Working Paper Series 619
(ISSN 1211-3298)

Formation of Market Beliefs in the Oil Market

Stanislav Anatolyev
Sergei Seleznev
Veronika Selezneva

CERGE-EI
Prague, June 2018

ISBN 978-80-7343-426-7 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-469-3 (Národohospodářský ústav AV ČR, v. v. i.)

Formation of Market Beliefs in the Oil Market*

Stanislav Anatolyev — CERGE-EI and NES

Sergei Seleznev — INECO Capital Ltd

Veronika Selezneva† — CERGE-EI

June 8, 2018

Abstract

We characterize formation of market beliefs in the oil market by providing a complete characterization of the market reaction to oil inventory surprises. We utilize the unique sequential nature of inventory announcements to identify inventory shocks. We estimate an AR-ARCH-MEM model of the joint dynamics of returns, return volatilities and trading volumes around the announcements using high frequency data on oil futures contracts. Our model (i) handles illiquidity of long maturity contracts by accounting for trading inactivity, (ii) captures time varying trading intensity, and (iii) allows for structural changes in the dynamics and responses to news over time. We show (i) uniform formation of expectations across oil futures contracts with different maturities, (ii) a strong negative relation between inventories surprises and re-turns, (iii) no effect on the term premium, which suggests that inventory shocks are always considered to be permanent, and (iv) di_erentiation in the reaction of volume by maturity. We demonstrate how our results can be used to test theories of oil price determination and contribute to the debate on the recent oil glut.

Keywords: oil market, ultra high frequency data, trading intensity, futures returns, return volatility, inventory surprises, expectation formation

JEL classifications: C22, C32, C58, G12, G13

---

*Grant GA ČR 17-27567S from the Czech Science Foundation is gratefully acknowledged.
We thank James Hamilton and conference participants of CFE-2017, AMEF-2018, and ISCEF-2018 for helpful comments. Valentin Artemyev, Matěj Bělín and Mariia Kosar provided excellent research assistance.
†Corresponding author. Address: CERGE-EI, Politických vězňů 7, 11121 Prague 1, Czech Republic; e-mail veronika.selezneva@cerge-ei.cz.
CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, P.O. Box 882, 111 21 Prague 1, Czech Republic.

1 Introduction

The oil market is undergoing a major transformation driven by the shale boom. Oil traders need to learn a great deal about the productivity of the new technology, not to mention to understand its effect on the market. While ex post one can observe an increase in oil production, it is much harder to track changes in the market perception of supply and demand conditions. Unfortunately, there is little extant analysis of how market beliefs are formed and evolve over time. In this paper, we aim to fill the gap by offering a way to characterize the expectation formation process in the oil market.

The key element is oil inventories. Oil inventories reflect the supply and demand balance in the oil market, and thus are informative about the scarcity of the commodity not only today, but also in the future. Our contribution is a complete characterization of the oil market reaction to oil inventories announcements. Revisions of market expectations regarding oil scarcity can be partially gauged from the reaction of the oil futures returns to inventory news, while the adjustments of the term structure of futures prices are informative about the perceived persistence of shocks. Finally, as idiosyncratic trading behavior reflects revisions of individual beliefs in response to news, we can investigate the intensity of oil trading around an announcement in order to characterize disagreement among traders.

In this paper we offer a novel procedure to identify inventory shocks by utilizing the unique sequential nature of oil stock announcements. First, weekly estimates of crude oil inventories in the U.S. are provided by the U.S. Energy Information Administration (EIA) within the U.S. Department of Energy. The report is released at a pre-specified time each week and is closely followed by the media. What is less well known is that there is another reporting agency that collects and disseminates exactly the same information privately to its subscribers at a cost. That agency is an association of oil producers known as the American Petroleum Institute (API), and the report comes out one day before the EIA report. Historically, the API and EIA estimates tend to be similar to each other. So, even though the purchase of API information may be prohibitively costly for unsophisticated traders whose main interests lie outside the energy market, for professionals it represents an opportunity to learn the fundamentals before the rest of the market does.

The existence of two sequential announcements has two implications. On the negative side, it creates an identification issue, as neglecting the API information can lead to incorrect identification of inventory surprises and potentially distorts estimates of a market impact.[1] To overcome this issue, we assume a simple form of market expectation formation that the market's expected value of the EIA estimate of an inventory change is a weighted sum of the API estimate and the median forecast of the survey of professional analysts conducted by Bloomberg two days before the EIA announcement. We estimate the weight placed on the API signal along with other parameters of the model. On the positive side, the sequential news structure offers a unique opportunity to identify changes in the market value of early fundamental information. In particular, our results show that, since 2014, the overall market awareness and usage of API information has significantly increased, despite unchanged cost of subscription and unchanged precision of the API estimates.

Using the identified market inventory surprises, we can estimate the market impact of news. We develop a model of the joint high-frequency dynamics of futures returns, return volatilities

[1] In the academic literature, there are only a handful of papers that utilize API information to shape market expectations. One exception is Armstrong et al. (2017), which defines an inventory surprise as a difference between EIA and API reported values, but which neglects information in initial surveys. Another example is Ye and Karali (2016), which defines inventory shocks sequentially: an 'API shock' as the difference between the API announced change and the forecast, and an 'EIA shock' as the difference between EIA and API.

and trading volumes around the EIA announcements. The model is estimated using ultra high frequency (5-second time intervals) data on short and long maturity oil futures contracts on the WTI oil traded at the NYMEX. We use dynamic models because of the ultra high frequency nature of the data and because our interest lies not only in identification of an instantaneous impact of inventory surprises, but also in their dynamic impact on the market.

We specify dynamic conditional densities for volumes and returns (and return volatility). To shed light on the idiosyncratic trading behavior, we model time varying trading intensity using the multiplicative error model (MEM) of Engle (2002). To investigate the effects of oil inventories on returns, we use the AR-ARCH framework. In particular, we utilize a gaussian autoregressive conditional heteroskedasticity model, where the conditional variance follows exponential GARCH dynamics (Nelson, 1991). Even though the gaussian EGARCH model (partially) captures heavy tail behavior and may serve as a quasi-likelihood model for consistent estimation of the mean and volatility equations, we also pay special attention to the tails – as a robustness check, we utilize the Student's t distribution to capture the shape of the conditional density in the tails more accurately. One additional issue that needs to be addressed is illiquidity of futures contracts, especially those with expiration dates far in the future. We handle illiquidity by explicitly accounting for trading inactivity following the approach developed in Hautsch et al. (2013). Namely, we assign a discrete probability mass to the event of no trading, and we allow this probability to be time varying.

We abbreviate our model as an ARI-ARCHI-MEMI model, in which the additional letter 'I' signifies the presence of indicators for announcements. The inventory shocks identified according to our procedure appear in the model via these indicators. A typical procedure would be to assume a linear response function;[2] however we find such specification too restrictive. Hence, we deviate from the literature and project the universe of our surprises into a number of indicator functions, and then allow the response coefficients to differ in an unrestricted way. We distinguish 'large positive' surprises, 'uninformative announcements' , and 'large negative' surprises, i.e. we split inventory shocks into three groups.

To estimate the model parameters, we utilize the composite likelihood approach, i.e. we construct a product of the specified conditional densities. The maximum composite likelihood estimate of the model is the value of parameters that maximizes the composite conditional log-likelihood. It should be noted that the weight placed on the API signal is estimated along with the other parameters of the model, separately for each contract.

Our first main result is that contract maturity is irrelevant for identification of inventory surprises. We find that the estimated weight that the market puts on the API signal is the same for contracts with different maturities. We interpret this as evidence of uniformity of rules guiding the expectation formation process. In principle, one could worry about market segmentation, i.e. that each separate market attracts its own unique clientele looking for exposure at just one particular horizon. Different compositions of traders may form expectations in different ways, either due to differences in sophistication levels, or because of different access to the API information. However, our results suggest that one need not worry about segmentation at least as far as the *average market expectations* are concerned.

Our results indicate a strong negative relation between inventory shocks and returns. In particular, an unexpected decline of inventories by at least 0.5% caused the oil price to increase by 1.2 in 2016. That is to be expected, given that inventories reflect scarcity of oil in the market.

Our second main finding is the absence of any effect of inventories on the term premium. The entire futures curve shifts up or down when an announcement comes in. The term structure

[2] See Andersen et al. (2003).

adjustments depend on the perceived persistence of shocks. That is, the effects of temporary shocks should vanish with maturity. Our results show that expectations of future oil prices are revised uniformly, and by the same amount irrespective of maturity, which implies that the traders view inventory changes as mostly reflecting *permanent* or long lasting shocks that hit the market.[3]

To interpret our results, we turn to a canonical competitive storage theory, first developed in Deaton and Laroque (1992). As long as the non-negativity constraint does not bind, inventories serve to smooth out temporary demand and supply shocks. Oil flows in or out of storage until current and expected future prices are equalized (up to interest and storage costs). In other words, in equilibrium, speculative activity by owners of storage facilities transforms temporary demand and supply shocks into small but permanent ones. Our finding of a parallel shift of the term structure curve confirms that prediction.

However, the equalization of prices crucially depends on speculation activity to move oil across periods. When the spare capacity is near exhaustion, or when inventories have been depleted, speculation activity becomes limited. In this case, some temporary shocks can no longer be smoothed out, and the spot oil price strongly adjusts to clear the inelastic market. Hence, the theory predicts (i) stronger and asymmetric market reactions at times of extremely high or low inventories, and (ii) term structure adjustments at the shorter end of the curve. Our results indicate intensification of reaction when inventories are high, but asymmetry goes in the wrong direction: negative shocks induce larger price adjustments. More importantly, we find no term structure adjustments over the entire period.

The lack of any effect on the term spread when inventories are high is surprising and contradicts the conventional wisdom. Indeed, all recent episodes of high inventories in the oil market have been accompanied by a widening term premiums, especially at the shorter end of the term structure curve. Two examples are traditionally used as anecdotal evidence. In 2008, a negative demand shock created an abundance of oil and depressed spot oil prices. The term structure curve became upward sloping and especially steep at the short end. Similarly, in 2014, the market was hit by a positive supply shock due to rising shale oil production and by a negative demand shock due to slowdown of the Chinese economy. The market again experienced large term premiums, but over a much longer period of time. In both cases, a steep term structure curve was attributed to the presence of an excess supply of oil in the market. However, we find that when inventory news announcements came in, traders did not revise expectations accordingly.

Explicit modeling of trading intensity allows us to analyze idiosyncratic trading behavior of market participants around inventory announcements. We observe a significant immediate reaction of volumes to all announcements, including uninformative ones, in line with previous findings in the literature.[4] One explanation of why trading volumes can be decoupled from movements in returns is disagreement among traders – that is, agents disagree about the interpretation of public information or have heterogeneous priors. When news comes in, agents revise their beliefs in different ways. As a result, trading may occur as agents adjust their positions according to individual changes in their beliefs. Under this interpretation, our results not only reveal the presence of disagreement and its variation with contract maturity, but also indicate that disagreement varies over time, independently for each contract. Thus, even though the *average market expectations* are formed uniformly across contracts (as our previous results have shown), the dispersion of belief or

[3]The risk premium may also adjust in response to inventory news, but it is extremely unlikely that it would move in the opposite way to expectations revisions, by the exact amount needed to keep the term premium constant for all maturities over our sample of 7 years.

[4]Kandel and Pearson (1995) find a significant volume reaction to earning announcements even when announcement returns are zero. See also Banerjee and Kremer (2010).

revision of beliefs in response to news are different across contracts, even the two most liquid ones. We argue that one reason for the lack of uniformity in higher moments is each contract's unique exposure to financial innovation. New financial instruments such as exchange traded funds have become proliferate over the last decade. The flow of investment in these funds varies over time and reaches significant amounts. However, ETFs offer exposure to contracts with specific maturities; hence these funds are active only in specific markets. Thus, by attracting a unique composition of traders and users of specific strategies and accommodating their investment needs, the funds create market segmentation, the degree of which seems to vary over time.

Overall, we find that modeling trading probability is crucial to understand and measure the propagation of inventory surprises for long distant futures contracts. Not accounting for time-variability of trading activity significantly underestimates the total effect of the shocks. Moreover, our results indicate a significant change in the trading pattern in 2015. We show that, even though surprising announcements trigger a larger immediate trading response than uninformative news, the overall cumulative response of the trading volume happens to be lower for surprising announcements. People trade less when surprised.

Finally, we demonstrate how our results can be used to test theories of oil price determination. We focus on the recent oil glut in the US. By 2014, the US oil production had exploded, reaching almost 9 mln barrels per day, and was expected to grow even further. In July of 2014, the price of oil started to decline and gained momentum in November when OPEC refused to cut production. By the beginning of 2015, oil inventories had reached unprecedented levels and were interpreted as a sign of immense oversupply. Considerable attention has been paid to understanding exactly how the belief in oversupply was unfolding. Arezki and Blanchard (2015) argue for a structural shift in market expectations after the OPEC meeting in November of 2014. Before the meeting, the market had maintained the belief that the total oil production could still be capped, despite growing shale oil supplies. That is, OPEC producers were believed to be willing to adjust, or basically cut, their production and give way to shale oil producers, only in order to maintain high prices. After the meeting it became clear that OPEC producers would not sacrifice their market share. Hence, not only did global expectations regarding the future path of the oil supply have to be reconsidered, but the possibility of having an excess supply of oil in the nearest future became substantial.

Do we see any structural shifts in market expectations or changes in market responses to inventory news around the OPEC meeting? To parsimoniously allow for time variability of market response parameters, we utilize threshold autoregression (TAR) modeling. Before proceeding to estimation, however, we need to resolve one issue. Both daily intensity of trading (say a fraction of 5-second intervals with zero trading volume) and daily volatility of returns indicate a distinct shift in the trading pattern occurring around the end of 2014. Hence, by trying to identify a structural change in the market response parameters we may wrongly pick up background changes in the overall trading pattern. To overcome this issue, we allow for two independent transitions: one for changes in the trading patterns, and the other for changes in the market response to news.

Not surprisingly, our results show an apparent break in the trading patterns which occurred around the first week of December of 2014. We believe it can be attributed to the explosion of money flowing into ETFs. However, our results do not indicate any clear shift in the market reaction to news. If anything, the break occurred in mid-February of 2015, i.e. *after* inventories reached extremely high levels. Hence, beliefs in oversupply tend to lag behind the actual inventory build up. In sum, our results do not confirm the theory of a structural shift in expectations after the OPEC meeting in November of 2014.

The remainder of this article is organized as follows. In the following section, we describe the relation of our research questions and findings to the previous literature. In Section 3 we provide institutional background information. We discuss formation of expectations in Section 4. The econometric framework is described in Section 5. The description of the data and estimation procedure is given in Section 6. Our main findings are presented in Section 7. Application of our results to the recent oil boom is discussed in Section 8. Section 9 concludes.

2 Relation to previous literature

Our paper is related to the literature on formation of beliefs, which typically relies on survey data, for example, see Coibion and Gorodnichenko (2015), Andrade and Le Bihan (2013), Gennaioli et al. (2016), Greenwood and Shleifer (2014), Mankiw et al. (2013). Our approach is complementary to these studies. We propose to use high frequency transaction data including trading volumes around major oil announcements to characterize the expectation formation process in the oil market. A similar approach is taken in Bollerslev et al. (2018), in which high frequency data are used to estimate the volume-volatility elasticity around macroeconomic news announcements. The relationship between trading intensity and volatility can be related to disagreement among market participants; see Kandel and Pearson (1995) and Banerjee and Kremer (2010). However, we use dynamic models because our interest lies not only in identification of any instantaneous impact of inventory surprises, but also in their dynamic impact, as this provides more information about formation of beliefs.

Our paper is also related to the literature on the role of inventories in shaping the behavior of commodity prices. We contribute to it in two ways. As a minor extension, we notice the existence of an upper limit of inventories, while the literature has been exclusively focused on the non-negativity constraint. The oil storage capacity is limited, and generally believed to have been close to exhaustion recently at least twice: in 2009 and again in 2015. More importantly, we offer a direct way to test the classic inventory theory. Existing empirical evidence is mostly based on indirect evidence and is somewhat mixed. Hamilton (2009) analyzes the time series behavior of oil prices and tests whether oil price changes are permanent and unpredictable, and finds no evidence against random walk behavior. However, the power of such tests in a small sample can be limited. Deaton and Laroque (1996) investigate agricultural commodities and find that, contrary to the model predictions, autocorrelation of prices in the data is as pronounced at high prices as it is at low prices. Moreover, autocorrelation is too high to be accounted for by a simple competitive storage model. Another strand of empirical literature uses futures prices to extract expectations of oil prices in the future. Using a large cross-section of commodity futures, Gorton et al. (2013) show that high levels of inventories are associated with an upward sloping futures curve. However, the presence of a time-varying risk premium potentially depending on the level of inventories complicates the interpretation of results. The risk premium can be estimated, but the estimates are extremely sensitive to model assumptions,[5] as the results of Baumeister and Kilian (2017) suggest. The errors introduced by estimation of the risk premium mask the predictions of the storage model.

