How does the financial market update beliefs about the real economy? Evidence from the oil market*

Stanislav Anatolyev — CERGE-EI and NES

Sergei Seleznev — INECO Capital Ltd

Veronika Selezneva† — CERGE-EI

August 17, 2020

Abstract

Oil traders find it challenging to process all information and choose which sources to follow. Inventories represent a perfect source, as they provide important information regarding real agents' intertemporal decisions and can easily be observed in real time. However, inventories do not contain full information about the state of the oil market. We show that financial traders fail to acquire additional information and treat inventories as a sufficient statistic. As a result, the financial market fails to distinguish realized shocks from news shocks, and treats all shocks as persistent. To confirm this hypothesis, we use oil inventory announcements to identify market inventory surprises; and we estimate a model of the joint dynamics of returns and return volatilities around announcements using high frequency data on oil futures contracts with short and long maturities.

Keywords: oil market, ultra high frequency data, futures returns, inventory surprises, term structure curve, theory of storage

JEL classifications: C32, C58, G12, G13, G14

*GACR Grant 17-27567S from the Czech Science Foundation is gratefully acknowledged. We are grateful to the Co-Editor and three anonymous referees for insightful suggestions and constructive critique, and Christiane Baumeister and James Hamilton for useful discussions. We also thank conference participants of CFE-2017, AMEF-2018, ISCEF-2018, CEF-2018, CEMA-2018, IAAE-2018, and EEA-ESEM 2018 for helpful comments. Valentin Artemyev, Matěj Bělín and Mariia Kosar provided excellent research assistance.

†Corresponding author. Address: CERGE-EI, Politickych vězňů 7, 11121 Prague 1, Czech Republic; e-mail veronika.selezneva@cerge-ei.cz.

1 Introduction

Clearly, the financial market does not have perfect information about the state of the real economy. It is challenging, however, to characterize the processes through which information is updated and document specific informational frictions. We use the unique structure of the oil market to investigate how the financial market updates beliefs about the shocks that hit the real economy.

What is unique about the oil market is the dynamic behavior of oil inventories.[1] Inventories provide important information about real agents' intertemporal decisions. The U.S. Energy Information Administration (EIA) aggregates dispersed information about thousands of individual inventory decisions and makes public official announcements every week, which are closely monitored by the financial market. However, we argue that changes in inventories, taken alone, cannot uniquely identify the shock that hit the oil market. The futures market reaction to inventory news reveals whether financial traders acquire additional information to distinguish the shocks, or simply treat inventory changes as a sufficient statistic.

The real agents with access to storage use inventories to optimally respond to different shocks. Shocks that can trigger inventory response, may include not only the shocks that affect current market conditions (we call them *realized shocks*), but also shocks that change agents' expectations about future economic fundamentals without affecting current fundamentals (*news shocks*).[2] The key mechanism is that oil should flow in or out of storage until current and expected future oil prices are equalized (up to all the costs and a risk premium). However, equalization of prices requires inventories to be unconstrained. When the spare capacity is near exhaustion or when the inventories have been depleted, temporary shocks can no longer be smoothed out, and the prices have to adjust. We rely on this intuition to develop a number of empirical predictions on the joint behavior of oil prices and inventories in response to various shocks.

[1] Other commodities could be studied in this way as well, including natural gas, metals, some agricultural commodities etc. However, the oil futures market is more liquid and thus better serves our purposes.

[2] The theory of storage was developed by Deaton & Laroque (1992), Pindyck (1990), Fama & French (1987), Ng & Pirrong (1994), Nielsen & Schwartz (2004), Alquist & Kilian (2010), Kilian & Murphy (2014), Wen (2005).

