

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Learning in the Oil Futures Markets: Evidence and Macroeconomic Implications

Sylvain Leduc
Federal Reserve Bank of San Francisco

Kevin Moran
Universite Laval

Robert J. Vigfusson
Federal Reserve Board of Governors

October 2020

Working Paper 2020-33

<https://www.frbsf.org/economic-research/publications/working-papers/2020/33/>

Suggested citation:

Leduc, Sylvain, Kevin Moran, Robert J. Vigfusson. 2020. "Learning in the Oil Futures Markets: Evidence and Macroeconomic Implications," Federal Reserve Bank of San Francisco Working Paper 2020-33. <https://doi.org/10.24148/wp2020-33>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Learning in the Oil Futures Markets: Evidence and Macroeconomic Implications*

Sylvain Leduc[†] Kevin Moran[‡] Robert J. Vigfusson[§]

October 1, 2020

Abstract

Using expectations embodied in oil futures prices, we examine how expectations are formed and how they affect the macroeconomic transmission of shocks. We show that an empirical framework in which investors form expectations by learning about the persistence of oil-price movements successfully replicates the fluctuations in oil-price futures since the late 1990s. We then embed this learning mechanism in a model with oil usage and storage. Estimating the model, we document that an increase in the persistence of TFP-driven fluctuations in oil demand largely account for investors' perceptions that oil-price movements became increasingly permanent during the 2000s before declining thereafter. We show that the presence of learning alters the macroeconomic impact of shocks, making the responses time-dependent and conditional on the views of economic agents about the shocks' likely persistence.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco, the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System. We thank seminar participants at the Federal Reserve Board, the Federal Reserve Bank of Boston, the Federal Reserve Bank of Cleveland, the Swiss National Bank, the Bank of Canada, Larefi (University of Bordeaux), and conference participants at the 2017 Federal Reserve System Energy Meeting, the 2017 CEBRA Annual Meeting, the BI Norwegian Business School and the 2015 NBER meeting on the Economics of Commodity Markets. In particular thanks to Ron Alquist, Benjamin Johannsen, Matt Smith, Hilde Bjornland, and Fabio Canova. Comments to robert.j.vigfusson@frb.gov.

[†]Federal Reserve Bank of San Francisco (sylvain.leduc@sf.frb.org).

[‡]Université Laval (kmoran@ecn.ulaval.ca).

[§]Federal Reserve Board (robert.j.vigfusson@frb.gov)

1 Introduction

Expectations are central in macroeconomics. They affect the response of price and wage inflation to shocks and impact the transmission of shocks more broadly. Whereas the literature has often looked at expectations formation in surveys (see Coibion, Gorodnichenko, and Kamdar (2018) for a recent overview), we look at the expectations embodied in futures markets. We focus on the oil market in part because of its importance for economic activity.

The futures market appears to have only slowly recognized that events since the early 2000s—from the growing importance of China in the world economy to the onset of the financial crisis and the shale revolution—would radically change the outlook for oil prices. Whereas the resulting forecast errors over this period have often been attributed to speculation or to time-varying risk, this paper provides an explanation of the movements in oil-price futures based on expectations formed through learning. We then show, using an estimated dynamic, stochastic, general equilibrium (DSGE) model, that this learning process has important macroeconomic implications. In particular, responses to shocks become time-dependent and conditional on the views of economic agents about the shocks' likely persistence.

Our analysis indicates that the developments in the oil futures markets since the late 1990s are largely consistent with investors gradually learning about the persistence of underlying shocks. Using the Kalman filter to infer the permanent and transitory components of shocks to spot oil prices, we show that this simple form of learning can replicate the observed behavior of futures prices.¹ Our estimates suggest a rising contribution of permanent shocks to the variance of oil prices from 2002 to 2008. Our analysis further indicates that investors perceived the late-2014 drop in oil prices as somewhat transitory. Since then, our estimates suggest some additional slight decline in the perceived persistence of shocks.

We then embed this gradual learning mechanism in a DSGE model in which oil is storable and used in production to impute the underlying source of shocks for the changes in investors' beliefs. We estimate the model focusing on three scenarios that

¹The Kalman filter has the advantage of being a straightforward approach to model learning but is also somewhat restrictive, assuming for example constant model parameters. In robustness exercises, we present results for two alternatives that allow for time-varying parameters: a constant-gain learning model and a variant of Stock and Watson's unobserved components model (2007) with stochastic volatility, estimated via the particle filter. We find that the baseline results derived using the Kalman filter are little changed under either alternative.

are consistent with our empirical results and capture perceptions about the persistence of oil price movements during key episodes: in the early 2000s, when changes in oil prices were largely thought to be transitory; in the late 2000s, when oil-price changes were expected to be much more persistent; and in the mid-2010s during the oil-price collapse. We document that an increase in the persistence of total factor productivity (TFP) driven fluctuations in oil demand largely accounts for investors' perceptions that oil-price movements became increasingly permanent during the 2000s, before declining thereafter. In contrast, we estimate that oil-supply shocks were mostly perceived as being temporary over this time period.

Using our estimated model, we also conduct some counterfactual experiments to assess the macroeconomic impact of learning about the persistence of shocks underlying oil-price movements. We show that misperceiving the persistence of shocks has large macroeconomic effects. For instance, failing to recognize the increase in the persistence of TFP-driven fluctuations in oil demand during the 2000s substantially increases the volatility of output and investment by raising the volatility of labor supply. In addition, this misperception lowers the volatility of the spot and futures prices of oil, as greater reliance on inventories is used to smooth the shocks in this case. We also show that the impact of learning partly occurs through storage, since dynamic, forward-looking inventory management decisions are importantly influenced by expectations.

Our paper thus highlights the importance of learning about the persistence of shocks underlying oil-price movements for the macroeconomy. Our finding that the perceived importance of TFP-induced oil-demand fluctuations rose during the 2000s is in line with Kilian's (2009) assessment of the period. For instance, he shows that increases in the global demand for industrial commodities, a shock with large and immediate impacts on oil prices, were the main source of oil-price increases in the mid-2000s. Our work complements this perspective by noting that demand and supply decompositions of oil shocks should also consider the role of perceptions about their likely persistence.

Our emphasis on investors learning about the persistence of shocks underlying oil-price fluctuations partly complements the work of Jin (2017), who examines the role of temporary and permanent shocks in driving movements in the spot and futures prices of oil and in improving forecasting accuracy. Moreover, by highlighting learning, we provide an alternative perspective to the literature emphasizing the role of time-varying risk premia embedded in futures prices. In addition, our approach relates to the analysis of Milani (2009) who examines the role of learning in altering the impact

of oil shocks on the macroeconomy using a standard New Keynesian model with oil used in production and directly consumed by households. Emphasizing constant-gain learning, he shows that learning contributed to a more muted impact of oil shocks since the 1970s, as agents' perceptions of the macroeconomic effects of these shocks changed over time, partly driven by the declining oil share in output. Similarly, Jin (2017) examines the importance of expectations shocks in driving commodity price dynamics, finding a limited role for the period between 1987 and 2014. We add to these papers by emphasizing the importance of learning about the persistence of shocks and by introducing oil inventories, which can help smooth transitory fluctuations.

The rest of the paper is organized as follows. After describing movements in oil prices over the past three decades in more detail, we lay out our empirical framework, estimating a time-series model of permanent and temporary shocks for oil prices. We also report results for extensions to our baseline empirical model. We then develop and estimate a DSGE model with oil usage and storage to examine the macroeconomic consequences of learning about the persistence of oil market developments, calibrating the learning process to match our empirical findings from 2003, 2007, and 2014. We conclude in the last section.

1.1 Related literature

Our paper's emphasis on how learning defines the relationship between spot and futures prices is an alternative perspective relative to most of the oil-price literature, which instead focuses on the role of time-varying risk premiums embedded in the futures markets.² For instance, Hamilton and Wu (2014) argue that increased participation by index-fund investing has reduced oil futures premiums 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 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.

²This literature is large and growing. Some of the many papers discussing this issue include Irwin and Sanders (2012); Fattouh, Kilian and Mahadeva (2013); Kilian and Murphy (2014); Alquist and Gervais (2013); and Singleton (2014).

This literature attributes developments in the futures markets to the increased financialization of commodity markets or to speculators' risk appetite. Relatedly, risk-premia variations in commodity markets are claimed to have important influence on futures prices. For instance, Baumeister and Kilian (2015) show such variations to be important during episodes like the height of the oil-price boom in the beginning of 2008. In addition Hitzemann (2016) presents a production-model view of asset and commodity markets incorporating long-run risk similar to Croce (2014); he shows such a setting replicates many observed features in oil futures returns.

By contrast, we abstract from financial factors and risk premia and show that learning about the persistence of shocks, in the tradition of Brunner et al., 1980, can help track the evolution of oil futures prices since the late 1990s and alter the macroeconomic impact of oil shocks. Our work does not intend to determine whether a risk-premium versus learning explanation should be preferred. Instead, it highlights that learning is a plausible approach, which may be complementary to a risk-premium view of oil markets.

The macroeconomic model in our paper is related to a broader macroeconomic literature on learning that focuses on situations when shocks have both persistent and transitory components not observable separately. A filtering mechanism helps economic agents disentangle both types gradually, even if early periods following a shock may be accompanied by expectation errors.³ This type of learning is related but conceptually distinct from the one originating from the work of Evans and Honkapohja, (2001), who use least-squares learning using realized economic outcomes to ascertain the general form of the economy's evolution (Milani (2009) uses this framework to study the changing relationship between oil prices and the macroeconomy since the 1970s). In our macroeconomic model (but not in our empirical section), all model parameters are known, but learning is necessary because, upon realization, the persistence of a shock affecting the economy is unknown; in such a setting, futures markets can provide valuable information about the learning process.⁴ Overall, these strands of literature argue that

³This literature is exemplified by contributions such as Erceg and Levin (2003), Schorfheide (2005), Andolfatto et al. (2008) or Saijo (2017).

⁴Our emphasis on the signaling value of futures markets for expectations is also compatible with Kellogg's result (2014), who shows that conditioning on oil futures better explains drilling decisions than assuming oil prices follow a random walk. Yet another form of learning is related to the "news shock" literature of Beaudry and Portier (2004), according to which information received in current periods may represent imperfect signals about future shocks (see Jin 2017 for an application to oil markets).

information included in expectations data should be incorporated in macroeconomic frameworks, notably via imperfect information and learning; see Coibion et al. (2018) for an exposition of this research program, as well as Slobodyan and Wouters (2012), Fornero and Kirchner (2014) or Ormeno and Molnar (2015) for recent contributions.⁵

Finally, our paper also complements the work of Alquist and Kilian (2010), who emphasize the presence of a convenience yield associated with oil inventories and its role 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 convenience yield may underlie the poor predictive performance of oil-price futures. We add to this literature by highlighting the interaction between learning and oil storage.

2 Oil prices over the past three decades

We start our analysis by presenting some evidence of the oil market's evolving views regarding the persistence of the shocks affecting the world economy. To do so, consider the movements in the spot and futures prices of oil since the early 1990s, depicted in the three panels of Figure 1. In each panel, the solid black line shows the evolution of the spot price, while the dotted red lines depict the path of futures prices at several given dates in the sample. The futures prices are quotes as reported by NYMEX. During the 1990s (top panel), the spot price tended to gyrate around fairly stable oil-price futures, suggesting that market participants viewed economic developments affecting the oil market as mainly temporary. Underlying shocks would tend to move the spot price of oil, at times substantially, but futures prices would indicate an expected return to roughly \$18 per barrel, about the average spot price during that period. These patterns suggest that market participants did not view the disturbances affecting oil markets as persistent enough to substantially alter their long-term view of oil prices, and these views were substantially correct during the 1990s.

However, the middle panel of Figure 1 indicates that the relationship between spot and futures prices changed during the early 2000s. Between 2000 and 2008, the spot price of oil rose steadily, from just above \$25 per barrel to more than \$140 per barrel. In

⁵ Sockin and Xiong (2015) also emphasize the importance of informational frictions in commodity markets using an equilibrium model with asymmetric idiosyncratic information among traders, where commodity prices provide imperfect signals about future economic conditions.

contrast, oil-price futures remained initially low, consistently fluctuating below \$20 per barrel until 2002. Then, oil-price futures started a gradual rise, increasing 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 2000-03, market participants perceived movements in spot prices to remain mostly temporary, as indeed had been the case throughout the 1990s. However, after being consistently surprised by the persistent rise in spot prices, market participants reassessed their views, placing more weight on the possibility that these steady increases in oil prices would be long-lasting.

According to this interpretation, by the time the oil market reached its peak in the spring of 2008, market participants largely expected the movements in spot prices to be highly persistent so that the price would remain at about \$135 per barrel over the next five years, as indicated by the futures curves at that time. The graph also suggests that, in the fall of 2008, the global financial crisis led to a reassessment of the long-run equilibrium price of oil. Indeed, the far-dated futures prices declined from roughly \$140 per barrel in the spring of 2008 to roughly \$60 per barrel by the end of that year.

Lastly, the evolution of spot and futures prices since 2010 is shown in the bottom panel of the figure. Between 2010 and 2014, the fluctuations in prices had more in common with the 1990s. That is, market participants appear to have perceived most fluctuations in the spot price of oil to be transitory, with the long-run futures price remaining fairly stable despite significant movements in the spot price. However, the oil market changed dramatically in mid-2014 when the spot price collapsed by roughly 50 percent.⁶ As in the early 2000s, this price change was initially perceived as being mostly temporary, with futures curves rising back to a long-run price of about \$80 per barrel during most of 2014 and only gradually adapting to the lower prices. By the end of 2015, this assessment was substantially changed, as far-dated futures prices only reached about \$55 per barrel. The next few years seemed broadly consistent with these forecasts. Although prices climbed back to near \$70 per barrel in late 2018, they retraced that increase and ended 2019 at \$60 per barrel.

⁶Further discussion of this decline can be found in Baumeister and Kilian (2016).

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. By design, we adopt a straightforward approach that highlights our main point and accordingly abstracts from many features of the oil market. Specifically, we postulate that spot oil prices are the result of movements in permanent and transitory components 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 re-estimate the model parameters as the data sample increases. We estimate our benchmark model using spot prices only. We adopt this approach to show that the learning process goes a long way in capturing the evolution of futures prices. We then extend our benchmark model to allow for an alternative form of learning that places less weight on the past than the present. We then augment our estimation by considering spot and futures prices jointly in our estimation, comparing estimates from the Kalman filter to those from a nonlinear model that explicitly allows for time-varying parameters.

3.1 A simple model

Consider the following linear process relating the spot price of oil p_t (expressed in logs) to a permanent component, $e_t^{\mathcal{P}}$, and a stationary one, e_t^{τ} :

$$p_t = e_t^{\mathcal{P}} + e_t^{\tau}, \tag{1}$$

Schwartz and Smith (2000) modeled oil prices using this assumption but, unlike in the current paper, they did not re-estimate the model coefficients as the sample increased. The permanent component is modeled as a random walk without drift:

$$e_t^{\mathcal{P}} = e_{t-1}^{\mathcal{P}} + v_t, \tag{2}$$

where v_t is independently and identically normally distributed with mean zero and constant variance σ_p^2 .⁷ For its part, the temporary component is assumed to follow the AR(1) process

$$e_t^{\tau} = \phi_{\tau} e_{t-1}^{\tau} + \varepsilon_t, \tag{3}$$

⁷Results are similar if the model allows for a drift parameter in $e_t^{\mathcal{P}}$.

where ε_t is also independently and identically normally distributed, with mean zero and constant variance σ_τ^2 , and where $|\phi_\tau| < 1$.

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

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

In contrast, absent full information about the current levels of $e_t^{\mathcal{P}}$ and e_t^τ , the (log) futures price will be based on the best forecasts given past values of p_t :

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

Note that the level of the futures price $F_{t,k}$ does not translate directly from this expression because of terms related to Jensen's inequality. The correct value for $F_{t,k}$ is instead

$$F_{t,k} = E_t \exp \left(e_t^{\mathcal{P}} + \phi_\tau^k e_t^\tau + \sum_{j=1}^k v_{t+j} + \sum_{j=0}^k \phi_\tau^{k-j} \varepsilon_{t+j} | \{p_{t-i}\}_{i=1}^t \right). \quad (6)$$

In assessing our model's performance, we compare actual values of futures prices, $F_{t,k}$, with estimated (fitted) values computed with our estimated parameters and model simulations using (6).

In our benchmark specification, we first estimate the model using spot prices to forecast the levels of the futures price. However, we complement this approach below by estimating the model using both the spot price and the futures price. In that case, we include an additional component, ζ_t , in the equation of the futures price:

$$f_{t,k} = E_t p_{t+k} + \zeta_t. \quad (7)$$

This component creates a wedge between the futures price and the expected price and could be interpreted as capturing either measurement error or a risk premium. For convenience, we will refer to it as a risk premium.

We assume that ζ_t also follows an AR(1) process

$$\zeta_t = \phi_\xi \zeta_{t-1} + \varepsilon_{\zeta,t}. \quad (8)$$

When we estimate this model using the Kalman filter, we assume that $\varepsilon_{\zeta,t}$ is inde-

pently and identically normally distributed with mean zero and constant variance σ_ζ^2 .