Our paper complements these studies. We directly estimate the term structure response to exogenous supply or demand shocks identified by inventory news. Although we do make an assumption about the risk premium, we only require it to be constant in a narrow window around

[5]A number of papers study the determination of the risk premium in a general equilibrium model with inventories subject to non-negativity constraints; see an illustrative model in Gorton et al. (2013) and an estimated structural model in Baker (2016). As with any structural approach it is subject to the joint hypothesis problem.

the inventory announcement, which we believe is a less restrictive assumption. Our results are consistent with the model predictions in normal times, but at times of high inventories we find that traders do not revise their expectations according to the predictions of a standard competitive storage theory.

The effect of oil announcements on returns and return volatility has been examined in the earlier literature. Our study differs in a number of important ways: (i) identification of inventory surprises, (ii) focus of study, (iii) modeling approach, including usage of high frequency data and trading volume data, and (iv) a longer and more recent coverage period.

First, we offer a novel procedure to identify market surprises. Most studies define a market surprise as the difference between the reported EIA estimate and the median survey forecast. Given general public interest in oil inventories, various surveys are available that directly ask about agents' expectations of EIA announced changes, including the surveys conducted by Reuters and utilized in Bu (2014), Bloomberg, used by Halova Wolfe and Rosenman (2014), Halova et al. (2014), Miao et al. (2018), and Platt's. By assumption, the weight placed on an API signal equals zero. There are only a handful of papers that utilize API information to shape market expectations. One exception is Armstrong et al. (2017), which defines an inventory surprise as the difference between EIA and API reported values, but neglects information in initial surveys, thus assuming the API weight to be unity. Another exception is Ye and Karali (2016), which defines inventory shocks sequentially: an 'API shock' as the difference between the API announced change and the forecast, and an 'EIA shock' as the difference between EIA and API. Both shocks are included in their regression model, assuming that the market may react differently to two consecutive positive surprises than to two alternating signals. This approach can be viewed as a reduced way to account for misspecified beliefs. In our paper, we directly model formation of market expectation, and allow the API weight to be unrestricted. We estimate it along with other parameters of the model, separately for each contract and each time period.

Second, our goal is to provide a complete characterization of the market impact of news. Most studies look only at the impact on returns and use only the nearby futures contract. Bu (2014) and Halova Wolfe and Rosenman (2014) also investigate the effect of news on return volatility. However, there are almost no studies of the effects of announcements on the term structure of futures prices.[6] One exception is Miao et al. (2018). Six continuous contracts are analyzed in this paper, and the results indicate slow weakening of the magnitude of the price response with maturity, a 1% increase in inventories decreases the price of the first month contract by 0.552%, the price of the second contract by 0.541%, and for the sixth contract – by 0.343%. In contrast, we do not find that the longer maturity contracts display smaller responses. The difference in the results can be attributed to our different identification of inventory surprises, usage of high frequency data instead of daily returns, and coverage of a more recent time period. To the best of our knowledge, our study is the first to analyze the impact of inventory shocks on trading intensity.

Third, in contrast to other studies on news impact (see Andersen et al. (2003)), we use ultra high frequency (5-seconds) data, which calls for different econometric tools, those that are used to deal with ultra-high frequency financial data. In particular, we analyze trading volumes, which are characterized by non-negativity, high volatility and the presence of zeros due to non-trading. Also, we analyze return volatility along with returns themselves, and hence employ models from the GARCH literature (also used by Bu (2014)). The literature on the impact of inventory shocks most frequently uses daily returns, see Bu (2014) and Miao et al. (2018). Halova et al. (2014)

[6] Although, there is a literature that documents reaction of the term structure of interest rates to various news announcements, see, for example, Gurkaynak et al. (2005) and Faust et al. (2007).

calculates continuously compounded returns in an intraday event window surrounding the EIA announcement.[7] One exception is Ye and Karali (2016) who work with 5-minute returns and use the methodology developed in Andersen et al. (2003).

Finally, our sample covers 7 years, from 2010 to 2016.[8] To the best of our knowledge, our study is the only one that covers the shale oil boom, the proliferation of ETFs, as well as one of the most dramatic oil price collapses in recent history.

Our paper also contributes to the debate on the recent oil glut in the US. It remains an open question whether oversupply of oil was the main determinant of the decline in the oil price in 2014; see, for example, Baumeister and Kilian (2016), Baumeister and Hamilton (2017), Baffes et al. (2015), Fantazzini (2016). What is even more controversial is exactly how beliefs regarding oversupply were developing in 2014. Arezki and Blanchard (2015) argue in favor of a structural shift in market expectations, following OPEC's announcement in November 2014 of their intention to maintain the production level despite growing shale oil supplies. Baumeister and Kilian (2016) use a four-variable vector autoregressive forecasting model for the real price of oil to show that more than half of the decline in the price of oil was predictable in real time as of June 2014. However, this does not necessarily imply that the market incorporated that information into its beliefs. Our approach allows us to characterize the time evolution of market expectations.

3 Institutional background

Weekly estimates of crude oil inventories in the U.S. are provided by the U.S. Energy Information Administration (EIA), a statistical and analytical agency within the U.S. Department of Energy. Any company which carries or stores more than 1000 barrels of oil may be selected into the EIA weekly sample based on a procedure that assures coverage of 90 percent of the market. Typically the sample includes gathering and pipeline companies, and storers of crude oil. The selected firms are required to report the end-of-week amount of oil in their storage facilities. On the following Wednesday, a summary report is released in the form of an EIA publication, the *Weekly Petroleum Status Report*. The report becomes available to the public at 10:30am Eastern time and is closely followed by the media.

However, what is less known is that there is an alternative reporting agency that collects and disseminates information about oil stocks on a weekly basis privately to its subscribers. An association of oil producers known as the American Petroleum Institute (API) surveys energy firms using exactly the same weekly survey forms that the EIA uses. While reporting to the EIA is mandatory, reporting to the API is voluntary, but despite this, the association claims its coverage is close to 90% of the industry. The API releases the data in the *Weekly Statistical Bulletin* on Tuesdays at 4:30 pm Eastern time, the day before the official EIA announcement. In contrast to the publicly observable EIA report, access to the API requires a costly subscription available only through Thompson Reuters. Thus, for less sophisticated traders and traders whose main interests are outside the energy market, the purchase of API information may be prohibitively costly. However, for more specialized traders, API reports represent quite valuable information, as historically, API and EIA estimates tend to be close to each other. Discrepancies are believed to occur due to different procedures utilized to estimate the remaining 10% of the market.

[7]Halova et al. (2014) also address the issue of measurement error in inventories changes.

[8]For comparison, Bu (2014) covers 2006 to 2011; Ye and Karali (2016) – mid-2012 to 2013; Halova Wolfe and Rosenman (2014) – mid-2003 to 2010; Halova et al. (2014) – mid-2003 to mid-2012; Miao et al. (2018) – mid-2003 to 2011.

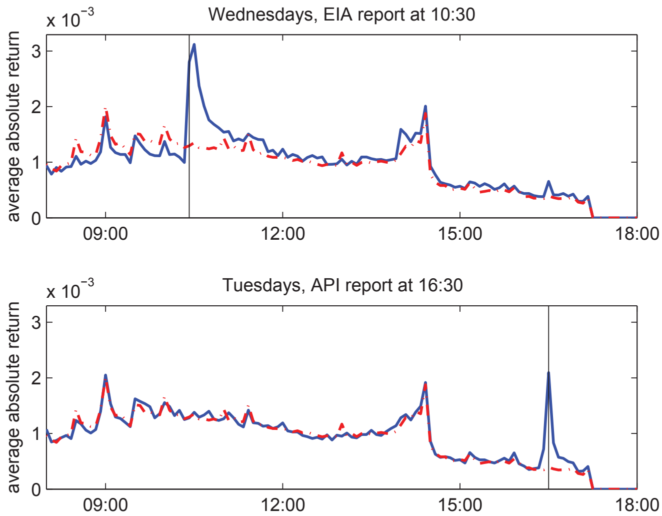

Figure 1: Market reaction to oil inventory announcements.

To get a feeling for whether the market follows oil inventories announcements, we use a simple procedure to estimate the market response to each announcement. We split each trading day into 5 minute intervals. Then we take, say, an interval from 9:00 to 9:05 am and use intraday data on the oil futures contract nearest to expiration to calculate an absolute return over this interval separately for each trading day. Finally, we take the average across trading days over the entire sample from 2010 to 2016. Figure 1 shows the results. The red line stands for non-report days which are Monday, Thursday, and Friday. We observe a typical daily pattern of trading, namely, some spikes when the open outcry on the oil market begins at 9 am and closes at 2:30 pm Eastern time. The blue lines correspond to report days: the top panel represents Wednesdays when an EIA report is released, while the bottom panel stands for Tuesdays and API reports. We see that for most part the blue and red line coincide, but at times of report releases (10:30 and 4:30), we observe considerable spikes in absolute returns, meaning that the market reacts quite strongly on average to inventory news. Moreover, we see that the market reaction to API reports (bottom panel) is comparable in magnitude to the market reaction to an EIA release. A slightly muted reaction to API reports can be partially attributed to its restricted access, or alternatively, to the late time of a release, as API reports come out after trading hours. This allows us to conclude that the market generally cares about API releases, and thus neglecting this announcement might be a significant factor in identification of surprises and may introduce a significant distortion in the estimates of market impact.

Finally, to get a sense of a magnitude of the effect, Figure 3 displays scatter plots of cumulative returns for the front month contract over a period of 1 minute around the announcements against our identified EIA surprises (whose identification will be discussed later). We break down the sample into three time periods, from 2010 to 2011, from 2012 to the end of June of 2014, and finally from July 2014 to the end of 2016. A negative relation can be clearly seen. Positive inventory surprises tend to depress oil prices, whereas negative surprises push prices up.Figure 3 also suggests that the strength of these effects evolves over time: the effects became much stronger in later years.

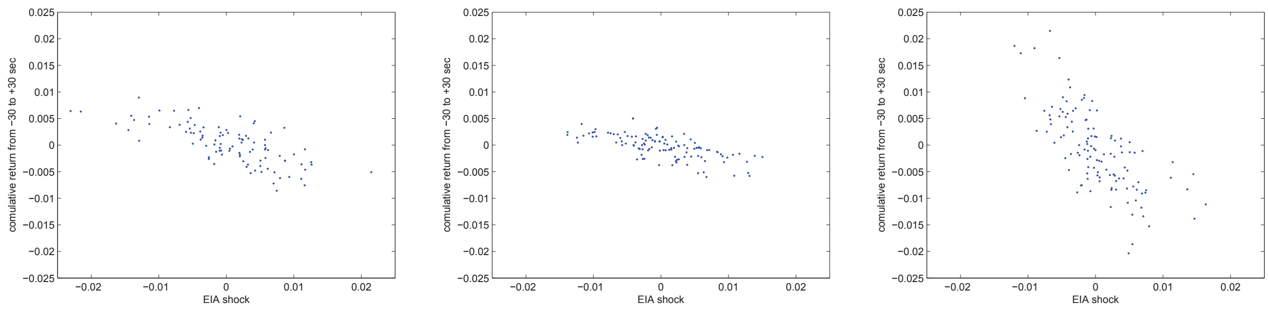

Figure 2: Cumulative returns on the front month contract over 1 minute around EIA announcements (± 30 seconds) against our identified EIA surprises. Left panel: 2010–2011; middle panel: 2012–mid-2014; right panel: mid-2014–2016.

4 Formation of expectations and identification of market surprises

We use EIA announcements to identify inventory surprises and study how they affect the market. This requires identification of inventory surprises and specification of response functions. In both aspects our approach differs from those in the literature.

For identification of inventory *surprises*, we need to estimate the *expected* value of a change in oil inventories. Ideally, we would like to observe this market expectation as close to the moment of the news announcement as possible. One way could be to estimate a simple time series model of weekly inventory changes, and use it to make the forecast.[9] However, it has become common in the literature to use surveys of professional forecasters to proxy for market expectations.[10] Given general public interest in oil inventories, various surveys are available that directly ask about agents' expectations of EIA announced changes, including the surveys conducted by Reuters,[11] Bloomberg,[12] and Platt's. The summary statistics, such as median forecasts, are typically released on Mondays. The pairwise correlations between the three median forecasts are well above 0.97, which suggests that the information content of these three surveys is the same. Thus, we decided to make use only of the Bloomberg median forecast, as the most commonly used in the literature. But, in contrast to other studies, we will also use the numbers released by the American Petroleum Institute.

Before we proceed, let us have a look at the data. The left panel of Figure 3 shows yearly pairwise correlations of EIA announced inventories with the API, as well as BBG median forecasts. The precision of the Bloomberg survey forecast fluctuates quite a lot over time, but no apparent trend is visible. In contrast, the precision of the API signal seems to be improving over time; the correlation increases from 0.6 in 2010 to over 0.8 in 2016. Thus, we should expect that the market would rely more on API information over time.

Expectations formation process Formally, denote by $\Delta \mathrm{Inv}^{EIA}$ the change of inventories (normalized by the total oil inventories for the prior week) to be released by the EIA. We will assume

[9] See Roesch and Schmidbauer (2011).
[10] See Andersen et al. (2003).
[11] See Bu (2014).
[12] See Halova Wolfe and Rosenman (2014), Halova et al. (2014), Miao et al. (2018).

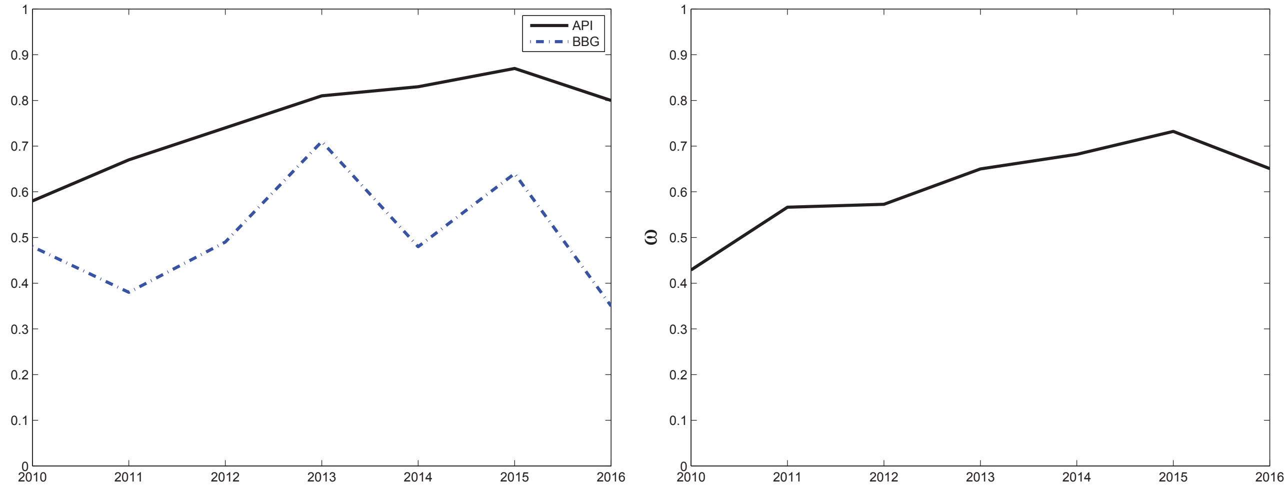

Figure 3: Left panel: correlation of EIA and API estimates of inventories changes (solid line); of EIA estimates and Bloomberg median forecast (dashed line).
Right panel: Optimal weight to be placed on the API signal, ω in each year, under the assumption of independent measurement errors in API and BBG, normally distributed signals, and a representative agent with access to API information.

that the market forms expectations based on realized values of two observable signals: an oil stocks number released by the API, which we denote as ΔInv^{API}, where the time subscript is dropped for simplicity, and the median *forecast* of the survey of professional analysts conducted by Bloomberg, which we denote as ΔInv^{BBG}, where conditioning is on the information set of professional forecasters.[13] For simplicity, for the most part we maintain the assumption of linearity; in other words, we assume that the market puts weight ω on the API signal and weight $(1-\omega)$ on the BBG signal, where $\omega \in [0,1]$. So the market expectation of the oil stocks change is given by

$$E[\Delta\text{Inv}^{EIA}|\Delta\text{Inv}^{BBG}, \Delta\text{Inv}^{API}] = \omega\Delta\text{Inv}^{API} + (1-\omega)\Delta\text{Inv}^{BBG}.$$

To provide intuition, we compute optimal ω for a simple case. Namely, we abstract from the complexities of aggregation of individual beliefs, and solve the optimal signal extraction problem of a single agent with access to API information, under an assumption of independent measurement errors in the API and BBG signals, normally distributed inventory changes and signals (see appendixA). In this case, linearity is known to be optimal, and the optimal weight depends on the signal to noise ratios in the BBG and API signals. We can use the data to calculate these signal-to-noise ratios and find the optimal weight ω. The right panel of Figure 3 shows the results, namely the optimal ωs calculated separately for each year under the assumptions made above. Not surprisingly, the better quality API information is reflected in higher weight place upon it over time. The optimal weight ω monotonically increases from 0.4 in 2010 to about 0.7 over the last three years.