If the financial traders achieve full information following inventory announcement, we should observe a large variety of market responses that conform with our theoretical predictions. In contrast, if the financial traders do not acquire all relevant information, the joint dynamics of inventories and returns may lack certain variation that can be detected and analyzed. In particular, if information acquisition is costly, the financial agents may optimally decide to not distinguish the shock by type (realized and news shocks), as theoretically predicted by Mackowiak et al. (2018), or by persistence.[3]

To estimate the market response to inventory surprises, we use weekly oil inventory announcements provided by the EIA and released at a pre-specified time. We carefully model market expectations about the upcoming inventory announcement using a standard Bloomberg survey forecast and the American Petroleum Institute (API) inventory report released one day prior to the EIA report. To estimate the market impact of news, we develop a model of the joint high-frequency dynamics of futures returns and return volatilities around the EIA announcements. The model is estimated using ultra high frequency (5-second time intervals) data on short and long maturity oil futures contracts on the WTI oil traded at the NYMEX. To account for illiquidity of long term futures contracts, we model trading inactivity directly following Hautsch et al. (2013). For the benchmark case, we estimate the model for the two periods of unconstrained inventories, 2010-2014 and 2017-2019, and for a period of constrained inventories, 2015-2016. As a separate exercise, we identify the exact moment of both transitions.

Our results indicate a strong negative link between inventory surprises and returns.[4] Moreover, the negative comovement of returns and surprises is observed in overwhelming 94% of large surprises. The predominance of the negative relationship suggests that the market treats all inventory changes as mainly reflecting realized shocks, rather than the shocks to expectations of future oil market conditions (which lead to a positive relationship). Our results are consistent with the theoretical prediction of Mackowiak et al. (2018) that the financial traders do not distinguish the shocks by type. Interestingly, in a different setting, Crouzet & Oh (2016) do not find a negative

[3] Mackowiak et al. (2018) argue that it can be optimal for rationally inattentive agents not to distinguish carefully between current increases in productivity and news about future increases in productivity.

[4] This result is robust and confirms previous findings of Bu (2014), Halova Wolfe & Rosenman (2014), Halova et al. (2014), and Miao et al. (2018).

comovement of sales and finished-goods inventories in the data, and interpret this as evidence of the insignificance of news shocks in business cycle fluctuations. We offer an alternative explanation for this result. We also carefully discuss the evidence of the presence and significance of news shocks in the oil market.

Our second main finding is a lack of any effect of inventory surprises on the term spread on the shorter end of the curve. When inventories are unconstrained, this result is consistent with the behavior of the real agents. However, the lack of term structure adjustments when inventories are high is a striking result and contradicts conventional wisdom. Indeed, the near-term prices are expected to respond stronger to shocks when inventories are constrained, thus diverging from the more distant end of the curve. 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. The co-occurrence of high inventories and a steep term structure curve is a robust and general phenomenon, observed for many commodities; see Gorton et al. (2012). Despite that, when inventory news announcements come, traders do not revise their expectations accordingly. Our results do indicate a muted response of longer maturity contracts to inventory news when inventories are constrained. However, the reaction is still stronger than in normal times, while one would expect the opposite.

We argue that the lack of term premium adjustments (or the overreaction of longer maturity contracts) is observed because financial agents do not acquire enough information to identify the persistence of shocks precisely. Intuitively, agents are not accustomed to doing this. During normal times, no term structure adjustments are needed, and thus traders may rationally develop a habit of disregarding the persistence, and trading as if all shocks were permanent. When suddenly inventories become constrained, traders fail to adapt. Such history dependence is a known feature of dynamic information acquisition, discussed by Mackowiak et al. (2018).

It could also be that the lack of term structure adjustments reflects an unawareness of the financial sectors of the constrained state of inventories. However, we believe that this is unlikely to be the case. Our transition result indicates that the market reaction to inventory surprises abruptly intensified in February 2015, exactly when inventories spiked. Moreover, the media had been following the evolution of inventories quite closely.[5]

Finally, our results also show negative asymmetry when inventories are high. The market responds more strongly to negative inventory surprises, though the theory predicts the opposite. This result may suggest that financial traders do not fully comprehend what the state of constrained inventories implies for equilibrium prices.

