# Formation of Market Beliefs in the Oil Market

Stanislav Anatolyev*, Sergei Seleznev, and Veronika Selezneva†

February 15, 2018

**Abstract**

We characterize the formation of market beliefs in the oil market by providing a complete characterization of the market reaction to oil inventory surprises. We utilize a 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 response to news over time. We show (i) uniform formation of expectations across oil futures contracts with different maturities; (ii) strong negative relation between inventories surprises and returns confirming involuntary inventory interpretation of oil stocks; (iii) no effect on the term premium, suggesting that inventory shocks are considered to be permanent; (iv) differential reaction of volumes 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**: trading intensity, futures returns, return volatility, ultra high frequency data, oil market, inventory surprises, expectation formation

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

*CERGE-EI, Czech Republic & New Economic School, Russia

†CERGE-EI, Czech Republic, Veronika.Selezneva@cerge-ei.cz. The Czech Science Foundation grant GACR 17-27567S is kindly acknowledged.

# 1 Introduction

The oil market is undergoing a major transformation driven by the shale boom. Oil traders are required to learn a great deal about the productivity of 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 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 supply and demand balance in the oil market, and thus are informative about the scarcity of a 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 scarcity of oil 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 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 the announcement, in order to characterize disagreement among traders.

In this paper we offer a novel procedure to identify inventory shocks by utilizing a unique sequential nature of oil stocks announcements. First, the 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 followed heavily 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 close to each other. So, even though the purchase of API information may be prohibitively costly for unsophisticated traders whose main interests are 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 leads to incorrect identification of inventory surprises and potentially distorts the estimates of the market impact[1]. To overcome this issue, we assume a simple form of market expectation formation, namely that the market's expected value of the EIA estimate of 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 chance to identify changes in 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

[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 inventory surprise as the difference between EIA and API reported values, but neglects information in initial surveys. Another example is Ye and Karali (2016), which defines inventory shocks sequentially: 'API shock' as the difference between the API announced change and the forecast, and 'EIA shock' as the difference between EIA and API.

increased, despite an unchanged cost of subscription and unchanged precision of the API estimates.

Using the market inventory surprises identified, we can estimate the market impact of news. We are interested in the determination of the oil futures prices along the term structure curve, as well as in understanding how exactly news about oil market fundamentals is incorporated in prices. Thus, we develop a model of the joint high-frequency dynamics of futures returns, return volatilities and trading volumes around the EIA announcements. The model is estimated using high frequency data (5-second time intervals) on short and long maturity oil futures contracts on WTI oil traded at the NYMEX. We use dynamic models because of ultra-high frequency data and because our interest lies not only in identification of instantaneous impact of inventory surprises, but also their dynamic impact on future returns, return volatilities, and trading volumes.

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 standard AR-ARCH framework. In particular, we utilize the gaussian (G)ARCH model, and the conditional variance follows EGARCH dynamics (Nelson, 1991). 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 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 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, where the additional letter 'I' signifies the presence of indicators for the announcements. The inventory shocks identified according to our procedure, appear in the model via these indicators. A typical way 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 response coefficients to differ in an unrestricted way. In the benchmark case we only distinguish 'large positive' surprises, 'uninformative announcements' , and 'large negative' surprises, i.e. we split inventory shocks into just three groups.

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

Our first main result is irrelevance of contract maturity 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 expec-

[2]See Andersen et al. (2003)

tation formation process. In principle, 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 a different way, either due to differences in sophistication levels, or because of differential access to API information. However, our results suggest that we shouldn't worry about segmentation at least as far as *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% , would cause the oil price to increase by 1.2% in 2016. The negative sign of the relationship shows that the market interprets fluctuations in oil stocks as fluctuations in *involuntary* inventories, which is oil that producers cannot sell and has to be put in storage. Say, for example, that when a surprisingly large amount of oil is stored, the market adjusts the spot oil price downwards in hopes of clearing the spot market. From a theoretical perspective, however, it is ambiguous how the spot price should move in response to inventory surprises. That is because oil can also be stored voluntarily, say by outsiders for speculative purposes, only to be released in the future at a higher expected price. Alternatively, producers may keep some inventory as a buffer stock for future demand shocks. So if it is extra *speculative* oil that surprises the market, than the spot price doesn't have to decrease. It can even increase, because speculative oil might continue to be taken away from the market.

Our second main result is the absence of any effect of inventories on the term premium. The entire futures curve shifts up or down when an announcement comes. The term structure adjustments depend on the perceived persistence of shocks. That is, the effects of temporary shocks should vanish with maturity. However, our results show that the expectations of future oil prices are revised uniformly, by the same amount irrespective of maturity, which implies that the traders view inventory changes as mostly reflecting *permanent* or long lasting shocks[3].

The lack of any effect on the term spread is surprising, as it contradicts 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 use as anecdotal evidence. In 2008 a negative demand shock created an abundance of oil and depressed the 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 a large term premium, but over a much longer period of time. So in both cases the observed steep term structure curve was attributed to the presence of excess supply of oil on the market.

Therefore, if the term premium is to be explained by an excess supply of oil on the market, we should observe a strong effect of inventory surprises on the time spread, especially at the shorter end of the term structure curve. It is exactly when inventory news come out that the market receives fundamental information about current scarcity of oil, and revises expectations of future market conditions. Moreover, the link should be especially pronounced at times of extremely high inventories, when the spare capacity is near exhaustion and speculation activity that could smooth

[3]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.

out the shocks is limited[4]. However, our results do not indicate the existence of such a relationship.

Explicit modeling of trading intensity allows us to analyze idiosyncratic trading behavior of market participants around inventory announcements. We observe significant immediate reaction of volumes to all announcements, including uninformative ones. Volume reaction to an uninformative announcement confirms previous findings in the literature[5]. One explanation offered to explain 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 in different ways. As a result, trade is possible, as agents adjust their positions according to individual changes in beliefs. Under this interpretation, our results not only reveal the presence of disagreement and its variation across maturity, but also indicate that disagreement varies over time, independently for each contract. Thus, even though *average market expectation*s are formed uniformly across contracts (as our previous results have shown), the dispersion of belief or beliefs' revisions 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, proliferated in 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 on 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 of trading probability is crucial for understanding and measuring 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 trading patter after 2015. We show that even though surprising announcements trigger larger immediate trading response relative to uninformative news, the overall cumulative volume response happens to be lower for surprising announcements.

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 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 declined to cut production. By the beginning of 2015, oil inventories reached unprecedented levels and were interpreted as a sign of immense oversupply. Considerable attention has been paid to understanding how exactly oversupply beliefs unwound. Arezki and Blanchard (2015) argue in favor of a structural shift in market expectations after the OPEC meeting in November of 2014. Before the meeting, the market maintained the belief that total oil production could still be capped, despite growing shale oil supplies. That is, OPEC producers were believed to be willing to adjust, basically cut, production and give way to shale oil producers, only to maintain high prices. After the meeting

[4]Facing low oil prices, owners of storage facilities have an incentive to take away oil from a satiated market, put it in storage, and release in the future once oil price normalizes. Such speculative operation at the same time prevents current oil prices from falling too much, and tends to decrease futures prices as more oil is brought to the market in the future. However, once storage space is exhausted, speculative activity stops.

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

it became clear that OPEC producers would not sacrifice their market share. Hence, not only the global expectations of the future path of the oil supply had to be reconsidered, but the possibility of having an excess supply of oil in the nearest future became substantial.

So do we see any structural breaks in market expectations, aka changes in market responses to inventory news, around the OPEC meeting? To parsimoniously allow for time variability of market response parameters, we utilize a threshold autoregression modeling approach (TAR). But before we can proceed to estimation, we need to solve one issue. Both daily intensity of trading (say a fraction of 5-sec intervals with zero trading volume) and daily volatility of returns indicate a distinct shift in the trading pattern happening around the end of 2014. Hence, by trying to identify a structural change in market response parameters, we may wrongly pick up background changes in the overall trading pattern. To overcome that issue, we allow for two independent transitions: one for changes in trading pattern, and the other for changes in market response to news.

Not surprisingly, our results show an apparent break in the trading pattern which occurred around the first week December of 2014. We believe it can be attributed to the explosion in money flowing into ETFs. However, our results do not indicate any clear shift in market reaction to news. If anything, the break happened in the middle of February of 2015, i.e. *after* inventories reached extremely high levels. Hence, oversupply beliefs 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 the article is organized as follows. In the following section, we provide institutional background information. We discuss formation of expectations in Section 3. The econometric framework is described in Section 4. The description of the data and estimation procedure is given in Section 5. Our main findings are presented in Section 6. Application of our results to the recent oil boom is discussed in Section 7. In Section 8, we relate our work to the existing literature. Section 9 concludes.

# 2 Institutional background

The weekly estimates of crude oil inventories in the U.S. are provided by the U.S. Energy Information Administration (EIA), which is the 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. On the following Wednesday a summary report is released in the form of the EIA publication, Weekly Petroleum Status Report. The report becomes available to the public at 10:30 am (Eastern Time) and is heavily followed by the media.

However, what is less know is the fact 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 the energy firms using exactly the same weekly survey forms that EIA uses. While reporting to EIA is mandatory, reporting to API is voluntary, but despite that the association claims its coverage

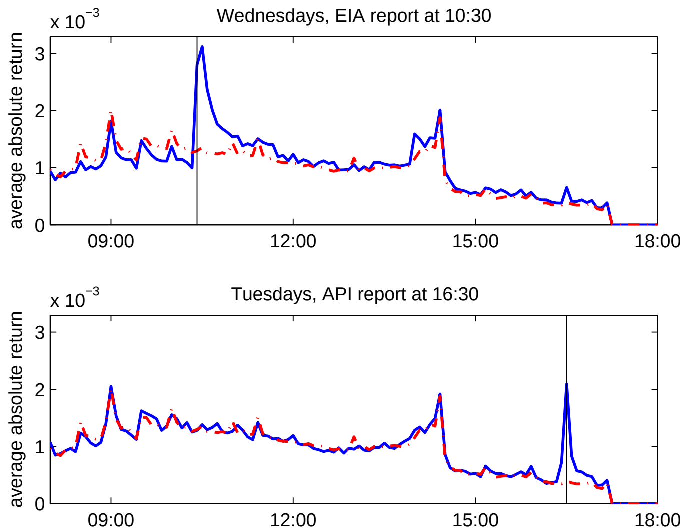


Figure 1: Market reaction to oil inventory announcements.

to be close to 90% of the industry. The API releases the data in a Weekly Statistical Bulletin on Tuesdays at 4:30 pm (Easter Time), the day before the official announcement. In contrast to publicly observable EIA report, access to API requires costly subscription available only through Thompson Reuters. Thus, for less sophisticated traders and traders with main interests outside the energy market, the purchase of API information may be prohibitively costly. But for the more specialized traders, API reports represent quite valuable information, as historically, API and EIA estimates tend to be close to each other. The discrepancy is believed to happen due to different procedures utilized to estimate the remaining 10% of the market.