In general, the weight placed on API information by the market as a whole could also depend on the overall access of investors to API reports. Of course, if markets are efficient, it is enough for just one trader to observe and trade on private information for it to be fully revealed in prices. Hence, every market participant in principle should be able to extract that information from observable

[13] Here we are working with expectations, however in the empirical work we will use the median, rather than the average forecast.

movements in the prices. In reality, information percolation may be slow, especially given that API reports come long after main trading hours end. The only problem with this explanation is the fact that the highlights of the API report, at least in terms of total crude oil stocks, tend to be published by major independent news providers, and thus can be freely accessed 5 minutes after the release of the API report. An alternative explanation may be related to heterogeneity in sophistication. Some traders, especially those who do not specialize in oil trading, may be unaware of the existence of this additional private source of information. Thus, we can expect the estimated weight placed on API to be lower than that displayed in Figure 3.

Identification of market inventory surprises Given ω, the market surprise is defined as follows:

$$x = \Delta\text{Inv}^{EIA} - E[\Delta\text{Inv}^{EIA}|\Delta\text{Inv}^{BBG}, \Delta\text{Inv}^{API}].$$

In other words, surprises represent *unexpected* by the market changes in inventories. A positive value of x implies that the market underestimated the change in inventories.

Response function Finally, we need to specify the response function. For example, how should the returns react to a market surprise of x? A usual approach would be to follow Andersen et al. (2003) and assume a linear response function. However, such specification may be too restrictive. Thus, we decided to deviate from the literature and project the universe of our surprises into a number of indicator functions, and then allow the response coefficients to differ in an unrestricted way. We distinguish large positive and negative surprises; in other words, we split all surprises into three bins according to the following rule: the market surprise of x is called

$$\begin{aligned} \text{'a positive suprise'}, &\quad \text{if } x > \bar{x} \\ \text{'an uninformative announcement'}, &\quad \text{if } -\bar{x} \leq x \leq \bar{x} \\ \text{'a negative suprise'}, &\quad \text{if } x < -\bar{x} \end{aligned}$$

where the threshold $\bar{x}$ is to be estimated.

Uniformity of expectation formation process Our data contain a number of futures contracts of different expiry dates. In principle, each single futures contract represents a separate market with its own distinct expectation formation process. That is, one could worry about market segmentation, when each separate market attracts its own unique clientele looking for exposure at just one particular horizon. And if, for example, shorter contracts in general attract less sophisticated traders or traders without access to API reports, we should see smaller weights placed on API reports for contracts with shorter maturities.

Ideally we would like to estimate $\{\omega, \bar{x}\}$ separately for each contract (and each year if there are concerns about stationarity). Unfortunately, we do not have sufficient numbers of inventory surprises per year to estimate all these parameters with reasonable precision. Thus, we proceed as follows. For the benchmark case, we assume uniform formation of expectations at the two most liquid markets – that is, we assume that $\{\omega, \bar{x}\}$ are the same for the two futures contracts with the closest expiration dates, while all other parameters may differ.[14] To check if our procedure is consistent with the data, we also estimate $\{\omega, \bar{x}\}$ separately for each contract and then perform a

[14]Estimation of these other parameters draws abundant information mostly from the dynamics of the variables modeled in no-surprise periods.

	F_1	F_2	F_3	F_4
% of intervals with zero volumes	22	55	82	91

Table 1: Trading activity with maturity measure as a fraction of 5-second intervals with no trading volumes.

likelihood ratio test. If the null hypothesis is not rejected, we will argue that the evidence is in favor of uniformity of the expectation formation process.

5 Econometric framework

In this section we develop an econometric model and outline our estimation procedure. In what follows we describe our model of joint high frequency dynamics of returns, return volatility and trading volumes around announcements. We abbreviate our joint model as ARI-ARCHI-MEMI, in which the additional letter 'I' signifies the presence of announcement indicators. We will show how market inventory surprises identified according to our procedure outlined above appear in our model.

Trading volume To describe the evolution of trading volumes, we use the multiplicative error model (MEM) of Engle (2002); see also Engle and Russell (1998). Denote by V_t the trading volume over 5-second interval t. Then,

$$V_t = \psi_t \varepsilon_t,$$

where $\psi_t = E_{t-1} V_t$ stands for the conditional mean of V_t based on information available up to period $t-1$. One may assume that $\varepsilon_t \sim i.i.d.\, D(1)$, where $D(1)$ is a specific distribution with non-negative support and mean unity.

However, we need a model flexible enough to be applied to both short and long maturity oil futures contracts. The problem with long maturity contracts is their illiquidity, hence we observe substantial periods with no recorded transactions. To illustrate the severity of this issue, Table 1 shows the fraction of 5-second intervals with zero trading volume over the entire sample, separately for each contract. We see that, while for the futures contract nearest to expiration we have only about 20% of intervals with no trading, this number increases to more than 90% for the four-month contract. Clearly, trading activity decreases dramatically with maturity.

Thus, we aim to build a model that can handle illiquidity of long maturity contracts by explicitly accounting for trading inactivity. We follow the approach developed in Hautsch et al. (2013). Namely, we assign a discrete probability mass to the event of no trading, and we allow this probability to be time varying. Formally, we assume that with probability π_t, the shock ε_t is drawn from a continuous distribution with a strictly positive support distribution and density f, say, whereas with the opposite probability $1-\pi_t$, we have $\varepsilon_t = 0$, which corresponds to the absence of any trading in period t. For most of the analysis we will use the generalized gamma distribution[15]

[15] See Lancaster(1990). The generalized gamma distribution has 3 parameters, a, m and λ_0 and

$$f(t) = \frac{a\lambda_0^{am} t^{am-1} \exp(-(\lambda_0 t)^a)}{\Gamma(m)}.$$

normalized so that $E\varepsilon_t = 1$; thus we need $E[\varepsilon_t|\varepsilon_t > 0] = \pi_t^{-1}$. After normalization, the density is

$$f(\varepsilon_t) = \frac{a(\pi_t \xi)^{am} \varepsilon_t^{am-1} \exp(-(\pi_t \xi \varepsilon_t)^a)}{\Gamma(m)},$$

where $\xi = \Gamma(m + a^{-1}) / \Gamma(m)$.

The conditional mean ψ_t is modeled as follows:

$$\psi_t = w + \sum_{k=1}^{q} \alpha_k V_{t-k} + \sum_{k=1}^{q^0} \alpha_k^0 \mathbb{I}_{\{V_{t-k}=0\}} + \sum_{k=1}^{p} \beta_k \psi_{t-k} + \sum_{type} c_v^{type} I_t^{type}.$$

In addition to standard lagged volumes and lagged conditional means, we also add indicators of the absence of trading in previous periods. The coefficients of primary interest are c_v^{type} that stand for the response of trading volume to inventory surprises.

Returns and volatility To investigate the effects of oil inventories on returns, we use the AR-ARCH framework augmented for trading inactivity. That is, with probability $1 - \pi_t$, no trading occurs, and thus the return is equal to zero: $r_t = 0$. With the opposite probability π_t, trading occurs, and the return r_t is drawn from the gaussian distribution with the conditional mean with the following dynamics:

$$\mu_t = \mu + \sum_{k=1}^{q_r} \rho_k r_{t-k} + \sum_{k=1}^{q_r^0} \rho_k^0 \mathbb{I}_{\{V_{t-k}=0\}} + \sum_{type} c_r^{type} \sigma_{D(t),S} I_t^{type},$$

and the conditional variance with the following EGARCH (Nelson (1991)) dynamics:

$$\ln \sigma_t^2 = \omega + \phi \ln \sigma_{D(t),S}^2 + \sum_{k=1}^{p_\sigma} \tau_k \ln \sigma_{t-k}^2 + \sum_{k=1}^{q_{\sigma,1}} \delta_k \eta_{t-k} + \sum_{k=1}^{q_{\sigma,2}} \gamma_k |\eta_{t-k}| + \sum_{type} c_\sigma^{type} I_t^{type},$$

where $\eta_t = r_t / \sigma_t$ are standardized returns. The EGARCH equation has a number of advantages among ARCH models with leverage (see Rodriguez and Ruiz (2012)), one of which is positiveness of conditional variances by construction. We include in the right hand side the daily level of volatility, $\sigma_{D(t),S}^2$, for day $D(t)$ to account for slow moving changes in the volatility. We will calculate $\sigma_{D(t),S}^2$ as filtered realized volatility for day $D(t)$; see Figure 17.[16,17]

A few things are worth discussing. First, the reaction of the conditional mean of returns to news is normalized by filtered realized volatility. The purpose is to simplify the comparison of the results over time. Without normalization we would not be able to interpret intensification of returns reaction as being linked to higher importance of inventory news. This is misleading, as the market could have become more sensitive to *all kinds* of news, not just to inventories. To disentangle fluctuations in overall market sensitivity to news from fluctuations in market response to inventory news, we suggest normalization by the level of volatility. That seems natural, given

[16] For each announcement day, we calculate realized variances using a range of sampling frequencies from 5 sec to 10 minutes, and then take an average. Volatility signature plots for most days are flat in that region of frequencies.

[17] We use the Hodrick–Prescott filter with a smoothing parameter of 1600. Alternative smoothing techniques yield similar results.

that volatility increases when prices respond more strongly to each news arrival.[18] Intuitively, $\sigma^2_{D(t),S}$ picks up the common component, whereas the coefficients c_r^{type} reflect a 'cleaned' market response to inventories news, beyond and above the common component.[19]

Second, even though the gaussian (G)ARCH model (partially) captures heavy tail behavior and may serve as a quasi-likelihood model for consistent estimation of the volatility equation, we also pay attention to conditional tails and utilize Student's t distribution to capture the shape of the conditional density in the tails more accurately. That is, when trading occurs, the return r_t is drawn from Student's t distribution with the shape parameter, degrees of freedom, defined by ν, with the conditional mean and the conditional variance defined as before. The results are similar and are presented in appendix B.

Time varying probability Following Hautsch et al. (2013), we assume that the probability of inactivity may also vary over time. We define

$$h_t = \ln \frac{\pi_t}{1-\pi_t},$$

and propose the following specification for the time varying probability of inactivity:

$$h_t = \omega_h + \varkappa_{D(t)} + \sum_{k=1}^{p_h} \zeta_k h_{t-k} + \sum_{k=1}^{q_h} \xi_k \mathbb{I}_{\{V_{t-k}>0\}} + c_\pi \mathbb{I}_{\{t=t^\star\}},$$

where, again, $\varkappa_{D(t)}$ is the daily component assumed to pick up any changes in average probability of trading over time.

Composite likelihood Thus far, we have specified conditional densities for volumes and returns (and return volatility). Even though this is not a complete model for the joint volumes–returns distributional dynamics, it is sufficient for the purposes of estimation of the parameters of our interest. Thus, we utilize the composite likelihood approach (Varin et al., 2011), i.e. we take the product of the specified conditional densities. The composite log-likelihood function for our model is given, up to a parameter-free constant term, by

$$L = \sum_{t=1}^{T} \left\{ \mathbb{I}_{\{V_t>0\}} \ln \pi_t + \mathbb{I}_{\{V_t=0\}} \ln(1-\pi_t) \right\} + \sum_{t=1}^{T} \mathbb{I}_{\{V_t>0\}} (L_t^{vol} + L_t^{ret}), \tag{1}$$

where

$$L_t^{vol} = \ln a + (am-1)\ln V_t - \left(\pi_t \xi \frac{V_t}{\psi_t}\right)^a - am\left(\ln \psi_t - \ln(\pi_t \xi)\right) - \ln \Gamma(m),$$

and

$$L_t^{ret} = -\frac{1}{2}\ln \sigma_t^2 - \frac{(r_t-\mu_t)^2}{2\sigma_t^2},$$

[18] Of course, alternatively, volatility can increase simply due to more more frequent news arrivals, which is exactly the logic behind (G)ARCH approach to volatility modeling. Thus, we can overestimate an increase in market sensitivity to news.

[19] We use filtered realized volatility, because the realized variance itself is extremely volatile. Although it is reasonable to attribute changes in volatility regimes to changes in market responsiveness to news, it would be hard to argue that high frequency fluctuations in RV also reflect changes in responsiveness.

in the benchmark case, and

$$L_t^{ret} = \ln\Gamma\left(\frac{\nu+1}{2}\right) - \ln\sqrt{\nu-2} - \ln\Gamma\left(\frac{\nu}{2}\right) - \frac{\nu+1}{2}\ln\left(1+\frac{(r_t-\mu_t)^2}{(\nu-2)\sigma_t^2}\right) - \frac{1}{2}\ln\sigma_t^2,$$

when we account for fat tails.

The maximum composite likelihood estimate of the model is the value of the parameters that maximizes the conditional composite log-likelihood.

Testing uniformity of formation expectation process Equation 1 specifies the composite log-likelihood function for a single futures contract. The market inventory surprises appear in this function via indicators for the announcements. Identification of these surprises depends on ω and $\bar{x}$, parameters of the expectation formation process, in a general case unique to the market for this specific futures contract.

As outlined above, we would like to perform a test of whether $\{\omega, \bar{x}\}$ indeed differ significantly across the markets. We use a simple LR test, where the joint log-likelihood function is defined as a sum of composite log-likelihood functions for the first and second month contracts as specified in 1.

Impulse responses Explicit modeling of trading volumes allows us to analyze the *evolution* of trading intensity in response to news. To characterize exactly how people trade, we calculate two types of volume impulse responses.

1. Our first measure is the difference between expected volumes following an uninformative announcement versus following no announcement at all:

$$IRF_h^0 = \underbrace{E\left[V_{t^\star+h}|I_{t^\star}^0 = 1, h_{t^\star-1}\right]}_{\text{uninformative}} - \underbrace{E\left[V_{t^\star+h}|I_{t^\star}^+ = I_{t^\star}^- = I_{t^\star}^0 = 0, h_{t^\star-1}\right]}_{\text{no announcement}}.$$

   where $t = t^\star$ denotes the time of the announcement, and $h_{t^\star-1}$ denotes the history up to the time right before the announcement. Thus, IRF_h^0 measures extra trading due to the announcement itself.

2. Our second measure is the difference between expected volumes following the arrival of an inventory surprise relative to the arrival of an uninformative announcement:

$$IRF_h^\pm = \underbrace{E\left[V_{t^\star+h}|I_{t^\star}^\pm = 1, h_{t^\star-1}\right]}_{\textit{surprising}} - \underbrace{E\left[V_{t^\star+h}|I_{t^\star}^0 = 1, h_{t^\star-1}\right]}_{\textit{uninformative}}.$$

   Thus, IRF_h^+ measures extra trading due to the arrival of a positive inventory surprise.

We have chosen to define the IRFs in the way we do for two reasons. First, our model is not linear. Non-linearity implies that history matters, and does not allow us to shut down all other shocks which occur after the announcement to prevent distortion of propagation properties. As a result, traditional impulse response functions cannot be used in our case. Second, we cannot calculate the generalized IRFs, for example, $E\left[V_{t^\star+h}|I_{t^\star}^+ = 1, h_{t^\star-1}\right] - E\left[V_{t^\star+h}|h_{t^\star-1}\right]$.[20] In our setting, an announcement always occurs at $t = t^\star$, that is, $I_{t^\star}^+ = I_{t^\star}^- = I_{t^\star}^0 = 1$; what is uncertain is the type

[20]See Koop et al. (1996).

of surprise. But, as we do not model inventory shocks, we cannot assign probabilities to these events. Therefore, our measures are somewhere in between traditional and generalized measures. Moreover, when we calculate IRF_h^0, we must rely on our model to reasonably capture the dynamics of the system that would prevail in the absence of any inventory announcement, which makes it a sort of counterfactual exercise.