Finally, the relative importance of permanent shocks can be assessed using a simple statistic derived using the formula for the change in the spot price of oil. To this end, note that

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

which implies that the variance of Δp_t can be expressed as

$$\sigma_{\Delta p}^2 = \sigma_{\mathcal{P}}^2 + \frac{2}{(1 + \phi_\tau)} \sigma_\tau^2,$$

and the fraction of $\sigma_{\Delta p}^2$ due to permanent shocks, in turn, is the following expression:

$$\frac{\sigma_{\mathcal{P}}^2}{\sigma_{\mathcal{P}}^2 + \frac{2}{(1 + \phi_\tau)} \sigma_\tau^2}. \quad (9)$$

In the model section below, we use this equation to calibrate the relative standard deviations of permanent and transitory shocks.

3.2 Learning

We assume that market participants use the Kalman filter to form expectations of future oil prices. Following the standard textbook treatment (Hamilton, 1994), define ξ_t as the unobserved state vector of the benchmark model above, comprising the permanent and temporary components: $\xi_t = \left(e_t^{\mathcal{P}} \quad e_t^\tau \quad e_{t-1}^{\mathcal{P}} \quad e_{t-1}^\tau \right)'$. Given values for the model's parameters, $\Gamma \equiv [\sigma_{\mathcal{P}}^2, \sigma_\tau^2, \phi_\tau]$, the Kalman filtering equation relates how observed variables (here, the change in the spot oil price Δp_t) respond to changes in the unobserved state vector ξ_t .

The equations for the dynamics of the observed variable Δp_t are thus given by the following system:

$$\begin{aligned} \Delta p_t &= H \xi_t; \\ \xi_t &= F \xi_{t-1} + \eta_t, \end{aligned} \quad (10)$$

where F and H are matrices of parameters and $\eta_t = \left(v_t \quad \varepsilon_t \quad 0 \quad 0 \right)$ are the shocks. In

turn, the unobserved state vector evolves according to the following standard equation:

$$\xi_{t|t} = \xi_{t|t-1} + P_{t|t-1}H(H'P_{t|t-1}H)^{-1}(\Delta p_t - H'\xi_{t|t-1}), \quad (11)$$

given initial estimates of $\xi_{t|t-1}$ and $P_{t|t-1}$, where the forecast error $(\Delta p_t - H'\xi_{t|t-1})$ is used to update estimates of the permanent and transitory components via the term $P_{t|t-1}H(H'P_{t|t-1}H)^{-1}(\Delta p_t - H'\xi_{t|t-1})$. The matrix $P_{t|t-1}$, which depends on the values of σ_p^2 and σ_τ^2 , governs the extent to which a given surprise is judged to have arisen from the permanent or the transitory components.

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

Our discussion so far has assumed that the parameter values were known with certainty, but in practice we will need to estimate the model parameters $\Gamma = [\sigma_p^2, \sigma_\tau^2, \phi_\tau]$ to derive forecasts of future oil prices. We assume that market participants estimate the model's parameters using the standard likelihood function:

$$LL(\Gamma) = - \sum_{t=1}^T \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_t\| + (\Delta p_t - E \Delta p_t)' V_t^{-1} (\Delta p_t - E \Delta p_t) \right), \quad (12)$$

where $V_t = H'P_{t|t-1}H$ is the variance of the prediction errors. Below, we also provide results obtained using an alternative likelihood function putting more weight on recent observations.

4 Benchmark results

In this section, we present our benchmark model 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 benchmark 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 solely from previous spot prices may be suboptimal, our emphasis on univariate methods has the benefit of simplicity and shares similarities with the learning algorithm used in the monetary policy literature (see, for instance, Orphanides and van Norden, 2005, and Primiceri, 2006). Using the price of West Texas Intermediate (WTI) crude oil from 1980:Q1 to 2019:Q4, we construct the model-implied estimates of the two-year-ahead futures prices and compare them with the actual futures price data.⁸

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 $t - 1$ to estimate the model parameters $\Gamma = [\sigma_{\mathcal{P}}^2, \sigma_{\tau}^2, \phi_{\tau}]$ using the likelihood in (12). Second, we apply the Kalman filter with the estimated model parameters and observed prices through time t to get estimates of the unobserved permanent and temporary components $e_t^{\mathcal{P}}$ and e_t^{τ} . In the third step, we use the estimated $e_t^{\mathcal{P}}$, e_t^{τ} and Γ to construct simulated values for $F_{t,k}$.

We are particularly interested in the behavior of futures prices since the late 1990s. As such, we first estimate the model from 1980:Q1 to 1998:Q4 and start calculating futures prices from this period on, using an expanding window of data. Thus, the fitted futures prices at the beginning of 2000Q1, say, are calculated using the model estimated over the period from 1980Q1 to 1999Q4 while the longer sample 1980:Q1–2004:Q3 would be used to estimate the model and compute the forecast of futures prices for 2004:Q4. In order to make this most comparable to a real-time exercise, our modeled two-year-ahead futures price is compared against the actual two-year-ahead futures price, measured using the closing quote from NYMEX for WTI crude at the end of the first week of the following quarter (typically near the 7th of the month). For instance, the futures price computed using information about the spot price through the third quarter of 2015 is compared against the closing futures price for October 2, 2015.

One concern is whether there is sufficient activity in these far-dated contracts to provide useful information. One relevant metric to assess the market’s liquidity is the

⁸We begin the sample in 1980 when U.S. oil production was deregulated. From 1986 onward, we use the WTI prices reported by the Energy Information Administration (EIA), while for earlier observations we use those reported in Alquist et al. (2013). Our measure of the quarterly price is the average daily price during the last month of the quarter. We use reported prices rather than the inflation-adjusted price. Given the modest level of overall price inflation volatility during this time period, inflation-adjusting our oil price forecasts would imply relatively modest effects on outcomes, whereas we gain a great deal of clarity by using reported prices. In particular, we avoid assuming that we know the inflation forecast tacitly embedded in the market’s oil price forecast.

fraction of total open interest that has a duration greater than 18 months. Before 2003, these further-dated futures contracts accounted for roughly 10 percent of open interest. However, their importance rose substantially between 2004 and 2013, averaging just under 20 percent of total open interest. With the decline in oil prices in mid-2014, far-dated futures contracts went back down to about 10 percent of total open interest. Although subject to fluctuations, the liquidity of the far-dated futures market appears sufficient to provide relevant price information.

Turning to our results, Figure 2 presents the behavior of the model’s parameter estimates over the expanding estimation window. In each of the figure’s three charts, solid black lines report the estimated coefficients while the gray intervals are two-standard-deviation confidence intervals. The results are broadly in line with the narrative of Figure 1. First, Figure 2 shows that the estimated variance coefficients do vary considerably as the estimation sample period expands and are fairly precisely estimated. Further, the estimated value of $\sigma_{\mathcal{P}}$ (middle panel) is low for the initial sample ending in 1998:Q4, implying that market participants then perceived oil prices to be mainly driven by temporary factors. As more data from the 2000s are included in the estimation sample, the estimated standard deviation of the permanent component steadily increases, peaking in the second quarter of 2008, before the financial meltdown and global recession. In contrast, the estimated value of σ_{τ} (right panel) broadly follows the opposite pattern. Interestingly, the estimated autoregressive coefficient of the temporary component, ϕ_{τ} , fluctuates less during the estimation process, remaining around a value of 0.7 throughout.

The evolution of these model parameters can also be assessed by considering the contribution of permanent shocks to the variance of Δp_t , which is reported in Figure 3. The figure shows that this contribution is very low early in the sample, with permanent shocks accounting for just 10 percent of $var(\Delta p_t)$ in the early 2000s, so that the temporary shocks were correspondingly perceived to be the main drivers of changes in spot prices. These results are very much in line with our assessment of Figure 1, where the futures curves show transitory deviations from a long-term price during the 1990s and early 2000s. However, the estimated contribution of permanent shocks steadily rises between 2002 and the first half of 2008, a point at which it accounts for more than 60 percent of $var(\Delta p_t)$. Thereafter, the sharp fall in oil prices in the last half of 2008 results in declines in that contribution to around the 35 percent mark, where it still stood as the price of oil collapsed in June 2014. This perception has generally continued

through the end of our sample, although the perceived role of permanent shocks rose slightly in early 2015.

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 (11) for each quarter from 1999:Q1 onward. Figure 4 illustrates the evolution of the estimated and actual two-year-ahead futures prices, alongside that of the spot price of oil. Comparing the actual futures prices with our model-implied futures price, the figure shows that our baseline model successfully replicates the evolution of futures prices.⁹ As with actual futures prices, note that our estimated futures prices remain well below the observed spot prices in the early 2000s, suggesting again that market participants viewed the underlying factors driving spot prices up to be mostly transitory. By the mid-2000s, our estimated futures prices move closely together with the spot price: in line with the rising estimated contribution of the permanent component to the variance in oil-price changes (Figure 3), changes in the spot price of oil are by then perceived as being mostly permanent and thus rapidly reflected in futures prices. As the financial crisis intensified in mid-2008, the spot price of oil falls rapidly, but this decline is much more pronounced than the fall in the actual futures price, which is well captured by our estimated value. Our model then tracks well the actual futures price between 2010 and the end of 2012. A gap developed afterward, as our predicted futures kept rising while the actual futures price was relatively stable; this gap was largely erased by mid-2014.

To assess the fit of our model, Figure 5 reports the cumulative mean-squared error (MSE) of fitting the futures price using the implied value from our Kalman filtering approach relative to using the previous quarter’s spot price. As is clear from the figure, the Kalman filtering approach is much better at fitting the futures price than the spot price throughout the 2000s, reaching a value as low as 0.2 and always remaining below 0.6. As visible in the figure, the model’s performance temporarily declines somewhat in 2013-2014, but the overall performance is still considerably better than the results using the spot price. Furthermore, as described below, alternative models such as the

⁹ The gray bands indicate the confidence interval, which is defined as the following set:

$$\left\{ E_t \left(p_{t+s} \mid \{p_{t-i}\}_{i=t}^1, \tilde{\Gamma}_{t-1} \right) \mid \left(\Gamma_{t-1} - \tilde{\Gamma}_{t-1} \right)' \hat{W} \left(\Gamma_{t-1} \right)^{-1} \left(\Gamma_{t-1} - \tilde{\Gamma}_{t-1} \right) \leq 6.25 \right\},$$

where $\tilde{\Gamma}_{t-1}$ is the maximum likelihood vector of parameter estimates and $\hat{W} \left(\Gamma_{t-1} \right)$ is the corresponding estimated variance-covariance matrix. The critical value of 6.25 is chosen as the 90th percentile of the chi-squared distribution with three degrees of freedom.

particle filter can improve the fit of our learning framework along that dimension.

Overall, our results indicate that learning about the persistence of underlying shocks helps capture movements in oil-price futures. Although our model is simple and only uses information from past movements in the spot price of oil, it accounts reasonably well for the fluctuations in oil-price futures over the past 15 years.

In the next three subsections, we consider extensions to our benchmark approach. We first examine the possibility that market participants may be concerned with structural breaks and thus examine a form of learning that places relatively more weight on recent observations. We then add futures prices to our estimation model. We estimate this augmented framework using the Kalman filter and then extend this model to formally address time-varying parameters. Overall, we find that the results from all of these extensions are similar to those from our benchmark model.

In a final subsection, we then close this empirical section by documenting that the movements in oil-price expectations uncovered by our approach are in line with those recovered from survey data of professional forecasters.

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 breaks and may choose to weight recent observations relatively more than distant ones. As a result, we now consider the possibility that market participants use a modified likelihood function, which is in the spirit of the recursive least squares algorithm in Evans and Honkapohja (2001).¹⁰ To this end, the likelihood function now becomes

$$LL_T = (1 - \chi_T) LL_{T-1} - \chi_T \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_t\| + (\Delta p_t - E \Delta p_t) V_t^{-1} (\Delta p_t - E \Delta p_t) \right). \quad (13)$$

If $\chi_t = \frac{1}{t}$, all observations have the same weight and the approach is equivalent to the standard likelihood function used above. In contrast, if χ_t is a constant, recent observations are then more important than lagged observations (in the learning literature, this approach is referred to as constant-gain learning, see Appendix D for details). In particular, for a data set of T observations, the first observation contributes $\prod_{t=2}^T (1 - \chi_t) \chi_1$,

¹⁰At first glance, our likelihood may not look like the algorithm in Evans and Honkapohja (2001). However, as shown in Appendix D, the two approaches are equivalent.

whereas the most recent observation, observed at time T), obtains a much greater weight of χ_T .

Different values for χ_T may be considered to conduct this exercise. As a starting point, we use information from the literature on learning and monetary policy to parameterize the gain. We set χ_T to 2 percent based on the value reported by Orphanides and Williams (2007) who estimate the (constant) gain that best fits the inflation forecasts from the Survey of Professional Forecasters. This value implies that an observation eight years in the past gets only half as much weight in the likelihood as the current observation. Figure 6 compares our baseline results with the estimate of the two-year-ahead futures price under the constant-gain filter when $\chi_T = 0.02$. The figure shows that discounting past observations at this rate generally leads to worse forecasts. Notably, the constant-gain filter significantly overpredicts oil-price futures from 2004 onward (relative to our baseline model) by predicting a peak of \$180 per barrel in the second quarter of 2008, well above the actual peak value. We also searched for the value of χ_T that results in the best model fit. For two-year-ahead futures, $\chi_T = 1.5\text{percent}$ provides the best match, a value similar to the one used by Primiceri (2006) in a model of U.S. inflation in which policymakers learn about the natural rate of unemployment using a constant-gain algorithm. However, the constant-gain estimation continues to overpredict the run-up in futures prices between 2004 and 2008 even with this optimized value (full results available on request).

4.2 Adding futures prices to the estimation

In this section, we add futures prices to the estimation equation and examine whether our benchmark results are robust to this additional information. We thus estimate our model using both spot and futures prices with the three components. Two of them, the permanent component, e_t^P , and the stationary one, e_t^T are the same as in the previous section. As previously discussed, the third component, ζ_t , captures fluctuations in a risk premium.

For futures prices, we use one-year-ahead futures. Due to data availability, our estimation period now begins in the first quarter of 1989, but still ends in the fourth quarter of 2019. Figure 7 reports the estimated model coefficients for the temporary and permanent components. Although both the sample period and data have changed, the resulting estimated model coefficients are very similar to those reported by our

benchmark model. Figure 8 reports the estimated standard deviation of ζ_t relative to the estimated values of σ_τ and $\sigma_{\mathcal{P}}$ for this two-variable model. Given that these estimated values of σ_τ and $\sigma_{\mathcal{P}}$ are very similar to the estimates from our benchmark model, it is perhaps not surprising that the estimated standard deviation of ζ_t is only a small share of σ_τ and $\sigma_{\mathcal{P}}$. Consistent with Schwartz and Smith (2000), having a permanent component, $e_t^{\mathcal{P}}$, and a stationary one, e_t^τ is enough to match the data. The third component is redundant and unnecessary.

To further illustrate this result, consider Table C, which reports the estimated coefficients for the full sample period in the first column. The table also reports the results from two other specifications that restrict respectively the value of the variance of the permanent component, $e_t^{\mathcal{P}}$, and the value of stationary one, e_t^τ to zero, whereas all other coefficients are estimated freely. The table shows that under either restriction the risk premium becomes important. However, as shown in the bottom rows of the table, the fit of these restricted models is poor relative to the unrestricted framework and one would reject these models using a standard likelihood ratio test at conventional critical values. These results suggest that using a framework of temporary and permanent shocks under learning is sufficient to broadly capture movements in oil futures prices.¹¹

4.3 Time-varying parameters: Particle filter

The three-step procedure used to derive our baseline results has the benefit of simplicity, but it also presents some potentially important limitations. It notably assumes that the model’s parameters are constant and, as a result, investors’ learning about the importance of temporary and permanent shocks is restricted by the requirement that parameter values fit the entire sample rather than just recent observations.

In this section, we address this limitation by assessing the robustness of our baseline results to a more general learning process. In particular, we consider a variant of Stock and Watson’s (2007) unobserved component model with stochastic volatility, which allows for time variation in Γ . As in the previous section, we estimate using both spot and year-ahead futures prices. Likewise, we include all three components $e_t^{\mathcal{P}}$, e_{t-1}^τ , and ζ_t , but with assumptions that the standard deviations of their innovations are themselves random variables. The estimation is conducted using the particle filter

¹¹This result contrasts with the findings of Jin (2017), who conducts a similar exercise, although without learning. Nonetheless, the estimated risk premium in her framework is relatively small (at most \$5 per barrel), which we view as largely in line with our finding.

as described in Creal (2012): Appendix A describes this alternative model and the estimation procedure in greater detail.

Figure 9 compares the expected futures prices using the Kalman filter model to those from the particle filter. Overall, the figure shows that both models have very similar predictions, especially during the 2002–05 period. In 2006 and 2007, the particle-filter-implied futures estimates are somewhat higher than futures price data, but the differences are relatively slight. In addition, the particle filter differs more from the spot price since 2013.