To get a feeling of whether the market follows oil inventories announcements, we use a simple procedure to estimate the market response to each of the announcements. We split each trading day into 5 minutes interval. Then we take, say, an interval from 9:00 to 9:05 am and use intraday data on nearest to expiration oil futures contract 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 opens at 9 am and closes at 2:30 pm (Easter Time). The blue lines correspond to report days: the top panel represents Wednesdays when EIA report is released, whereas 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 to inventory news on average. Moreover, we see that market reaction to API (bottom panel) is comparable in magnitude to the market reaction to EIA release. Slightly muted reaction to API reports, may be partially attributed to restricted access, or alternatively, to the late time of a release, as API reports come out after trading hours. That allows us to conclude that the market generally cares about API releases, and thus neglecting this announcement might be a big deal in identification of surprises and may introduce a significant distortion in the estimates of market impact.

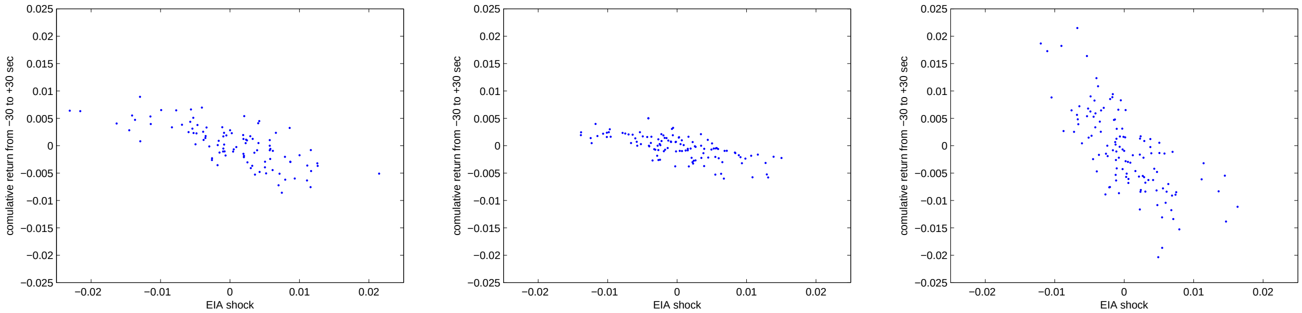


Figure 2: Cumulative returns on the front month contract over 1 minute around EIA announcement (±30 seconds) against our identified EIA surprises; Left panel: 2010 - 2011; middle panel: 2012 - mid 2014; right panel: mid 2014 - 2016.

Finally, to get a sense of the magnitude of the effect, Figure 2 displays scatter plot of cumulative returns for the front month contract over a period of 1 minute around the announcement against our identified EIA surprises, identification of which will be discussed later. We break down the sample into 3 time periods. A negative relation can be clearly seen. Positive inventory surprises tend to depress oil prices, whereas negative surprises push prices up. Figure 2 also suggests that the strength of the effects evolves over time. We break down the sample into 3 time periods, from 2010 to 2011, from 2012 - to end of June of 2014, and finally from July 2014 to the end of 2016. The scatter plot suggests that the effects became much stronger in the later years.

# 3 Formation of expectations and identification of market surprises

We would like to use EIA announcements to identify inventory surprises and study their effect on the market. That requires identification of inventory surprises and specification of response functions. In both aspects our approach differs from the literature.

For identification of inventory *surprises*, we need to estimate the *expected* value of the change in oil inventories. Ideally, we would like to observe this market expectation as close to the moment of 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[6]. However, it has become common in the literature to use surveys of professional forecasters to proxy for the market expectations[7]. Given general public interest to oil inventories, various surveys are available that directly ask agents' expectations of EIA annouced changes, including the surveys conducted by Reuters[8], Bloomberg[9], 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 informational 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

[6]See Roesch and Schmidbauer (2011)
[7]See Andersen et al. (2003)
[8]See Bu (2014)
[9]See Halova Wolfe and Rosenman (2014), Halova et al. (2014), Miao et al. (2018).

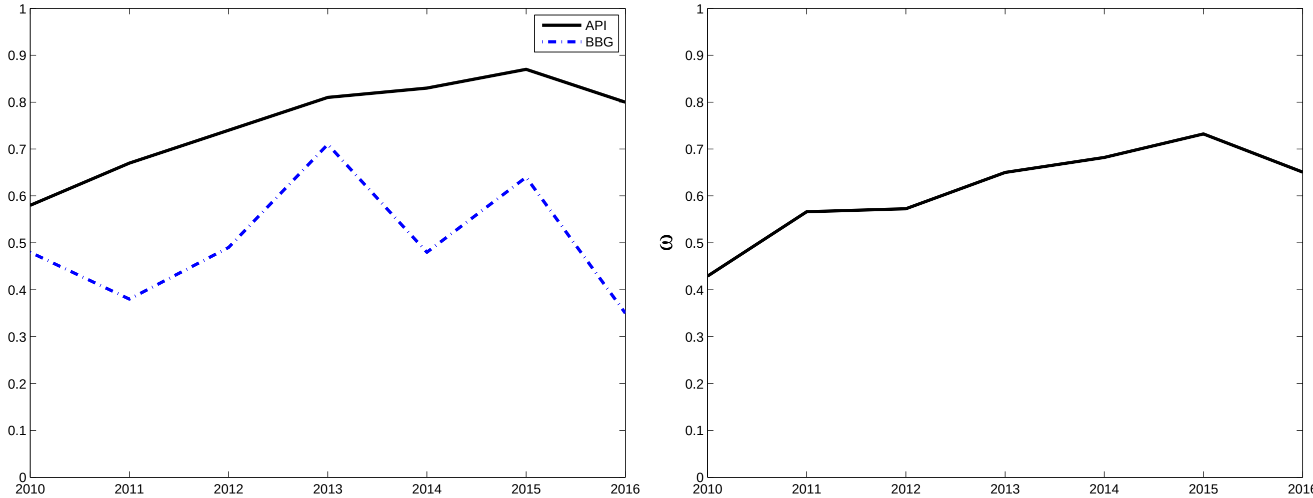


Figure 3: Left panel: Correlation of EIA and API estimates of inventories changes (solid line); of EIA estimates and Bloomberg median forecast (dashed line).
Right panel: Optimal weight to be placed on API signal, $\omega$ in each year, under the assumption of independent measurement errors in API and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

contrast to other papers we will also use the numbers released by the American Petroleum Institute.

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

**Expectations formation process** Formally, denote by $\Delta Inv^{EIA}$ the change of inventories (normalized by the total oil inventories as for last week) to be released by EIA. We will assume that the market forms expectation based on realized values of two observable signals: oil stocks number released by API, which we denote as$\Delta Inv^{API}$ where time subscript is dropped for simplicity, and the median *forecast* of the survey of professional analysts conducted by Bloomberg, denoted as $E_{BBG}[\Delta Inv^{EIA}]$, where conditioning is done over information set of professional forecasters[10] . For simplicity for the most part we maintain the assumption of linearity, in other words, we assume that the market puts $\omega$ weight on API signal, and $(1-\omega)$ weight on BBG signal, where $\omega \in [0,1]$. So market expectation of oil stocks change is given by

$$E_{market}[\Delta Inv^{EIA}] = (1-\omega)E_{BBG}[\Delta Inv^{EIA}] + \omega \Delta Inv^{API}.$$

To provide intuition, we calculate optimal omega for a simple case. Namely, we abstract from the complexities of the aggregation of individual beliefs, and solve optimal signal extraction problem of a single agent with access to API information, and under assumption of independent

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

measurement errors in API and BBG signals, normally distributed inventory changes and signals (see appendixA). In such a case linearity is known to be optimal, and optimal weight depends on signal to noise ratio in BBG and API signals. We can use the data to calculate these signal to noise ratios and find the optimal weight $\omega$. Right panel of Figure 3 shows the results, namely the optimal $\omega$s calculated separately for each year under the assumptions made above. Not surprisingly, better quality of API information gets reflected in higher weight that one should place on it over time. The optimal weight $\omega$ monotonically increases from 0.4 in 2010 to about 0.7 in the last three years.

In general, the weight placed by the market as a whole on API information could also depend on 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 movements in the prices. In reality, information percolation might be slow, especially given that API reports come long after main trading hours. The only problem with this explanation is the fact that the highlights of the API report, at least 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 API report. Alternative explanation can be related to heterogeneity in sophistication. Some traders especially those who don't specialize on oil trading might be unaware of the existence of this additional private source of information. Thus, we can expect estimated weight placed on API to be lower than displayed on Figure 3.

**Identification of market inventory surprises** Given $\omega$ the market surprise is defined as follows:

$$x = \Delta Inv^{EIA} - E_{market}[\Delta Inv^{EIA}].$$

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

**Response function** Now we are left with specifying the response function. How should, for example, returns react in response to a market surprise of $x$? Usual approach[11] would be to assume a linear response function. However, such a 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 response coefficients to differ in an unrestricted way. The benchmark case we only distinguish large positive and negative surprise, in other words we split all the surprises into 3 bins according to the following: the market surprise of $x$ is called

$$\begin{aligned}
&\text{'a positive suprise', denoted as } I^+ = 1, \text{if } x > \bar{x} \\
&\text{'an uninformative announcement', } I^0 = 1, \text{if } -\bar{x} \le x \le \bar{x} \\
&\text{'a negative suprise', } I^- = 1, \text{if } x < -\bar{x}
\end{aligned}$$

where threshold $\bar{x}$ is to be estimated. Notice also that for each weighting scheme, for each chosen $\omega$, which is the weight placed on API, the identified shocks are going to be different, and will

[11] See for example Andersen et al. (2003).

deliver a different value of log likelihood function. Thus, we search for $\{\omega, \bar{x}\}$ along with other parameters.

**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 could see smaller weight 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 we are afraid of stationarity issues). Unfortunately, we don't have that many inventory surprises per year to have enough information 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. To check if our procedure makes sense, we also estimate $\{\omega, \bar{x}\}$ separately for each contract and then perform a simple LR test. If the null hypothesis is not rejected, we will argue that evidence is in favor of uniformity of expectation formation process.

# 4 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, where the additional letter 'I' signifies the presence of indicators for the announcements. We will show how the 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 a 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 nono-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 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 if for the nearest to expiration futures contract we have only about 20% of no trading intervals, this number increases to more than 90% for the four-month contract. Clearly, the trading activity decreases dramatically with maturity.