Time variation in market responses

In general, market responses to news may vary over time. One way to capture time variation would be to split the sample by calendar year and estimate the model separately for each year. That approach is attractive for its simplicity, and for most part of this paper that is exactly what we do. However, following this approach, we will not be able to pin down the exact moment of a change, nor can we assess the speed of adjustment. This requires us to model the transition dynamics directly.

To parsimoniously model the evolution of some parameters of our model, we utilize a threshold autoregression model (TAR). We denote the time variable by d, which in our case represents a particular announcement week. The parameter $x(d)$ varies over time according to

$$x(d) = x_0 + (x_1 - x_0)G(d, d^\star, \delta),$$

where time variability is driven by a non decreasing transition function $G(d, d^\star, \delta)$ that lies between zero and one. When G goes from zero to one, $x(d)$ changes from x_0 to x_1 reflecting the transition. The arguments of the transition function, apart from the time index, are the threshold $d^\star$ and a vector of parameters δ, that control the timing and speed of adjustment respectively.

The simplest possible transition function is the indicator function: $G(d, d^\star) = \mathbb{I}_{d \geq d^\star}$. This specification assumes a structural break in week $d^\star$. In principle, we could allow all parameters to jump simultaneously and search for the optimal timing of the break. This approach, however, is problematic, as the oil market may be subject to other structural changes at the same time. One potential source of a 'background' change could be a dramatic growth of exchange traded funds tracking oil futures prices. An investment flow to ETFs could alter the composition of trading strategies utilized. Thus, as our dynamic model is direct reflection of oil traders' behavior, any changes in the strategies inevitably lead to changes in model parameters. As a result, a simplistic approach might pick up the wrong transition.[21]

To disentangle changes in parameters related to proliferation of new trading strategies, from the changes in market perception of inventory news, we could allow some parameters to follow a more complicated dynamics:

$$x(d) = x_0 + (x_1 - x_0)G_I(d, d_I^\star, \delta_I) + (x_2 - x_1)G_H(d, d_H^\star, \delta_H),$$

where $G_I(d, d_I^\star, \delta_I)$ and $G_H(d, d_H^\star, \delta_H)$ define two separate transition dynamics: *I-transition* corresponds to market perception of inventory news and *H-transition* corresponds to shifts in the composition of trading strategies. The identification comes from differential exposure of the model parameters to the sources of time variation, and potentially from different speeds at which the changes occur. In our notation the coefficients that govern market reaction to inventory news, $\{c_v^{type}, c_r^{type}, c_\sigma^{type}\}$, would follow I-transition only. Similarly, all parameters in the probability equation, $\{\omega_h, \zeta_k, \xi_k\}$, all parameters in the conditional mean equation for trading volumes,

[21]See appendix D for more details.

$\{w, \alpha_k, \alpha_k^0, \beta_k\}$, and all coefficients in the conditional mean equation for returns, $\{\mu, \rho_k, \rho_k^0\}$ and coefficients in the conditional variance equation, $\{\omega, \delta_k, \gamma_k, \tau_k, \phi\}$, will all follow H-transition only.

We will perform two exercises. The first employs the indicator transition functions, $G_I(d, d_I^\star) = I_{\{d \geq d_I^\star\}}$ and $G_H(d, d_H^\star) = I_{\{d \geq d_H^\star\}}$ and we search over all possible combinations of $(d_I^\star, d_H^\star)$ in the time range we are interested in. In the second exercise, we keep the same $G_H(d, d_H^\star)$, but replace $G_I(d, d_I^\star)$ with

$$G_I(d, d_I^\star, \delta) = \frac{1}{1 + e^{-\delta_1 (d - d_I^\star)}} \mathbb{I}_{\{d < d_I^\star\}} + \frac{1}{1 + e^{-\delta_2 (d - d_I^\star)}} \mathbb{I}_{\{d \geq d_I^\star\}}.$$

This specification allows for different speeds of adjustments before and after the threshold. After having fixed the threshold at $d_I^\star$, we estimate $\delta = (\delta_1, \delta_2)$.

6 Data and estimation

Data In this paper, we utilize changes in weekly U.S. ending stocks excluding SPR and including lease stocks of crude oil as published by the Energy Information Administration.[22] For identification of market surprises we use the estimates of inventory changes published by the American Petroleum Institute, as well as the median consensus forecast from the Bloomberg survey of analysts.

The high frequency data on WTI oil futures traded at NYMEX (CME group) was obtained from TickData. Our sample covers the period from 2010 to 2016[23] and after cleaning contains 337 announcement days. We focus on one hour around EIA announcements, that is, from 10 to 11 am. We believe that the one hour long interval is long enough to obtain a reasonable estimate of the parameters of our dynamic model. The sampling is done at a 5-second frequency, which yields 720 data points for each announcement day.

Dealing with futures contracts Before we proceed to estimation, we need to deal with a number of issues specific to the futures market. One such feature is expiration. At any moment in time, NYMEX offers a set of contracts that differ by delivery month. The expiry dates range from one month up to nine years in the future, thus constituting more than 100 contracts at any given moment. Trading in the current delivery month ceases on the third business day prior to the twenty-fifth calendar day of the month preceding the delivery month (for example, the last trading day of a February-2018 contract is January 22).[24] Fixed expiration dates mean that the maturity of a contract varies over time, as does the open interest and trading volume. Hence, we should be concerned about stationarity; after all, our dynamic properties of futures prices and volumes reflect the composition of trading strategies utilized at the moment. We follow a standard approach in the literature and rely on a rolling procedure to create continuous futures contracts. In particular, we replace the expiring contract with the next one on the 5th day of each month. Thus, the maturity of what we call the first month contract in our sample ranges from 5 weeks to 2 weeks, when the soon-to-expire contract is replaced with the next one, and a new maturity

[22] The EIA terminated publication of this series in September 2016. For the last few months, we use the weekly U.S. Ending Stocks excluding SPR and add the most recent available data on lease stocks. Lease stocks have been relatively stable with a range of 30.6 mln barrels to 33.1 mln barrels from January 2014 through June 2016.

[23] We only have data on API releases from 2010.

[24] See contract specifications at the CME website.

cycle commences. As long as oil traders follow the same procedure and shift their trading from one market to another at about the same time, stationarity concerns are alleviated.

The estimation of market responses for long-maturity futures contracts requires modification of our approach. Even though our model is specifically designed to handle a certain level of illiquidity, once we move beyond the first three futures contracts, liquidity drops considerably and precludes any meaningful analysis. Fortunately, we find that certain contracts maintain reasonable liquidity throughout the year: December and June contracts (CLZ and CLM). Disproportional interest in these contracts perhaps reflects the convenience that mid year and end of the year expiry brings to reporting of hedging procedures. Alternatively, it may simply reflect certain market coordination over time. But if we plan to utilize contracts with fixed maturity dates, rather than continuous contracts, we must reconsider the stationarity issue. If we stick with a particular contract, then necessarily its maturity would vary substantially. For example, if we analyze the trading of a contract that expires in December of 2016, its maturity would range from 23 months in January of 2015 to just 12 months in December of 2015. It is highly unlikely that a two-year contract is traded in exactly the same way as a one-year contract, and thus we do not have stationarity.

One way to mitigate this concern would be to restrict the sample to only a few months, e.g., to consider only three months from January to March of 2015, and thus limit the variation in maturity to 23-21 months. However, we would have to significantly sacrifice the sample size. To compensate, we combine the same quarters of different years into one sample. That is, we combine 2010, 2011, and 2012; then 2013 and 2014; and finally 2015 and 2016. This choice of sample partition is driven by our previous estimation results that show similarity of estimated coefficients across these years. Thus, we restrict changes in parameters over time, but in return we partially account for changes in parameters across maturities. Most importantly, this procedure allows us to make predictions for the longer end of the term structure curve by using contracts with sufficient liquidity.

Normalization The second issue is normalization of trading volumes. Our model is not rich enough to account for potential day-to-day changes in the composition of traders, or strategies utilized, or intensity of trading. For example, trading may be slow in August, or on the days before major holidays. Alternatively, oil traders' awareness of commodities (as of asset class) may increase over time, and thus market characteristics might not be constant. Thus, we decided to normalize trading volumes to eliminate any differences between trading days. For each announcement day we calculate the total volume accumulated from 8 am until 6 pm, i.e. over 10 trading hours. Then we calculate the mean daily volume, and for each day we can now find the volume ratio as volume traded within this day to the mean daily traded volume. Finally, high frequency trading volumes are normalized by this ratio. Formally, denote by $v_{t,d}$ the trading volume observed over a 5-second interval t in day d. Total volume traded this day is given by $v_d = \sum_{t=1}^{N} v_{t,d}$, where $N = 7200$ is number of data points. The mean daily volume equals $\bar{v}_d = D^{-1}\sum_{d=1}^{D} v_d$, where $D = 337$ is number of days, and thus, the volume ratio equals $R_d = v_d/\bar{v}_d$. Finally, normalized volumes are given by $\tilde{v}_{t,d} = v_{t,d}/R_d$. The total volumes traded over one hour around the announcement can still differ from day to day (we normalized total daily volumes), but these differences can now be attributed to the direct effect of the news.

Estimation We utilize numerical methods to solve for the optimal parameter values. As a part of our optimization routine, we utilize a simple grid to search over possible values of ω and $\bar{x}$. For most parts of the paper, we split the sample by a calendar year and focus on the two most liquid

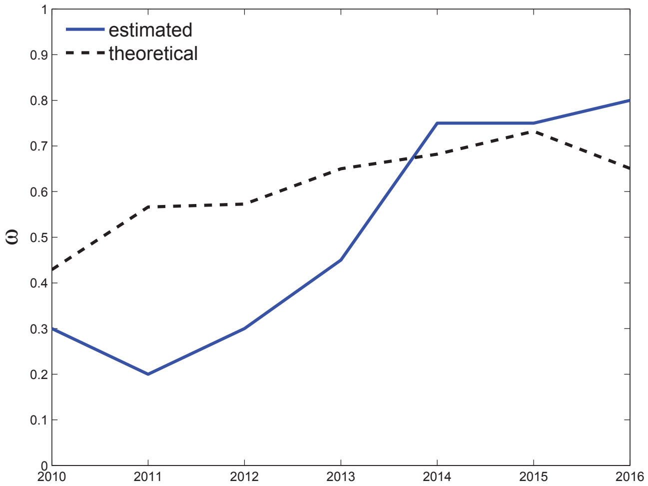

Figure 4: Estimated weight ω of the API signal. The solid blue line corresponds to ω estimated jointly for the first two most liquid contracts. The black dotted line displays the value of ω which represents a solution to the signal extraction problem.

futures contracts with the nearest expiration dates. We consider other contracts and alternative partitions of the sample when we study the inventory effects on the term structure of futures prices.

7 Results

Before characterizing the market response to inventory news, we analyze the expectation formation process in the oil futures market.

7.1 Irrelevance of maturity in expectation formation process

Figure 4 displays the estimated value of ω, the weight that the market places on API reports for the first month futures contract. We can see that the weight increases from 0.3 in 2010 to 0.8 in 2016, suggesting an increase in the importance of API information. To some extent that can be attributed to a better accuracy of API reports in recent years. To illustrate the improvement, the dashed black line depicts ω_{theory}, the optimal weight calculated based on the observed quality of API reports and BBG signals. The upward trend can clearly be seen. However, the lines do not always coincide, the market used to significantly underweight API signals in the early years, and only from 2014 do we see convergence. To some extent, ω can also be interpreted as a fraction of people willing to pay for API reports. Given that the cost of subscription has not changed significantly, we can conclude that it was the value of being informed that increased in 2014.

The second observation that we make is that the estimated values of ω happen to be the same for the first two most liquid futures contracts for all years except 2016.[25] Moreover, when we repeated the exercise for the United States Oil Fund, the exchange traded fund that tracks oil futures prices, and then for two oil producing companies, Exxon Mobil and Chevron, we obtained the same estimates. This is an important result, as it reassures us that we identify the market surprises correctly. Moreover, we can interpret our findings as evidence of the uniformity of rules traders

[25] In 2016, the estimated value of ω for the first futures contract is 0.80, and for the second – 0.95. The LR test performed using composite likelihood functions rejects the null hypothesis with a p-value of 0.0007.

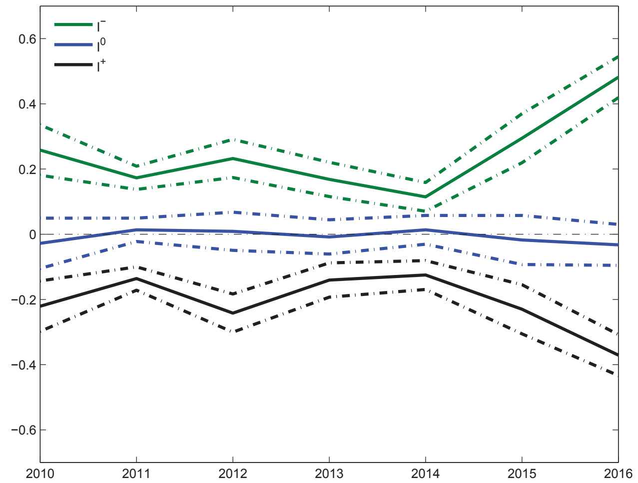

Figure 5: Return reaction to inventory surprises. The solid line shows estimated return coefficients on inventory surprises for the front month contract: c_r^+ in black corresponds to positive surprises, c_r^- in green corresponds to negative surprises, and c_r^0 in blue – to uninformative announcements. Dashed lines represent 95% confidence intervals.

follow when forming expectations. As discussed, one could worry about market segmentation, namely that each separate market attracts its own unique clientele looking for exposure at just one particular horizon. Different compositions of traders may form expectations in different ways, either due to differences in their sophistication levels, or because of differential access to API information. However, our results suggest that this is actually not the case, at least for the most liquid futures contracts.

Finally, the estimated value of the threshold $\bar{x}$ that defines the split of market surprises into categories was equal to 0.5% and we did not find any significant variation across years or contracts.

7.2 Market response to oil inventories

In this section we characterize the oil market response to inventory news. We present and discuss the estimated values of $\{c_v^{type}, c_r^{type}, c_\sigma^{type}, c_\pi^{type}\}_{type}$ that reflect, respectively, trading volumes, returns, and return volatility reaction to unexpected oil inventories changes.

Negative relation between inventories and returns

Our results indicate a strong negative link between oil inventory surprises and returns. To illustrate this link, Figure 5 displays the return coefficients on inventory surprises for the front month contract, $\{c_r^+, c_r^0, c_r^-\}$, and the 95% confidence intervals. As expected, we observe a clear negative relationship. The coefficients corresponding to large inventory surprises are all statistically significantly different from zero. As for the magnitude of the effects, we find that an unexpected decline of inventories by 0.5% or more in 2016 would cause the oil price to immediately increase by about 1.2%.[26] The effect is lower for other years, but is still quantitatively large.

A theory of competitive storage in its modern form was developed in Deaton and Laroque (1992). As long as the non-negativity constraint does not bind, inventories serve to smooth out

[26] To calculate the average effect, we need to multiply c_r^{type} by the average level of filtered realized volatility.

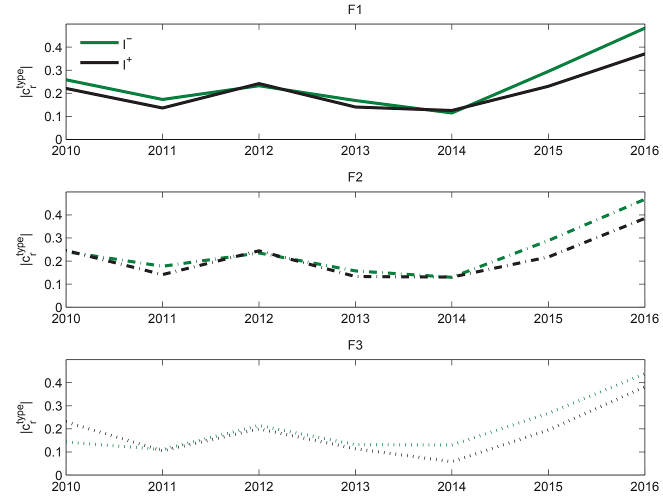

Figure 6: Absolute values of return responses to inventory news. Each panel corresponds to one of the three contracts with the nearest expiration dates.

temporary demand and supply shocks. Oil flows in or out of storage until current and expected future prices are equalized (up to interest and storage costs). Facing a positive supply or a negative demand shock, owners of inventories move oil out of the abundant market and store it for the future, equalizing current and expected oil prices. As a result, the current spot price needs only to decrease slightly to clear the inelastic market. Thus, a negative link between inventory surprises and returns is to be expected.

