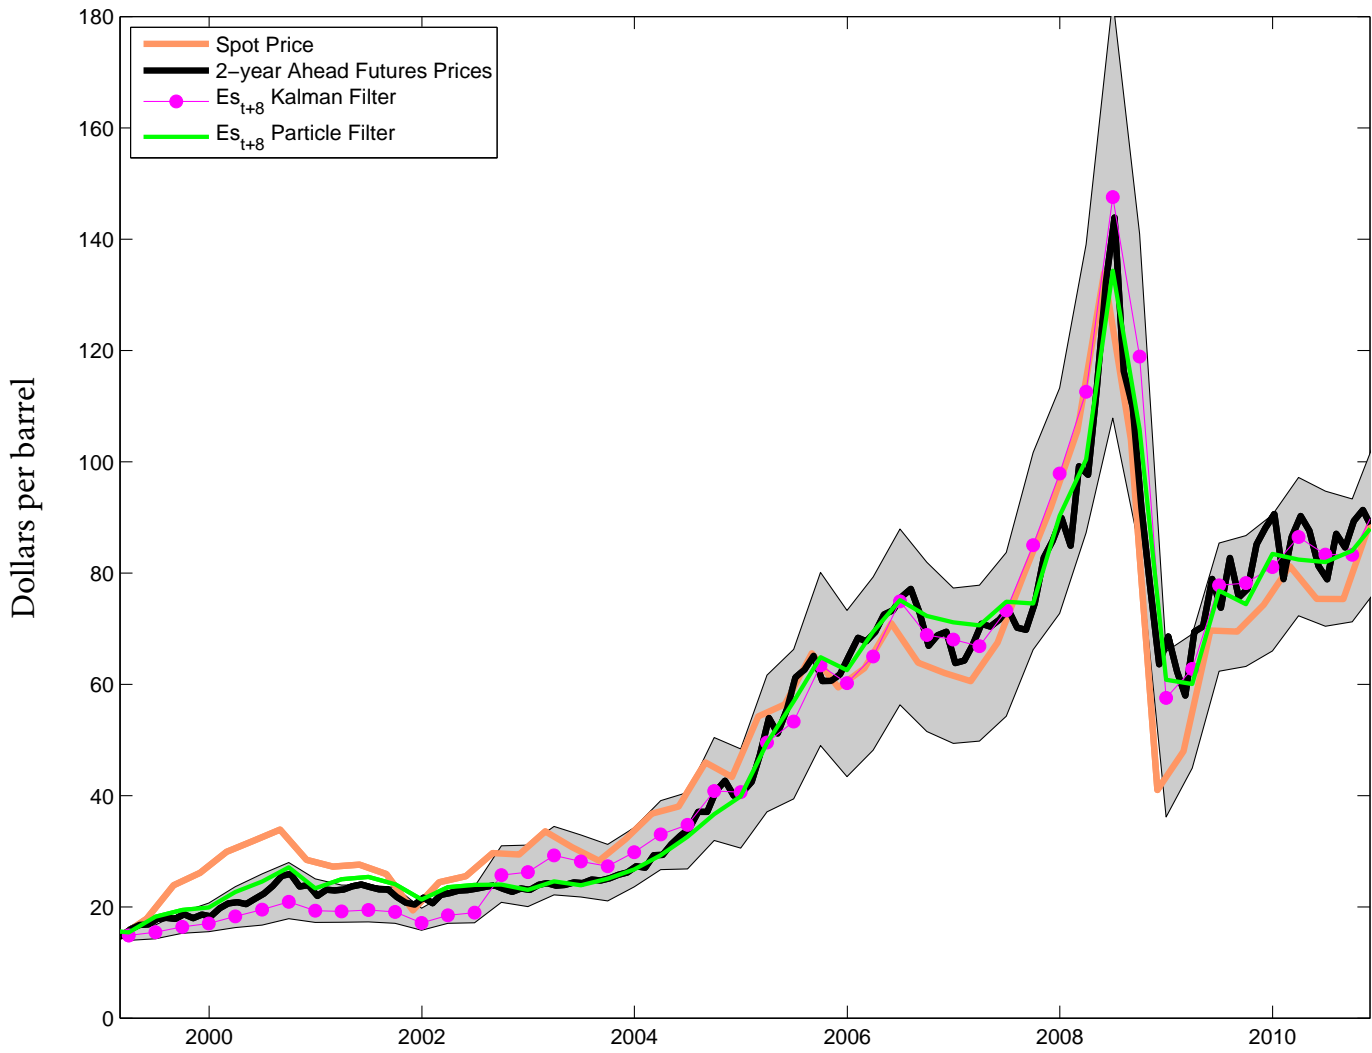


Figure 10: Estimated Risk Premium in Oil Market

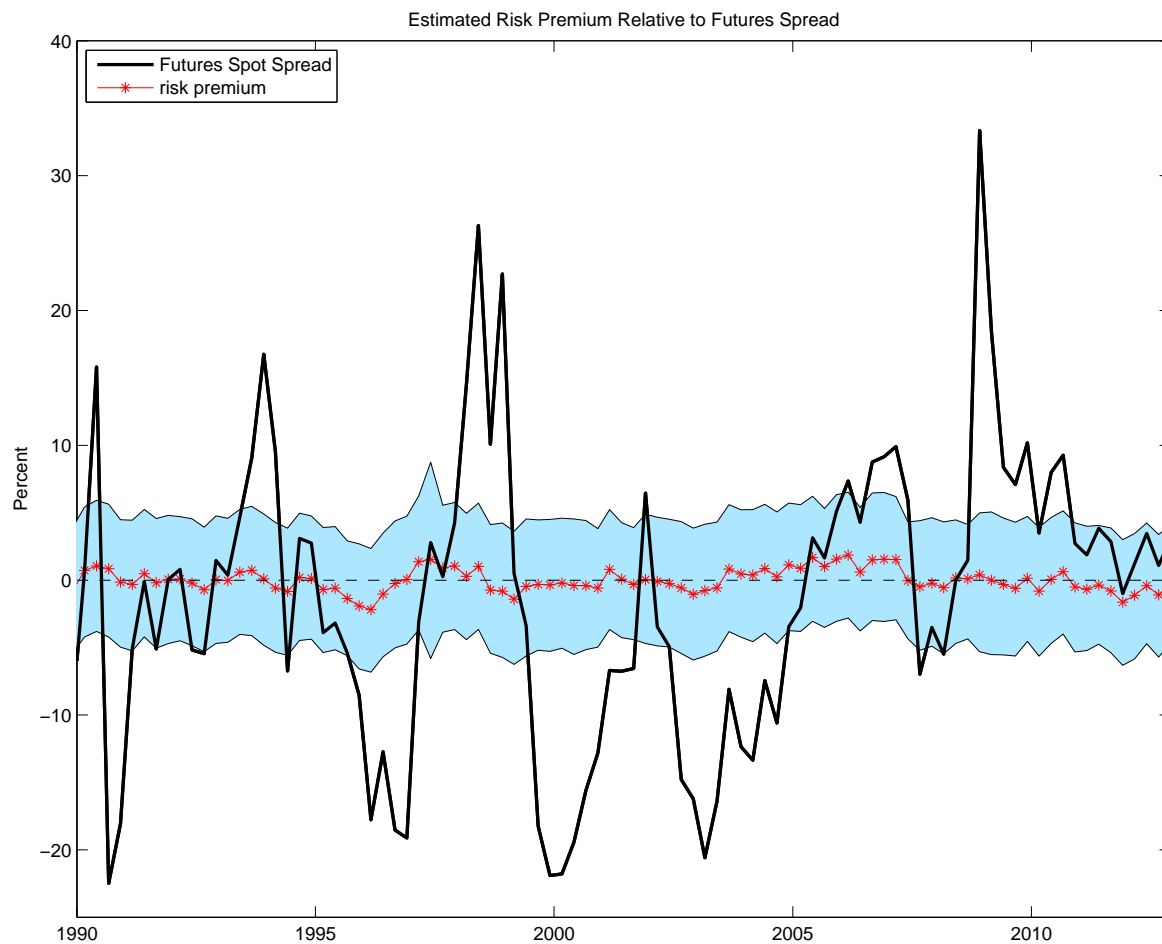


Figure 11: Estimated Futures Prices Since 2010




Figure 12: Impact of Transitory and Permanent Oil Demand Shocks Under Complete Information