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

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

Thus, we aim to build a model that would 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 $\pi_t$, the shock $\varepsilon_t$ 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 the most part of the analysis we will use the generalized gamma distribution[12] normalized so that $E\varepsilon_t = 1$; thus we need $E\left[\varepsilon_t | \varepsilon_t > 0\right] = \pi_t^{-1}$. 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 $\xi = \dfrac{\Gamma\left(m + a^{-1}\right)}{\Gamma\left(m\right)}$.

The conditional mean $\psi_t$ is modeled as follows:

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

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

**Returns and volatility** To investigate the effects of oil inventories on returns, we use the standard 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 $\pi_t$, trading occurs, and the return $r_t$ is drawn from the gaussian distribution with the conditional mean with the following dynamics

$$\mu_t = \mu + \sum_{k=1}^{q_r} \rho_k r_{t-k} + \sum_{k=1}^{q_r^0} \rho_k^0 I_{\{V_{t-k}=0\}} + \sum_{type} c_r^{type} I_t^{type},$$

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

[12] See Lancaster(1990). The generalized gamma distribution has 3 parameters, $a, m,$and $\lambda_0$ and

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

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

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 $D(t)$.

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 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 $\nu$, 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:

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

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

**Composite likelihood** So 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 parameters of our interest. Thus, we utilize the composite likelihood approach, i.e. we simply 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

$$L = \sum_{t=1}^{T} \left\{ I_{\{V_t>0\}} \ln \pi_t + I_{\{V_t=0\}} \ln(1-\pi_t) \right\} + \sum_{t=1}^{T} I_{\{V_t>0\}} (L_t^{volume} + L_t^{return}), \tag{1}$$

where

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

and

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

in the benchmark case, and

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

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 $\omega$ and $\bar{x}$ , parameters of the expectation formation process, in a general case unique to the market for this specific futures contract.

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

**Impulse responses** Explicit modeling of trading volumes allows us to analyze the *evolution* of trading intensity in response to news. We are interested in comparing the intensity of trading after a surprising announcement relative to an uninformative announcement. The measure that we use is the impulse response function defined as

$$\Delta^+(h_{t^\star-1}) = E\left[V_{t^\star+h}|I_{t^\star}^+ = 1, h_{t^\star-1}\right] - E\left[V_{t^\star+h}|I_{t^\star}^0 = 1, h_{t^\star-1}\right]$$

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 comes. Similarly, we define $\Delta^-(h_{t^\star-1})$.

This particular form or definition of the IRF is chosen for two reasons. First, our model is not linear. Non-linearity not only implies that history matters, but also doesn't allow us to shut down all other shocks which occur after the announcement, as it can distort the propagation properties. Therefore, traditional impulse response functions can't be used in our case. However, we also can't utilize what became known as the generalized IRF: $E\left[V_{t^\star+h}|I_{t^\star}^+ = 1, h_{t^\star-1}\right] - E\left[V_{t^\star+h}|h_{t^\star-1}\right]$[13]. In our model an announcement always happens at $t = t^\star$, $I_{t=t^\star}^+ = I_{t=t^\star}^- = I_{t=t^\star}^0 = 1$, what is uncertain is the type of surprise. That is, whether the market faces a positive inventory surprise shock, $I_{t=t^\star}^+ = 1$, a negative shock , $I_{t=t^\star}^- = 1$, or an uninformative announcement: $I_{t=t^\star}^0 = 1$. We don't

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

model inventory shocks, thus, we can't assign probabilities to these events. Therefore, our measure is something in between traditional and generalized measures.

However, we are still interested in estimating the time it takes for an announcement shock to be traded away. To get a rough estimate of propagation of the shock, we also look at

$$IRF^{type}(h^{t^\star-1}) = E\left[V_{t^\star+h}|I^{type}_{t^\star}=1, h_{t^\star-1}\right] - E\left[V_{t^\star+h}|I^{+}_{t=t^\star}=I^{-}_{t=t^\star}=I^{0}_{t=t^\star}=0, h_{t^\star-1}\right],$$

while hoping that the dynamics of the system in the absence of the announcement is reasonably captured by our model.

Given complexity of our model, we use simulations to calculate the expectations. Given history-dependence we investigate different combinations of initial values.

# 5 Data and estimation

## Data

In this paper we use the announced changes in weekly U.S. ending stocks excluding SPR and including lease stock of crude oil as published by the Energy Information Administration[14]. For identification of market surprises we use the corresponding changes in oil inventories as announced 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[15] and contains 337 announcement days. We focus on an hour around EIA announcements, that is, from 10 to 11 am. We believe that one hour long interval is long enough to get a reasonable estimate of the parameters of our dynamic model. The sampling is done at 5-second frequency, which leaves us with 720 data points for each announcement day.

## Data preparation

Before we proceed to estimation, we need to deal with a number of issues specific to the futures market. First unique feature is the expiration of contracts. At any moment in time, the exchange offers a number of contracts, that differ by so called "delivery month". The expiry dates stretch out from one month up to nine years forward, thus constituting more than 100 contracts at any time. 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, last trading day of February-2018 contract is January 22)[16]. What matters for us is that crude oil futures contracts are not constant maturity contracts, as that questions stationarity. To overcome this issue we follow a standard approach in the literature and utilize basic rolling procedures to create continuous

[14]The EIA terminated publishing of this series in September of 2016. For the last few month we use the weekly U.S. Ending Stocks excluding SPR and add the last available data on lease stock (about 30 mln barrels, and is a slowly moving variable).

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

[16]See contract specifications at CME website.

futures contracts. In particular, we replace the expiring contract with the next one on the 5th day of each month. We believe that the traders in the market follow the same procedure, that is, move away from trading the expiring contract to trading the next one at around the same time. As our dynamic model reflects the composition of currently utilized trading strategies, the stationarity should be achieved.

Second issue is normalization of trading volumes. Our model is not rich enough to account for potential day-to-day changes in composition of traders, or strategies utilized, or intensity of trading. For example, the trading may slow down in August, or before the major holidays. Thus, we decided to normalize trading volumes to eliminate any differences across days. For each announcement day we calculate the total volume traded 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 over this day to the mean daily traded volume. Finally, high frequency trading volumes can be normalized by this ratio. Formally, denote by $v_{t,d}$ the trading volume observed over a 5-sec interval $t$ in day $d$. Total volume traded this day is given by $v_d = \sum_{t=1}^{N} v_{t,d}$, where $N = 7200$ data points. Mean daily volume equals to $\bar{v}_d = \frac{1}{D}\sum_{d=1}^{D} v_d$, where $D = 337$ days, and thus, the volume ratio equals to $R_d = \frac{v_d}{\bar{v}_d}$ . Finally, normalized volumes are given by $\tilde{v}_{t,d} = \frac{v_{t,d}}{R_d}$. The total volumes traded over one hour around the announcement can still differ from day to day (we normalized total daily volumes), but these difference 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 optimization routine, we utilize a simple grid to search over possible values of $\omega$ and $\bar{x}$. The estimation of all other parameters is done for each contract separately. For most of the paper we will present the estimation result for the first 4 continuous futures contracts. Finally, we split the entire sample by a calendar year.

## 6 Results

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

### Formation of expectations

Our first main result is irrelevance of maturity in expectation formation process. We find that estimated values of the weight $\omega$ put on API signal and threshold $\bar{x}$ do not differ across the two contracts for all years. We interpret this as the evidence of uniformity of rules traders follow while forming expectations. As we discussed before, 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 composition of traders may form expectations in a different way, either due to differences in sophistication levels, or because of differential access to API information. However, our results suggest that this is actually not the case, we see no differences in the

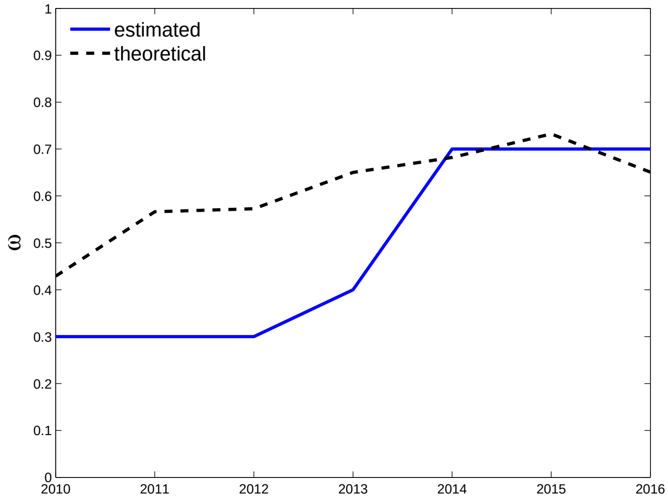


Figure 4: Estimated weight $\omega$ of API signal. Solid blue line corresponds to estimated omega jointly for the first two most liquid contracts. The black dotted line displays the value $\omega$ which represents a solution to the simple signal extraction problem.

expectation formation process at the aggregate level.

Next, we analyze the estimated values of $\omega$, the weight traders place on API reports while forming expectations. The estimates are presented at the left panel of Figure 4. We can observe a clear upward trend: the weight increases from 0.3 in 2010 to 0.7 in 2016. That is perhaps not surprising, as we have seen that accuracy of API information has been increasing, while the precision of BBG numbers remained the same. For comparison, we also plot our calculated series, $\omega_{theory}$, which solves a simple signal extraction problem. Even though the overall trend is similar, the estimated omegas remain constant until 2012 and sharply jump in 2014, rather than smoothly adjust as $\omega_{theory}$ does. The market used to significantly underweight API signals in early years. Something must have spurred public interest to the oil market in 2014, and brought greater awareness of and wider access to the API reports.

The estimated value of the threshold $\bar{x}$ found to be almost always the same across years and contracts and equal to 0.5%. Value of $\bar{x}$ defines how large the unexpected inventory shock should be in order to be considered as a surprise, either positive or negative.