Time variation and asymmetry in market responses

Figure 5 also indicates much stronger market reaction to inventory news since 2014. In absolute terms, the difference in magnitudes is striking: the returns reaction is 2.5 times larger in 2016 than in 2014. The normalization of returns reaction by volatility allows as to attribute this increase to the greater importance of inventory information, and not just higher overall market sensitivity to news.[27].

From a theoretical perspective, the intensification of market reaction may seem natural. Early in 2015, oil inventories reached unprecedented levels; see Figure 13. The positive supply and negative demand shocks were so large that the market became oversupplied and oil producers simply could not market the entire supply of crude and much of it had to be stored. When the space in storage facilities ran out, speculative activity became problematic. Oil speculators could not respond to any additional temporary positive supply or negative demand shocks by taking extra oil off the market and putting it in storage. Thus, the smoothing mechanism that could prevent spot prices from decreasing too much became limited, if not entirely shut down. Therefore, at times of high inventories, we should expect to see intensified market reaction to unexpected positive inventory surprises, which is exactly what we find.

The problem, however, is that the market reaction to negative inventory surprises increased

[27] To illustrate that fluctuations in market sensitivity to news may be substantial, Figure 17 shows realized variance estimated using high-frequency returns on the front-month futures contracts. We can see a distinct change in the volatility regime at the end of 2014, when volatility increased dramatically and remained high for a long period of time. So, to proceed, we normalize the return reaction by the filtered and smoothed version of the estimated realized variance.

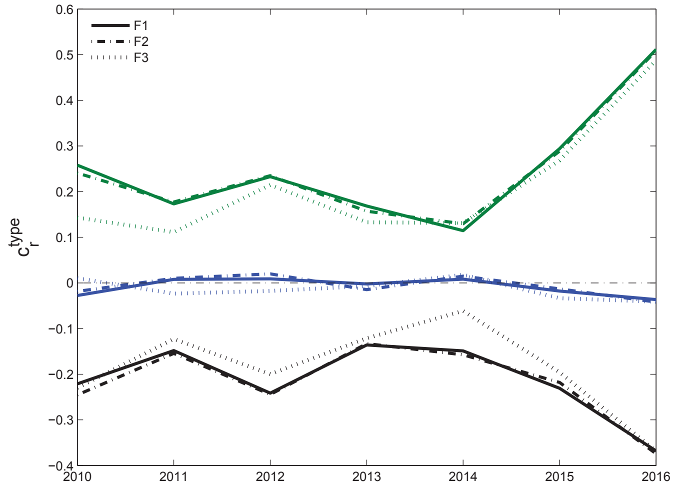

Figure 7: Return reactions to inventory surprises for the first three futures contracts.

even more than the market reaction to positive inventory surprise! For the sake of visibility, Figure 6 plots the absolute values of estimated coefficients that govern returns reaction to news. Until 2014, there is no systematic difference between market responses to negative and positive news. However, in the last two years, the market reacts much more strongly to negative inventory surprises. That is problematic, because a negative supply or a positive demand shock could still be counteracted by speculators releasing some of the oil into the spot market, which should have limited the oil price responses. The standard inventory theory fails to explain the joint time variation and asymmetry pattern.[28]

Alternatively, the market *perception* of inventories surprises may have changed in 2015. The oil futures market experienced a significant change in the composition of oil traders. At the end of 2014, investment in oil exchange traded funds increased dramatically. ETFs offer a simple alternative to trading oil futures, and thus tend to bring less sophisticated traders into the market. It is only natural to assume that traders differ in knowledge and experience, and thus may form beliefs in different ways. As the composition of traders changed, so may have the overall market perception of inventories. Indeed, news of oil inventories have been continuously hitting the wires, and all major news outlets regularly publish highlights of the EIA reports. Given that real time information about the oil market is scarce, rookie oil traders could have overestimated the importance of inventory news, and failed to understand the implications of a binding constraint. As a result, we observe stronger and asymmetric market reactions to inventory surprises.

Effects of inventories of the term structure curve

In this section we present one of the main results of the paper. We investigate whether inventory announcements alter the futures term structure curve.[29] The term structure adjustments depend on the perceived persistence of shocks. If the shock is perceived as temporary, its effect should

[28]The evidence of asymmetry in the literature is mixed. Miao et al. (2018) find no evidence of an asymmetric impact of inventories on futures prices. Similarly, Ye and Karali (2016) document that the impact of positive and negative crude oil inventory shocks are not statistically different from each other. In contrast, Bu (2014) finds that negative shocks cause stronger movements of the market, whereas Halova et al. (2014) investigates cumulative returns and argues that positive inventory surprises tend to be followed by larger price moves.

[29]The comparison of market impact across maturities is possible as identified inventory surprises are the same.

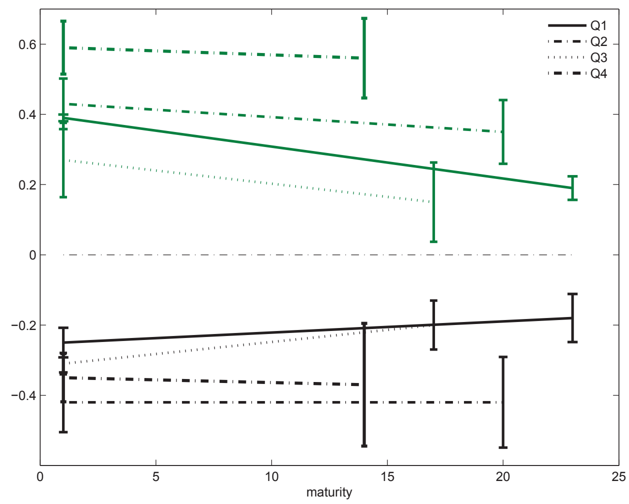

Figure 8: The effects of inventory shocks on long term contracts 2015-2016. The return responses to positive inventory news are in black, and to negative in green. Each line represents the estimation results for a particular quarter, and consists of two points corresponding to two contracts, one of which is the first month contract (maturity equals one month), and the other is the CLZ contract that expires the next year in December (maturity ranges from 23 months to 12 months). Maturity is calculated at the beginning of the period, and is measured in months.

vanish with maturity, and thus we should not see any movements of long term futures prices. As estimation of our model for long term contracts requires a slight modification of our approach, we present the results separately for the short and long end of the term structure curve.

We start with the short end. Figure 7 displays the estimated coefficients c_r^{type} for the first three futures contracts by maturity. The estimates for the first and second month futures contracts coincide, implying that the prices of both contracts adjust by exactly the same amount in response to inventory news. Hence, the term premium remains constant at the short end of the curve. The dotted line corresponding to the third month futures contract slightly diverges at times, but clearly follows the same adjustments.

The estimates for the longest possible contract that expires the next year in December and for the subsample that covers 2015 and 2016 are plotted at Figure 8. Each lines contains just two points, representing the shortest and the longest parts of the term structure curve as estimated separately for each quarter. All lines bend towards zero at the longer end, implying that the price of the next year's December contract is less reactive to news relative to the price of the front month contract. However, the difference is small; for many quarters the curve is flat or almost flat. The only sample in which we observe strong weakening of the market response to news with maturity is the sample of the first quarters of 2015 and 2016. The results for alternative partitions of the sample, as well as for other long term contracts, including the current year December contract, and the next year June contract, are similar and presented in Appendix C.

In sum, our results indicate an absence of any effect of inventory news on the term premium on the short end of the curve and only a weak effect for contracts with longer maturities. In other words, the entire term structure curve moves *in parallel* in response to inventory news. Under a reasonable assumption of constant risk premium over the event window, the shift in futures prices can be attributed to uniform *revision of traders expectations* of future oil prices.[30] We interpret

[30]The risk premium may also adjust in response to inventory news, but it is extremely unlikely that it would

this as follows: the traders view most inventory changes as reflecting *permanent* or at least long lasting shocks that hit the market.

Our results provide evidence in support of Hamilton's (2009) conjecture that, in the presence of inventories, all oil price movements have to be permanent and unpredictable. If oil prices are expected to rise, owners of storage facilities have an incentive to purchase oil, put it in storage, and sell it later at a profit. Such speculative transactions simultaneously push up spot oil prices, and tend to decrease futures prices, as more oil is brought to the market in the future. Hence, inventories serve to smooth out demand and supply shocks over time. As a result, every shock that hits the market ends up moving the entire futures curve either up or down, with no adjustments of the term structure. Our results are consistent with the theory.

However, the equalization of prices crucially depends on speculation activity to move oil across periods. When spare capacity is near exhaustion, or when inventories have been depleted, speculation activity becomes limited. In this case, temporary shocks in the direction of the binding constraint can no longer be smoothed out, and the spot oil price strongly adjusts to clear the inelastic market. The above narrative suggests that when inventories are binding, we should see term structure adjustments. In particular, we would expect to see the effects of positive inventory news on the term structure since 2015. However, our results do not indicate this. Even at times of binding inventories, the futures prices move uniformly in response to news.

The lack of any effect on the term spread when inventories are high is surprising and contradicts the conventional wisdom. Indeed, all recent episodes of high inventories in the oil market have been accompanied by a widening term premium, especially at the shorter end of the term structure curve. Two examples are traditionally used as anecdotal evidence. In 2008, a negative demand shock created an abundance of oil and depressed spot oil prices. The term structure curve became upward sloping and especially steep at the short end. Similarly, in 2014, the market was hit by a positive supply shock due to rising shale oil production and by a negative demand shock due to the slowdown of the Chinese economy. The market again experienced large term premiums, but over a much longer period of time. In both cases, a steep term structure curve was attributed to the presence of an excess supply of oil in the market. However, we find that when inventory news announcements came, traders did not revise expectations accordingly. We believe that our results justify the search for alternative mechanisms of the term premium determination.[31]

7.3 Volume response

Explicit modeling of trading intensity allows us to analyze idiosyncratic trading behavior of market participants around inventory announcements. In contrast to other papers, we not only analyze the instantaneous impact of inventory news, but also characterize trading patterns following the event.

In our model, on impact, trading volume can spike due to jumps in both the conditional mean of volumes, ψ_t, and the probability of trading, π_t. The magnitude of these jumps is defined by the coefficients, $\{c_v^{type}, c_\pi\}$. The propagation, however, will depend not only on the dynamic properties of volumes and probability of trading, but also on the response of the conditional variance to news, driven by $\{c_\sigma^{type}\}$, and it subsequent evolution.

move in the opposite way to expectations revisions, by the exact amount needed to keep the term premium constant for all maturities over our sample of 7 years.

[31]See Selezneva (2015) for an example of such a mechanism.

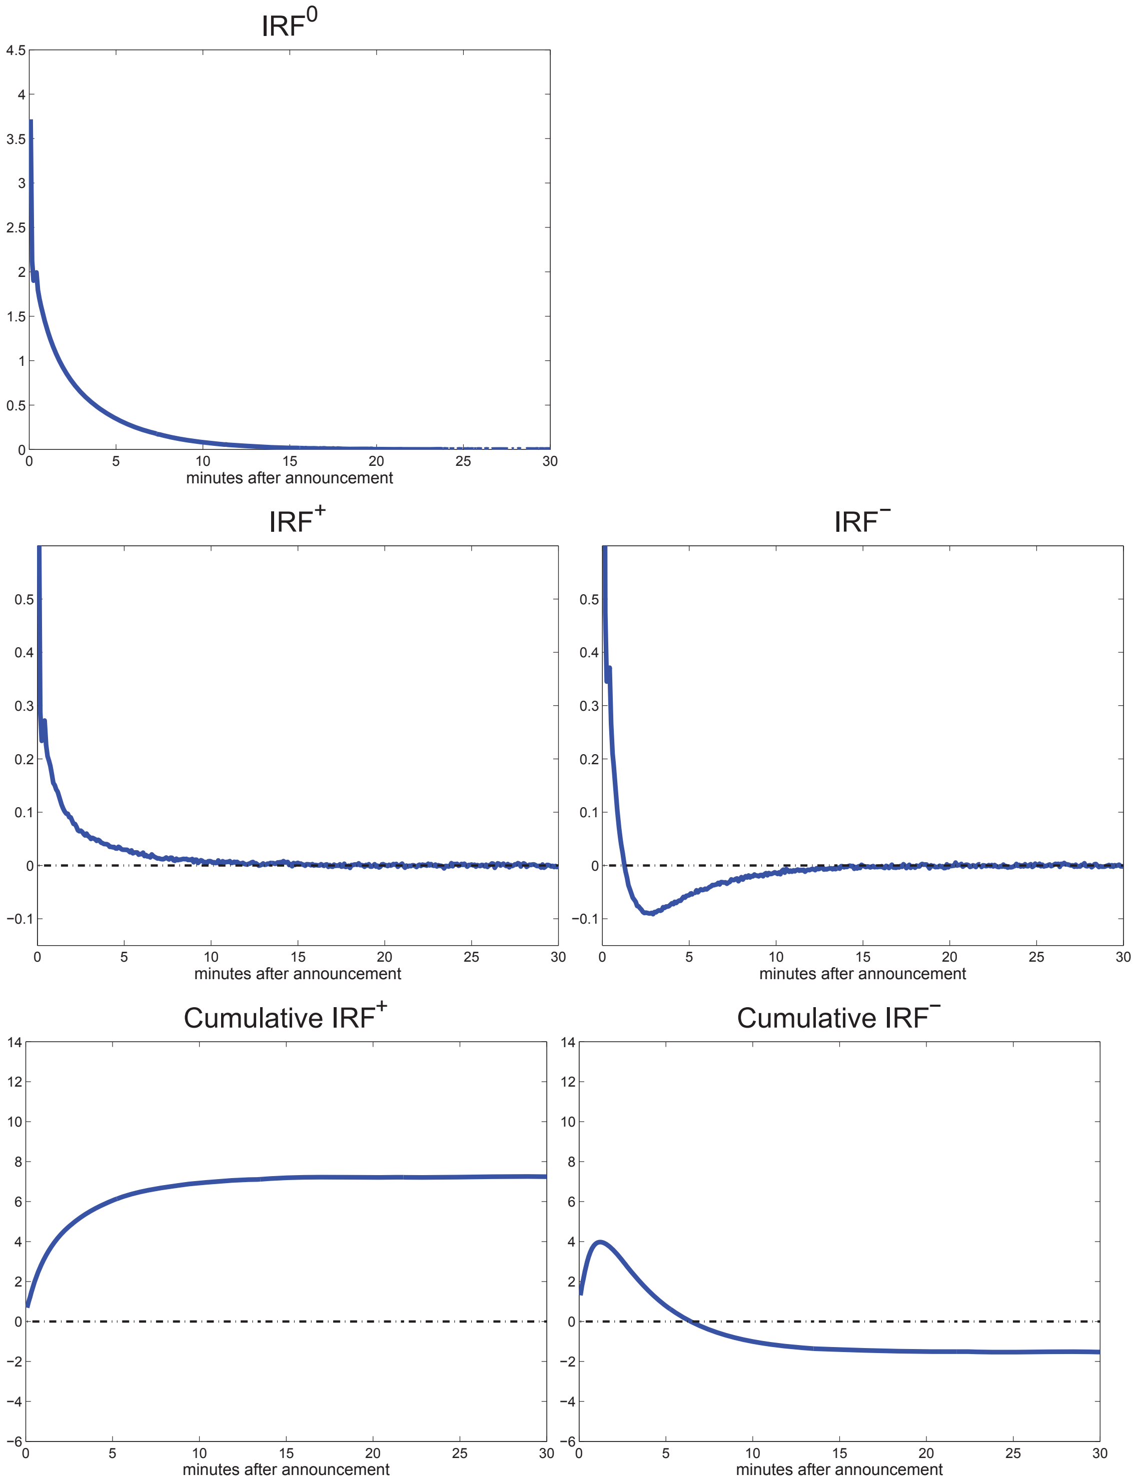

Figure 9: Volume impulse response functions, estimates for 2010, front-month futures contract.
Top: Volume response to uninformative announcement.
Middle: Extra trading volumes following a positive (negative) inventory shock relative to uninformative announcements.
Bottom: Extra cumulative trading volumes following a positive (negative) inventory shock relative to uninformative announcement.