4.4 Professional forecasters

Our baseline model extracts oil-price expectations, from which we construct an estimated futures price path. Our results indicate that the estimated two-year-ahead futures price tracks the actual one reasonably well. Since we did not find an important role for a risk premium, our results suggest that changes in price expectations broadly capture the behavior of futures prices. In this section, we examine whether our model-based oil price expectations match direct measures from survey data. We therefore compare our estimated one-year-ahead futures prices using our benchmark model estimates to the spot price expectations from professional forecasters compiled by Bloomberg (see also Bianchi and Piana, 2016, for an analysis of this survey).

Each quarter, this survey compiles the one-year-ahead price expectations of around 100 forecasters. Using these individual forecasts, we construct the range of one-year-ahead expected spot price, which is depicted in Figure 10.¹² The figure shows that our estimated futures prices typically fall within the range of individual forecasts, suggesting that spot price expectations derived from our approach are broadly in line with those of market participants.

5 A DSGE model with learning

In the previous sections, we showed both that investors’ perceptions of the persistence of oil prices changed over the past 30 years and that these changes are captured well by

¹²The range of professional forecasts could indicate that professional forecasters use different models to predict future oil prices. To keep the analysis tractable, we abstract from modeling this heterogeneity. Similarly, as we note below, we also abstract from possible differences in the way that consumers and professional forecasters may predict future oil prices.

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 macroeconomic impact of shocks that affect the outlook for oil prices, using a DSGE model with oil usage and oil storage similar to those developed in Leduc and Sill (2004), Unalmis et al. (2012), Arseneau and Leduc (2013) or Hitzemann (2016). In contrast to this literature however, the model is constructed and solved to allow for agents' learning about the persistence of shocks. Specifically, we assume that economic agents are subject to the same information constraint faced by investors in the empirical models developed above: observed shocks may be permanent or transitory and agents use a filtering mechanism to perform that decomposition. We thus abstract from the possibility that households form expectations about future oil prices in a different manner than market participants, as suggested, for instance, by Baumeister and Kilian (2016). As such, our learning mechanism is similar to those employed by Erceg and Levin (2003), Gourinchas and Tornell (2004), Schorfheide (2005), Andolfatto et al. (2008), or Saijo (2017) to study the macroeconomic implications of learning about the persistence of shocks.¹³

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. One important aspect of our approach is the use of an oil storage model. This framework is a priori appealing for our purpose, since storage directly links current oil prices to expectations of their future values. We assume that risk-neutral speculators make the decisions related to storage, but results are similar if we follow Arseneau and Leduc (2013) and assume that households hold the oil inventories, or Hitzemann (2016) where firms hold oil inventories but evaluate future profits using households' stochastic discount factors. As is typical in the rational expectations storage literature, speculation in inventory holdings allows agents to smooth temporary volatility in the oil market. The production side of the model is composed of firms producing a consumption good using labor, capital, and oil. Using our model, we estimate and analyze three scenarios that capture salient features of our empirical results at important points during the past two decades.

¹³Our work is related to but conceptually different from the literature that arose from Evans and Honkapohja (2001) which considers that the general form of the data generating process for the economy's evolution under rational expectations is known but where exact parameter values for it are established via least-squares learning using realized economic outcomes (Milani, 2009, provides an application to the oil markets). By contrast, our contribution puts forth a view whereby imperfect signals about future shocks are the key challenge when forming expectations. Thus futures markets, which embody expectations in the model, provide crucial signals about agents' learning.

5.1 Households

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

$$U_t = \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\sigma}}{1-\sigma} - \eta \frac{n_t^{1+\varphi}}{1+\varphi} \right] \quad (14)$$

where β denotes the subjective discount factor, σ is the coefficient of relative risk aversion, η affects the disutility of labor, and φ controls the Hicksian labor supply elasticity. In turn, the composite consumption good is itself the following CES combination of final goods consumption, c_t^g , and oil consumption, c_t^o :

$$c_t = \left[(1 - \omega_c)^{1/\xi_c} c_t^g{}^{\frac{\xi_c-1}{\xi_c}} + \omega_c^{1/\xi_c} c_t^o{}^{\frac{\xi_c-1}{\xi_c}} \right]^{\frac{\xi_c}{\xi_c-1}}, \quad (15)$$

with ξ_c denoting the elasticity of substitution between final goods and oil consumption and where ω_c represents the share of oil in the consumption composite.

The price of the composite good, P_t^c , is given by the standard expression:

$$P_t^c = \left[(1 - \omega_c) P_t^g + \omega_c P_t^o{}^{1-\xi_c} \right]^{\frac{1}{1-\xi_c}}, \quad (16)$$

where P_t^g and P_t^o are the price of final goods and oil, respectively. Using the price of final goods as the numéraire, this equation expresses the relative price of the composite good, p_t^c , as a function of the relative price of oil p_t^o :

$$p_t^c = \left[(1 - \omega_c) + \omega_c p_t^o{}^{1-\xi_c} \right]^{\frac{1}{1-\xi_c}}. \quad (17)$$

Households supply labor and capital services to firms producing the final good, and the associated income from these two activities are $w_t n_t$ and $r_t k_t$, with the real wage and rental rate of capital denoted by w_t and r_t , respectively.

Households hold two types of assets: risk-free bonds, capital, and oil inventories. First, we assume that they can purchase a risk-free bond, b_{t+1} yielding a rate of return R_t . Next, capital k_t accumulates according to the usual law of motion

$$k_{t+1} = i_t + (1 - \delta)k_t, \quad (18)$$

where i_t denotes investment and δ represents the depreciation rate of capital.

Households also receive an exogenous and stochastic endowment of oil each period, \bar{o}_t , which they can sell to firms on the spot market for oil. We assume that the endowment \bar{o}_t is affected by both persistent and transitory components, as in:

$$\bar{o}_t = \bar{o} \varepsilon_{o,t}^P \varepsilon_{o,t}^\tau, \quad (19)$$

where the persistent and transitory components $\varepsilon_{o,t}^P$ and $\varepsilon_{o,t}^\tau$ follow AR(1) processes:

$$\log(\varepsilon_{o,t}^P) = \rho_o^P \log(\varepsilon_{o,t-1}^P) + u_{o,t}^P, \quad (20)$$

$$\log(\varepsilon_{o,t}^\tau) = \rho_o^\tau \log(\varepsilon_{o,t-1}^\tau) + u_{o,t}^\tau. \quad (21)$$

In this context, the household's optimization problem is to choose sequences of c_t , n_t , k_{t+1} , and b_{t+1} to maximize (14) subject to (18) and an infinite sequence of flow budget constraints given by:

$$p_t^c c_t + b_{t+1} + i_t = w_t n_t + r_t k_t + R_t b_t + p_t^o \bar{o}_t. \quad (22)$$

5.2 Production

Final Goods

Firms combine capital, labor, and oil inputs to produce the final good y_t using the production function

$$y_t = z_t \left[(1 - \omega_y)^{1/\xi_y} v a_t^{\frac{\xi_y - 1}{\xi_y}} + \omega_y^{1/\xi_y} o_t^{\frac{\xi_y - 1}{\xi_y}} \right]^{\frac{\xi_y}{\xi_y - 1}}, \quad (23)$$

where o_t represents the oil input, ω_y is the share of oil in final output, and ξ_y is the elasticity of substitution between energy and value-added $v a_t$, which itself is the following CES combination of capital and labor:

$$v a_t = \left[(1 - \omega_{va})^{1/\xi_{va}} k_t^{\frac{\xi_{va} - 1}{\xi_{va}}} + \omega_{va}^{1/\xi_{va}} n_t^{\frac{\xi_{va} - 1}{\xi_{va}}} \right]^{\frac{\xi_{va}}{\xi_{va} - 1}}, \quad (24)$$

where ω_{va} is the share of labor in value added. The neutral technology shock, z_t , is affected by both persistent and transitory components, as captured by the following

expression:

$$z_t = \bar{z} \varepsilon_{z,t}^P \varepsilon_{z,t}^\tau, \quad (25)$$

where the persistent and transitory components $\varepsilon_{z,t}^P$ and $\varepsilon_{z,t}^\tau$ follow the AR(1) processes

$$\log(\varepsilon_{z,t}^P) = \rho_z^P \log(\varepsilon_{z,t-1}^P) + u_{z,t}^P; \quad (26)$$

$$\log(\varepsilon_{z,t}^\tau) = \rho_z^\tau \log(\varepsilon_{z,t-1}^\tau) + u_{z,t}^\tau. \quad (27)$$

5.3 Storage

We assume that there is a continuum of competitive risk-neutral speculators who buy and sell oil in the spot market. In particular, they can purchase s_{t+1} units of oil to store until the next period. Holding inventories entails a per-unit cost $\phi(s_{t+1})$ in terms of oil, with $\phi'(s_{t+1}) > 0$.¹⁴ Because speculators 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, as in Unalmis et al (2012).¹⁵ We verify that this condition is met in our simulations below.

Speculators choose the amount of oil to store to maximize expected profits given by

$$E_t(p_{t+1}^o) s_{t+1} - R_t p_t^o s_{t+1} (1 + \phi(s_{t+1})). \quad (28)$$

The optimal demand for oil inventories is characterized by the following:

$$R_t p_t^o (1 + \phi(s_{t+1}) + \phi'(s_{t+1}) s_{t+1}) = E_t(p_{t+1}^o). \quad (29)$$

This expression states that speculators will accumulate oil inventories up until the marginal cost of holding one additional unit, inclusive of the opportunity cost of investing in the risk-free asset and of the cost of storage, is equal to the expected gain from holding oil for one period and selling it at next period's spot price. In turn, we can define $f_{t+k|t}$ as the futures price for delivery for any future date $t+k$. Arbitrage among speculators will be such that $f_{t+1|t} \equiv E_t p_{t+1}^o$.

¹⁴Our results are quantitatively similar if the adjustment cost is instead denoted in terms of output.

¹⁵See Williams and Wright (1991) and Arseneau and Leduc (2013) for partial and general equilibrium analyses directly tackling the non-negativity constraints on inventories.

5.4 Learning

We contrast two different information structures when solving the model. In the first, agents have full information and directly observe the components $\varepsilon_{o,t}^P$ and $\varepsilon_{o,t}^\tau$ of the shock z_t^o in equation (19) and the components $\varepsilon_{z,t}^P$ and $\varepsilon_{z,t}^\tau$ of the shock z_t in equation (25). In the second, agents observe z_t^o and z_t and know all model parameters, but they cannot distinguish the permanent and transitory components of the shocks. Consistent with our empirical results above, we assume that they use the Kalman filter to gradually learn about them. Initially, agents' perception is driven by the degree to which they see TFP-induced fluctuations in oil demand and oil-supply shocks as temporary or permanent. Under this information structure, a permanent shock occurring under a scenario in which they are relatively rare will thus only gradually be incorporated into economic agents' expectations. The appendix describes how we extend a standard solution method for rational expectations models to allow Kalman filtering of persistent and transitory shocks.

5.4.1 Equilibrium

Taking as given the exogenous variables z_t and \bar{o}_t , the equilibrium of the model is a sequence of $\{y_t, va_t, c_t, c_t^g, c_t^o, n_t, k_t, o_t, s_t, w_t, p_t^c, r_t\}$ that satisfy the household optimality conditions; the optimality conditions for firms producing final and consumption goods; **the optimality condition for speculators**; the bond market clearing condition, as well as the oil-market clearing condition $(1 + \phi(s_{t+1}))s_{t+1} - s_t + c_t^o + o_t = \bar{o}_t$ and the resource constraint $c_t^g + i_t = y_t$.

5.5 Calibration and estimation results

We adopt a flexible empirical strategy in which we calibrate some parameters and estimate others using a simulated method of moments. We calibrate the model to a quarterly frequency and accordingly set the discount factor β and the depreciation rate δ to 0.99 and 0.025, respectively. The parameter σ , controlling the degree of intertemporal substitution in consumption, is 2, and the Hicksian elasticity of labor supply φ is 1. The scaling parameter η is set in order for the labor input in the economy's steady state to be 1.

We calibrate the parameters related to the production and storage of oil. Following Unalms et al. (2012), we adopt the storage cost $\phi(s) = \kappa + \frac{\phi}{2}s$, where the constant

κ can be interpreted as a convenience yield which we set to 5 percent and where ϕ is calibrated such that the steady-state stock of oil stored as a ratio of total (quarterly) output is 50 percent, also as in Unalmis et al. (2012).

We calibrate the production function parameters as follows. We set $\omega_{va} = 0.66$ and $\xi_{va} = 1.0$, which matches the values used in standard business cycle models where substitution between capital and labor in production is unity and the capital share in final output is around one-third. We set $\omega_y = \omega_c = 0.0585$ and $\xi_y = \xi_c = 0.1$, which ensures that the elasticity of substitution between oil and other inputs is small, at 0.1, and that the steady-state value of oil used in production and in composite consumption is about 4 percent, the values also used in Bodenstein, Erceg, and Guerrieri (2011).

We set ρ_z^P and ρ_o^P arbitrarily close to 1, such that the permanent components of the technology and oil-supply shocks are (near) random walks. We then fix σ_z^τ to an arbitrary value (1 percent), since only the relative importance of transitory versus permanent shocks matter for the analysis.

The remaining parameters are estimated using a simulated method of moments approach. We thus choose the parameters of the technology shock process, ρ_z^τ , σ_z^P and those of the oil-supply shock process, ρ_o^τ , σ_o^P , σ_o^τ to match the following five moments: The volatility of oil inventories (relative to the volatility of real GDP), the correlation between oil prices and oil inventories, the correlation between oil prices and real GDP, the autocorrelation of GDP, and the fraction of the volatility of the changes in spot prices explained by permanent shocks from our Kalman filter estimate. The first four moments are calculated after filtering the log of each series, while the last moment uses the log of the spot price of oil.

In matching these moments, we consider three scenarios, labeled *Early 2000s*, *Late 2000s*, and *Mid-2010s*. For the *Early 2000s*, we conduct the moment matching exercise over the period 1980 to 2002. We then use data from 1980 to 2008 to estimate the parameters under the *Late 2000s* scenario, while for the *Mid 2010s* scenario, we use the data sample from 1980 to 2014.¹⁶

Table C reports the results for the three scenarios. Focusing on the moments to match, the first panel shows that, for the *Early 2000s* scenario, the perceived importance of permanent shocks in oil price fluctuations was small at around 10 percent.

¹⁶For each historical episode, estimates are arrived at as follows. Starting for initial guesses, the Simplex method is used to minimize the discrepancies between targeted moments (conditional on the episode) and those in simulated data, which consist of 500 sample economies of 250 periods' each.

It also indicates that the volatility of oil inventories was about six times that of real GDP. Although oil prices are roughly acyclical and not correlated with GDP, oil inventories are substantially negatively correlated with oil prices. This negative correlation is consistent with inventories increasing when prices are temporarily low and decreasing when inventories are temporarily high. Across the three scenarios and consistent with our empirical results, we observe an increase in the perceived importance of permanent shocks in driving oil price fluctuations to about 50 percent during the late 2000s, before decreasing somewhat in the mid 2010s. The correlation between oil prices and real GDP also rises to about 30 percent in the mid-2010s, while inventories become slightly more negatively correlated with oil prices. GDP is highly autocorrelated under each scenario.

The table shows that our matching exercise can capture the broad patterns in the data. For instance, it is consistent with a rise in the perceived importance of permanent shocks as drivers of oil price fluctuations during the 2000s and the slight decline in the mid-2010s. Our estimation also captures the decline in the volatility of inventories over time, although the estimated volatility for the *Early 2000s* scenario is somewhat higher than in the data. Oil prices in our model become increasingly more procyclical across the three scenarios, while inventories are negatively correlated with spot prices as in the data. Overall, we view the resulting simulated moments as broadly consistent with the observed empirical evidence.

Table 3 then reports the estimated values of the five parameters driving the shock processes across the different scenarios. Comparing the *Early 2000s* to the *Late 2000s* scenarios, the striking change in the estimated parameter values is the large increase in the estimated standard deviation of the permanent component of TFP shocks, which rises from 0.1 percent in the early 2000s to 0.6 percent in the late 2000s. We also obtain a similar, though more modest, increase in the estimated volatility of the permanent component of oil-supply shocks. In contrast, the estimated volatility of the transitory component of TFP shocks declines only slightly across the two scenarios. The estimated persistence of transitory shocks to either TFP or oil supply are roughly constant. As we move to the *Mid 2010s* scenario, the estimated volatility of the permanent components of TFP shocks declines substantially, while the volatility of temporary TFP shocks is estimated to have risen.

An alternative way to synthesize these results is by considering the implications for the gain in the Kalman filter. When a shock hits the economy, the gain determines

how agents interpret the source of the disturbance, i.e., whether it is due to TFP-induced fluctuations in oil demand or the supply of oil, and its persistence, i.e., whether temporary or permanent. Table C reports the changes in the gains across the three scenarios. First, it shows that oil-supply shocks are viewed as being mostly transitory under each scenario. As a result, a realized permanent oil-supply shock will be highly unexpected and learning about its persistence will be gradual. In contrast, the table also indicates that agents view TFP shocks very differently in the three scenarios. Although these innovations are largely considered transitory in the early 2000s, agents' perceptions dramatically changed in the late 2000s, viewing TFP shocks as largely permanent. By the mid-2010s, agents placed an equal probability that TFP shocks were either transitory or permanent.