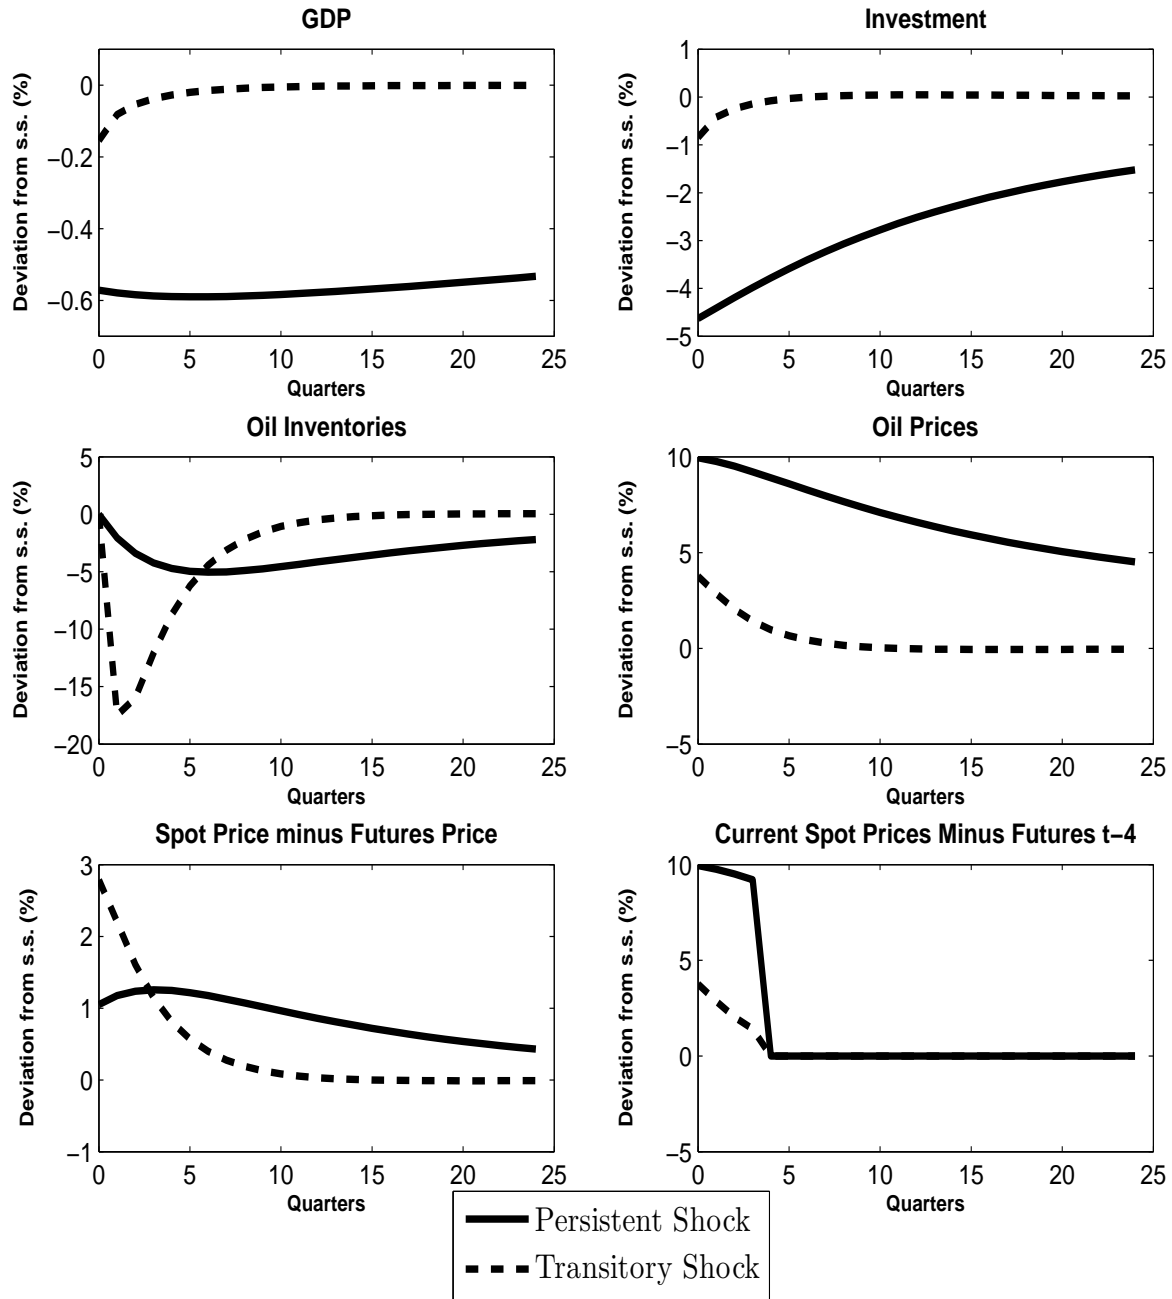


Figure 13: Impact of a Permanent Oil Demand Shock: 2003 Scenario Information