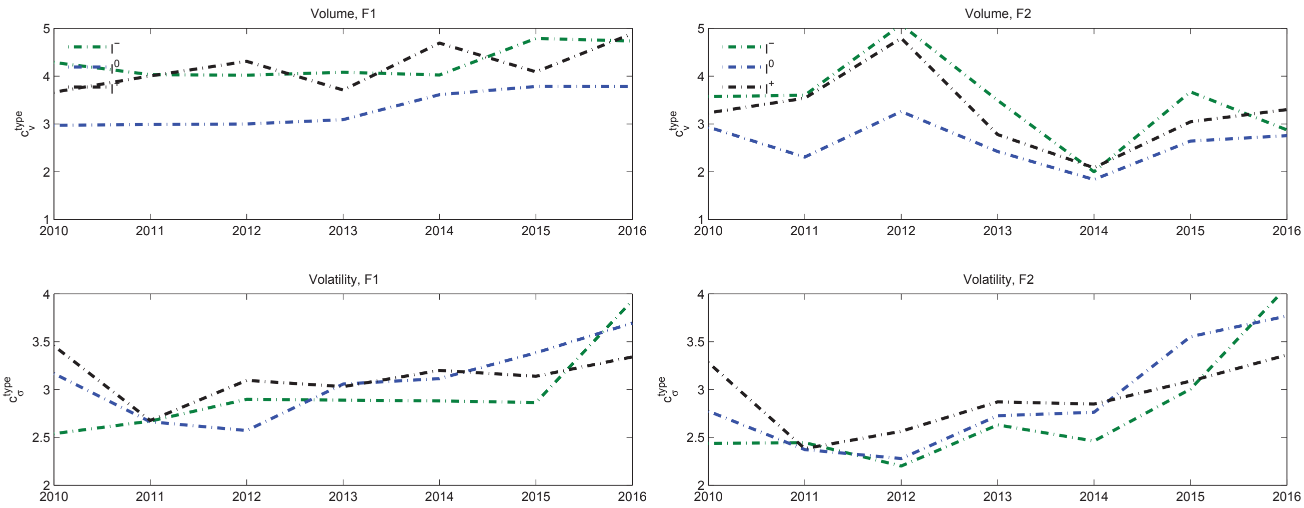

Figure 10: Trading volume and volatility of returns reactions to inventory surprises by a type of surprise. First month futures contract appears on the left, and second month futures contract appears on the right.

Trading patterns in the oil market To give a general idea of how people trade in the oil market, Figure 9 draws the volume impulse response functions for the front month contract for 2010. The top picture shows extra trading upon the arrival of the inventory report, IRF_h^0. We observe a strong immediate reaction followed by a gradual decay afterwards. Quantitatively, our results indicate that it takes about 15 minutes for an uninformative 'announcement shock' to be fully traded away.

The next two figures show additional trading volumes due to arrival of an inventory surprise, either positive, on the left, IRF_h^+, or negative, on the right, IRF_h^-. Clearly, the pictures suggest the existence of two completely different trading regimes.

The average market reaction to positive news in 2010 reflects what can be called a typical pattern. The arrival of an inventory surprise immediately intensifies trading even further (on top of what is predicted by IRF_h^0). The extra trading activity remains uniformly positive, and on average becomes negligible 10 minutes after the event. When we integrate the trading intensity over time, as in the bottom left figure, we see that people trade *more* when they are surprised by a positive inventory shock. That result is intuitively appealing.

In contrast, the average market reaction to negative news in 2010 is puzzling. Even though, on impact, an inventory shock intensifies trading activity, after less than a minute it drops dramatically, and even becomes negative. Moreover, the cumulative trading volume ends up lower following an inventory surprise. People trade *less* when they are surprised by a negative inventory shock.

Do people trade differently across years? How typical are the trading patterns that we found?

As for immediate responses to inventory news, our results draw a robust picture. We always document (i) intensification of trading activity irrespective of announcement type, and (ii) higher trading following surprises. The point estimates of $\{c_v^{type}, c_\sigma^{type}\}$ for the two most liquid contracts are presented in Figure 10. There is a strong reaction of volumes to uninformative announcements, depicted by the blue lines on the top two pictures. Even when the market correctly predicts

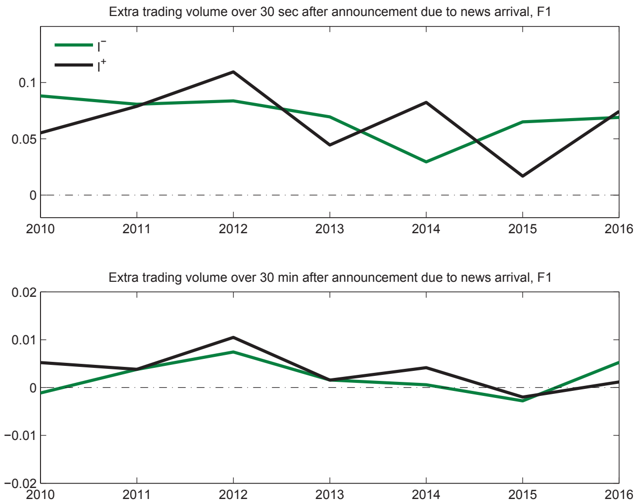

Figure 11: Extra trading volume due to the arrival of new information: the difference between cumulative (log) trading volumes following a large inventory shock and following an uninformative event, by year, %.

inventory change and prices do not adjust, we still see intensification of trading activity. Instances of excessive trading volumes not accompanied by movements in prices have been documented in the literature before, thus our results reconfirm previous findings.[32] However, when the market is surprised, we see even more active trading: green and black lines representing market responses to negative and positive surprises are uniformly above the blue line corresponding to uninformative events. A similar pattern is observed for the reaction of conditional variance and probability of trading to news.[33] Finally, in contrast to returns responses, the volume responses are symmetric: large negative and large positive surprises cause lead to similar trading intensity.

However, we find considerable variation in dynamic trading patterns over years. Figure 11 summarizes the results. We decided not to display $IRF_h^{\pm}$ for each year; instead we show two consecutive snapshots of extra cumulative volumes triggered by the arrival of inventory surprise, taken 30 seconds and again 30 minutes after the report. Black (green) lines correspond to positive (negative) shocks. When we consider the short-term reactions, for all years and for both contracts, the lines are above zero, which reflects the previous finding that news arrivals intensify trading on impact. In contrast, the second row shows the differences across years; in particular, the year 2015 stands out, as surprising inventory shocks triggered lower overall trading. The figure is designed to illustrate the importance of distinguishing immediate trading reactions from dynamic trading.

From the modeling perspective, it is the effect of inventory news on conditional variance of returns that triggers the differences across years. When uninformative events cause larger jumps in conditional variance, we observe more intensive trading over a longer period of time, as conditional variance is persistent. Overall, we find that modeling trading probability is crucial to understand and measure the propagation of inventory surprises. Not accounting for time-variability of trading

[32]See Kandel and Pearson (1995) and Banerjee and Kremer (2010)

[33]For the first futures contract the probability of trading is already close to unity around EIA news reports. Our estimates of $\{c_\pi\}$ for the second futures contract imply that the probability of tradings jumps all the way to unity. For instance, in early years, the probability of trading happens to be on average about 0.6 before the event, but jumps to 1 immediately after the announcement. Similarly, in 2016, the probability of activity increases from 0.9 to 1. The results are available upon request.

activity significantly underestimates the total effect of the shocks.

Interpretation of results: implications for disagreement and learning Our results for the dynamic behavior of trading volumes can be summarized as follows: (i) on impact, trading activity intensifies irrespective of announcement type, but it is (ii) higher following surprises. In addition, we document that typically (iii) trading activity is uniformly larger over the first half hour after the arrival of an inventory surprise. However, sometimes, as in 2015, we observe the reverse, namely that (iv) the trading activity quickly drops and cumulative trading ends up being lower.

One explanation for why trading volumes can be decoupled from movements in returns is disagreement among traders – that is, agents disagree about the interpretation of public information or have heterogeneous priors. When news comes, agents revise their beliefs individually, potentially even in opposite directions. As a result, agents can trade and adjust their asset positions according to their individual changes in beliefs.

Banerjee and Kremer (2010) develop a dynamic model of trade by assuming that investors differ in their interpretation of public signals. When a public signal arrives, it initially causes a large disagreement and leads to idiosyncratic trade. The more investors disagree, the larger would be the immediate volume response. But as investors' beliefs converge, volume gradually decays to normal levels.

The framework developed by Banerjee and Kremer (2010) takes the degree of disagreement as exogenous, and thus puts no restrictions on its initial value, nor on its evolution after the event. Given this flexibility, most of the empirical facts that we document, including reactions to uninformative news and gradual dissipation of trading activity following an event (or autocorrelation of volumes), can be easily generated by their model. However, the remaining two facts are harder to replicate and require some extra very strong assumptions on the degree of disagreement. First, to account for stronger trading in response to large inventory surprises, once has to assume that oil traders have more extreme interpretations following large inventory shocks. Second, to generate the atypical trading behavior that we observe in 2015, one has to assume occasional *faster convergence of beliefs* following an inventory surprise.

The question remains whether a more microfounded model can generate such specific behavior of disagreement endogenously. Naturally, disagreement among traders should be driven by their information acquisition and information processing behavior. Some of the identified properties of disagreement, such as stronger reactions to inventory surprises, may arise from the state-contingent information acquisition. The realization of a large inventory shock forces agents to allocate more attention to process the news, and thus potentially simultaneously triggers larger initial trading and faster belief convergence.

Applied theory literature has recently made considerable efforts to study dynamic information acquisition. For example, Banerjee and Breon-Drish (2017) consider information acquisition as a real option. In their model, the information acquisition decision is a cutoff rule; the agent chooses to acquire information only when public uncertainty reaches a threshold. When uncertainty about the asset payoff is sufficiently high, the expected profits from being informed are sufficiently large. In our case, the arrival of an inventory surprise can increase the uncertainty enough to trigger information processing, which in turns intensifies trading and simultaneously speeds up the convergence of beliefs.

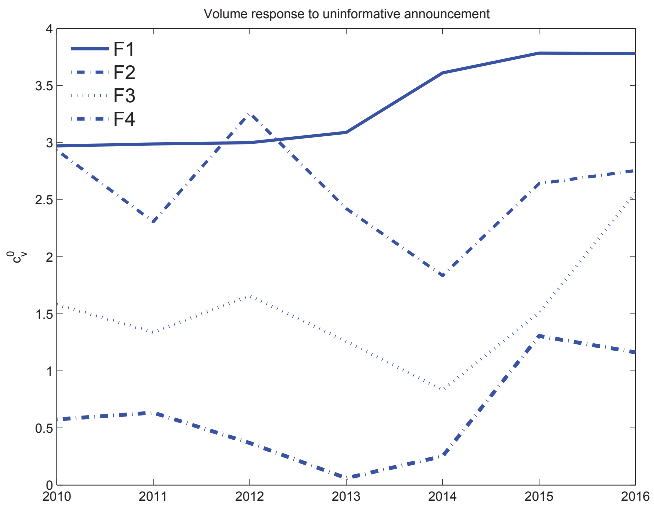

Figure 12: Volume responses to uninformative announcements by contract maturity.

Trading intensity by contract maturity Finally, we compare the trading patterns of contracts with different maturities. Figure 12 plots the volume response to uninformative announcements for the first four futures contracts. Volume reaction clearly varies over time. We also see that the first futures contract stands out from the rest. While the volume response of the first futures contract monotonically increases over time, speeding up in 2014, the second month contract shows more variation and a clear downward trend from 2012 to 2014. The same pattern is observed for the third and forth month contracts. Differential volume reaction suggests some form of market segmentation. That is, even though *average market expectations* are formed uniformly across contracts (as our previous results have shown), revisions of individual beliefs in response to news are unique for each maturity and evolve uniquely over time. If markets attract different clientele we might expect differences in information acquisition and processing, and thus differential immediate volume reactions, as well as differential speeds of volume decaying after the public signal.

We argue that one reason for the lack of uniformity in higher moments is each contract's unique exposure to financial innovation. New financial instruments, such as exchange traded funds, proliferated in the last decade. The flow of investment in these funds varies over time and reaches significant amounts. However, ETFs offer exposure only to a limited set of contracts, usually the most liquid ones. For example, one of the largest ETFs is the United States Oil Fund, which invests in the first futures contract, and rolls over the entire portfolio about two weeks before the expiration. Thus, by attracting a unique composition of traders and users of specific strategies and accommodating their investment needs, the funds facilitate market segmentation. Moreover, as investment flows to such an instrument fluctuate quite substantially over time, the degree of market segmentation may vary over time as well.

8 Understanding the recent glut of oil

In this section we illustrate how our methodology can be applied to test existing oil pricing theories. In particular, we investigate the development of oversupply fears in the US.

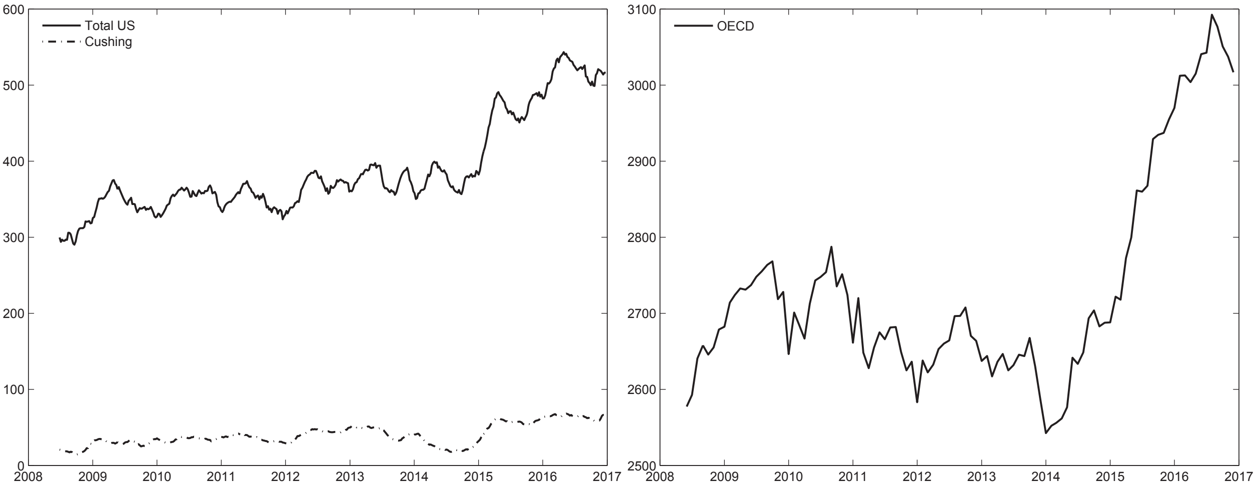

Figure 13: Left panel: Weekly U.S. ending stocks of crude oil excluding SPR (mln barrels, solid line). Weekly Cushing, Oklahoma ending stocks of crude oil excluding SPR (mln barrels, dashed line).
Right panel: Monthly OECD commercial crude oil and other liquids inventory (mln barrels).

Structural shift in oversupply beliefs in 2014

First, let us briefly describe the chronology of events. By 2014, US oil production exploded, reaching almost 9 mln barrels per day, and it was expected to grow even further. In addition, struggling economies in China and Europe raised doubts that it would be possible to maintain the same pace of demand growth as before, despite the fact that the oil prices remained high. However, in July of 2014, the price of oil started to decline. By November the oil price had fallen by 30%, and the term structure of oil prices became upward sloping. On November 27, 2014, OPEC announced their decision to maintain production levels, and soon afterwards the price of oil crashed even further. By the beginning of 2015, oil inventories reached unprecedented levels and were interpreted as a sign of an immense oil oversupply (see figure 13).