Overall, our procedure highlights a substantial rise in the persistence of TFP shocks that drove a persistent increase in economic activity and the demand for oil, which households and investors learned over time. The emphasis on the important of an oil-demand channel, driven by an aggregate shock, is broadly in line, for instance, with the analysis in Kilian (2009).

6 Macro effects of learning

We now examine the response of the economy to a permanent shock, $\varepsilon_{z,t}^o$, to oil supply and to TFP-driven oil demand (i.e., a permanent shock to TFP, $\varepsilon_{z,t}^P$, which boosts the demand for all inputs, including oil). We consider shocks whose amplitude entails a 10 percent increase in oil prices under full information and compare the effects of the shocks under full information to those arising when agents do not directly observe the shocks and must instead infer their persistence via Kalman filtering. Under this learning mechanism, agents gradually reassess their views regarding the shock's persistence over time, which in turn informs their forward-looking decisions.

To demonstrate the model's mechanism, Figures 11 first contrast the transmission of permanent and transitory oil-supply shocks, while 12 shows the response of the economy to a permanent and a transitory TFP-driven increase in oil demand, assuming full information. This comparison helps provide intuition for the effects of learning, since the consequence of learning can be interpreted as an endogenous convex combination between the impact of permanent and transitory shocks.

Figure 11 shows that the permanent oil-supply shock lowers GDP, investment, and

consumption, with oil inventories declining following the rise in the price of oil. A transitory shock has similar qualitative effects, although the responses are muted because of the larger inventory decline, which mutes the rise in oil prices. In turn, the oil-price basis (i.e., the difference between the spot and futures price) increases more following the transitory oil-supply increase than in response to a permanent one.

As shown in Figure 12, the relative impact of transitory and permanent fluctuations in TFP-driven increase in oil demand differs from those arising from the supply side of the oil market. For instance, the transitory shock amplifies the GDP and investment responses compared to those following a permanent shock. This partly reflects a more pronounced increase in labor supply (not shown), which boosts the rate of return on capital and investment relatively more. It also results from a decline in oil inventories that mutes the rise in oil prices. In line with the greater increase in income, consumption rises relatively more following a permanent TFP shock. Still, the oil-price basis rises following a transitory shock, while it declines in response to a permanent one.

We now examine the effect of learning on the economic responses following permanent shocks to the supply of oil and to TFP-driven increases in the demand for oil. We compare the full-information responses to those arising under learning in the *Early 2000s* and *Late 2000s* scenarios. Recall that the shocks are misperceived to be temporary as estimated under our *Early 2000s* scenario and are more rapidly identified as being permanent under the *Late 2000s* scenario.

The impact of learning following an oil-supply shock is described in Figure 13. The figure shows that, when agents misperceive the shock's persistence as being transitory, the decline in GDP is substantially smaller, by about one-half under the *Early 2000s* scenario, than when agents have full information. The responses of investment and consumption are similarly more muted. This partly reflects the more substantial decline in inventories and the associated smaller rise in oil prices that occur under incomplete information. In addition, the spot price of oil rises persistently above the one-year-ahead futures price, a feature observed empirically during the early 2000s.

Figure 14 describes the responses to a TFP-driven increase in oil demand under learning and full information. It shows that when agents misperceived the permanent shock to be transitory, the effects on GDP and investment are amplified. Under the *Early 2000s* scenario, when agents put little weight on the possibility that shocks are permanent, these effects are particularly strong and persistent. Boosted by the increase in economic activity, the misperceived shock ultimately leads to a substantial rise in oil

prices and in the oil-price basis, again consistent with developments in the early 2000s shown in Figure 1. In line with our empirical evidence and estimated parameters, agents perceived permanent shocks to be much more likely under the *Late 2000s* scenario, resulting in more muted differences with the full information responses. In addition, spot prices do not rise as much above futures prices in this case.

To complement these impulse responses, we highlight the importance of learning for macroeconomic volatility through counterfactual simulations. In the empirical section, we documented a gradual revision in the perceived persistence of oil price fluctuations. We documented that this learning process helped account, for instance, for the lack of movements in futures prices in the early 2000s despite large increases in spot prices. To examine the impact of this gradual learning process for the macroeconomy, we examine a polar case in which agents maintain their priors despite changes in the structural shocks. The first column of Table 5 reports the counterfactual output and oil price volatilities if economic agents had kept their priors constant in the face of changing shock properties during the 2000s, thus missing the increased persistence of TFP-induced oil-demand fluctuations. To do so, we simulate the economy with the estimated shock processes under the *Late 2000s* scenario, assuming, however, that agents perceived oil-price movements to be as temporary as under the *Early 2000s* scenario. We report the volatilities of this counterfactual as a ratio to the volatility under our benchmark model with learning under the *Late 2000s* scenario. The results first indicate that futures prices experience a decline in volatility of roughly 15 percent, as fluctuations are too likely to be misperceived as temporary. Importantly, it also reports a substantial impact of output volatility, which increases by nearly 40 percent.

The second column of the table reports the results of a similar experiment for the *Mid 2010s* episode. In this case, we simulate the economy with the estimated shock processes under the *Mid 2010s* scenario, assuming, however, that agents perceive oil-price movements to be as persistent as under the *Late 2000s* scenario. We report the volatilities of this counterfactual exercise as a ratio to the volatility under our benchmark model with learning under the *Mid 2010s* scenario. In contrast to the previous counterfactual, oil prices are now misperceived to be too persistent. As a result, this misperception leads to a 15-percent decline in output volatility, while futures prices fluctuate more.¹⁷

¹⁷In addition to the expectations channel that we emphasize, the empirical effects of shocks to the demand and supply of oil during the mid-2010s is further altered by the rise in shale gas production in

The impact of learning on macroeconomic outcomes also stem from an interaction between agents' expectations, which embed their views about the shocks' persistence, and the storage capabilities of the economy, which allow the impact of transitory shocks to be smoothed. Figure 15 shows the importance of this interaction. The figure contrasts the full-information and learning responses in the economy with storage (left column of the figure) with those arising when agents are unable to store oil (right column). Following an oil-supply shock, the figure shows that the ability to store oil is particularly important for dampening the fall in output during the quarters immediately following the shock: the left-hand panels show that GDP falls substantially less when inventories are available to smooth the effects of the shocks and the rise in oil prices is concurrently more modest. A similar pattern occurs following a TFP-driven increase in oil demand, though the presence of inventories mutes the impact on GDP to a lesser extent in this case. Similarly, the last of column of Table 6 reports volatility ratios for GDP and oil prices in the benchmark model relative to those in the model without storage. It shows that introducing storage reduced output volatility by 8 percent, while oil price futures are 13 percent less volatile in this case. Overall, learning partly impacts the macroeconomic transmission of shocks through its effects on forward-looking storage decisions.

7 Conclusion

Our analysis highlights that learning about the persistence of shocks underlying oil-price movements is consistent with developments in the oil market and has important macroeconomic implications. We first provided an analysis of movements in oil-price futures since 2000 based on learning. We showed that a simple model with unobserved components in which investors must form beliefs about the persistence of changes in oil prices accounts well for the fluctuations in oil-price futures. Our 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 volatile period during which oil prices reached historic highs, the model predicts a level of futures prices broadly in line with the data. Our estimates suggest that through learning investors revised up their estimate of the contribution of permanent

the United States. As emphasized by Baumeister and Kilian (2015), this change affected the impact of oil-price changes on the U.S. economy via changes to its effects on U.S. investment.

shocks to the variance of oil prices over this period. Similarly, our results suggest that throughout 2014 investors perceived the dramatic decline in oil prices as somewhat temporary. Importantly, we show the importance of learning for oil markets is robust to allowing for the presence of a risk premium and for time-varying volatility.

Using a DSGE model in which oil is storable and used in production, we estimated that changes in TFP-driven fluctuations in oil demand largely accounted for the changes in investors' perceptions. In addition, we showed that misperceptions about the persistence of shocks underlying oil price movements had large macroeconomic effects. Because of learning, failing to recognize the increased persistence of TFP-driven fluctuations in oil demand substantially increases the volatility of output and oil prices. As such, our analysis highlights expectations formation and the importance of learning about the persistence of shocks underlying oil-price movements for macroeconomic fluctuations.

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), pp. 539-573.
- Alquist, Ron, and Olivier Gervais, 2013. "The Role of Financial Speculation in Driving the Price of Crude Oil," *The Energy Journal* 34(3), pp. 35-54.
- Alquist Ron, Lutz Kilian and Robert J. Vigfusson, 2013. "Forecasting the Price of Oil," in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting* 2, pp. 427-507.
- Andolfatto, David, Scott Hendry, and Kevin Moran, 2008. "Are inflation expectations rational?" *Journal of Monetary Economics* 55(2), pp. 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 an Application to Crude Oil," Manuscript.
- Baumeister, Christiane and Lutz Kilian, 2016. "Understanding the Decline in the Price of Oil since June 2014," *Journal of the Association of Environmental and Resource Economists*, 3(1), pp. 131-158.
- Beaudry, P. and F. Portier, 2004. "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics* 51, pp. 1183-1216.
- Bianchi, Daniele and Jacopo Piana, 2016. "Expected Spot Prices and the Dynamics of Commodity Risk Premia," Manuscript.
- Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri, 2011. "Oil shocks and external adjustment," *Journal of International Economics* 83(2), pp. 168-184.

Buyuksahin, B., M. Haigh, J. Harris, J., J. Overdahl, and M. Robe 2008. *Fundamentals, Trader Activity and Derivative Pricing*, Manuscript.

K. Brunner, Cukierman, A., and A. Meltzer. 1980 Stagflation, persistent unemployment and the permanence of economic shocks. *Journal of Monetary Economics*, 6(1), pp. 467-492.

O. Coibion, Y. Gorodnichenko and R. Kamdar. 2018 The Formation of Expectations, Inflation and the Phillips Curve. *Journal of Economic Literature*, 56(4), pp 1447-91.

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

Croce, M. M. 2014. Long-run productivity risk. A new hope for production-based asset pricing? *Journal of Monetary Economics*, 66(1), pp. 13-31.

Erceg, Christopher and Andrew T. Levin, 2003. "Imperfect credibility and inflation persistence," *Journal of Monetary Economics* 50(4), pp. 915-944.

Evans, George W. and Seppo Honkapohja, 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton,

Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva, 2013. "The Role of Speculation in Oil Markets: What Have We Learned So Far?" *Energy Journal* 34(3), pp. 7-33.

Fornero, Jorge and Markus Kirchner, 2014. "Learning About Commodity Cycles and Saving- Investment Dynamics in a Commodity-Exporting Economy," Working Papers Central Bank of Chile 727, Central Bank of Chile.

Hamilton, James D., 1994. *Time Series Analysis*. Princeton University Press, Princeton,

Hamilton, James D. and Jing Cynthia Wu, 2014. "Risk Premia in Crude Oil Futures Prices," *Journal of International Money and Finance* 42(C), pp. 9-37.

Hitzemann, Steffen, 2016. "Macroeconomic Fluctuations, Oil Supply Shocks, and Equilibrium Oil Futures Prices," Manuscript, Rutgers Business School. September 2016.

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

Jin, Xin, 2017. “The Role of Market Expectations in Commodity Price Dynamics: Evidence from Oil Data,” University of Aberdeen Discussion Paper in Economics No. 17-7. September 2017.

Kellogg, Ryan, 2014. “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review* 104(6), pp. 1698-1734.

Kilian, Lutz, 2009. “Not All Oil Prices Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review* 99(3), pp. 1053-1069.

Kilian, Lutz, and Christiane Baumeister, 2016. “Lower Oil Prices and the U.S. Economy: Is This Time Different?,” *Brookings Papers on Economic Activity* Fall 2016, pp. 287-357.

Kilian, Lutz, and Bruce Hicks, 2013. “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003–2008?” *Journal of Forecasting* 32, pp. 385–394.

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(3), pp. 454-478.

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

Leduc, Sylvain and Keith Sill, 2004. “A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns,” *Journal of Monetary Economics* 51(4), pp. 781-808.

Milani, Fabio, 2009. “Expectations, learning, and the changing relationship between oil prices and the macroeconomy,” *Energy Economics* 31(6), pp. 827-837.

Ormeno, Arturo and Krisztina Molnar, 2015. “Using Survey Data of Inflation Expectations in the Estimation of Learning and Rational Expectations Models,” *Journal of Money, Credit and Banking* 47(4), pp. 673-699.

Orphanides, Athanasios and Simon van Norden, 2005. “The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time,” *Journal of Money, Credit and Banking* 37(3), pp. 583-601.

- Orphanides, Athanasios and John C. Williams, 2007. "Robust monetary policy with imperfect knowledge," *Journal of Monetary Economics* 54(5), pp. 1406-1435
- Primiceri, Giorgio E., 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy," *The Quarterly Journal of Economics* 121(3), pp. 867-901.
- Saijo, Hikaru, 2017. "The uncertainty multiplier and business cycles," *Journal of Economic Dynamics and Control* 78, pp. 1-25
- Schorfheide, Frank, 2005. "Learning and monetary policy shifts," *Review of Economic Dynamics* 8(2), pp. 392-419
- Schwartz, Eduardo, and James E. Smith 2000. "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science* 46(7), pp. 893-911.
- Singleton, Kenneth J. 2014. "Investor Flows and the 2008 Boom/Bust in Oil Prices," *Management Science* 60(2), pp. 300-318.
- Slobodyan, Sergey and Raf Wouters, 2012. "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," *American Economic Journal: Macroeconomics* 4(2), pp. 65-101.
- Sockin, M. and W. Xiong, 2015. "Informational Frictions and Commodity Markets," *Journal of Finance* 70(5), pp. 2063-2097.
- Stock, J. H. and Mark Watson, 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39(1), pp. 3-33.
- Unalmis, Deren, Ibrahim Unalmis, 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.

A Particle Filter

A.1 Model and estimation

In this section, we generalize from our estimated baseline model and formally model time variation in the learning process. In particular, we consider a variant of Stock and Watson's (2007) unobserved components model with stochastic volatility, in which we introduce an additional temporary shock to the level of oil prices. The model is thus similar to our baseline framework but allows for time variation in the parameter vector Γ . Specifically, the log spot price of oil is, as before,

$$p_t = e_t^{\mathcal{P}} + e_t^{\tau}, \quad (30)$$

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

$$e_t^{\mathcal{P}} = e_{t-1}^{\mathcal{P}} + v_t. \quad (31)$$

In contrast to our baseline framework however, we assume that the disturbance v_t is Gaussian with time-varying variance $\sigma_{\mathcal{P},t}^2$: $v_t \sim N(0, \sigma_{\mathcal{P},t}^2)$. Moreover, we postulate that $\sigma_{\mathcal{P},t}^2$ evolves according to

$$\ln \sigma_{\mathcal{P},t}^2 = \ln \sigma_{\mathcal{P},t-1}^2 + \xi_{\mathcal{P},t}, \quad (32)$$

where $\xi_{\mathcal{P},t}$ is a Gaussian disturbance with zero mean and variance $\sigma_{\xi_{\mathcal{P}}}^2$.

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

$$e_t^{\tau} = \phi_{\tau} e_{t-1}^{\tau} + \varepsilon_t; \quad (33)$$

although we fix ϕ_{τ} at a constant 0.7, the variance of $\varepsilon_t \sim N(0, \sigma_{\tau,t}^2)$ is time-varying, with the following equation

$$\ln \sigma_{\tau,t}^2 = \ln \sigma_{\tau,t-1}^2 + \xi_{\tau,t}, \quad (34)$$

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

Following Schwartz and Smith (2000), we further allow for the presence of measurement error in the pricing of oil futures, which could capture errors in reporting or

deviations between our model’s fit and observed prices. We accordingly specify

$$f_t^k = E p_{t+k} + \xi_t, \quad (35)$$

where the measurement error term ξ_t is assumed to be independently and identically normally distributed with zero mean and time-varying variance $\sigma_{\xi,t}^2$, which evolves according to

$$\ln \sigma_{\xi,t}^2 = \ln \sigma_{\xi,t-1}^2 + \xi_{\Phi,t}, \quad (36)$$

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

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

$$\Delta p_t = v_t + (\phi_\tau - 1) e_{t-1}^\tau + \varepsilon_t; \quad (37)$$

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

$$f_t^k - p_t = (\phi_\tau^k - 1) (\phi_\tau e_{t-1}^\tau + \varepsilon_t) + \xi_t; \quad (38)$$

as well as (31), (32), (33), (34), and (36). To better discipline the particle filter, we complement the use of the price of West Texas Intermediate oil used for estimating our baseline model with the nine-month-ahead futures price. We then use the model’s estimates to forecast two-year-ahead futures prices. Because we are limited by the availability of nine-month-ahead futures contracts, our estimation period begins in 1989:Q1. The sample still ends in 2015:Q4. The estimation is conducted using the particle filter as described in Creal (2012), which we describe in the next section.

A.2 Filtering

This appendix briefly describes our use of the particle filter, which draws on the survey by Creal (2012). A more thorough and advanced treatment is available in Creal (2012).