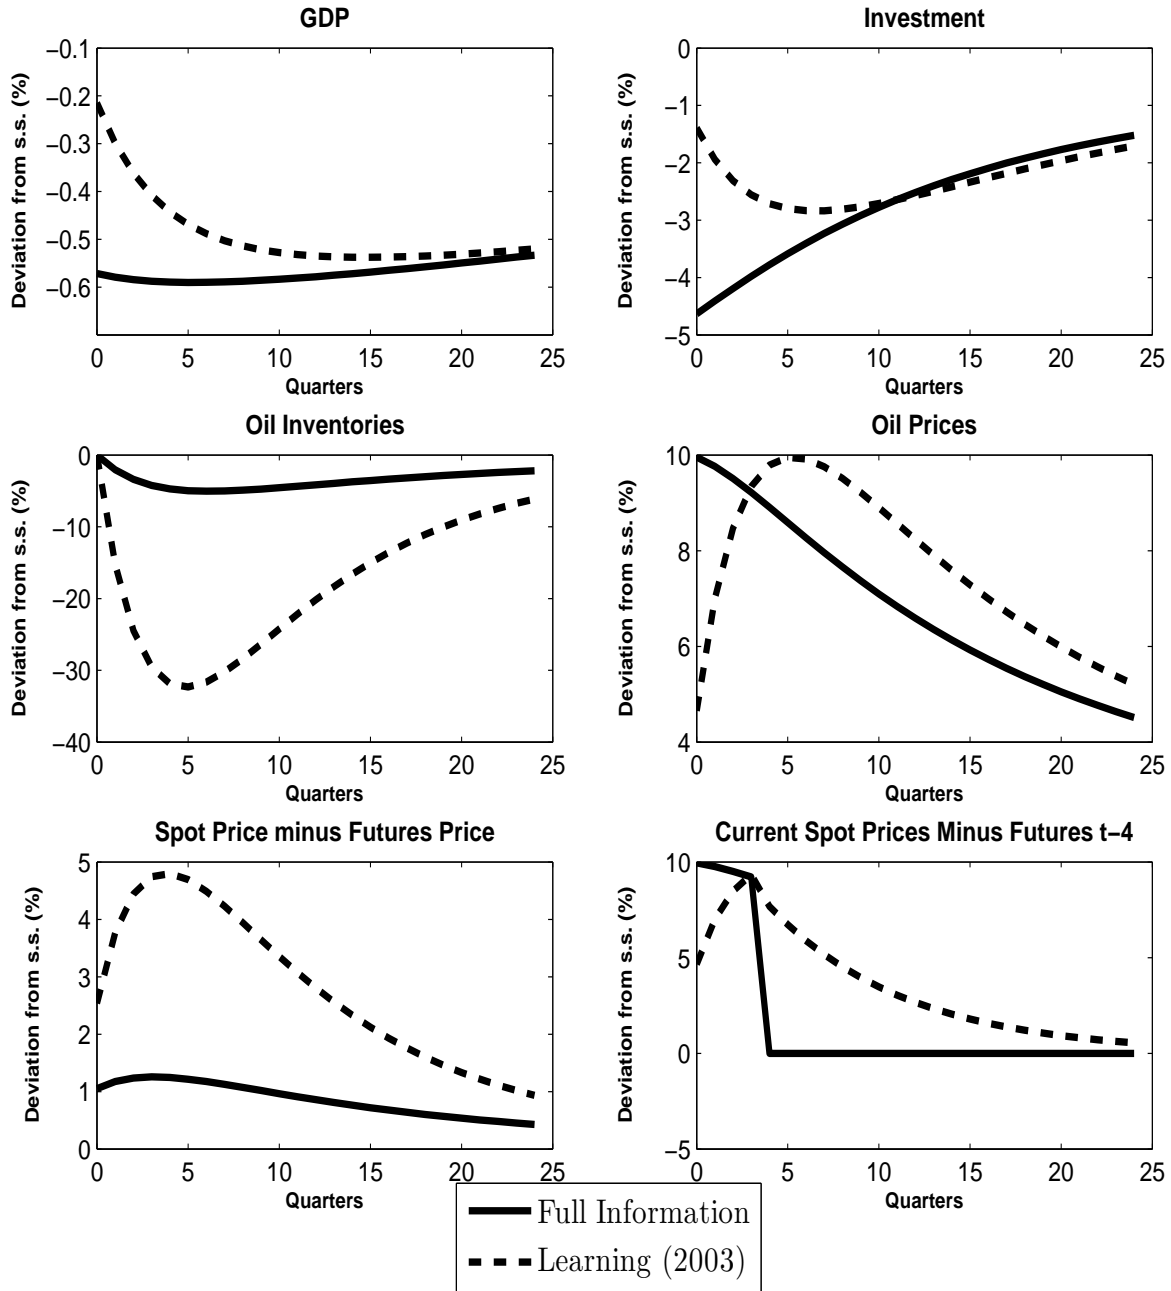


Figure 14: Impact of a Permanent Oil Demand Shock: 2007 Scenario Information