Failure to identify the true underlying reasons for changes in inventories can have long-term consequences. One example is an oil futures liquidity shock related to the rollover of the United States Oil Fund. Mou (2010) and Selezneva (2015) argue that the rollover practice of such funds could temporarily steepen the term structure curve. In response, owners of oil storage facilities engaged in physical arbitrage, further increasing oil inventories.[6] However, financial traders interpreted the rise in inventories as reflecting additional already realized positive supply or negative demand shock, and thus put additional downward pressure on the current and expected oil prices. The lower current prices lead to additional increases in oil stored, and the spiral continued. Although the mispricing can be only temporal, the resulting additional increase in oil inventories is more long-lived, as it affects the oil supply in both current and future periods. Our results thus suggest that the decreases in the price of oil in 2008/9 and 2015 could have been amplified by the inability of the financial market to correctly identify the underlying reasons for inventory changes.

It is common in the literature to use survey data to test for the presence of information frictions.[7] Our approach is complementary to these studies. We provide a framework that integrates inventories data and ultra high frequency futures prices data into an analysis of the formation of beliefs of financial traders. Our paper is also related to Li (2019), who tests the rational inattention assumption by comparing bond price underreactions to default-relevant and interest rate news.

Our study is not the first to estimate the market impact of oil inventory announcements. However, we are the first to conduct the analysis and interpret the results through the prism of the theory. First, we argue that is it crucial to distinguish periods of constrained and unconstrained inventories, though none of the existing studies do.[8] We also emphasize the importance of the upper

[5]See, for example, https://www.reuters.com/search/news?blob=EIA+inventory+report.

[6]Ederington et al. (2020) empirically document the effect of price spreads on crude oil inventories.

[7]See Froot (1989), Coibion & Gorodnichenko (2012), Andrade & Le Bihan (2013), Mankiw et al. (2013), Greenwood & Shleifer (2014), Gennaioli et al. (2016).

[8]See Bu (2014), Halova Wolfe & Rosenman (2014), Halova et al. (2014), Miao et al. (2018).

bound on inventory storage, whereas earlier literature mainly focuses on the lower bound following Deaton & Laroque (1992). Second, we investigate longer maturity contracts, while existing studies mostly focus on the nearby futures contract. One exception is Miao et al. (2018), who analyze the first six continuous contracts and documents weakening of the magnitude of the price response with maturity, but these results cannot be meaningfully interpreted without distinguishing periods of constrained and unconstrained inventories.

Methodologically, our study also differs in a number of important ways. Most importantly, we sharpen the identification of market surprises by estimating the weights placed by the market on survey information and on API information. Most studies only use various median survey forecast: Reuters, Bloomberg, or Platt's.[9] One exception is Armstrong et al. (2017) who define an inventory surprise as the difference between EIA and API reported values, but neglect information in initial surveys. Another exception is Ye & Karali (2016) who define inventory shocks sequentially and include both shocks in the regression model. In this paper, we directly model the formation of market expectations, and estimate the weights on different signals along with the other parameters of the model.

To the best of our knowledge, our study is the first to use trading intensity data to account for illiquidity of long maturity futures contracts. The literature on the impact of inventory shocks most frequently uses daily returns; see Bu (2014) and Miao et al. (2018). Halova et al. (2014) calculate continuously compounded returns in an intraday event window surrounding the EIA announcement.[10] One exception is Ye & Karali (2016), who work with 5-minute returns and use the methodology developed in Andersen et al. (2003). Finally, our sample covers 10 years, from 2010 to 2019, thus including the shale oil boom, the proliferation of ETFs, and one of the most dramatic oil price collapses in the recent history.[11]

The remainder of this article is organized as follows. In Section 2 we briefly outline the theory of storage and use the theory to formulate testable implications. Our empirical methodology is described in Section 3. Our main findings are presented in Section 4. We discuss our findings and

[9] See footnote 8; see also Crego (2020) who estimates the effect of EIA announcements on the stock market.

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

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

conclude in Section 5. An online Appendix contains auxiliary results and discussions.