Given the observed data y_t

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

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

$$x_t = \left(\ln \sigma_p^2 \quad \ln \sigma_{\tau,t}^2 \quad \ln \sigma_{\xi,t}^2 \quad v_t \quad e_{t-1}^\tau \quad \varepsilon_t \quad \xi_t \right).$$

To construct estimates of x_t , we will attempt to draw out the distribution. We begin with a set of model parameters θ (below these are estimated to maximize likelihood):

$$\theta = \left(\sigma_{\xi p}^2 \quad \sigma_{\xi \tau}^2 \quad \sigma_{\xi \oplus}^2 \quad \phi_\tau \right).$$

Now start with a set of initial particles x_0^i , where $i = 1$ to N , which are drawn randomly from a distribution conditional on θ . For the stationary variables, we draw from the unconditional distribution. For the nonstationary variables, we choose random variables that are reasonable approximations.

$$\begin{aligned} \ln \sigma_{\tau,t}^2 &= v_t^\tau \quad \text{where } v_t^\tau \sim N(0, \sigma_\tau^2) \\ \ln \sigma_{p,t}^2 &= v_t^p \quad \text{where } v_t^p \sim N(0, \sigma_p^2). \end{aligned}$$

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 distributed with the conditional mean given

above. The variance covariance matrix is

$$N \left(\begin{pmatrix} \Delta p_t - [v_t + (\phi_\tau - 1) e_{t-1}^\tau] \\ f_t^k - p_t - [(\phi_\tau^k - 1) \phi_\tau e_{t-1}^\tau] \end{pmatrix}, \begin{pmatrix} \sigma_{\tau,t}^2 & \phi^k \sigma_{\tau,t}^2 \\ \phi^k \sigma_{\tau,t}^2 & \phi^{2k} \sigma_{\tau,t}^2 + \sigma_{\xi,t}^2 \end{pmatrix} \right)$$

We define

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

Note that, given the model structure here, these weights w_t^i are simpler than the general expression for w_t^i given in Creal. We normalize the weights

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

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

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

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.

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

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

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

We then move on to the next observation.

B Model Solution

We extend King and Watson's (2002) solution method to allow Kalman filtering of persistent and transitory shocks. Denote dynamic (d_t) and flow (f_t) variables as functions of exogenous expected future variables x_t , as follows (King and Watson, Sec. 4):

$$E_t d_{t+1} = W d_t + E_t [\Psi_d(F) x_t]; \quad (39)$$

$$f_t = -K d_t - E_t [\Psi_f(F) x_t]; \quad (40)$$

with $\Psi_d(F)$ and $\Psi_f(F)$ matrix polynomials in the forward operator (i.e, $F x_t = x_{t+1}$). The dynamic variables d_t can be further separated into nonpredetermined and predetermined variables labeled λ_t and k_t , respectively, so that $d_t = [\lambda_t k_t]'$.

King and Watson use the following decomposition of the matrix W in (39):

$$V_u W = \mu V_u, \quad (41)$$

where μ is a lower-triangular matrix with the unstable eigenvalues of W on its diagonal. Next, they define $u_t \equiv V_u d_t$, or, since $d_t = [\lambda_t k_t]'$,

$$u_t = V_{u\lambda} \lambda_t + V_{uk} k_t. \quad (42)$$

Finally, they apply V_u to (39) and use the definition of u_t and (39) to obtain:

$$E_t u_{t+1} = \mu u_t + V_u \Psi_d(F) E_t(x_t). \quad (43)$$

Meanwhile, the exogenous variables x_t evolve as

$$x_t = \Theta \xi_t, \quad (44)$$

$$\xi_t = \rho \xi_{t-1} + \theta \eta_t, \quad \eta_t \sim (0, Q); \quad (45)$$

where η_t is a martingale difference sequence.

B.1 Solution when shocks in ξ_t are observed

For reference, we first describe how King and Watson (2002) solve the model when all shocks ξ_t are observed. The system (44)–(45) is used to evaluate all expressions

depending on expected future values of x_t . Notably, the last part of expression (43) becomes

$$\begin{aligned}
V_u \Psi_d(F) E_t(x_t) &= V_u \Psi_{d,0} x_t + V_u \Psi_{d,1} E_t(x_{t+1}) + V_u \Psi_{d,2} E_t(x_{t+2}) + \dots \\
&= V_u [\Psi_{d,0} \Theta + \Psi_{d,1} \Theta \rho + \Psi_{d,2} \Theta \rho^2 + \dots] \xi_t \\
&\equiv \varphi_u \xi_t;
\end{aligned} \tag{46}$$

while the last part of (40) becomes

$$\begin{aligned}
\Psi_f(F) E_t(x_t) &= \Psi_{f,0} x_t + \Psi_{f,1} E_t(x_{t+1}) + \Psi_{f,2} E_t(x_{t+2}) + \dots \\
&= [\Psi_{f,0} \Theta + \Psi_{f,1} \Theta \rho + \Psi_{f,2} \Theta \rho^2 + \dots] \xi_t \\
&\equiv \varphi_f \xi_t.
\end{aligned} \tag{47}$$

Eigenvalues of μ are unstable so (43) can be solved forward, equation by equation and yield

$$u_t = \nu \xi_t, \tag{48}$$

where ν is a function of the coefficients in φ_u , in μ , and in ρ . Once u_t is solved as a function of ξ_t , the remainder is straightforward: using (42) allows us to compute λ_t as a function of the predetermined variables k_t and the exogenous shocks ξ_t :

$$\lambda_t = -V_{u\lambda}^{-1} V_{uk} k_t + V_{u\lambda}^{-1} \nu \xi_t, \tag{49}$$

and the stable part of (39) allows us to compute the dynamic evolution of predetermined variables k_t . Finally, (40) is used to solve for f_t , and the complete solution reads as

$$\begin{aligned}
\begin{bmatrix} f_t \\ \lambda_t \\ k_t \\ x_t \end{bmatrix} &= \begin{bmatrix} \Pi_{fk} & \Pi_{f\xi} \\ \Pi_{\lambda k} & \Pi_{\lambda\xi} \\ I & 0 \\ 0 & \Theta \end{bmatrix} \begin{bmatrix} k_t \\ \xi_t \end{bmatrix} \\
\begin{bmatrix} k_{t+1} \\ \xi_{t+1} \end{bmatrix} &= \begin{bmatrix} M_{kk} & M_{k\xi} \\ 0 & \rho \end{bmatrix} \begin{bmatrix} k_t \\ \xi_t \end{bmatrix} + \begin{bmatrix} 0 \\ \theta \end{bmatrix} \eta_{t+1}.
\end{aligned}$$

B.2 Solution when x_t is observed but not ξ_t

Consider now that only the composite shock x_t in (44) is observable and that its decomposition into components of ξ_t is not. Knowledge about (44)–(45) can still be used to infer probable values for ξ_t through the application of the Kalman filter on the observed values of x_t . In turn, these inferences can be used to compute $E_t[x_{t+h}]$ for any h . To this end, denote the best estimate of ξ_t based on information up to time t as $\widehat{\xi}_t|t$ and the best linear forecast for ξ_{t+1} as $\widehat{\xi}_{t+1}|t$. Further, $P_{t+1}|t$ is the mean-squared error of that forecast. The following recursions obtain for these quantities:

$$\widehat{\xi}_t|t = \widehat{\xi}_t|t-1 + K_t \left(x_t - \Theta \widehat{\xi}_t|t-1 \right); \quad (50)$$

$$\widehat{\xi}_{t+1}|t = \rho \widehat{\xi}_t|t = \rho \widehat{\xi}_t|t-1 + \rho K_t \left(x_t - \Theta \widehat{\xi}_t|t-1 \right); \quad (51)$$

$$K_t = P_t|t-1 Q' (Q P_t|t-1 Q')^{-1}; \quad (52)$$

$$P_{t+1}|t = (\rho - K_t \Theta) P_t|t-1 (\rho' - \Theta' K_t') + Q; \quad (53)$$

Before proceeding, notice that (50) can be rewritten

$$\begin{aligned} \widehat{\xi}_t|t &= K_t \Theta \widehat{\xi}_t + [I - K_t \Theta] \widehat{\xi}_t|t-1; \\ &= A \widehat{\xi}_t + B \widehat{\xi}_t|t-1 \end{aligned} \quad (54)$$

We can now use (50)–(53) in computations involving expectations of future values of exogenous shocks above: (46) and (47) become

$$\begin{aligned} V_u \Psi_d(F) E_t(x_t) &= V_u \Psi_{d,0} x_t + V_u \Psi_{d,1} E_t(x_{t+1}) + V_u \Psi_{d,2} E_t(x_{t+2}) + \dots \\ &= V_u \Psi_{d,0} \Theta \xi_t + V_u \Psi_{d,1} \Theta \widehat{\xi}_{t+1}|t + V_u \Psi_{d,2} \Theta \widehat{\xi}_{t+2}|t + \dots \\ &= V_u \Psi_{d,0} \Theta \xi_t + V_u \Psi_{d,1} \Theta \rho \widehat{\xi}_t|t + V_u \Psi_{d,2} \Theta \rho^2 \widehat{\xi}_t|t + \dots \\ &\equiv \varphi_{d,\xi} \xi_t + \varphi_{d,\widehat{\xi}} \widehat{\xi}_t|t; \end{aligned} \quad (55)$$

$$\begin{aligned} \Psi_f(F) E_t(x_t) &= \Psi_{f,0} x_t + \Psi_{f,1} E_t(x_{t+1}) + \Psi_{f,2} E_t(x_{t+2}) + \dots \\ &= \Psi_{f,0} \Theta \xi_t + \Psi_{f,1} \Theta \widehat{\xi}_{t+1}|t + \Psi_{f,2} \Theta \widehat{\xi}_{t+2}|t + \dots \\ &= \Psi_{f,0} \Theta \xi_t + \Psi_{f,1} \Theta \rho \widehat{\xi}_t|t + \Psi_{f,2} \Theta \rho^2 \widehat{\xi}_t|t + \dots \\ &\equiv \varphi_{f,\xi} \xi_t + \varphi_{f,\widehat{\xi}} \widehat{\xi}_t|t. \end{aligned} \quad (56)$$

Again, use (54) and (55) to help solve (43) forward, equation by equation to get

$$u_t = \nu_\xi \xi_t + \nu_{\widehat{\xi}} \widehat{\xi}_t|_{t-1}. \quad (57)$$

The solution for u_t allows us to express λ_t as a function of predetermined variables k_t and exogenous shocks ξ_t , as follows:

$$\lambda_t = -V_{u\lambda}^{-1} V_{uk} k_t + V_{u\lambda}^{-1} \nu_\xi \xi_t + V_{u\lambda}^{-1} \nu_{\widehat{\xi}} \widehat{\xi}_t|_{t-1}. \quad (58)$$

Further the “stable” part of (39) is again used to find a dynamic solution for the evolution of the predetermined variables k_t ; and (40) can then be used to find the solution for f_t . In the end, the complete solution has the following form:

$$\begin{bmatrix} f_t \\ \lambda_t \\ k_t \\ x_t \end{bmatrix} = \begin{bmatrix} \Pi_{fk} & \Pi_{f\xi} & \Pi_{f\widehat{\xi}} \\ \Pi_{\lambda k} & \Pi_{\lambda\xi} & \Pi_{\lambda\widehat{\xi}} \\ I & 0 & 0 \\ 0 & \Theta & 0 \end{bmatrix} \begin{bmatrix} k_t \\ \xi_t \\ \widehat{\xi}_t \end{bmatrix}$$

$$\begin{bmatrix} k_{t+1} \\ \xi_{t+1} \\ \widehat{\xi}_{t+1} \end{bmatrix} = \begin{bmatrix} M_{kk} & M_{k\xi} & M_{k\widehat{\xi}} \\ 0 & \rho & 0 \\ 0 & \rho A & \rho B \end{bmatrix} \begin{bmatrix} k_t \\ \xi_t \\ \widehat{\xi}_t \end{bmatrix} + \begin{bmatrix} 0 \\ \theta \\ 0 \end{bmatrix} \eta_{t+1}.$$

Note that, in effect, we have added $\widehat{\xi}_t|_t$ as an additional vector of dynamic variables to the original solution and taken into account its impact on the other variables through the matrices $\Pi_{f\widehat{\xi}}$, $\Pi_{\lambda\widehat{\xi}}$ and $M_{k\widehat{\xi}}$. Further details about the solution method can be obtained from the authors.

C Constant-gain learning: Derivation

This appendix demonstrates that for estimating a linear regression that maximizing the following likelihood function

$$LL_T = (1 - \chi_T) LL_{T-1} - \chi_T \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_t\| + (p_t - Ep_t) V_t^{-1} (p_t - Ep_t) \right) \quad (59)$$

is equivalent to least squares learning as described in Evans and Honkapohja (2001).

In this section, we adopt the notation of Evans and Honkapohja. For the following linear model

$$y_i = c'x_i + \varepsilon,$$

the recursive least squares algorithm in Evans and Honkapohja is defined as

$$\begin{aligned} c_t &= c_{t-1} + \kappa_t R_t^{-1} x_{t-1} \left(y_t - c'_{t-1} x_{t-1} \right); \\ R_t &= R_{t-1} + \kappa_t \left(x_{t-1} x'_{t-1} - R_{t-1} \right). \end{aligned}$$

Note that the algorithm is first defined as

$$\kappa_t = t^{-1}$$

which is standard OLS, but a generalization with a constant gain is also discussed, viz,

$$\kappa_t = \kappa.$$

We need to prove that the solution to maximizing the recursive likelihood function is equivalent to this recursive least squares algorithm. As such, we first defined the likelihood at time 1 as

$$LL_1 = -\chi_1 \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_1\| + (p_1 - Ep_1) V_1^{-1} (p_1 - Ep_1) \right),$$

we can show that our equation is equivalent to the following weighted likelihood function:

$$LL_T = -\sum_{t=1}^T \prod_{j=t}^T (1 - \chi_j) \chi_t \left(\frac{1}{2} \ln 2\pi + 0.5 \log \|V_1\| + (p_t - Ep_t) V_t^{-1} (p_t - Ep_t) \right)$$

Now, we first change the notation to match Evans and Honkapohja

$$Ep_t = c'x_t$$

and that

$$V_t = V$$

and V is known. The standard approach is that maximizing the weighted likelihood

LL_T is equivalent to minimizing the sum of squared residuals

$$SS_T = \sum_{t=1}^T \prod_{j=t}^T (1 - \kappa_j) \kappa_t (y_t - c'x_t) (y_t - c'x_t)$$

where we have replaced p_t and χ_j to match the notation in Evans and Honkapohja.

The solution to the minimization problem is c_T where

$$\begin{aligned} c_T &= \arg \min SS_T \\ &= \arg \min \sum_{t=1}^T \prod_{j=t}^T (1 - \kappa_j) \kappa_t (y_t - c'x_t) (y_t - c'x_t) \end{aligned}$$

Now, we need to establish that this c_T is equivalent to the solution of the algorithm in Evans and Honkapohja (2001)

$$\begin{aligned} c_t &= c_{t-1} + \kappa_t R_t^{-1} x_{t-1} \left(y_t - \phi'_{t-1} x_{t-1} \right), \\ R_t &= R_{t-1} + \kappa_t \left(x_{t-1} x'_{t-1} - R_{t-1} \right). \end{aligned}$$

For ease of notation, we assume that y_t and x_t are scalars but the proof where x is a vector would be straightforward. Note that by definition the solution to our recursive least squares solution is calculated using the well-known weighted least squares solution

$$\begin{aligned} c_T &= \arg \min SS_T \\ &= \frac{\sum_{t=1}^T \prod_{j=t}^T (1 - \kappa_j) \kappa_t x_t y_t}{\sum_{t=1}^T \prod_{j=t}^T (1 - \kappa_j) \kappa_t x_t^2}. \end{aligned}$$

To prove equivalence, we first define

$$R_1 = \kappa_1 x_1^2$$

then we can show that

$$\begin{aligned}
R_T &= \sum_{t=1}^T \prod_{j=t}^T (1 - \kappa_j) \kappa_t x_t^2 \\
&= (1 - \kappa_T) \sum_{t=1}^{T-1} \prod_{j=t}^{T-1} (1 - \kappa_j) \kappa_t x_t^2 + \kappa_T x_T^2 \\
&= R_{T-1} + \kappa_T (x_T^2 - R_{T-1}),
\end{aligned}$$

which is the second equation in Evans and Honkapohja (2001). Next we have that

$$\begin{aligned}
C_T &= \frac{(1 - \kappa_T) \sum_{t=1}^{T-1} \prod_{j=t}^{T-1} (1 - \kappa_j) \kappa_t x_t y_t + \kappa_T x_T y_T}{R_T} \\
&= C_{T-1} \frac{R_{T-1}}{R_T} + \kappa_T R_T^{-1} x_T y_T \\
&= C_{T-1} \frac{(1 - \kappa_T) R_{T-1}}{R_T} + \kappa_T R_T^{-1} x_T y_T \\
&= C_{T-1} \frac{(1 - \kappa_T) R_{T-1} - \kappa_T x_T^2}{R_T} + \kappa_T R_T^{-1} x_T y_T \\
&= C_{T-1} + \kappa_T R_T^{-1} x_T (y_T - x_T C_{T-1})
\end{aligned}$$

which is the first equation in Evans and Honkapohja (2001), completing our proof.