## Market response to oil inventories

We are ready to characterize the oil market response to inventory news. We will 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 it, Figure 5 displays the return coefficients on inventory surprises for the front month contract, $\{c_r^+, c_r^0, c_r^-\}$, as well as the 95% confidence intervals. As expected, a clear negative relationship is

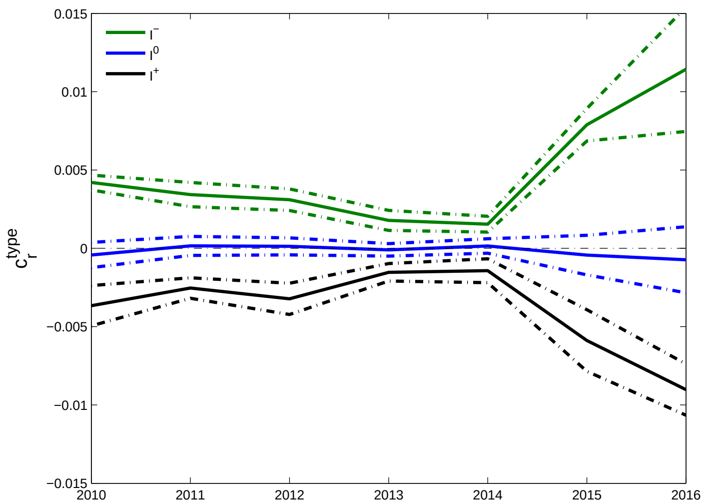


Figure 5: Return reaction to inventory surprises. Solid line - estimated return coefficients on inventory surprises for the front month contract: $c_r^+$in black corresponds to positive surprises, $c_r^-$in green corresponds to negative surprises, and $c_r^0$ in blue - to uninformative announcements.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

observed. All coefficients are statistically significantly different from zero. As for the magnitude of the effect, we find that an unexpected decline of inventories 0.5% or more in 2016, would cause the oil price to immediately increase by about 1.2%. The effect is lower for the other years, but is still quantitatively very large.

From a theoretical perspective, however, it is ambiguous how the spot price should move in response to inventory surprises. The direction in general depends on the nature of this inventories shock, or more precisely of its perceived nature. To interpret a negative relation between inventory surprises and announcement returns, we need to understand what kind of information the market extracts by looking at inventories. Inventory accumulation can be broadly distinguished between voluntary and involuntary. First, inventories may consist of oil unsold in equilibrium due to oil market frictions. Producers simply can't market the crude and thus have to put it in storage. Alternatively, oil can be stored voluntary, say by outsiders for speculative purposes in order to be released in the future at a higher expected price, or by producers as a buffer stock for future demand shocks.

The distinction between voluntary and involuntary oil inventories is critical for the direction of the market reaction to shocks. If the inventory announcement reveals surprisingly large amount of *unsold* oil in storage, than the market receives a signal that the spot oil price needs to adjust downwards in order to clear the spot market[17]. However, if it is extra *speculative* oil that surprises the market, than the spot price doesn't have to decrease. It can even increase, because speculative oil might continue to be taken away from the market. [18]

[17]The futures prices also need to decrease, because market tightness is likely to be persistent, and also because currently large inventories will eventually be brought back to the market in the future.

[18]As for the futures prices, the reaction can be ambiguous, because on the one hand extra oil is transferred into the future, but at the same time speculators might have some private information that suggests higher spot prices in the future which should be incorporated into prices.

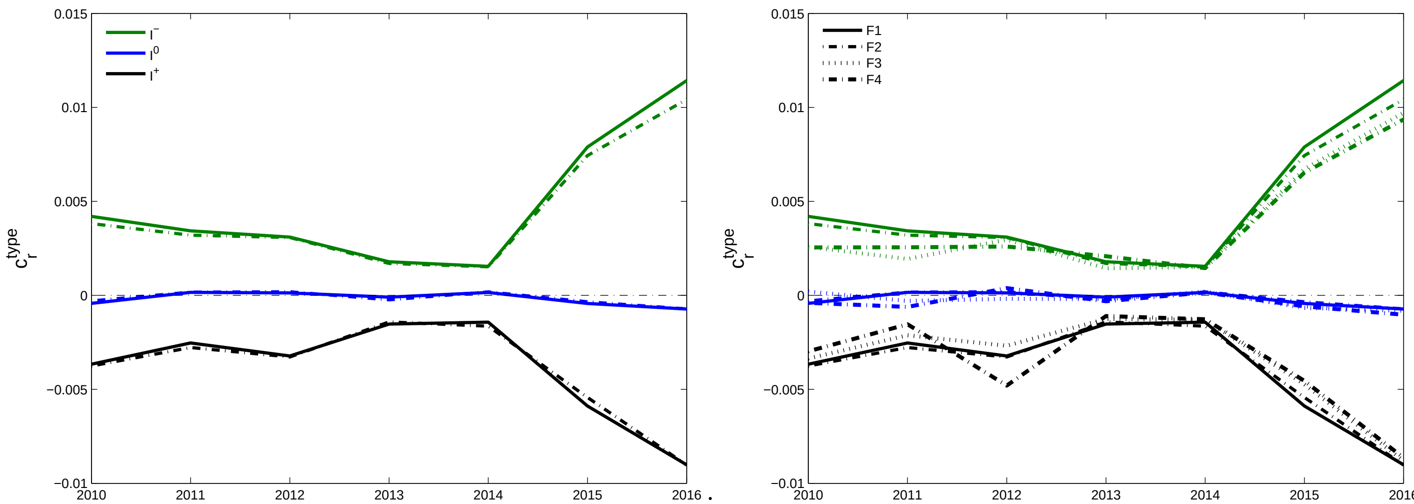


Figure 6: Left panel: Return reaction to inventory surprises by contract. Solid lines correspond to the point estimates of $\{c_r^+, c_r^0, c_r^-\}$ for the front-month contract, dashed line - for the second-month contract.
Right panel: Return reaction to inventory surprises for the first 4 contracts.

Thus, depending on the *perceived* nature of inventories, the prices can move either way in response to announced inventory changes. But our estimation results show a clear negative relation. Later we will see that the entire forward curve follows the movements of the spot price. The negative sign of the relationship shows that the market interprets fluctuations in oil stocks as fluctuations in *involuntary* inventories.

### Effects of inventories of the term structure curve

In this section we investigate whether inventory announcements affect the term premium, i.e. the difference between more distant and the nearest futures prices. The term structure adjustments depend on the perceived persistence of shocks. That is, the effects of temporary shocks should vanish with maturity. The comparison across maturities is meaningful because we have shown that expectation formation process is uniform, i.e. identified surprises are the same.

Our second main result is the absence of any effect of inventories on the term premium. Let's first focus on the shortest end of the term structure curve. Figure 6 displays the estimated coefficients $c_r^{type}$ for the first and second-month contracts. We see that the prices of both contracts adjust by exactly the same amount in response to all news, which implies that the term premium remains constant.

The lack of any effect on the term spread is surprising, as it contradicts 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. In 2008 a negative demand shock created an abundance of oil and depressed the 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 a large term premium, but over a much longer period of time.

Thus, the steep term structure curve tends to be attributed to the presence of excess supply of oil on the market. High levels of inventories are viewed as a sign of an oversupplied crude oil market (according to our previous results), and industry is believed to produce more oil than can be consumed. But such an imbalance is typically believed to be relatively short-lived. As a result, excess supply puts downward pressure on the spot oil price, but less so on the futures prices, and so the spot price becomes greatly discounted relative to the futures prices, i.e. we observe the term premium.

Therefore, if the term premium is to be explained by an excess supply of oil on the market, we should observe a strong effect of inventory surprises on the time spread, especially at the shorter end of the term structure curve. It is exactly when inventory news come out that the market receives fundamental information about current scarcity of oil, and revises expectations of future market conditions. Moreover, the link should be especially pronounced at times of extremely high inventories, when the spare capacity is near exhaustion and speculation activity that could smooth out the shocks is limited[19]. However, our results do not indicate the existence of such a relationship. We believe that our results justify the search for an alternative mechanism of the term premium determination.[20]

The results for the more distant futures contracts are similar and presented at Figure 6. Now we see that the entire futures curve moves in parallel in response to inventory news.

Before we can interpret this result, an important remark should be made. We document absence of any effect of inventory news on term premium and would like to attribute that to uniform *revision of traders expectations* of future oil prices. 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. Thus, Occam's razor logic leads us to argue that traders view most inventory changes as reflecting *permanent* or at least long lasting shocks that hit the market, and thus, he expectations of future oil prices are revised uniformly, by the same amount irrespective of maturity, and the entire futures curve simply slides up or down.

### Time variation in market responses

Our results so far indicate much stronger market reaction to inventory news in recent years. In absolute terms, the difference in magnitudes is striking: the returns reaction is 2.5 times larger in 2016, than in 2014 (see Figure 5). But can we interpret that as larger importance of inventory announcements in the last few years? Well, such interpretation might be misleading. What if the market have become more sensitive to *all kinds* of news, not just to inventories, that is, *relative* importance of inventories remained unchanged.

[19]Facing low oil prices, owners of storage facilities have an incentive to take away oil from a satiated market, put it in storage, and release in the future once oil price normalizes. Such speculative operation at the same time prevents current oil prices from falling too much, and tends to decrease futures prices as more oil is brought to the market in the future. However, once storage space is exhausted, speculative activity stops.

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

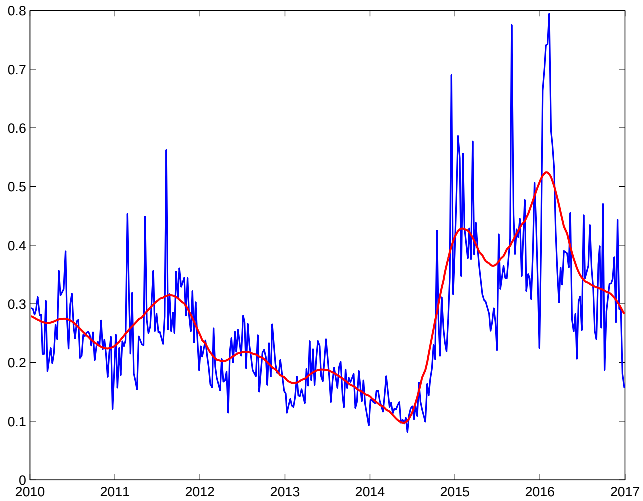


Figure 7: Realized variance (annualized). Blue line - daily realized variance of returns on the first futures contract. Red line represents smoothed filtered series (HP filter with standard smoothing parameter).

One sign that this can actually be the case is the fact that the oil market has become much more volatile in the last few years. Figure 7 displays time series behavior of realized variance, calculated for each announcement day using high frequency returns on the first futures contract[21]. A distinct change in volatility regime can be seen at the end of 2014. So if higher volatility reflects larger information flow, i.e. more frequent news arrivals[22], than our interpretation of the results is correct and no adjustments are needed. However, what if the amount of news remain the same, but the market becomes more sensitive to every news arrival, triggering higher volatility. In this case some sort of normalization is required.