2 What can be identified from joint behavior of inventories and futures prices?

In this section we show how the inventories data can be useful to characterize the beliefs of financial traders about the changes in the current and future state of the oil market. We begin by illustrating the role of inventories in oil price determination, then talk about the disaggregated nature of the oil market, and finally discuss the information aggregation role of the financial market. Rather than developing a full scale model of the interaction of the real and financial sectors, we aim to identify a number of robust empirical predictions that can be tested.

2.1 Role of inventories in oil price determination

The theory of storage governs the behavior of the real agents in response to shocks that hit the oil market. We begin by briefly outlining a toy model to illustrate the role of inventories in oil price determination. For simplicity, we consider a single isolated oil market and abstract from the existence of the financial market.

Illustrative example

Imagine a two period economy, $t = 1, 2$. There is a single risk neutral producer of oil. The oil production is exogenously given and equals to $q_t^s = 1 + \varepsilon_t$ in period t, where ε_t is a zero mean iid random variable with variance σ^2. This assumption is a limiting case of inelastic production. We abstract from modeling the demand side, and simply assume that the demand for oil is given by a downward sloping function $q_t^d(p_t) = 2 - p_t$.

Current vs. expected future prices Start from the case when oil is a non-storable commodity. Imagine that a positive supply shock hits the market, $\varepsilon_1 > 0$. When oil cannot be stored, this shock depresses the current spot price, $p_1 = 1 - \varepsilon_1$, but as shocks are independent, the expectation of the second period price is unaffected, and an upward sloping term structure curve is observed:

$$\text{Case 1: } p_1 = 1 - \varepsilon_1 < Ep_2 = 1.$$

The total expected profit conditional on the realization of a positive supply shock is $\pi(\varepsilon_1) = 2 - \varepsilon_1^2 - \sigma^2$.

Now imagine that oil is a storable commodity. Let us denote by x an amount of oil that the oil producer decides to put in storage after observing the supply shock.[12] In response to a positive supply shock, $\varepsilon_1 > 0$, the producer may put oil in storage, let us denote it by $x > 0$. Extra barrels oil stored away from the market, increase the current price, $p_1 = 1 - \varepsilon_1 + x$, while decrease the expected future price, $Ep_2 = 1 - x$.

The maximum total expected profit over the two periods is achieved at $x = \varepsilon_1/2$. By smoothing production optimally over time, the oil producer achieves a higher profit. One can think of inventories as of a mechanism that helps producers to transform a large one time supply shock ε_1 (a.k.a. temporary shock), into a smaller shock $\varepsilon_1/2$ that hits the market in both periods (a.k.a. permanent shock). In equilibrium, perfect equalization of current and expected future price of oil is observed:

$$\text{Case 2: } p_1 = Ep_2 = 1 - \frac{\varepsilon_1}{2}.$$

Imagine now that inventories cannot exceed $\varepsilon_1/4$, for example due to a capacity constraint. In this case, producers can only partially increase expected profits by storing maximum amount of oil allowed. As a result, perfect equalization of prices cannot be achieved:

$$\text{Case 3: } p_1 = 1 - \frac{3}{4}\varepsilon_1 < Ep_2 = 1 - \frac{\varepsilon_1}{4}.$$

We see that relative to no inventories case (case 1), the first period price is now higher, but it is lower than the price in the unconstrained case (case 2).

We can make our first main observation. As long as the inventory constraint does not bind,

[12] We do not put any restrictions on x. To avoid negative inventories stored in equilibrium, one can assume that the producer owns an initial stock of oil, $x_0 > 0$, and restrict the range of ε_t.

inventories serve to smooth out temporary shocks. Oil flows in or out of storage until current and expected future prices are equalized.[13] However, in the absence of inventories or when inventories are constrained, temporary shocks strongly affect the spot price.[14] More generally, when the storage capacity is limited, expected long term prices should be less sensitive to the current shocks. Only truly persistent shocks should move the expectations of the prices many periods ahead.