Table 1: **Bivariate Model Estimates**

	Unrestricted	Restricted	
		$e_t^r = 0$	$e_t^p = 0$
ϕ_τ	0.809 <i>0.020</i>	-	0.899 <i>0.008</i>
σ_τ	0.132 <i>0.013</i>	-	0.182 <i>0.012</i>
$\sigma_{\mathcal{P}}$	0.102 <i>0.007</i>	0.177 <i>0.011</i>	-
ϕ_ξ	-0.674 <i>0.626</i>	0.716 <i>0.063</i>	0.984 <i>0.014</i>
σ_ξ	0.003 <i>0.007</i>	0.075 <i>0.005</i>	0.047 <i>0.003</i>
Std Dev of ξ_t	0.004	0.108	0.263
Likelihood Ratio Test	-	132	27.6
95% Critical Value	-	5.99	3.84

Note: Observed variables are change in log front month oil futures quote and the log difference between one-year-ahead and front-month futures prices. Estimation results are for the sample period 1989:Q1 to 2019:Q4. Standard errors are reported in italics. Critical values are based on the chi-squared distribution with 2 and 1 degrees of freedom, respectively.

Table 2: **Model Estimation - Targeted and Simulated Moments**

Moment	Targeted	Simulated
<i>Panel A: Early 2000s</i>		
$\text{frac } \sigma_{\Delta_p}^2$	0.10	0.10
$\sigma(\text{Inv.})/\sigma(Y)$	6.00	7.01
$\rho(Y, Y_{-1})$	0.82	0.65
$\rho(\text{Inv.}, \text{Spot})$	-0.40	-0.47
$\rho(\text{Spot}, Y)$	-0.07	-0.07
<i>Panel B: Late 2000s</i>		
$\text{frac } \sigma_{\Delta_p}^2$	0.50	0.48
$\sigma(\text{Inv.})/\sigma(Y)$	5.00	4.67
$\rho(Y, Y_{-1})$	0.85	0.70
$\rho(\text{Inv.}, \text{Spot})$	-0.45	-0.44
$\rho(\text{Spot}, Y)$	0.10	0.10
<i>Panel C: Mid 2010s</i>		
$\text{frac } \sigma_{\Delta_p}^2$	0.30	0.30
$\sigma(\text{Inv.})/\sigma(Y)$	4.00	4.12
$\rho(Y, Y_{-1})$	0.87	0.70
$\rho(\text{Inv.}, \text{Spot})$	-0.50	-0.50
$\rho(\text{Spot}, Y)$	0.36	0.36

Note : $\text{frac } \sigma_{\Delta_p}^2$ refers to the fraction of the variance in spot price changes explained by permanent shocks, as per the Kalman filter mechanism described in the paper. *Inv.* refers to (US) Oil Inventories and *Spot* to the spot price of oil. Individual data series are logged and HP-filtered to produce the two relative volatilities and the two correlations.

Table 3: Estimated Parameter Values

Parameter	Estimate		
	<i>Early 2000s</i>	<i>Late 2000s</i>	<i>Mid 2010s</i>
σ_o^τ (s.d. trans. oil-supply shocks)	0.01	0.01	0.01
σ_o^P (s.d. perm. oil-supply shocks)	0.0007	0.0024	0.0013
ρ_o^τ (autocorr. trans. oil-supply shocks)	0.82	0.83	0.69
σ_{TFP}^τ (s.d. trans. oil-demand shocks)	0.0037	0.0026	0.0062
σ_{TFP}^P (s.d. perm. oil-demand shocks)	0.001	0.0064	0.0035
ρ_{TFP}^τ (autocorr. tran. oil-demand shocks)	0.85	0.85	0.93

Note : Estimated values for the standard deviation of permanent oil-supply shocks (σ_o^P), the autocorrelation of transitory oil-supply shocks (ρ_o^τ), the standard deviation of transitory oil-demand (TFP) shocks (σ_{TFP}^τ), the standard deviation of permanent oil-demand (TFP) shocks (σ_{TFP}^P), and the autocorrelation of transitory oil-demand (TFP) shocks (ρ_{TFP}^τ). Estimates are obtained by minimizing the distance between simulated and observed moments, via the Simplex algorithm.

Table 4: **Model Estimation - Implication for Kalman Filter Gain**

Shock	Attributed to Permanent Comp.	Attributed to Transitory Comp.
<i>Panel A: Early 2000s</i>		
Oil Supply	0.014	0.986
TFP	0.105	0.895
<i>Panel B: Late 2000s</i>		
Oil Supply	0.234	0.766
TFP	0.923	0.077
<i>Panel B: Mid 2010s</i>		
Oil Supply	0.134	0.867
TFP	0.490	0.510

Note : Implications for the Kalman filter's gain under learning according to the parameter estimates in Table ???. The table gives, for each historical episode, how much an unexpected innovation to oil supply or to TFP will be attributed to its permanent (left column) or transitory components (right column).

Table 5: Impact of Learning on Macroeconomic Volatility

Moment	<i>Misperceiving the Late 2000s</i>	<i>Misperceiving the Mid 2010s</i>
σ_Y	1.37	0.85
$\sigma_{futures}$	0.84	1.17

Note : Business cycle characteristics of simulated data. Table entries are standard deviations when using the previous episode's Kalman filter gain, *as a ratio* of the corresponding figure when using the current gain.

Table 6: Impact of Storage on Macroeconomic Volatility

Moment Simulated	<i>Episode</i>		
	<i>Early 2000s</i>	<i>Late 2000s</i>	<i>Mid 2010s</i>
σ_Y	0.92	0.95	0.97
σ_{spot}	0.50	0.61	0.48

Note : Business cycle characteristics of simulated data. Table entries are the standard deviations under the benchmark assumption of storage as a ratio of the corresponding value when abstracting from storage

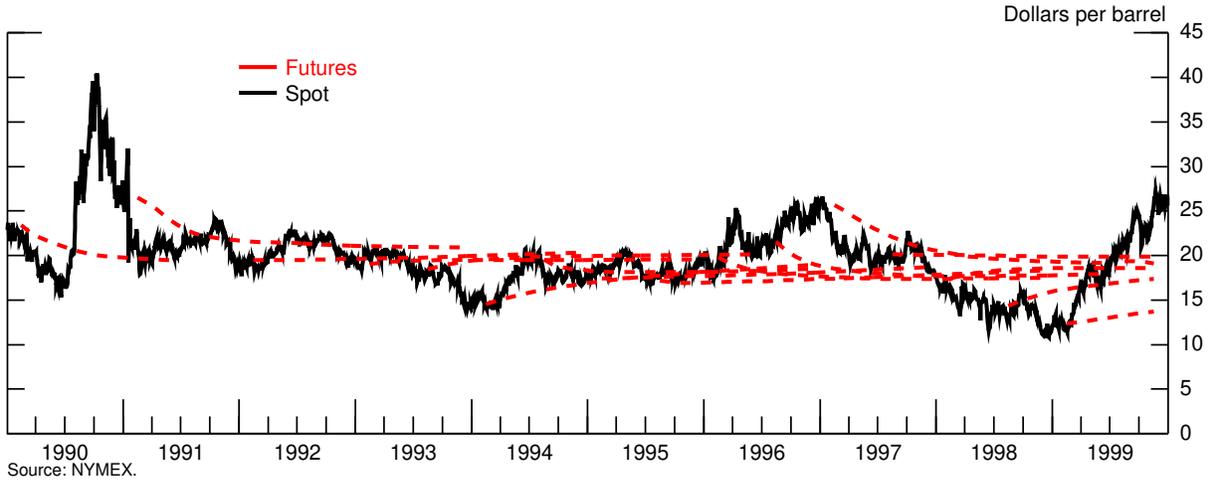
Figure 1: Oil Spot and Futures Prices

Exhibit 1

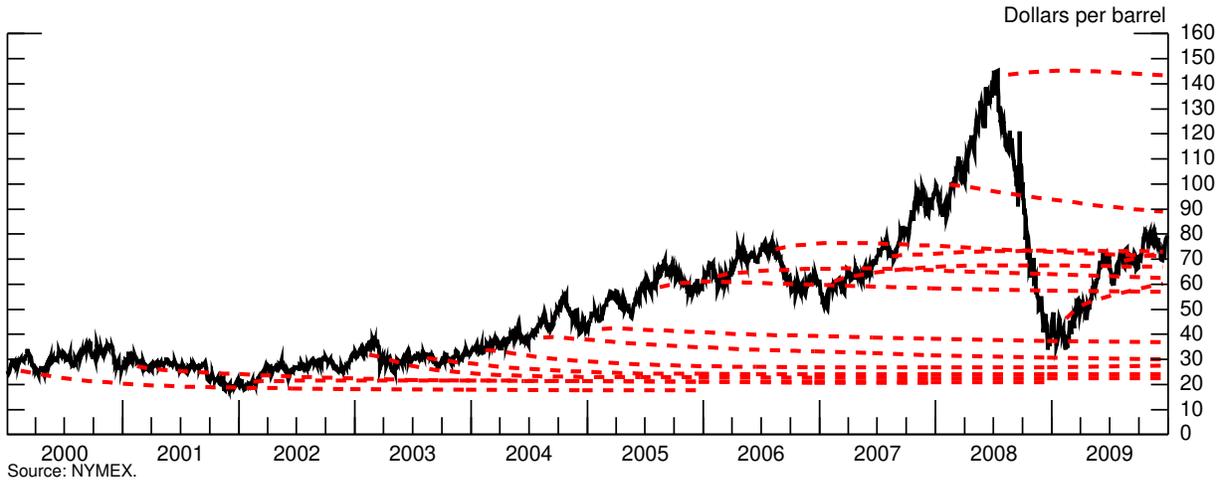
6/30/20

WTI Spot and Futures Prices

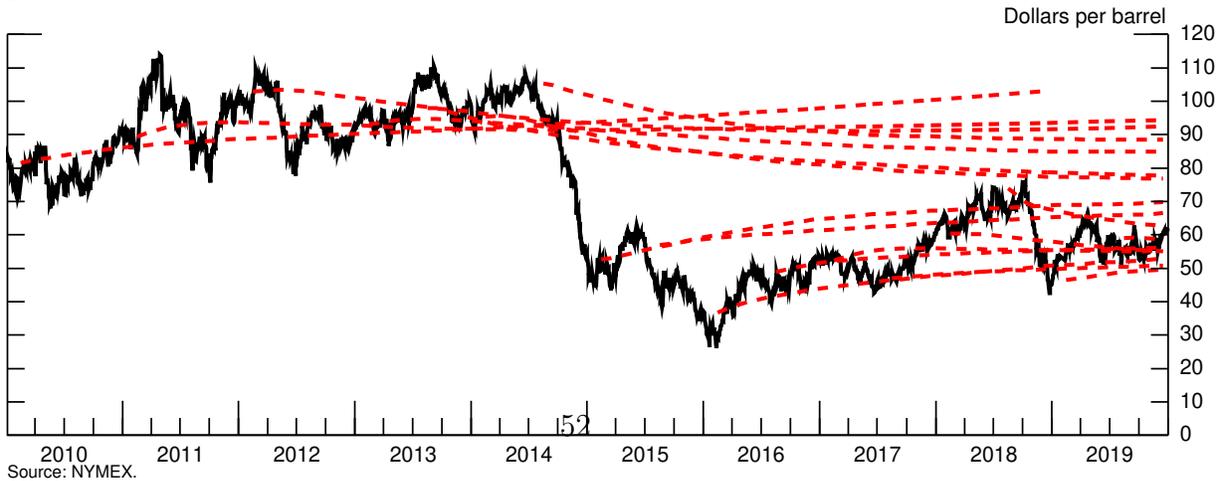
1990s



2000s

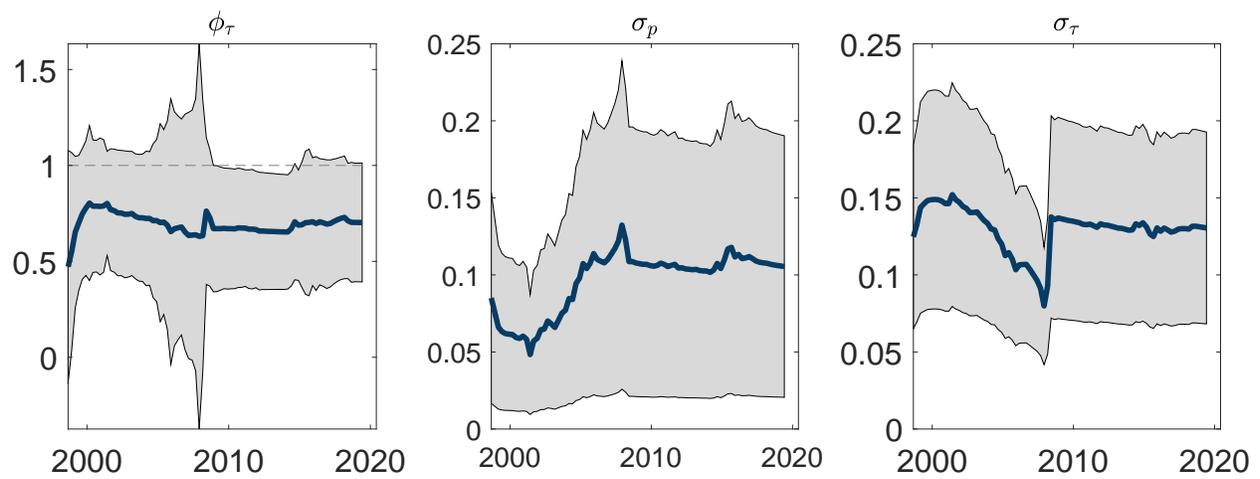


2010s



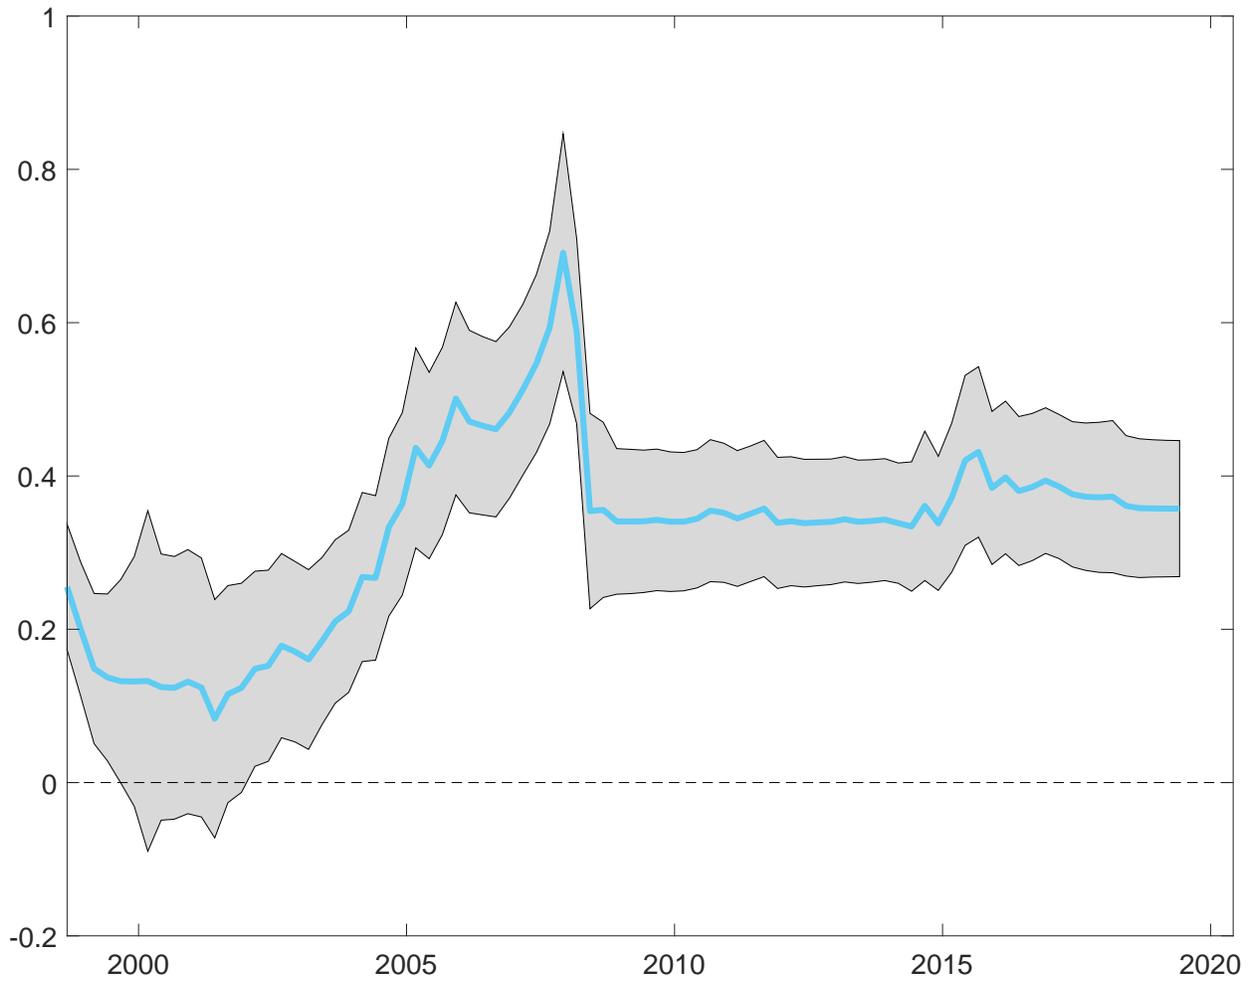
NOTES: Evolution of WTI spot prices (black lines) as well as paths for sequentially-dated futures prices (red): 1990s; 2000s and 2010s. All data reported from NYMEX.