One way to assess the *relative* importance of inventories would be to "normalize" the returns reaction to inventories by some measure of current market responsiveness to news. The obvious measure could be the current level of realized variance. However, as Figure 7 suggests, realized variance itself is extremely volatile, even within each regime we see quite a lot of fluctuations. 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. Thus, we aim to extract slow moving component of realized variance. We proceed using standard filtering procedure[23]. Finally, we adjust the return equation as follows:

$$\mu_t = \mu + \sum_{k=1}^{q_r} \rho_k r_{t-k} + \sum_{k=1}^{q_r^0} \rho_k^0 I_{\{V_{t-k}=0\}} + \sum_{type} c_{r,adj}^{type} \sigma_{D(t),S} I_t^{type},$$

where $\sigma_{D(t),S}$ is the extracted slow moving component of daily realized volatility. So if the market generally responses stronger to all kinds of news, the level of realized volatility would be

[21]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.

[22]That logic is formalized by (G)ARCH approach to volatility modeling.

[23]We use Hodrick - Prescott filter with smoothing parameter 1600. Alternative smoothing techniques yield similar results.

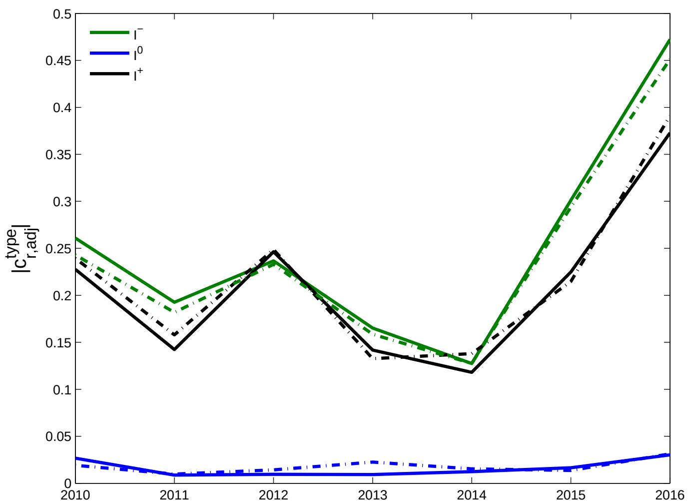


Figure 8: Normalized return reaction to inventory surprises for the first 2 contracts. Solid line - F1, dashed - F2.

larger[24], and we expect returns to react stronger to inventories news as well. That is a common component of the market response to inventories. Whereas, normalized coefficients $c^{type}_{r,adj}$ would reflect market sensitivity to specific, inventories news, beyond and above common component. And we can now compare 'normalized' responses, $c^{type}_{r,adj} - s$ over time.

Figure 8 compares normalized responses, $c^{type}_{r,adj}$. Even after normalization, the market reaction to inventory announcement is uniformly stronger in the last time period. But now we can attribute that to higher *relative* importance of inventory announcements after 2014.

Not surprisingly, certain events on the oil market could have attracted public attention to EIA inventory news. Indeed, in the summer of 2014 the oil price started what would be one of the fastest and deepest falls. Early in 2015 oil inventories reached unprecedented levels convincing the market participants in the presence of immense oversupply of oil. Since that oil inventories have been continuously hitting the news wires, and all major news outlets have been publishing highlights of the EIA reports. Therefore, it is natural to expect greater awareness about inventories and EIA announcements.

However, by itself greater awareness is not enough to justify larger market movements in response to the EIA releases. If markets are efficient simply increasing the number of traders following the news would not affect movements of prices. In reality, however, the traders differ in knowledge and experience, and thus may have different pre-announcement beliefs, as well as interpret the same announcement differently. Thus, if information processing and portfolio rebalancing are heterogeneous, greater awareness of EIA news and potential entrance of less sophisticated traders into the post announcement trading, could have amplified the market responses to news.

[24]In principle, we could estimate the model directly incorporating a general component proportional to daily volatility into the coefficient of return response to announcement...

**Asymmetry in responses**

In this section we explore the asymmetry in returns responses. In principle, various frictions can create differences in market responses to negative and positive inventory surprises. The most obvious reason for asymmetry is the existence of upper and lower bounds on inventories. Indeed, inventories can't be negative as well as can't exceed a certain threshold defined by the total amount of available space in storage facilities. Thus, if over a particular time period, the market is relatively close to one of the limits, one could expect that the change in inventories in the direction towards this limit may trigger larger responses, than the movement away from it.

Historically, the academic literature has been mostly exploring non-negativity constraint on inventories and the possibility of stock-outs.[25]. However, recent oversupply episodes in the US market attractive public attention to the upper limit on inventories. When the space in storage facilities is running out, physical arbitrage becomes problematic. Thus, if the market experiences a further negative demand shock that is presumed to be temporary, oil speculators cannot respond by taking extra oil from the market and putting it in storage, to sell it in the future. The smoothing mechanism that should prevent current spot price from fluctuating downwards is now limited if not entirely shut down. And thus, one may expect the effect of extra negative demand shock on price to be large. Whereas a positive demand shock can be counteracted by speculators releasing some of the oil into the spot market, which should limit the oil price spikes.

Figure 8 plots the absolute values of previously estimated normalized returns coefficients. We don't see any systematic pattern in returns responses in the earlier period. However, our results indicate asymmetry present in the last two years.

The evidence on 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 bigger price moves. We believe our results to be more accurate, as we more precisely identify inventory surprises by using information in the API reports, but we also allow the data to speak about the optimal weighting scheme, which may evolve over time. Moreover, we estimate a detailed dynamic model using high-frequency data, whereas the papers above either focus on daily returns, or calculate cumulative returns over an event window, thus we achieve a more accurate estimation of price jumps.

## Volume response

Explicit modeling of trading intensity allows us to analyze idiosyncratic trading behavior of market participants around inventory announcements. On impact, trading volumes spike due to jumps

[25]See Deaton and Laroque (1992), Coibion and Yuriy (2015), Wen (2005). If inventories are scarce, than further unexpected draws from inventories should trigger larger spikes in oil prices, that unexpected build ups. Theoretical justification of this argument can be found, for example, in Deaton and Laroque (1992). Figure 2 in that paper displays a nice illustrative example. When inventories are low, the price is defined by the inverse demand function and currently available amount of commodity as a state variable. If the inverse demand is a convex function, than a decrease in current inventories would cause a larger increase in price, than a decrease in price caused by the same amount of extra commodity in storage

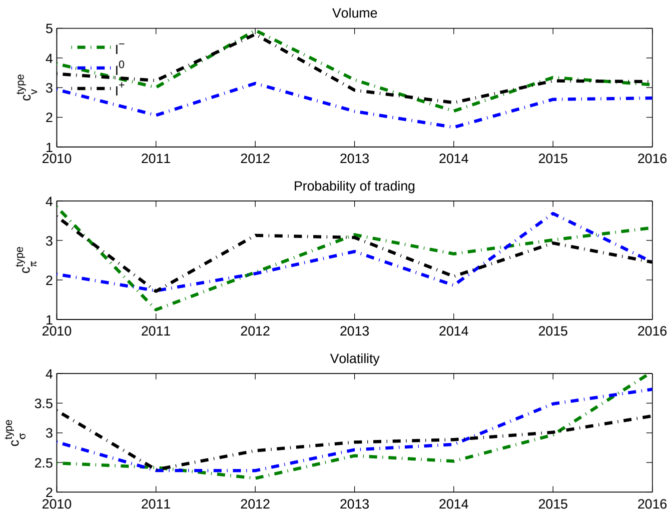


Figure 9: Trading volume, probability of trading, and volatility of returns reaction to inventory surprises by type of surprise.

in both the conditional mean of volumes, $\psi_t$, and the probability of trading, $\pi_t$. The magnitude of these jumps is defined by the indicators coefficients, $\{c_v^{type}, c_\pi^{type}\}$. The propagation, however, depends not only on the dynamic properties of volumes and probability of trading, but also on the evolution of conditional variance in response to news. Overall, we find that modeling trading probability is crucial for understanding and measuring propagation of inventory shocks for long distant futures contracts.

**Instantaneous effect** The point estimates of $c_v^{type}$ and of $c_\pi^0$ for the second contract are presented on the left panel of Figure 9. Both conditional mean of trading volumes and the probability of trading jump on impact. Our results show that probability of trading for the second-month contract jumps on impact from 0.9 to 1 in 2016, and from 0.6 to slightly less then 1 in earlier years.

Noticeably, we observe strong reaction of volumes to all announcements, including uninformative ones. That is, even when the market correctly predicts inventory change and no movements in prices are observed, we still see intensification of trading activity. Abnormal trading volumes without movements in prices have been documented in the literature before, thus our results reconfirm previous findings[26].

One explanation proposed in the literature 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, agents revise their beliefs in different ways. As a result, trade is possible, as agents adjust their positions according to individual changes in beliefs. However, when the market is surprised and average expectations need to adjust, that should trigger even larger response, and that is exactly what we observe immediately after the announcement.

Finally, in contrast to returns responses, we don't find any asymmetry, large negative and large positive surprises affect volumes in the same way.

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

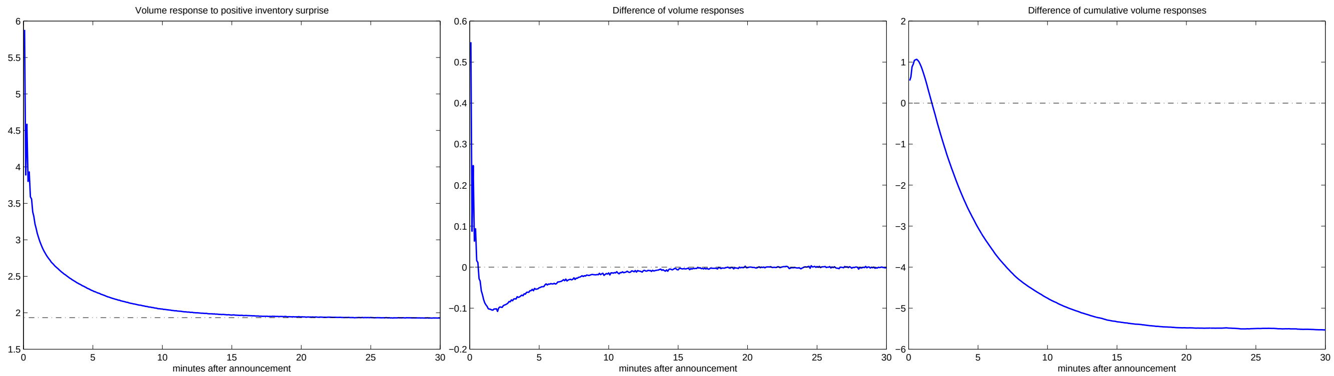


Figure 10: Left: Volume response to a large positive inventory shock.
Middle: The difference between trading volumes following a surprised announcement and trading volumes after an uninformative event.
Right: The difference between cumulative trading volumes following a surprised announcement and trading volumes after an uninformative event.
Estimated model using 2016 sample, second month futures contract.