Similar analysis, of course, can be conducted for negative ε_1. We then immediately see that when inventories are not constrained, the market response to positive and negative shocks is symmetric. In contrast, when inventories reach the capacity constraint, an additional positive ε_1 cannot be fully stored away, while a negative ε_1 can still be fully smoothed out by taking oil from existing storage. Thus, the price reaction is stronger to a positive shock, that is, a shock in the direction of currently binding constraint.[15] This is our second insight from the theory of storage.

News shocks vs. realized shocks The final prediction of the model compares the outcomes of the two different types of shocks. Note that in response to current shocks ε_1, inventories and prices move in the *opposite* direction. Indeed, oil flows into (out of) the storage to be sold at a relatively higher future (higher current) price, when a positive (negative) supply shock or negative (positive) demand shock hits the economy in the current period. This result holds irrespective of whether inventories are constrained or not. We will call shocks such as ε_1 as *realized shocks.*

Let us now investigate an arrival of a *news shock*. We assume that in period $t = 1$, the agent receives news about a negative supple shock *in the future*, in period $t = 2$. In our model that means that the agent observes ε_2 already in period $t = 1$.[16] To be specific, assume that the agent observes $\varepsilon_2 < 0$, and there are no other shocks today, $\varepsilon_1 = 0$.[17] The total expected profit over the

[13] So far we have abstracted from storage costs, this is a straightforward extension. Also, if agents are risk averse the main argument still holds. The only difference is that if, for example, prices are expected to be only a little bit higher in the future, risk averse agents would not find it worthwhile to put extra oil in storage because it is risky. Hence, equalization of prices is observed up to a risk premium.

[14] We do not model investments in storage capacity. Construction of additional tanks is a slow and expensive process. Moreover, most of the time capacity utilization is low, hence additional investment to storage capacity is likely to be unprofitable. If we look at the storage in Cushing, Oklahoma, its working storage capacity increased from 46 mln bbls in September 2010 to 77 mln bbl in March 2016 and remained constant afterwards.

[15] This effect is similar to consumers having a larger marginal propensity to consume out of unexpected wealth shocks, when they sit on the borrowing constraint.

[16] We could assume that the agent receives a noisy signal about ε_2, but that would not change the argument.

[17] We separately analyze news and realized shocks for the sake of exposition only. An arrival of two shocks simultaneously can be considered accordingly. The inventory response depends on the direction and relative size of

two periods equals $2 - x^2 - (\varepsilon_2 + x)^2$. The maximum of this function is achieved at $x = -\varepsilon_2/2 > 0$. The equilibrium price in the first period is $p_1 = 1 - \varepsilon_2/2 > 1$, and thus is equal to $p_2 = 1 - \varepsilon_2/2$.

Hence, an arrival of a negative *news* shock is associated with a build up inventories today and an increase in the equilibrium price; in other words, both inventories and prices move in the *same* direction. The result continues to hold when inventories are constrained. In this case, producers can only partially increase expected profits by storing a maximum amount of oil allowed, and thus the first period price still increases, but by a smaller amount that in an unconstrained case.[18]

Our final observation is the following. Realized shocks and news arrivals about future fundamentals lead to opposite predictions about the joint behavior of inventories and prices. We can use the data on the comovement of inventories and prices to identify the types of shocks that hit the market.

It should be noted that the first two predictions (about the term structure response and asymmetry) continue to hold for news shocks as well. We have already shown that when inventories are unconstrained, current and expected future oil prices are equalized and the market response is symmetric. When inventory capacity is limited, equalization of price may not be achieved. Indeed, a large negative ε_2 would make the expected future price higher, than the current price (see footnote 18). Finally, the current price reaction to a negative ε_2 shock that cannot be fully smoothed out is now *smaller* than the reaction to a positive ε_2 shock that can be smoothed out. This is because the only channel through which future shocks can affect current prices is via the equilibrium response of inventories. If inventories are constrained, the response can only be relatively small.