Figure 2: Estimates of Model Parameters: Baseline Model



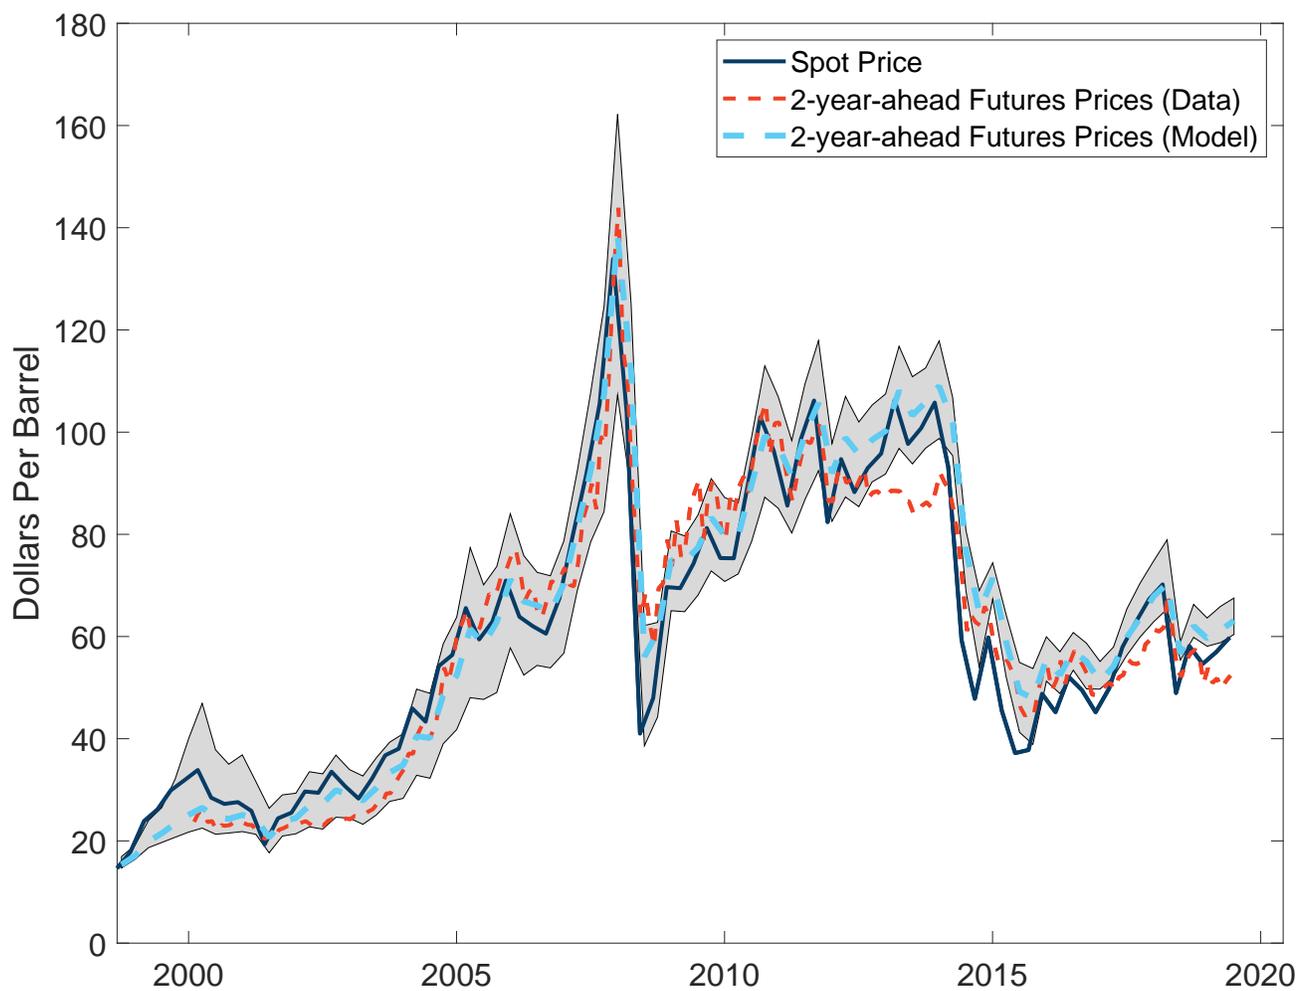
NOTES: Benchmark estimates for serial correlation of transitory shocks (ϕ) and volatility of permanent (σ_p) and transitory (σ_τ) shocks.

Figure 3: Estimated Importance of Permanent Shocks in $var(\Delta p_t)$



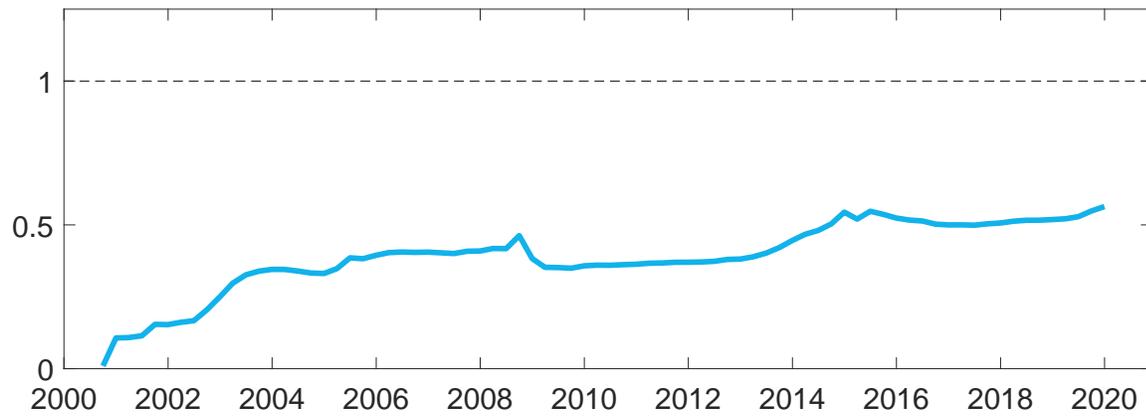
NOTES: Estimated fraction of log-spot price changes variability attributable to permanent shocks ($= \frac{\hat{\sigma}_p^2}{\hat{\sigma}_p^2 + \frac{2}{(1+\phi_\tau)} \hat{\sigma}_\tau^2}$)

Figure 4: Predicting Futures Prices



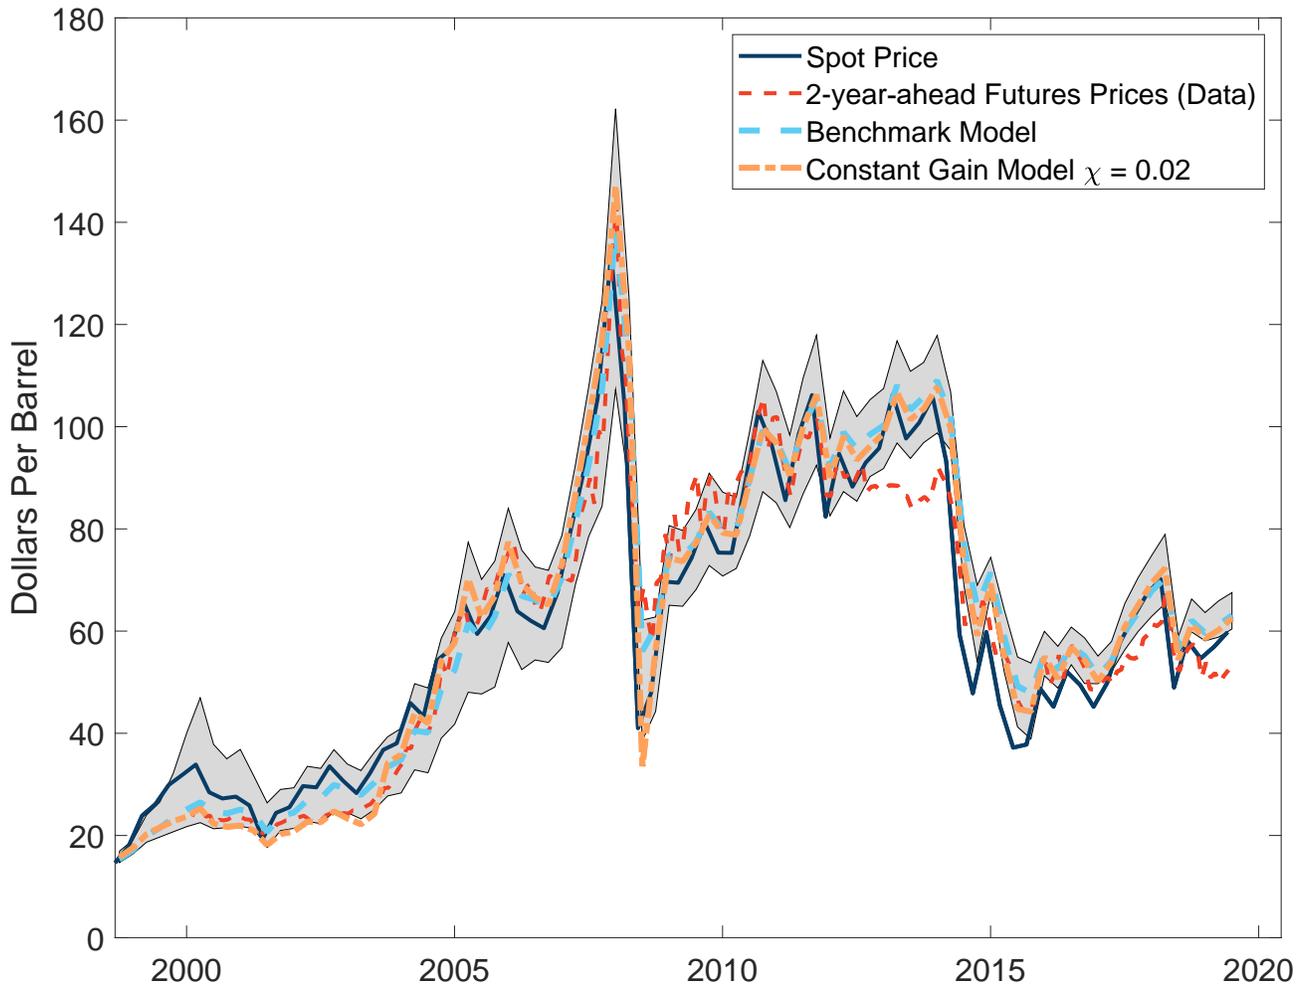
NOTES: Benchmark model estimates of two-year-ahead futures prices, along with observed futures and spot prices.

Figure 5: Assessment of Model Performance



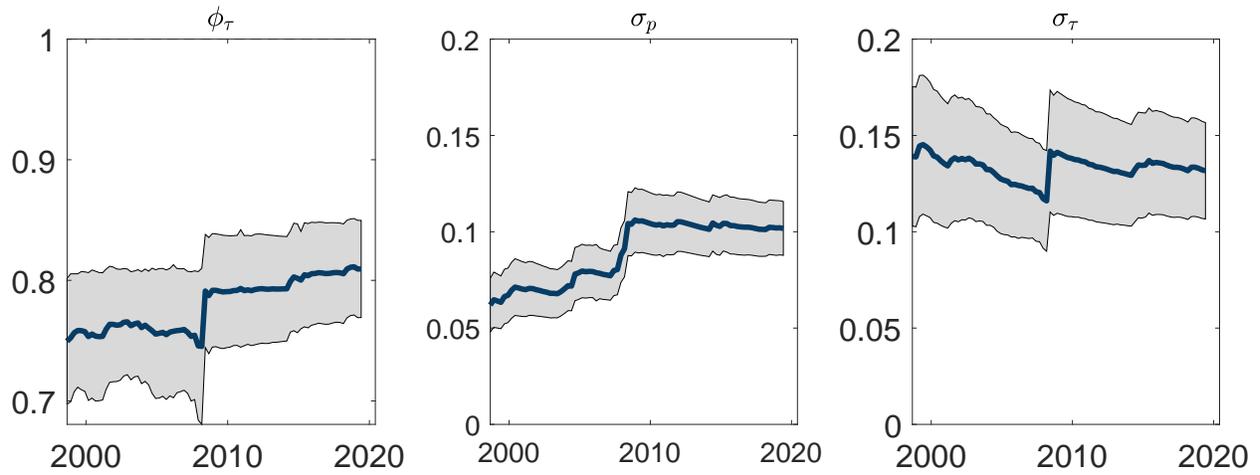
NOTES: Cumulative mean-squared error (MSE) of fitting two-year-ahead futures price using benchmark Kalman filter estimates, relative to MSE of random walk alternative.

Figure 6: Predicted Futures Prices: Constant-Gain Learning



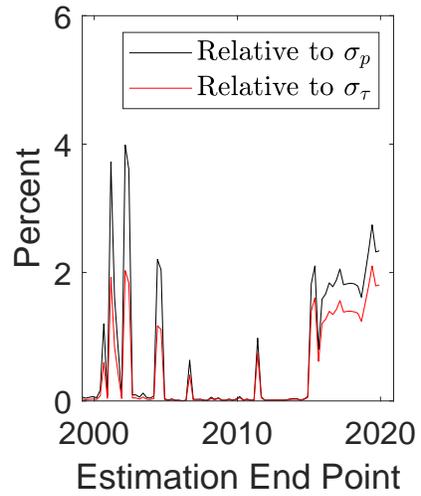
NOTES: Benchmark model versus constant-gain alternative estimates of two-year-ahead futures prices, along with observed futures and spot prices. Gray interval indicates 90 percent confidence interval around benchmark model's implied futures prices, constructed using a Monte Carlo simulation, under the assumption that estimated coefficients are normally distributed and with a correction for Jensen's inequality term.

Figure 7: Estimates of Model Parameters: Bivariate Model



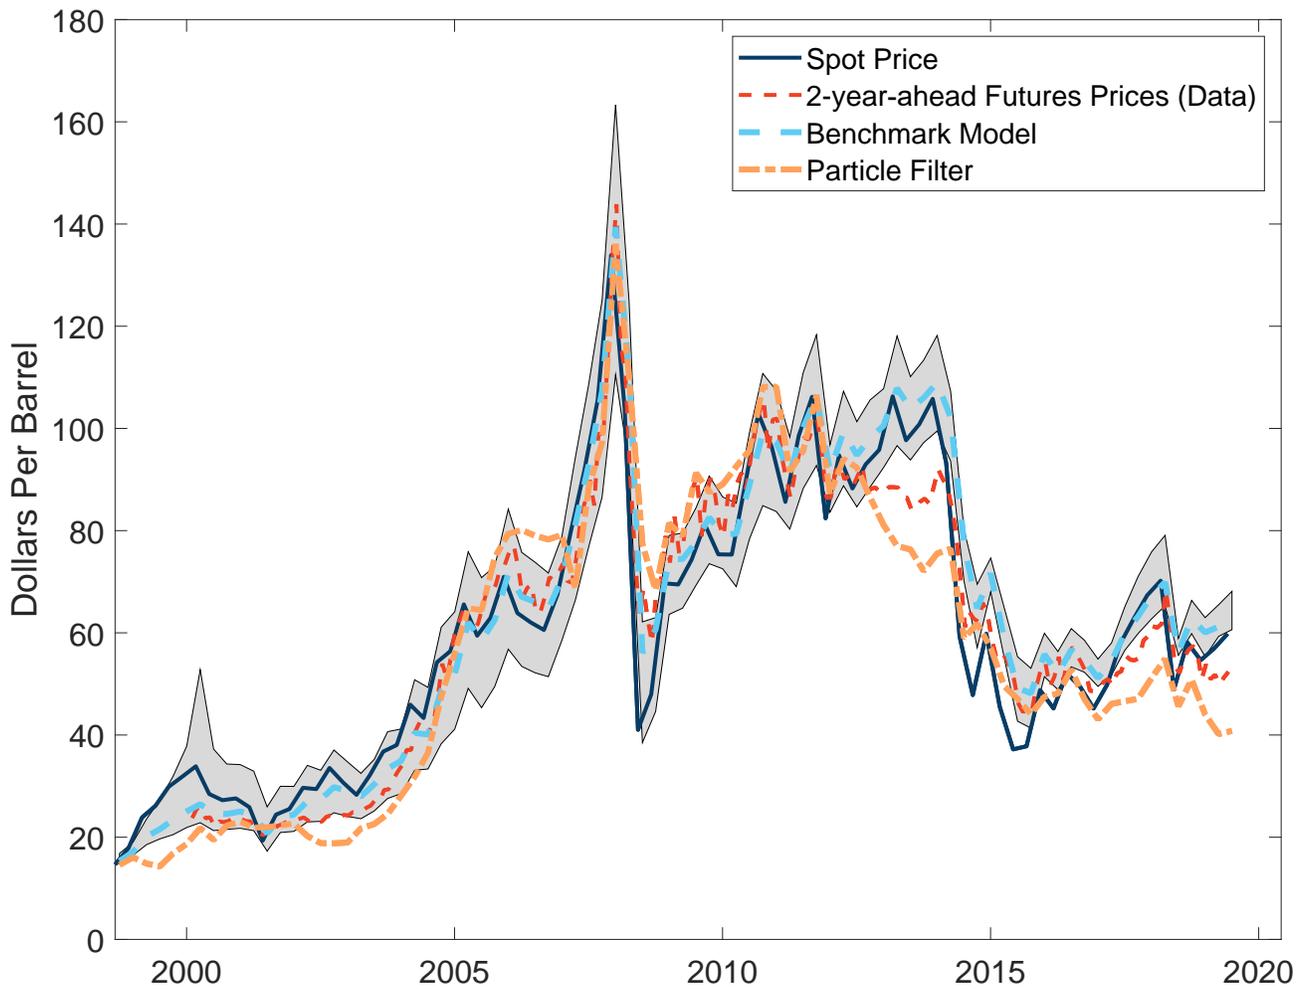
NOTES: Black lines report bivariate estimates for serial correlation of transitory shocks (ϕ) and volatility of permanent (σ_p) and transitory (σ_τ) shocks. Gray-shaded interval indicates two-standard deviation confidence intervals around those estimates, computed using the appropriate inverse Hessian. The bivariate estimation is for front-month and one-year-ahead futures prices over the sample period from 1990 through the end point indicated on the x-axis.