**Propagation of inventory shocks** What is novel in our paper is the analysis of the propagation of inventory shocks. Naturally, one would expect volumes to jump on impact and then gradually fall. And in addition, we expect more intensive trading following surprising announcement at any moment after the announcement. The second prediction, however, turns out not to be always correct.

Left side of Figure 10 shows volume response to a large positive inventory shock in 2016. Indeed, we see a spike in volumes on impact and a slow decay afterwords. The usage of dynamic model allows us to make quantitative predictions. Our results indicate that it takes about 17-18 minutes for an inventory shock to be fully traded away. That number is lower for earlier years, and in 2010 it takes less then 10 minutes.

The picture in the middle of Figure 10 compares the trading intensity response following a surprising announcement with trading after an uninformative event. Surprisingly, we observe a non monotone relationship. On impact we already have seen that surprising event triggers larger trading. However, after less then a minute, that is reversed and trading intensity actually becomes *lower* following an inventory surprise. Thus, large inventory shocks seem to be traded away faster. Moreover, when we integrate the volume responses over time (see the picture on the right), we find that overall trading is *smaller* following a surprising announcement.

Main propagation results are summarized in Figure 11 which draws the difference in cumulative volumes following a large positive inventory shock and following an uninformative event, for all years in our sample. First, we clearly see an existence of two different trading regimes. In the early years, it would be an arrival of new information that triggers larger overall trading. However, in 2015 and 2016 it is the reverse.

A second important result is the difference in time needed for the difference in trading activity to fully unwind. In 2011 and 2014 the extra trading due to arrival of new information is finished in less then 3 minutes after the announcement, that is, we see that the curve completely flattens

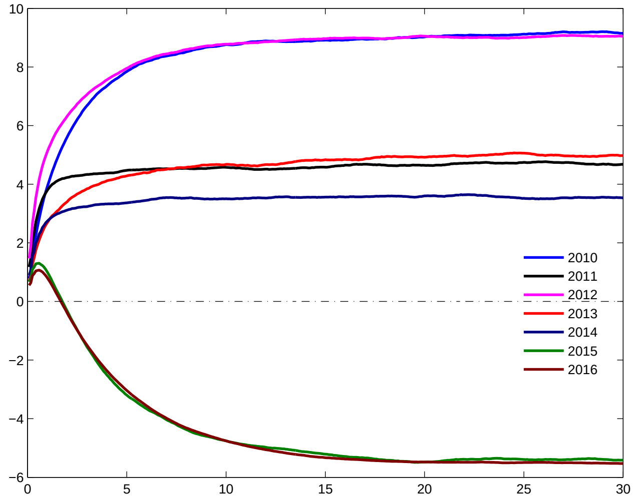


Figure 11: Extra trading due to arrival of new information: the difference between cumulative trading volumes following a large positive inventory shock and following an uninformative event, by year.

out. In contrast, in 2010, 2011, and 2012, the difference disappears after 10 minutes, whereas in 2015 and 2016 - only after 20 minutes. We interpret that as the time needed for individual market participants to process information and adjust the holdings correspondingly. In 2011 and 2014 the information is processed extremely fast, and after a few minutes, we don't observe any extra (or less) trading relative to uninformative announcement. However, in the last two years, the information processing takes much longer. Which is perhaps not surprising, given the extraordinary events happening in the oil market since the end of 2014.

Formally, it is the effect of inventory news on conditional variance of returns that triggers the difference. 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 of trading probability is crucial for understanding and measuring the propagation of inventory surprises for all futures contracts except for the nearest one. Not accounting for time-variability of trading activity significantly underestimates the total effect of the shocks, by distorting the propagation.

**Trading intensity by contract maturity** We document a difference between volume reactions of the first and second month futures contracts. Volume reaction to news, and thus potentially disagreement, varies over time, and it does so independently for each contract. To see that more clearly, we plot the volume response to uninformative announcements for both contracts on Figure 12. While the volume response of the first futures contract remains flat over the first 5 years, or even slightly increases, the second month contract shows more variation and a clear downward trend, from 2012 to 2014. Thus, even though *average market expectation*s are formed uniformly across contracts (as our previous results have shown), the dispersion of belief or beliefs' revisions in response to news are different across contracts, even the two most liquid ones. Second picture on the same figure plots average daily traded volumes for comparison, to show that there is no systematic difference in trading pattern, and it is the processing of inventory announcements that

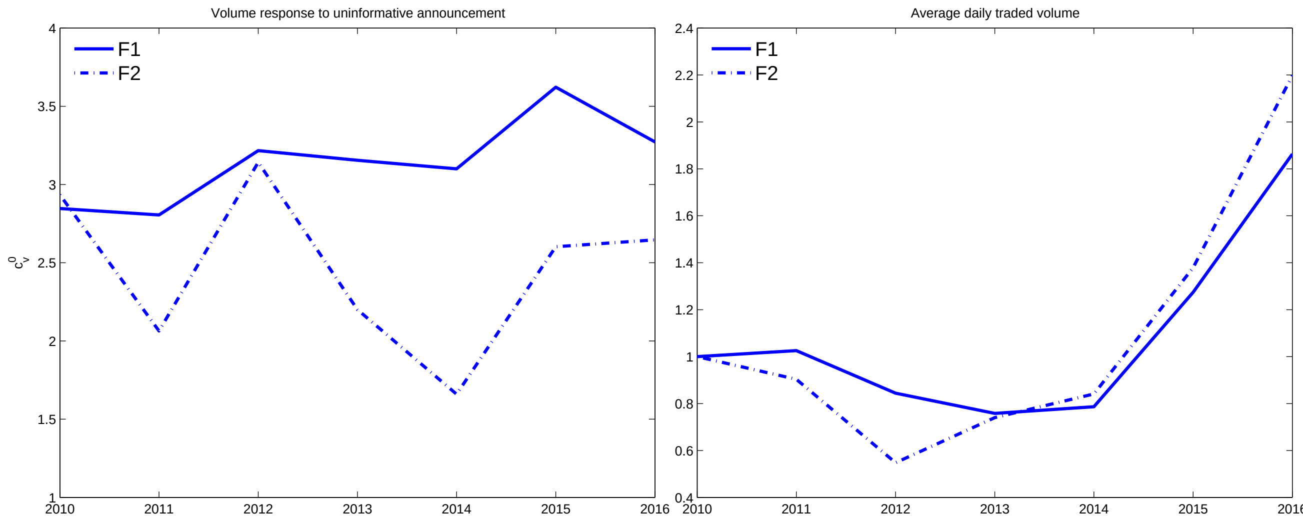


Figure 12: Trading pattern by contract maturity:
(on the left) Volume response to uninformative announcement.
(on the right) Average daily traded volumes.
Front-month contract,$F_1$−solid line, second month contract, $F_2$- dashed line.

is different.

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 to contracts with specific maturities; hence these funds are active only on 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.

### Summary of results

Our results can be summarized as follows. We document

- uniform formation of expectations;
- rising weight place on API estimates;
- strong negative relation between inventories surprises and returns;
- no effects on the term premium;
- differential volume effects by maturity;
- intensification of returns reaction a change in trading pattern after 2015.

# 7 Understanding the recent glut of oil

In this chapter we discuss how focusing on inventories and studying the market impact of inventory surprises improve our understanding of the oil glut.

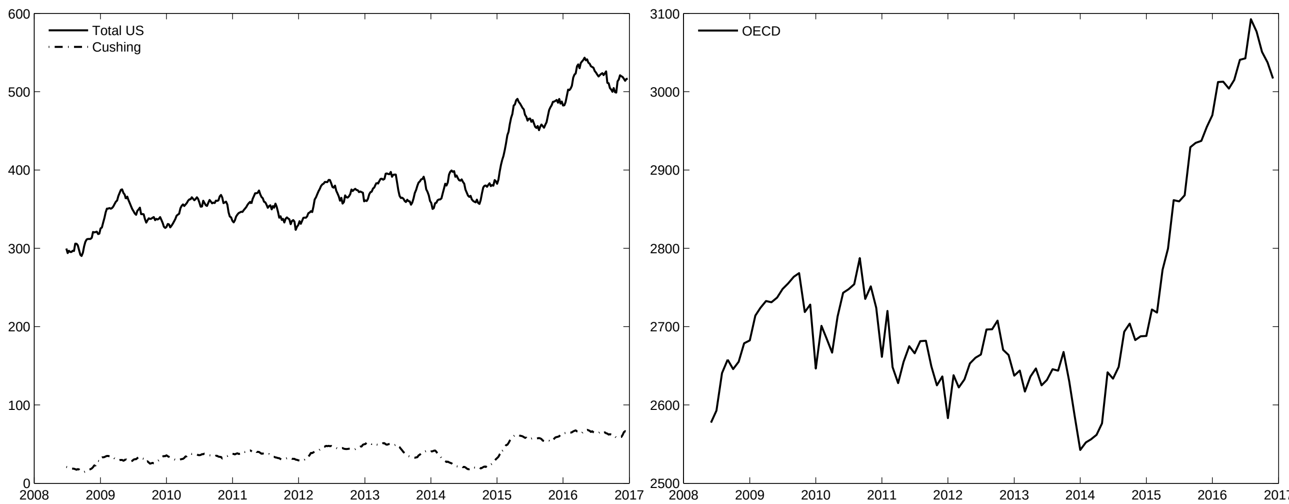


Figure 13: Left Panel: Weekly U.S. ending stocks excluding SPR of crude oil (mln barrels, solid line).Weekly Cushing, Oklahoma ending stocks excluding SPR of crude oil (mln barrels, dashed line).
Right Panel: Montly OECD commercial crude oil and other liquids inventory (mln barrels).

**The oil market before and after 2014**

Let us first describe the chronology of events. By 2014 US oil production exploded, reaching almost 9 mln barrels per day and was expected to grow even further. But at the same time struggling or slowing economies in China and Europe raised doubts that it would be possible to maintain the same pace of demand growth as before. In July of 2014 the price of oil started to decline, and gained momentum in November when OPEC declined to cut the production. By the beginning of 2015, oil inventories reached unprecedented levels and was interpreted as a sign of immense oil oversupply. In accordance with oversupply view, various agencies, including the International Energy Agency and EIA were reporting strong excess supply of oil. Very quickly the investment in oil drilling collapsed, give hopes that production cuts will follow and will finally clear the market. But the speed at which the production side could adjust was not certain. At any point of time, the traders were not sure whether the oil price had reached rock bottom or could continue falling. So while inventories continued climbing up, the fears that the market would no be able to absorb this excess supply intensified and further depressed the oil prices. Despite that the production side of the economy showed remarkable resilience, long-awaited production cuts were not happening. Month after month, the projections for a recovery in the oil price had been continuously pushed back.