Empirical predictions of theory of storage

Our toy model is extremely stylized; however, it captures the role of inventories in shaping the equilibrium behavior of oil prices. A more general model with endogenous convex production and uncertain demand generates similar empirical predictions; see Deaton & Laroque (1992) or Wen (2005).[19] We summarize these predictions below.

---

the shocks.

[18]If inventories cannot exceed $-\varepsilon_2/4$, then $x = -\varepsilon_2/4$, and $p_1 = 1 - \varepsilon_2/4$, $p_2 = 1 - 3\varepsilon_2/4$.

[19]In Appendix A.2, we discuss how our approach relates to the notion of the convenience yield.

Summary of empirical predictions If the real agent's behavior conforms with the theory of storage, then

1. Realized shocks lead to a negative correlation between changes in inventories and returns, while news shocks lead to a positive correlation.

2. When inventories are unconstrained, current and expected future oil prices are equalized. When inventories are constrained, the relative response of current and expected future oil prices is defined by the *persistence* of shocks and by the expected date of the shock (realized or news).

3. The lack of spare storage capacity is associated with an increase in the oil price sensitivity to shocks and asymmetry.

The last prediction requires a short discussion. In the model above we assume a linear demand function. In reality, the demand for oil is likely to become less elastic for higher volumes, as it becomes increasingly harder for the refineries to process an additional barrel of oil and market the products. Larger adjustments of the spot price may be needed to clear the already abundant market. Hence, when inventories are high and approach the capacity level, we may expect to see a larger reaction of the oil price to shocks, both positive and negative. Moreover, the convexity of the demand function creates an additional reason for asymmetry, as positive shocks require even larger adjustments of the spot price to clear the market than negative shocks.

2.2 From spot to futures prices: aggregation of dispersed information

The empirical predictions above have been derived for a single isolated oil market. In reality, the oil market represents a union of a large number of local markets spread all over the country. Recent technological advances in oil extraction have opened oil production fields in many isolated and distant locations such as North Dakota. The complexity of oil production and transportation networks (and also the diversity of oil grades) makes it infeasible for a centralized market place for oil trading to exist. There does not exist a single oil spot price, instead, each oil grade in each location has a separate price negotiated by local market participants. It is natural to expect some variation in local prices, as every local market is not only hit by common aggregate shocks, but may also experience changes in local conditions.

Each producer must form expectations about the dynamics of the local price to make optimal decisions. The inability to observe prices at all other locations, may prevent producers from acquiring full information about aggregate shocks. To clarify intuition, we further extend our toy model to introduce dispersed beliefs, see Appendix A.1.

The oil pricing can be greatly facilitated by the financial market, especially by the oil futures market. The financial market not only provides risk sharing opportunities, but also aggregates dispersed information.[20,21] Indeed, the trading activity of the real agents reflects their oil production and consumption decisions; moreover, the financial traders have expertise, technical possibilities, and incentives to acquire dispersed information.

In this paper, we investigate whether the financial market is able to perfectly identify the type, magnitude and persistence of the aggregate shock that hits the oil market. We suggest a reduced form analysis using public inventory announcements.

2.3 Understanding market response to public inventory announcements

Inventory announcements reveal aggregate information about *past inventory decisions* of a large number of the real agents. Clearly, individual inventory decisions are potentially observable. If the market could perfectly aggregate information and incorporate it into prices, such announcements would bring zero new information and would have absolutely no effect on the oil futures prices.

In contrast, under dispersed information, inventory announcements are informative; see Appendix A.1 for intuition on the role of inventory announcements and the connection between 'inventory' shocks and underlying demand and supply shocks that hit the real economy. Not being able to collect information about each and every local market, the financial traders may form biased beliefs, such as under- or overestimate the magnitude of a shock, incorrectly assess its persistence, misclassify news and realized shocks. The main question, however, is whether observing aggregate

[20]It has become a common practice to use the WTI futures price as a reference or a benchmark price, while each oil grade in each location is then traded at a local pre-specified price discount, which can be revised occasionally. See crude oil bulletins published by the Plain All American that specify spot prices for each location: https://www.plainsallamerican.com/customer-center/crude-oil-bulletins. The discounts are revised monthly.