Figure 8: Relative Importance the Risk Premium



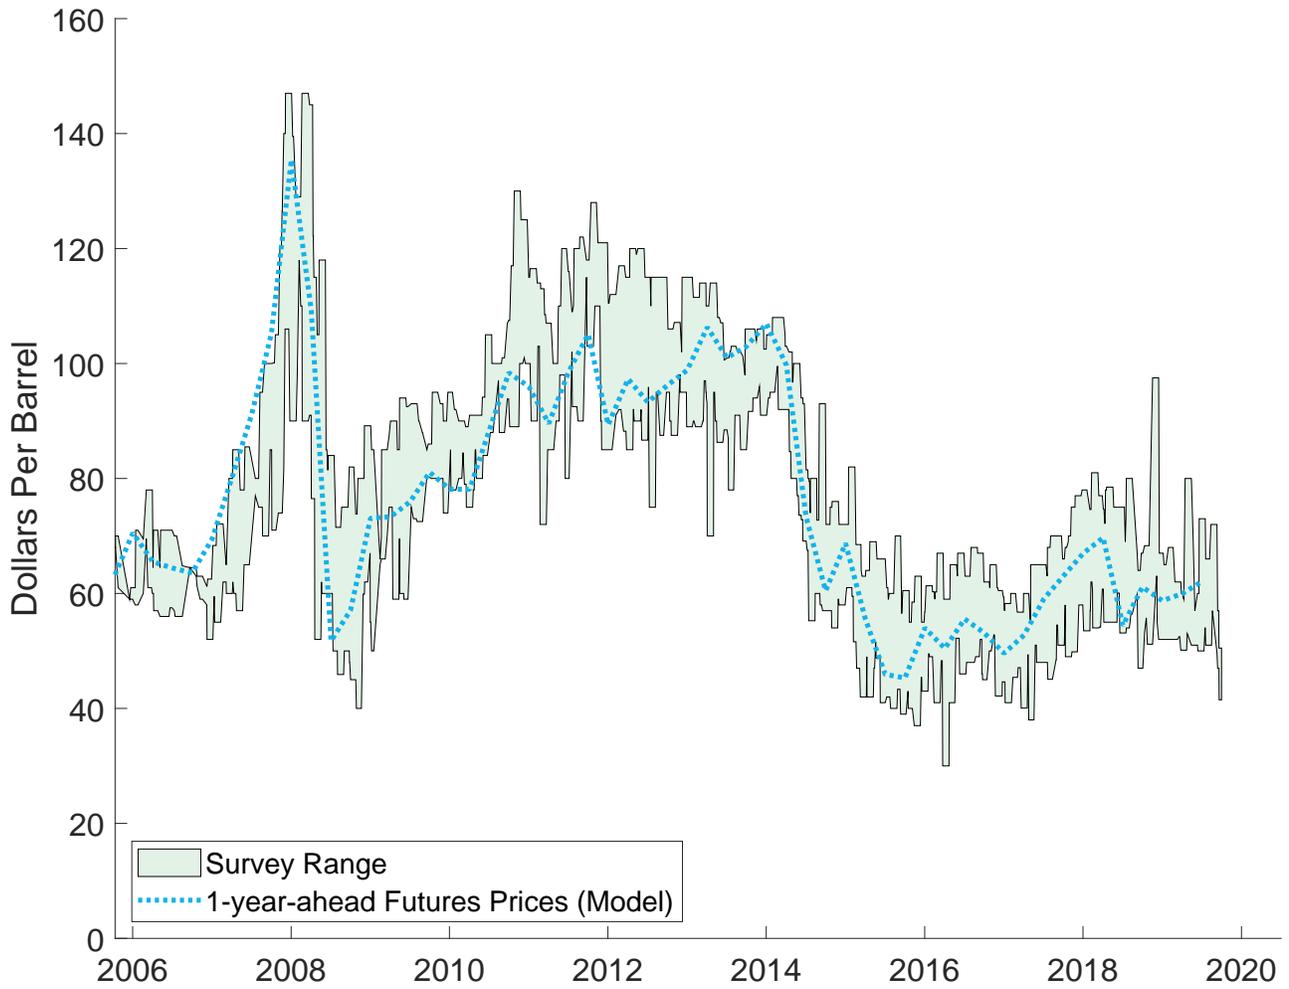
NOTES: The black line reports 100 times the ratio of the standard deviation of ξ to σ_p , the standard deviation of permanent shocks as estimated in our bivariate model. The red line reports the same statistic for ξ relative to σ_τ , the standard deviation of transitory shocks.

Figure 9: Predicted Futures Prices: Particle Filter



NOTES: Benchmark model versus particle-filter alternative estimates of two-year-ahead futures prices, along with observed futures and spot prices. Gray interval indicates 90 percent confidence interval around benchmark model's implied futures prices, constructed using a Monte Carlo simulation, under the assumption that estimated coefficients are normally distributed and with a correction for Jensen's inequality term.

Figure 10: Predicted Futures Prices and Market Expectations



NOTES: Benchmark model estimates of one-year-ahead futures prices versus range of Bloomberg survey expectations reported over the previous 30 days for one-year-ahead spot price.

Figure 11: Impact of Oil Supply Shocks: Full Information

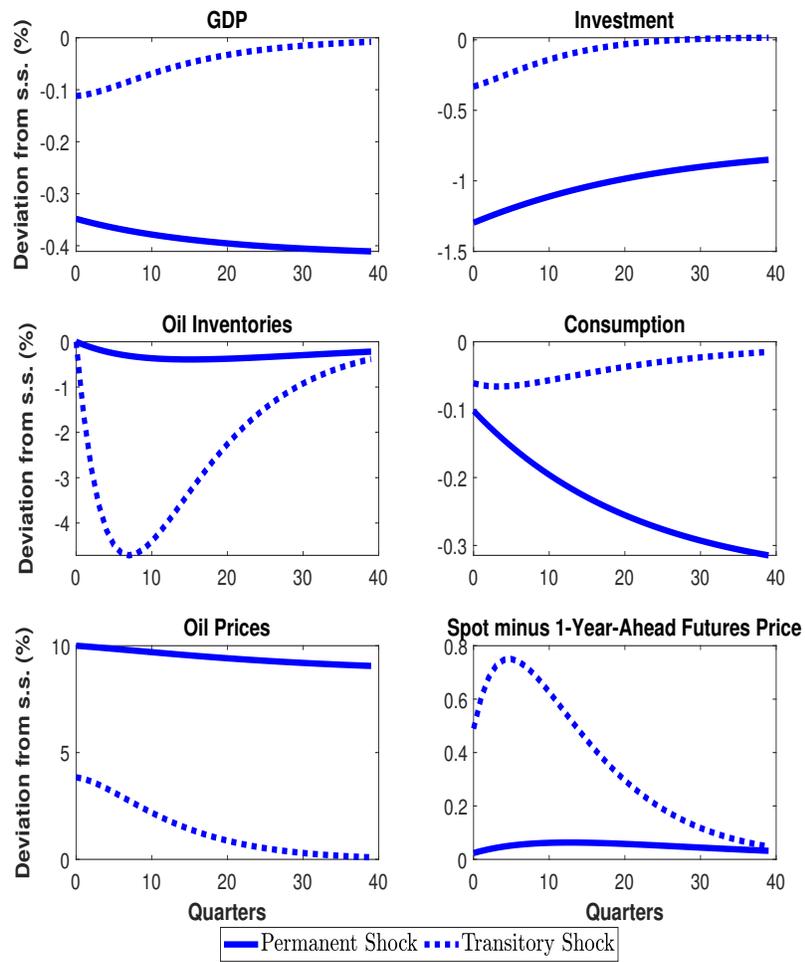


Figure 12: Impact of TFP Shocks under Full Information

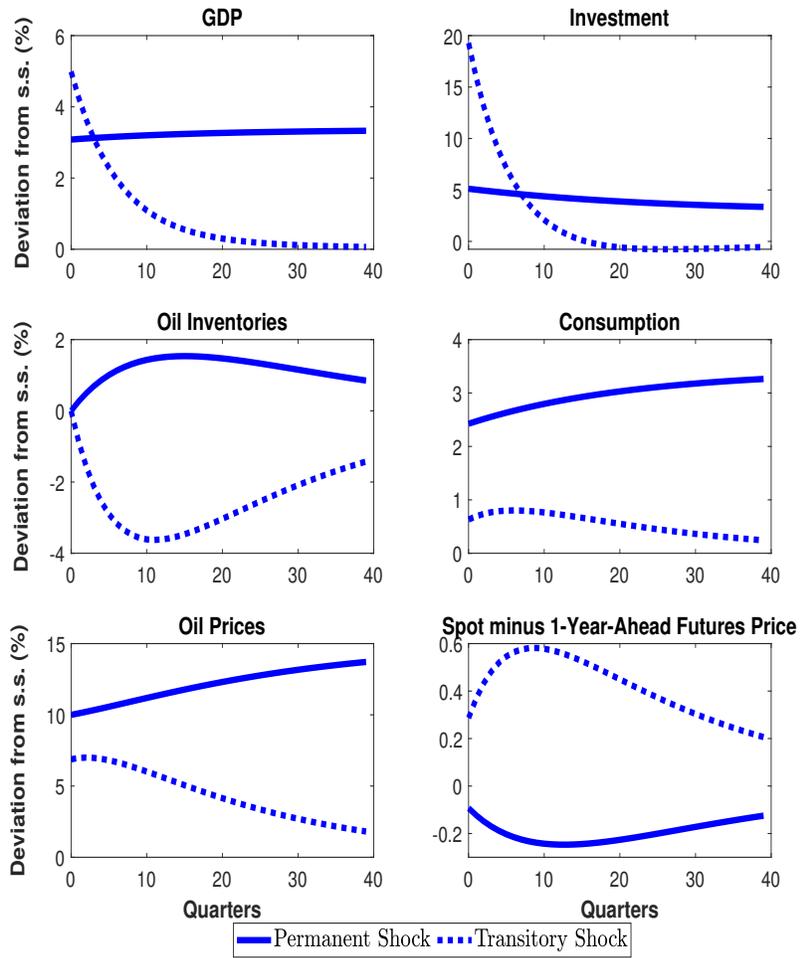


Figure 13: Impact of Oil Supply Shocks: Full Information versus Learning

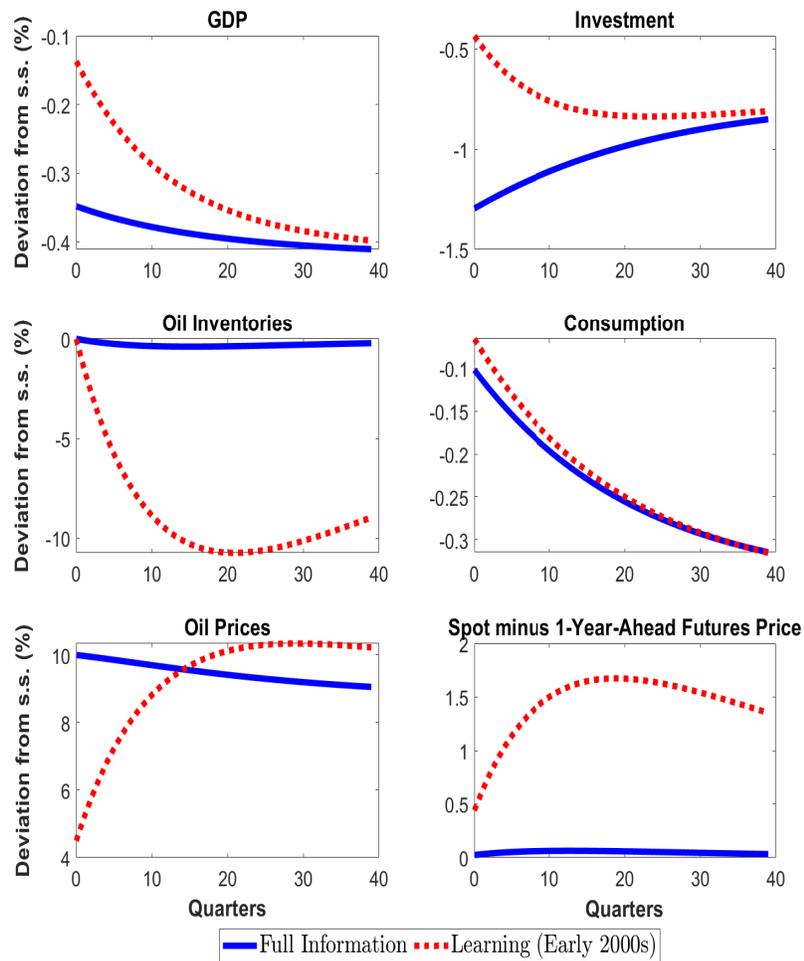


Figure 14: Impact of TFP Shocks: Full Information versus Learning

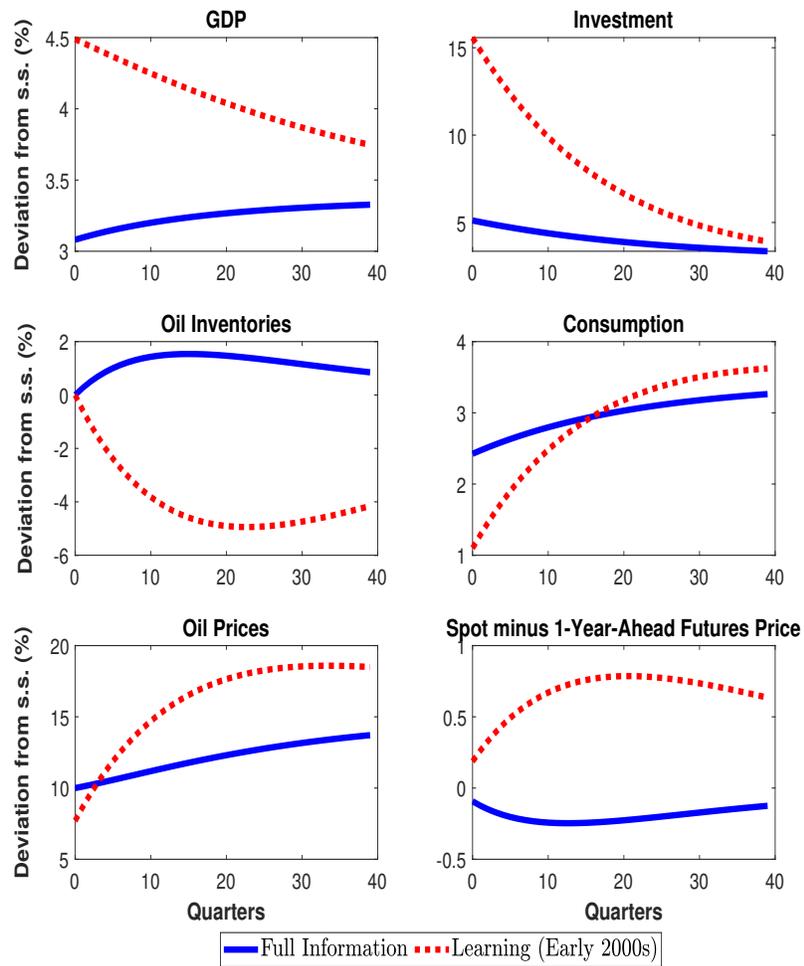


Figure 15: The Impact of Oil Inventories