Figure 13 shows evolution of oil inventories in the US. Surprisingly we don't observe any upward trend over the entire 2014, if anything, crude inventories have been decreasing for most of the year. It was not until early 2015 when inventories started to accumulate. Figure 13 also displays oil stocks in Cushing, Oklahoma, the main delivery point for oil futures contracts. Again for most of the year we observe decreasing stocks. Finally, right graph on Figure 13 displays crude stock in OECD countries relative to previous years average. Although inventories have been slowly rising over 2014, only in 2015 the average range has been exceeded. Turns it would be an oversimplification to attribute falling oil prices to oversupply of oil. However, plunging oil prices can still be explained by unwinding of oversupply *fears*.

So the key facts about the oil market in 2014 and 2015 are

- plunging oil prices since mid 2014;
- huge negative oil price reaction to OPEC announcement of no production cuts;
- flat oil inventories in 2014 and sharp increase in 2015.

Any consistent theory that aims to explain the collapse of oil prices in 2014 and subsequent events should be able to simultaneously account for these facts. Unfortunately, as we cannot directly observe market expectations, the story of particular evolution of oversupply fears matches the facts but cannot be refuted. Our goal is to bring another fact to the table that can be used to test such a theory.

**Evolution of market expectations**

We contribute to the literature by characterizing evolution of market reaction to inventory news. Changes in market reaction to news are connected with changes in market beliefs of scarcity of oil. When traders are more willing to attribute surprising increases of inventories to larger amounts of excess and unsold oil because they believe oil market to be oversupplied, market *expectations are revised stronger* in response to any given inventory shock, and thus we should observe larger market response to inventory shocks. By identifying shifts in market reaction to inventories, we can pin down structural changes in market belief, which should help us better understand the nature of the recent oil price collapse.

There are two competing theories behind the collapse of oil prices, which assume completely different behavior of market beliefs. First, 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 the overall oil production was expected to grow, the market maintained the belief that it would be OPEC producers who would adjust, basically would cut production to give way to shale oil producers. However, after the meeting in November, it became clear that OPEC producers were not willing to sacrifice their share of production. Hence, not only the global expectations of future path of oil supply had to be reconsidered, but the possibility of temporary oversupply in the nearest future became substantial. In other words, after the OPEC announcement, the market reevaluated the probability of transferring to oversupply state and found it considerable. In this case, ever since the announcement the market should have been constantly watching for signs of realized oversupply. But that implies that the link between oil inventories and perceived market tightness should have been intensified. Any extra barrel of oil in storage observed since after the announcement could have reflected the beginning of oversupply. Thus, the market response to inventory surprises, especially to positive surprise should have been sharply intensified since after the OPEC meeting.

Alternatively, there is a hypothesis of gradually unwinding oversupply fears. Partially based on Baumeister and Kilian (2016) result of partial predictability of oil price collapse, the proponents of this theory emphasize that ever since July of 2014, the market gradually became convinced in the presence of oversupply, thus justifying the oil price fall since the very beginning.

Thus, one theory assumes gradual revision of expectations, whereas the other conjectures a distinct structural shift in market expectations after the OPEC announcement. Our methodoly allows to estimate the evoltuion of market beliefs directly, and thus compare the two theories.

### Background changes

Our goal is to search for structural breaks in market reaction to inventory news. Hence, our model will be adjusted to allow some parameters, mainly the market reaction parameters, to vary over time.

The problem, however, is that there may be other structural changes happening on the oil futures market. One potential reason is proliferation of alternative investments, such as exchange traded funds. Flow of investment to ETFs tracking oil futures prices (such as USO) may have shifted the composition of trading strategies utilized at the oil futures market. And as parameters of our model reflect the set of strategies utilized, there may be another structural shift in this parameters. Thus, by allowing only a subset of parameters to adjust, we may wrongly pick up these 'background' structural changes.

One sign of this changes is an increasing in volatility since 2014 that we have depicted on Figure 7. Second reason to be worried about is a change in intensity of trading. Figure 14 shows the number of 5-sec intervals with zero traded volumes by contract over time. Again we see a distinct shift in intensity of trading at the end of 2014.

To overcome this problem, we will allow for two separate changes. First, we call it $I{-}transition$, corresponds to evolution of market reaction to inventory news, that is the one that we are interested in. The second one, called as $H - transition$, corresponds to time varying composition of trading strategies and intensity of trading.

### Econometric framework

To parsimoniously model the evolution of some parameters of our model, we utilize a threshold autoregression model (TAR). Denote time variable by $d$ , in our case it represents a particular announcement (week) in our sample. Formally, a 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 given by a non decreasing transition function $G(d, d^\star, \delta)$ that lies between zero and one and depends on the time index, which is given by a week $d$, threshold $d^\star$, and a vector of parameters $\delta$. By specifying different functions $G$ we specify a particular timing and speed of adjustments of market responses. Notice, that when $G$ goes from zero to one, $x(d)$ smoothly goes from $x_0$ or $x_1$, reflecting the transition. Parameters $\delta$ define how fast the transition occurs.

First, we perform a simple LR test for non linearity. So consider a linear form of time dependency, in other words the transition function is given by

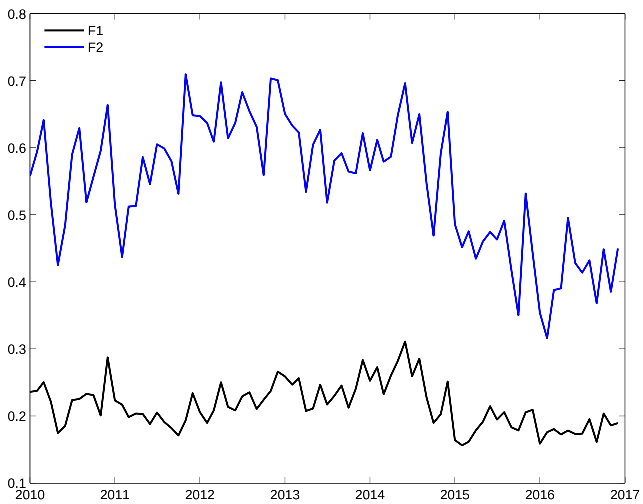


Figure 14: Fraction of 5-sec intervals with zero traded volumes. Blue line corresponds to the first-month contract, black line to the second month contract.

$$G^{linear}(d, d^\star) = d - d^\star$$

where $d^\star$can be chosen arbitrary. We can estimate parameters $\beta_x = x_1 - x_0$ and test whether all the betas are jointly equal to zero:

$$H_0 : \beta_{c_v^{type}} = \beta_{c_r^{type}} = ...\beta_{\tau_1} = 0$$

The null hypothesis can be tested with a LR test[27].

Once non linearity is established, we do a formal search for a distinct or instantaneous structural *break* in market responses to announcements. In this case the transition function is given by a "before-after" indicator:

$$G^{break}(d, d^\star) = I_{d \geq d^\star}.$$

As was mentioned before, we allow for two separate transition, $I - transition$ and $H - transition$, corresponding to changes in market response to news and to changes in trading pattern respectively. Thus, we will have two transition functions, $G_I^{break}(d, d_I^\star)$ and $G_H^{break}(d, d_H^\star)$, where we allow two breaks two happen at potentially different times, in other words $d_I^\star$ and $d_H^\star$ may be different.

For some parameters, such as constants, we will need to allow for adjustments following both transitions, in which case the parameter $x(d)$ varies over time according to

$$x(d) = x_0 + (x_1 - x_0)G_I^{break}(d, d_I^\star) + (x_2 - x_1)G_H^{break}(d, d_H^\star),$$

and thus we need to estimate three different levels.

[27]Technically, we use composite likelihood and not the real one. Thus, we can't use standard asymptotic results. However, we follow the applied literature and apply the test anyways.

### Estimation procedure

We need to decide which parameters vary over time according to one or the other transition dynamics. We group all parameters into 3 groups. In our notation $\{c_v^{type}, c_r^{type}, c_\sigma^{type}\}$ for all types of news will now be functions of time and will follow $I-transition$ only. Similarly, all parameters in probability equation, namely $\{\omega_h, \zeta_k, \xi_k\}$, all parameters but a constant in the conditional mean equation for trading volumes, namely $\{\alpha_k, \alpha_k^0, \beta_k\}$, and all coefficients but a constant in the conditional mean equation for returns: $\{\rho_k, \rho_k^0\}$ and some coefficients in the conditional variance equation, $\{\delta_k, \gamma_k\}$, will all follow $H-transition$ only. Finally, all the constants, $\{w, \mu, \omega\}$ as well as persistence parameters in the conditional variance, $\{\tau_k\}$ and daily level of variance, $\phi$ are allowed to follow both transitions.

We estimate the augmented model for each set of values of $d_I^\star$ and $d_H^\star$ corresponding to weeks from September of 2014 to the end of April of 2015.

### Results

We find that for both contracts the null hypothesis is rejected suggesting presence of time variability.

By comparing the values of composite log-likelihood we find that the break in trading pattern is clearly chosen to occur around the first week of December of 2014. We believe that this can be attributed to an immense money inflow to ETFs, such as the United States Oil Fund trading crude light oil futures with short maturities. However, our results do not indicate a clear break in market response to news. If anything, the break happened in the middle of February of 2015, *after* inventories reached extremely high levels. Hence, oversupply beliefs tend to lag behind 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.

# 8 Literature Review

Our paper is related to the literature on belief formation. Typically, these studies rely on survey expectation data to test expectation formation theories, see for example, Coibion and Yuriy (2015), Andrade and Le Bihan (2013), Gennaioli et al. (2016), or Mankiw et al. (2013). Our approach is complementary to these studies as we use high frequency transaction data, including trading volume, to characterize the expectation formation process. Closest to our paper is the recent work of Bollerslev et al. (2016) that proposes to use the intraday volume-volatility elasticity estimated using high-frequency data to measure disagreement among market participants. From a theoretical perspective, the paper builds on the seminal work of Kandel and Pearson (1995) that proposes a differences-of-opinion model to explain decoupling of trading volumes from movements in returns, see also Banerjee and Kremer (2010). We adopt the same view and interpret movements in trading volume in response to uninformative announcements as reflecting disagreement among the traders. However, we utilize a model of the joint high-frequency dynamics of futures returns, return volatilities and trading volumes around the event to estimate the impact of news, whereas the measure of Bollerslev et al. (2016) is based on the differences between the post- and pre-event levels of the instantaneous volume intensity and volatility, which are recovered non-parametrically using high-frequency data.