An oversupply of oil soon became considered as the main drive of the extreme movements of oil prices. But how exactly did oversupply fears develop? Arezki and Blanchard (2015) argue in favor of a structural shift in market expectations after the OPEC meeting in November of 2014. The growing oil production and potentially weakening demand had been observed long before the meeting and had been reflected in falling oil prices. But even though overall oil production was expected to grow, the market maintained the belief that OPEC producers would adjust, by cutting production to give way to shale oil producers. A glutted market on which producers would not be able to market their crude oil was assumed to be an impossibility. However, after the meeting in November, it became clear that OPEC producers were not willing to sacrifice their share of production. Hence, not only did the global expectations of the future path of the oil supply have to be reconsidered, but the possibility of temporary oversupply in the nearest future became substantial. Early in 2015, this oversupply actually occurred. However, the assumption was that any oversupply would be short-lived, which was reflected in the upward sloping forward curve. Basically any oversupply story follows more or less the same pattern.

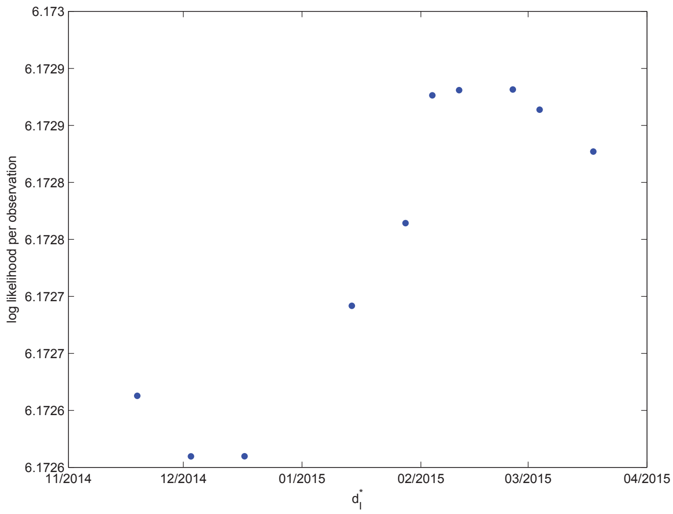

Figure 14: Identifying structural break in market reaction to news: the value of the composite log-likelihood function for each $d_I^\star$. The value of the other threshold $d_H^\star$ corresponds to December 3, 2014.

Testing the theory

In this section we use our methodology to test oil price theories. The crucial element of Arezki and Blanchard's (2015) narrative is a structural shift in market expectations after the OPEC meeting. The market began viewing oversupply as an actual possibility. Since the OPEC decision, the market had been constantly watching for signs of oversupply. Any unexpected barrel of oil in storage could signal the beginning of a period of excess supply. But if traders become more willing to attribute surprising increases in inventories to increased amounts of excess or unsold oil, market *expectations must be revised more strongly* in response to inventory shocks. Similarly, as inventories were rising, the speculative possibilities for smoothing shocks became limited. Thus, the market response to inventory surprises, especially to positive surprises, should have been sharply intensified after the OPEC meeting, reflecting a structural shift in expectations.

We search for a structural break in market reactions to news over the time period from September 2014 to the end of March 2015. For each combination of thresholds $(d_I^\star, d_H^\star)$ in the given range, we estimate our model and compare the values of the objective function. The overall sample covers 2013 to 2016.[34] Thus, we have approximately two years before and after the break period to perform a meaningful estimation, but not a long enough period to raise the issues of stationarity.

Our results clearly indicate that the break in the trading pattern, $d_H^\star$, occurred around the first week of December 2014. The time of the change is consistent with observed changes in volatility and intensity of trading as measured by the fraction of intervals with no trading, see Figures 17 and 18. We believe that the break in trading patterns can be attributed to immense monetary inflows to exchange traded funds, such as the United States Oil Fund, that track short-term crude oil futures contracts.

In contrast, the break in market response to news occurred much later. To illustrate this, Figure 14 shows the value of the composite log likelihood function for each value of the threshold $d_I^\star$. The maximum is achieved for the values of the threshold corresponding to the week of February

[34]For robustness, we repeat the exercise for a shorter sample, from 2014 to 2015; the results are similar and are available upon request.

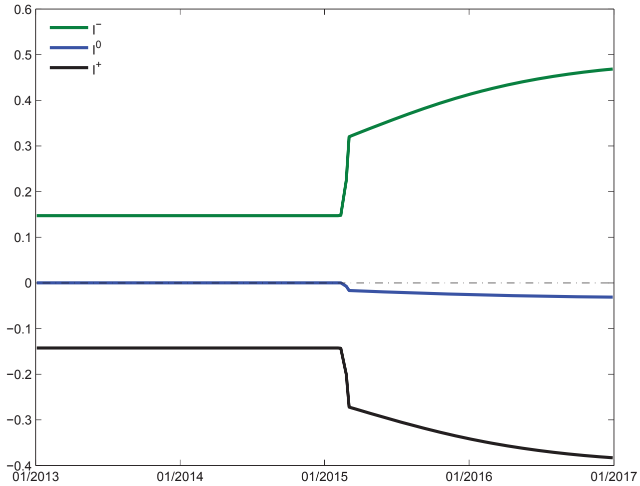

Figure 15: Evolution of returns responses to inventory news.

25, 2015.

To allow for more flexibility, as well as to determine more precisely the speed of adjustment, we utilized the smoothed specification of the transition function as specified in section 5. Figure 15 depicts the estimated evolution of returns reactions to inventory surprises. The returns reactions to news remained unchanged until the end of 2014. We find neither evidence of a structural shift in expectations around the OPEC meeting on November, 27 of 2014, nor gradual development of oversupply fears throughout 2014. Baumeister and Kilian (2016) use a VAR forecasting model for the real price of oil to show that more than half of the decline in the price of oil was predictable in real time as of June 2014. Of course, this does not necessarily imply that the market indeed incorporated that information to update beliefs, however, it suggests that evolution of oversupply beliefs could have been gradual. In contrast, our results indicate that oversupply fears developed abruptly in the middle of February of 2015, and importantly, *after* inventories reached extremely high levels.[35]

Alternative theory

When combined, our results challenge the conventional narrative of the evolution of oversupply beliefs in 2014. We have found an absence of the effects of inventories on the term structure in 2015, when the storage capacity was believed to be close to exhaustion. We found market responses to be asymmetric, but the asymmetry goes in the wrong direction relative to the predictions of standard inventory theory. Finally, we have revealed a market reluctance to believe in oversupply in 2014. What type of narrative would be consistent with our evidence?

One potential mechanism, which is typically overlooked, is the dramatic inflows of investment in ETFs tracking oil prices. By attracting a certain clientele with specific trading preferences, ETFs alter the dynamics of the system. Our results clearly pick that up. However, the effect of ETFs is not limited to alteration of trading patterns. In addition, when such funds become large enough, a significant effect of their operations on the term premium can be expected.[36] Thus, widening of

[35]We obtain similar results when we estimate even more flexible specification. The results are available upon request.

[36]See Mou, 2012, Hamilton and Wu, 2014 and Selezneva, 2015

the term structure curve from December of 2014 could have a financial nature, not directly related to oversupply. But as long as the distant futures prices are significantly higher that the spot price, physical arbitrage becomes profitable. The owners of storage facilities may purchase oil, put it in storage and sell in the future at a profit. By performing this speculative transaction, the traders build up oil inventories. Finally, large amounts of stored oil are wrongly are wrongly considered to be the basis of an immense oversupply.

The narrative above is a potential story. We do not go as far as to claim to develop a convincing explanation of the evolution of market beliefs in 2014, nor do we try to argue that oversupply was not present in 2014. Rather, we believe that more research is needed to improve our understanding of the oil market conditions in the US in those years, as well as to assess the vulnerability of the market to financial innovations.

9 Conclusion

The characterization of market beliefs is an important tool in understanding price formations. We provide a framework that integrates ultra high frequency data into characterization of market beliefs. Our approach can be used to distinguish between competing theories of price formation, as any consistent theory that aims to explain the behavior of prices inevitably assumes a certain evolution of market beliefs, which can then be tested against our empirical evidence.

Recently there has been substantial interest in dynamic information acquisition theory. However, empirical verification has been limited. The models are predominantly tested either using survey data or in an experimental setting. At the same time, the finance literature has long been developing the disagreement theory of trading volumes. We believe that there is a strong link between these two strands of the literature that is yet to be fully explored. Naturally, disagreement among traders is driven by information acquisition and the processing decisions of the traders, eventually reflected in trading volumes. Thus, high frequency modeling of trading volumes has the potential to bring new insights as well as to allow development of new empirical facts to challenge and shape the development of information theories. Our paper is a step in this direction.

Our study only examines the oil futures market. However, with very little adjustment, our methodology can be applied to other commodity markets, and we expect to see similar results. Using a large cross-section of commodity futures, Gorton et al. (2013) show that high levels of inventories are associated with an upward sloping futures curve. Our methodology and data allow us to directly test the mechanism potentially underlying this relationship. We believe that our results justify the search for an alternative mechanism of the term premium determination.[37]

We apply our methodology to understand the unfolding of beliefs in an oversupply of oil in 2014 and early 2015. However, our approach can also be used to understand the perceived market conditions in 2011, when a WTI-Brent spread occurred. Over many years, two oil benchmarks have reflected existing supply-demand conditions in the oil market. One, West Texas Intermediate (WTI) crude oil has long served as a primary global benchmark, a reference price for oil traders. However, in 2011, the ability of WTI to serve as a global benchmark was called into question, when WTI prices significantly diverged from another benchmark, Brent, representing the European market. Historically, the two benchmarks traded in line. However, in September of 2011 the WTI-Brent spread reached more than $25 a barrel. The anomaly was soon resolved, and the spread became a major manifestation of structural shifts in the oil market. The shale boom in North Dakota and the rapid growth of Canadian oil production flooded Cushing, Oklahoma. The lack

[37]See Selezneva (2015) for an example of such a mechanism.

of pipeline capacity to move oil south from Cushing to the Gulf Coast created a glut of oil in the Midwest and put significant downward pressure on the price of WTI. The oil glut explanation is generally accepted in the academic literature (e.g. Kilian (2016), Borenstein and Kellogg (2014), Kaminski (2014), Fattouh (2011), McRae (2015), Buyuksahin et al. (2013), Liu et al. (2015)) as well as by government agencies,[38] global providers of commodities information,[39,40] exchanges,[41,42] and by the media[43] and energy analysts.[44] Our approach can be used to understand the evolution of market beliefs in times of extraordinary events. We leave that for future research.

References

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review 93*(1), 38–62.

Andrade, P. and H. Le Bihan (2013). Inattentive professional forecasters. *Journal of Monetary Economics 60*(8), 967–982.

Arezki, R. and O. Blanchard (2015). The 2014 oil price slump: Seven key questions. VoxEU.

Armstrong, W., L. Cardella, and N. Sabah (2017). Information shocks and liquidity innovations. Manuscript, Texas Tech University.

Baffes, J., M. Kose, F. Ohnsorge, and M. Stocker (2015). The great plunge in oil: causes, consequences, and policy responses. World Bank Policy Research Note PRN/15/01.

Banerjee, S. and B. M. Breon-Drish (2017). Dynamic information acquisition and strategic trading. Manuscript, University of California, San Diego.

Banerjee, S. and I. Kremer (2010). Disagreement and learning: Dynamic patterns of trade. *Journal of Finance 65*(4), 1269–1302.

Baumeister, C. and J. D. Hamilton (2017). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. National Bureau of Economic Research working paper, No. w24167.

Baumeister, C. and L. Kilian (2016). Understanding the decline in the price of oil since june 2014. *Journal of Association of Environmental and Resource Economists 3*(1), 131–158.

Bollerslev, T., J. Li, and Y. Xue (2018). Volume, volatility and public news announcements. *Review of Economic Studies*, forthcoming.

[38] U.S. Enery Information Administration (EIA), 'This week in Petroleum', September 12, 2012.
[39] See Platts, Price group/Oil Division, 'New Crudes, New Markets' special report of March, 2013.
[40] Argus White Paper, 'LLS – The Emergence of a Secondary Benchmark', 2014
[41] Intercontinental Exchange(ICE), 'Oil Benchmarks: Implications & Opportunities', May, 2011
[42] CME Group, 'Trading WTI and Brent 101. CME Group as the exchange of choice', March 19, 2014.
[43] WSJ, 'Oil Pipeline Sale Pushes Up Prices', November 17, 2011 by Gold, Cummins, and Pleven.
FT Alphaville, "WTI's upcoming 'Keystone' problem", January 26, 2011 by I. Kamiska.
CNBC, 'Why Brent, Not US Crude, Is a Better Gauge for Oil Prices', March 9, 2011 by J. Cox.
[44] The Oil Drum, 'Why are WTI and brent Prices so Different', February 22, 2011 by Gail Tverberg

Borenstein, S. and R. Kellogg (2014). The incidence of an oil glut: Who benefits from cheap crude oil in the midwest? *Energy Journal 35* (1), 15–33.

Bu, H. (2014). Effect of inventory announcements on crude oil price volatility. *Energy Economics 45*, 485–494.

Buyuksahin, B., L. T. K., J. T. Moser, and M. A. Robe (2013). Physical markets, paper markets and the WTI-Brent spread. *Energy Journal 34* (3), 129–152.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105* (8), 2644–2678.

Engle, R. (2002). New frontiers for arch models. *Journal of Applied Econometrics 17* (5), 425–446.

Engle, R. and J. Russell (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica 66* (5), 1127–1162.

Fantazzini, D. (2016). The oil price crash in 2014/15: Was there a (negative) financial bubble? *Energy Policy 96*, 383–396.

Fattouh, B. (2011). *An anatomy of the crude oil pricing system.* Oxford Institute for Energy Studies.

Gennaioli, N., Y. Ma, and A. Shleifer (2016). Expectations and investment. *NBER Macroeconomics Annual 30* (1), 379–431.

Gorton, G. B., F. Hayashi, and G. Rouwenhorst (2013). The fundamentals of commodity futures returns. *Review of Finance 17* (1), 35–105.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies 27* (3), 714–746.

Halova, M., A. Kurov, and O. Kucher (2014). Noisy inventory announcements and energy prices. *Journal of Futures Markets 34* (10), 911–933.

Halova Wolfe, M. and R. Rosenman (2014). Bidirectional causality in oil and gas markets. *Energy Economics 42*, 325–331.

Hautsch, N., P. Malec, and M. Schienle (2013). Capturing the zero: a new class of zero-augmented distributions and multiplicative error processes. *Journal of Financial Econometrics 12* (1), 89–121.

Kaminski, V. (2014). The microstructure of the North American oil market. *Energy Economics 46*, S1–S10.

Kandel, E. and N. D. Pearson (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy 103* (4), 831–872.

Kilian, L. (2016). The impact of the shale oil revolution on U.S. oil and gasoline prices. Manuscript, University of Michigan.

Liu, W.-M., E. Schultz, and J. Swieringa (2015). Price dynamics in global crude oil markets. *Journal of Futures Markets 35* (2), 148–162.

Mankiw, N. G., R. Reis, and J. Wolfers (2013). Disagreement about inflation expectations. *NBER macroeconomics annual 18*, 209–248.

McRae, S. (2015). Vertical integration and price differentials in the US crude oil market. Manuscript, University of Michigan.

Miao, H., S. Ramchander, T. Wang, and J. Yang (2018). The impact of crude oil inventory announcements on prices: Evidence from derivatives markets. *Journal of Futures Markets 38*(1), 38–65.

Nelson, D. (1991, March). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–370.

Rodriguez, M. J. and E. Ruiz (2012). Revisiting several popular garch models with leverage effect: Differences and similarities. *Journal of Financial Econometrics 10*(4), 637–668.

Roesch, A. and H. Schmidbauer (2011). Crude oil spot prices and the market's perception of inventory news. Conference presentation, ISF 2011.

Selezneva, V. (2015). A side effect of financial innovation. Manuscript, CERGE-EI.

Varin, C., N. Reid, and D. Firth (2011). An overview of composite likelihood methods. *Statistica Sinica 21*, 5–42.

Ye, S. and B. Karali (2016). The informational content of inventory announcements: Intraday evidence from crude oil futures market. *Energy Economics 59*, 349–364.