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

First, we offer a novel procedure to identify inventory surprises. Most studies define inventory surprises as the differences between the reported EIA estimates and the median survey forecasts. Given general public interest to oil inventories, various surveys are available that directly ask agents' expectations of EIA annouced changes, including the surveys conducted by Reuters and utilized in See Bu (2014), Bloomberg used by Halova Wolfe and Rosenman (2014), Halova et al. (2014), Miao et al. (2018), and Platt's. There are only a handful of papers that utilize API information to shape market expectations. One exception is Armstrong et al. (2017) which defines inventory surprise as the difference between EIA and API reported values, but neglects information in initial surveys. Another exception is Ye and Karali (2016), which defines inventory shocks sequentially: 'API shock' as the difference between the API announced change and the forecast, and 'EIA shock' as the difference between EIA and API. Thus, in both papers it is implicitly assumed that API value is observed by all market participants and that right before the EIA release, the market expectation is exactly equals to API released value. However, Ye and Karali (2016) also argue that the market reaction might depend on the history shocks, and thus allow the market return to react differently to the same EIA shock depending on the previously realized API shock. The intuition is simple, the market might react differently to two consecutive positive surprises, than to the contradicting signals. However, such logic should be viewed as a reduced way to account for mispecified beliefs. Instead, we would maintain the assumption that only the magnitude of the surprise matters for the market reaction to EIA announcement, however we will carefully specify the formation of market expectation (and thus define surprises) building on a simple logic of an optimal signal extraction problem. In particular we will look for the optimal weighting scheme, and we compare the expectation formation process across contracts with different maturities.

Second, our goal is to provide a complete characterization of market impact of news. Most studies look only at the impact on oil futures returns, and use only the first-month futures contract. Bu (2014) and Halova Wolfe and Rosenman (2014) also investigate the effect of news on return volatility. But there are almost no papers that study the effects of the announcements on the term structure of futures prices. 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, our results do not show that the longer maturity contrast displays smaller responses. The difference in the results can be attributed to different identification of 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 in dealing with ultra-high frequency financial data. In particular, we analyze trading volumes which are characterized by non-negativity, high volatility and 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 impact of inventory shocks mostly uses even non lower frequency data, such as daily returns, see Bu (2014) and Miao et al. (2018). Halova et al. (2014)calculates continuously compounded returns in an intraday event window surrounding the EIA announcement[28]. One exception is Ye and Karali (2016) that works with 5-minute returns and uses the methodology developed in Andersen et al. (2003).

Finally, our sample covers 7 years, from 2010 to 2016. [29] 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 the recent history.

Our paper contributes to the debate on the recent oil glut in the US. It still remains an open question whether oversupply of oil was the main determinant of the decline in the price of oil 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 how exactly the oversupply beliefs unwound. Arezki and Blanchard (2015)argue in favor of a structural shift in market expectations, following the OPEC's announcement in November of 2014 to maintain production level despite growing shale 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, it doesn't necessarily imply that the market incorporated that information into beliefs. Our approach allows to directly test for structural breaks in expectations.

# 9 Conclusion

The characterization of market beliefs is an important tool in understanding of the price formation. We provide a framework that integrates ultra high frequency into characterizing of the 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 then can be tested against our empirical evidence.

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.[30]

We apply our methodology to understand the unwinding of the oversupply beliefs in 2014 and early 2015. However, our approach can also be used to understand the perceived market conditions in 2011, when WTI-Brent spread occurred. Over many years two oil benchmarks have reflected existing supply-demand conditions in the oil market. One of them, West Texas Intermediate (WTI) crude oil has long served as a primary global benchmark, a reference price for oil traders.

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

[29]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.

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

However, in 2011 the ability of WTI to serve as a global benchmark was called into question, when the price of WTI significantly diverged from another benchmark, Brent, representing 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. The lack of pipeline capacity to move oil south from Cushing to the Gulf Coast created the 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[31], global providers of commodities information[32,33], exchanges [34,35], and by the media[36] and energy analysts[37]. Our approach can be used to understand the evolution of market beliefs at times of these 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.

Baffes, J., M. Kose, F. Ohnsorge, and M. Stocker (2015). The great plunge in oil: causes, consequences, and policy responses. World Bank Group.

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. No. w24167. National Bureau of Economic Research.

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


[31] U.S. Enery Information Administration (EIA), 'This week in Petroleum', September 12, 2012.
[32] See Platts, Price group/Oil Division, 'New Crudes, New Markets' special report of March, 2013.
[33] Argus White Paper, 'LLS – The Emergence of a Secondary Benchmark', 2014
[34] Intercontinental Exchange(ICE), 'Oil Benchmarks: Implications & Opportunities', May, 2011
[35] CME Group, 'Trading WTI and Brent 101. CME Group as the exchange of choice', March 19, 2014.
[36] 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.
[37] The Oil Drum, 'Why are WTI and brent Prices so Different', February 22, 2011 by Gail Tverberg

Bollerslev, T., J. Li, and Y. Xue (2016). Volume, volatility and public news announcements. *The Review of Economic Studies*.

Borenstein, S. and R. Kellogg (2014, January). 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 G. Yuriy (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–78.

Deaton, A. and G. Laroque (1992). On the behaviour of commodity prices. *Review of Economic Studies 59*(1), 1–23.

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

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

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

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.

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.

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.

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.

Selezneva, V. (2015). A side effect of financial innovation.

Wen, Y. (2005, November). Understanding the inventory cycle. *Journal of Monetary Economics 52*(8).

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 Expectation formation process

What defines $\omega$? To see the intuition behind it$\omega$, let's first abstract from the complexities of market trading and heterogeneity of traders. Assume that there is only one agent that sequentially observes two signals, $BBG$ and $API$, and optimally weights them. Let's assume that EIA reveals the fundamental value without noise, whereas both the survey and the API report contain independent errors. Denote by $x$ the fundamental value of the change in inventories release by EIA. Let $z^{bbg}$ be the signal received (privately) by professional forecasters, whereas let $z^{api}$ be the signal provided by API. Let's assume that everything is distributed normally. So let

$$\begin{aligned}
&x \sim N(0, \sigma_x^2) \\
&z^{bbg} = x + \eta, \quad \eta \sim N(0, \sigma_\eta^2) \\
&z^{api} = x + \xi, \quad \xi \sim N(0, \sigma_\xi^2) \\
&\xi \perp \eta \perp x
\end{aligned}$$

Optimal signal extraction problem with normally distributed disturbances implies $E[x|z^{bbg}] = \dfrac{1}{1+\sigma_\eta^2/\sigma_x^2} z^{bbg}$ and

$$E[EIA|BBG, API] = (1-\omega)BBG + \omega API$$

where $BBG \equiv E[x|z^{bbg}]$ is the survey estimate [38] , $API \equiv z^{api}$ denotes the API estimate, and the weight placed on API should be given by

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

Intuition is straightforward. If $API$ signal is very precise, noise-to-signal ratio $\sigma_\xi^2/\sigma_x^2 \to 0$ goes to zero, while $\sigma_\eta^2 > 0$, then the weight placed on API increases, $\omega \to 1$, meaning that agents disregard imprecise information from the survey. Alternatively, if initial information revealed by the survey is very precise, $\sigma_\eta^2/\sigma_x^2 \to 0$ while $\sigma_\xi^2 > 0$, then $\omega \to 0$, and API signal is disregarded. Thus, higher $\omega$ reflects better quality of API signal relative to information available to professional forecasters.

# 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 $\nu$, with the conditional mean and the conditional variance defined as before.

[38] Here we are working with expectation, however in the data we will use the median, rather than the average forecast.

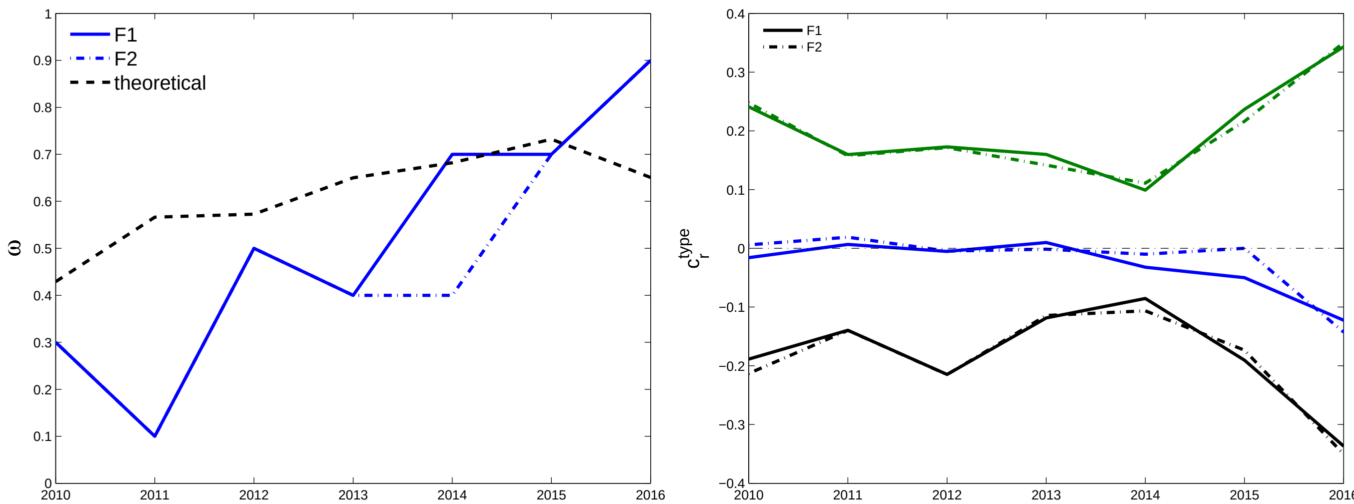


Figure 15: Capturing fat tails.
Left Panel: Estimated weight $\omega$ of API signal. Solid blue line corresponds to the nearest to expiration contract, dashed blue line corresponds to the second-month contract. The black dotted line displays the value $\omega$ which represents a solution to the simple signal extraction problem.
Right Panel: Normalized return reaction to inventory surprises for the first 2 contracts. Solid line - F1, dashed - F2.

The shape parameter,$\nu$, 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, expecially more pronounced for the long maturity contracts.

The estimated values of $\omega$are presented on Figure 15 separately for the first two contracts. In general the values are similar, however with fat tails the estimates of $\omega$ fluctuate around 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 all years 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 15 compares the normalized return reaction to news for the first two most liquid futures contracts. A strong negative relation is present (st.err are not presented on 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 15 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 $\omega$ 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.