A Optimality of linear market expectation function

What defines ω? To see the intuition behind it, we abstract from the complexities of market trading and heterogeneity of traders. Denote by $x = \Delta \text{Inv}^{EIA}$ the change in inventories to be released by the EIA. Assume that a representative professional forecaster sequentially observes two signals and optimally weights them. Let z^{BBG} be the initial signal received by the agent. Once he processes it, he delivers his expectation to Bloomberg to be published; we denote it as $\Delta \text{Inv}^{BBG} \equiv E[x|z^{BBG}]$. Finally, let z^{API} be the API estimate of the inventories, observed by the agent before the public announcement. Finally, assume that everything is distributed normally, and all disturbances and x are independent:

$$\begin{aligned} x &\sim N(0, \sigma_x^2), \\ z^{BBG} &= x + \eta, \;\; \eta \sim N(0, \sigma_\eta^2), \\ z^{API} &= x + \xi, \;\; \xi \sim N(0, \sigma_\xi^2). \end{aligned}$$

The optimal signal extraction problem with normally distributed disturbances implies

$$E[x|z^{BBG}] = \frac{1}{1 + \sigma_\eta^2/\sigma_x^2} z^{BBG}.$$

Similarly,

$$E[x|z^{BBG}, z^{API}] = \omega z^{API} + (1 - \omega) E[x|z^{BBG}],$$

where

$$\omega = \frac{1 - 1/(1 + \sigma_\eta^2/\sigma_x^2)}{1 - 1/(1 + \sigma_\eta^2/\sigma_x^2) + \sigma_\xi^2/\sigma_x^2}.$$

The intuition is straightforward. If the API estimate happens to be very precise, that is, if the noise-to-signal ratio in the API signal, σ_ξ^2/σ_x^2, is close to zero, while the survey contains the error, $\sigma_\eta^2/\sigma_x^2 > 0$, then the weight placed on API is close to one, $\omega \approx 1$. Agents disregard imprecise information from the survey. Alternatively, if the initial information received by professional forecasters is very precise, σ_η^2/σ_x^2 is close to zero, while the API estimate is noisy, $\sigma_\xi^2/\sigma_x^2 > 0$, then $\omega \approx 0$, and the API signal is disregarded. In sum, higher weight ω placed on the API signal reflects better quality API estimates relative to information available to professional forecasters.

Using our previous notation from the main text, the formula above can be rewritten as

$$E[\Delta \text{Inv}^{EIA} | \Delta \text{Inv}^{BBG}, \Delta \text{Inv}^{API}] = \omega \Delta \text{Inv}^{API} + (1 - \omega) \Delta \text{Inv}^{BBG}.$$

As we see, the linear specification for market expectation function turns out to be *optimal* under the assumed information structure.

B Capturing fat tails

In this section we pay attention to conditional tails and utilize the Student's t distribution to capture the shape of the conditional density in the tails more accurately. That is, when trading occurs, the return r_t is drawn from the Student's t distribution with the shape parameter, degrees of freedom, defined by ν, with the conditional mean and the conditional variance defined as before.

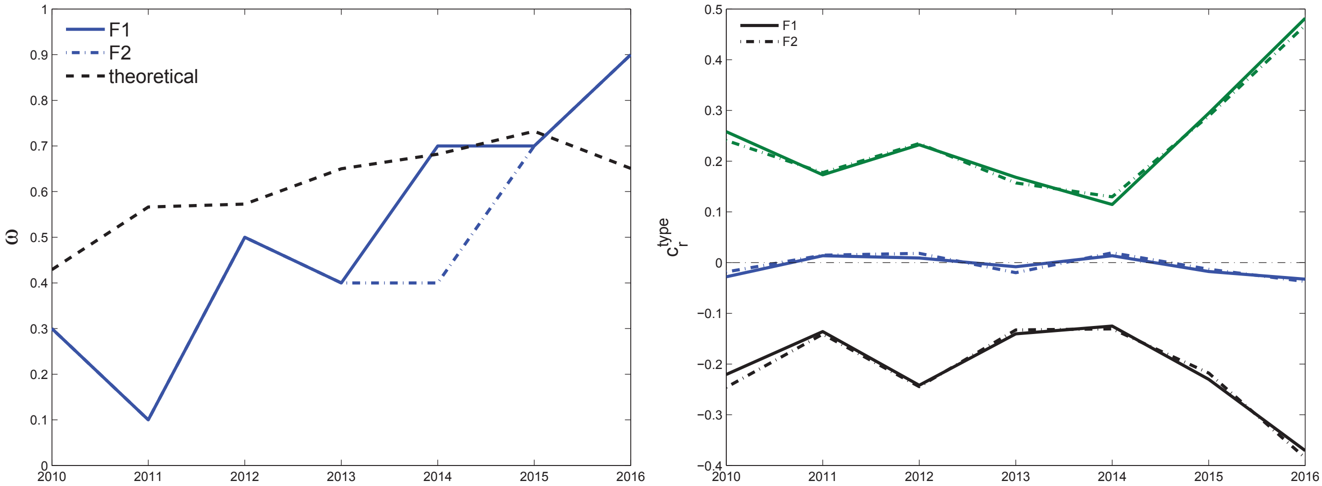

Figure 16: Capturing fat tails.
Left panel: Estimated weight ω of API signal. The solid blue line corresponds to the nearest to expiration contract, the dashed blue line corresponds to the second month contract. The black dotted line displays the value ω which represents a solution to the simple signal extraction problem. Right panel: Normalized return reaction to inventory surprises for the first two contracts: solid line – F1, dashed line – F2.

The shape parameter, ν, is estimated to be equal to 5.22 on average for the first month contract, 4.88 for the second, 3.61 for the third, and 2.82 for the fourth. Thus our results indicate fairly fat tails which are more pronounced for the long maturity contracts.

The estimated values of ω are presented in Figure 16 separately for the first two contracts. In general the values are similar, however with fat tails, the estimates of ω fluctuate more. For example, the highest value is equal to 0.9 in 2016. In 2011, we see a drop to 0.1, however the total composite function is almost entirely flat for this year for values of omega from 0.05 to 0.4. Finally, we cannot reject the hypothesis of uniform formation of expectations across contracts for any year except 2014. For this year, the LR test rejects the null hypothesis at 0.01% confidence level.

The market response results are generally similar. Figure 16 compares the normalized return reaction to news for the first two most liquid futures contracts. A strong negative relation is present (standard errors are not presented in the picture), and the lack of any effect on the term premium. The only thing that is new here is that we find negative prices movements in response to uninformative announcements, however, the estimates are not significantly different from zero.

However, Figure 16 also shows that when fat tails are accounted for, our results do not indicate asymmetry of returns responses. We also find that, if API information is not taken into consideration when modeling market expectations by fixing ω at zero, the asymmetry is also not revealed. Thus, these two assumptions seem to be critical to asymmetry results, which could explain why the evidence on asymmetry in the literature is mixed.

C Term structure effects of inventory news

Term structure effects of inventory news in 2015 and 2016 In this section we consider the subsample of 2015 and 2016, and split both calendar years into quarters. For each quarter of

2015, we analyze trading activity on the markets for the current December contract (CLZ 15), the next year June contract (CLM 16), and the next year December contract (CLZ 16). Similarly, for each quarter of 2016, we use the current December contract (CLZ 16), the next year June contract (CLM 17), and the next year December contract (CLZ 17). We then combine the first quarters of 2015 and 2016 into one sample and call it 'Q1', the second quarters into another sample called 'Q2' and so on.

The maturity of these contracts differ by quarter. The table below gives further information about maturities ranges. For example, it shows that when we consider the first quarters, we are dealing with contracts that have maturity ranging from 11 months to 9 months (the current December contracts), from 17 months to 15 months (the next year June contracts), and from 23 months to months (the current December contracts):

	Maturity at the beginning of the period in months			
	F_1	Current December	Next June	Next December
Q1	1	11	17	23
Q2	1	8	14	20
Q3	1	5	11	17
Q5	1	2	8	14

Our identification of market responses to inventory surprises requires the presence of a sufficient number of larger inventory surprises, which can be tricky given our short sample size in this exercise. The following table shows the number of large inventory surprises of each type per quarter:

	I^+	I^-
Q1	7	1
Q2	2	5
Q3	4	4
Q5	4	4

The returns reaction to inventory surprises are presented in the following table:

	F1	Current December*	Next June**	Next December
Panel A: Returns reaction to large positive inventory surprises (I^+)				
Q1	−0.25 (0.04)	−0.24 (0.03)	–	−0.18 (0.07)
Q2	−0.42 (0.09)	−0.20 (0.22)	–	−0.42 (0.13)
Q3	−0.31 (0.03)	−0.27 (0.03)	−0.23 (0.07)	−0.20 (0.07)
Q4	−0.35 (0.07)	–	−0.45 (0.07)	−0.37 (0.17)
Panel B: Returns reaction to large negative inventory surprises (I^-)				
Q1***	0.39 (0.01)	0.35 (0.01)	–	0.19 (0.03)
Q2	0.43 (0.07)	0.47 (0.10)	–	0.35 (0.09)
Q3	0.27 (0.11)	0.23 (0.09)	0.11 (0.21)	0.15 (0.11)
Q4	0.59 (0.08)	–	0.52 (0.14)	0.56 (0.11)

Standard errors in parenthesis. *In 'Q4' we don't consider current year December contract due to its approaching expiration. **In 'Q1' and 'Q2' the trading of the next year June contract is still very limited. ***Only one negative surprise was recorded over this time.

As we have seen, some of the quarters have very limited numbers of inventory surprises, so we repeat the calculations for the enlarged samples, in which we combine quarters:

	F1	Current December*	Next June**	Next December
Panel A: Returns reaction to large positive inventory surprises (I^+)				
Q1-Q2	−0.32 (0.03)	−0.26 (0.06)	−0.34 (0.08)	−0.27 (0.06)
Q3-Q4	−0.33 (0.03)	–	−0.44 (0.13)	−0.30 (0.11)
Panel B: Returns reaction to large negative inventory surprises (I^-)				
Q1-Q2	0.30 (0.05)	0.37 (0.08)	0.23 (0.19)	0.26 (0.06)
Q3-Q4	0.43 (0.08)	–	0.40 (0.12)	0.38 (0.09)

Standard errors in parenthesis. *In 'Q3'–'Q4' we do not consider the current year December contract due to its approaching expiration. **In 'Q1'–'Q2' the trading of the next year June contract is still very limited.

Term structure effects of inventory news in 2013–2014 The number of large inventory surprises of each type per quarter:

	I^+	I^-
Q1	4	5
Q2	4	4
Q3	6	6

The returns reaction to inventory surprises are presented in the following table:

	F1	Current December*
Panel A: Returns reaction to large positive inventory surprises (I^+)		
Q1	−0.08 (0.03)	−0.06 (0.04)
Q2	−0.12 (0.03)	−0.14 (0.02)
Q3	−0.07 (0.03)	−0.05 (0.03)
Panel B: Returns reaction to large negative inventory surprises (I^-)		
Q1	0.12 (0.05)	0.12 (0.06)
Q2	0.08 (0.02)	0.05 (0.02)
Q3	0.15 (0.03)	0.14 (0.03)

Standard errors in parenthesis. *In 'Q4' we do not consider the current year December contract due to its approaching expiration.

Term structure effects of inventory news in 2010–2012 Number of large inventory surprises of each type per quarter:

	I^+	I^-
Q1	13	3
Q2	9	4
Q3	9	11

The returns reaction to inventory surprises are presented in the following table:

	F1	Current December⋆
Panel A: Returns reaction to large positive inventory surprises (I^+)		
Q1	−0.21 (0.02)	−0.16 (0.03)
Q2	−0.20 (0.04)	−0.18 (0.03)
Q3	−0.23 (0.04)	−0.22 (0.05)
Panel B: Returns reaction to large negative inventory surprises (I^-)		
Q1	0.25 (0.03)	0.00 (0.01)
Q2	0.18 (0.06)	0.12 (0.06)
Q3	0.25 (0.03)	0.22 (0.00)

Standard errors in parenthesis. ⋆In 'Q4' we don't consider current year December contract due to its approaching expiration.

D Time variations in trading patterns

Two observations raise concerns about non stationarity, namely, time variation in oil price intensity of trading. First, Figure 17 shows realized variance estimated using high-frequency returns on the front-month futures contract. We can see a distinct change in the volatility regime that occurred at the end of 2014, when volatility increased dramatically and remained high for a long period of time. An increase in volatility from 2014 has been well documented; what is less known is the contemporaneous change in trading intensity in the futures market. To demonstrate this, we calculate the fraction of 5-second intervals with zero traded volumes and present the results for the first two futures contracts in Figure 18. The figure shows a dramatic drop in trading intensity in the last few years. If we take the second-month futures contract, the empirical probability of no-trading over 5-second intervals decreased from 0.6 in 2014, to 0.45 in 2015, and to 0.4 in 2016. The results suggest the presence of significant changes in the way oil futures are traded since 2014.

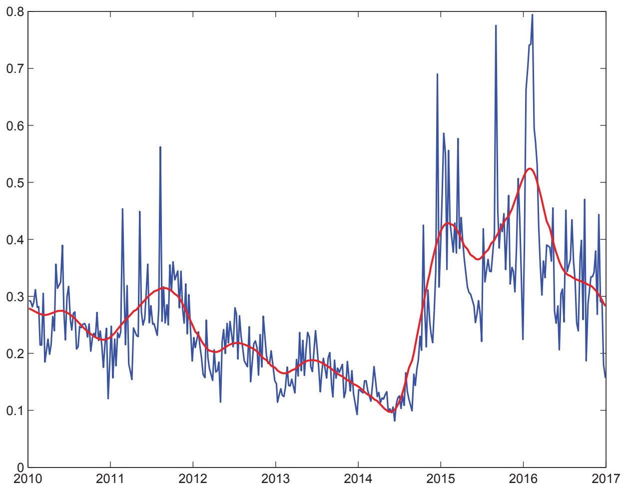

Figure 17: Realized variance, annualized. The blue line - daily realized variance of returns on the first futures contract. The red line represents a smoothed filtered series (HP filter with standard smoothing parameter).

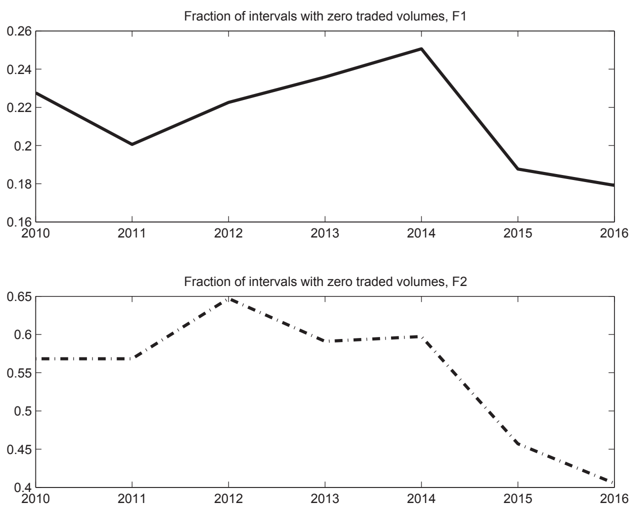

Figure 18: Time variation in probability of trading: fraction of 5-sec intervals with zero traded volumes.

Abstrakt

Charakterizujeme utváření tržních očekávání na ropném trhu kompletním popisem reakcí trhu na neočekávané změny zásob ropy. Využíváme jedinečnou sekvenční povahu zpráv o stavu zásob k identifikaci šoků. Odhadujeme AR-ARCH-ARCHMEM model sdružené dynamiky výnosů, volatility výnosů a objemu obchodů v době vydání zpráv s použitím vysokofrekvenčních dat futures kontraktů na ropu. Náš model (i) popisuje nízkou likviditu kontraktů s dlouhou dobou do splatnosti s použitím nízké obchodní aktivity, (ii) zachycuje kolísání intenzity obchodování v čase, a (iii) umožňuje strukturální změny v dynamice reakcí na nové informace v průběhu času. Ukazujeme (i) jednotnou tvorbu očekávání napříč kontrakty s různou splatností, (ii) silný negativní vztah mezi neočekávanými změnami zásob ropy a výnosy, (iii) žádný vliv na výnosovou strukturu, což naznačuje, že ropné šoky jsou vždy považovány za trvalé, a (iv) rozdílnou reakci obchodovaných objemů pro různé doby do splatnosti. Demonstrujeme, jakým způsobem mohou být naše výsledky použity k testování teorií o tvorbě ceny ropy, a přispíváme k debatě o nedávném nadbytku ropy.

Working Paper Series
ISSN 1211-3298
Registration No. (Ministry of Culture): E 19443

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: http://www.cerge-ei.cz

Editor: Byeongju Jeong

The paper is available online at http://www.cerge-ei.cz/publications/working_papers/.

ISBN 978-80-7343-426-7 (Univerzita Karlova, Centrum pro ekonomický výzkum
a doktorské studium)
ISBN 978-80-7344-469-3 (Národohospodářský ústav AV ČR, v. v. i.)