[21]See the recent paper Goldstein & Yang (2019) for an analysis of various roles of the financial market and a feedback effect on the real economy.

information is sufficient for the financial traders to perfectly identify the aggregate shock.

It should be noted that we will maintain an assumption that risk premium is not directly affected by information conveyed in the inventory reports. This is a standard assumption in the event study literature; see for example, Bianchi et al. (2019).[22] However, we will later show that allowing for time varying risk premium only makes our results stronger.

Full information is achieved after inventory announcement Our analysis is based on the following main insight. When aggregate inventory information fully completes the information set of the financial agents, the updates of their beliefs and the corresponding changes in the futures prices should *embed the empirical predictions* derived above. As a result, we should observe a large *variety of market responses*, especially diverse when inventories are constrained.

To illustrate, we show two examples of beliefs updates that can generate both positive and negative comovement of inventory surprises (unexpected part of total inventory change) and the futures price changes following the announcement.

1. Imagine that the oil market is hit by a one-period positive supply shock. The real agents increase total inventories by x barrels, decreasing the average spot price by y \$/bbl. Imagine that the financial market *underestimates the magnitude* of the shock by half. The financial agents believe that total inventories will increase by $0.5x$ barrels, and the average spot price decrease by $0.5y$ \$/bbl. Accordingly, the front month futures price decreases by $0.5y$ \$/bbl. Once the announcement is made, the market finds out that the change in total inventories is actually *higher than expected*, and the futures price *decreases* by additional $0.5y$ \$/bbl.

2. Imagine that the market receives news about a one-period negative supply shock in the next period. Total inventories increase by x barrels and the price increases by y \$/bbl. Imagine that the financial market *misclassifies the shock* as a positive current supply shock and *underestimates its magnitude*. It expects the total inventories to increase by $0.5x$ barrels and the average spot price to fall by $0.5y$ \$/bbl. The futures price decreases accordingly by $0.5y$ \$/bbl. Once the announcement is made, the market observes that the combined change in inventories is *higher than expected*. If

[22] We would like to emphasize, that we do not assume that risk premium is zero and thus the futures price equals expected future price. We assume that risk premium is constant in a short window around inventory announcement.

at that moment the market finally *learns the type of the shock*, then the futures price dramatically *increases* by $1.5y$ \$/bbl to offset the initial incorrect reaction.

Other combinations of aggregate shocks and mistakes reflected in pre-announcement beliefs can be considered in a similar way.

Full information is not achieved even after inventory announcement The lack of certain market responses could suggest that the market fails to identify certain shocks even *after observing the total inventory change.* We have seen already that inventories do not contain full information about the oil market.[23] Of course, the financial traders may *acquire additional information* from other sources. However, information acquisition is costly. The cost may be so large that the financial agents may *optimally* decide not to distinguish between shocks of certain types, leading to the following empirical predictions:

1. Not distinguishing shocks by type (news vs realized) leads to the same observed sign of the correlation between inventory changes and returns.

2. Not distinguishing shocks by persistence leads to parallel shifts of the term structure curve in response to all inventory surprises even when inventories are constrained.

The first prediction follows the theoretical result of Mackowiak et al. (2018) who solve the problem of rationally inattentive agents facing news shocks in a standard macro setting. Intuitively, if the realized shocks are more prevalent than news shocks, then rationally inattentive agents would treat all shocks as realized and react accordingly.

Similarly, the financial agents may not acquire enough information to precisely identify the persistence of shocks. When inventories are not constrained, no term structure adjustments are needed as inventories smooth out all the shocks. As a result, traders may rationally develop a habit of disregarding the persistence, and trading as if all the shocks were permanent. When suddenly inventories become constrained, traders may take time or fail to adapt.[24]

[23] The same increase in inventories may reflect an arrival of a current positive supply shock or a future negative supply shock.

[24] One may also wonder if the financial traders are aware of the status of inventories. However, as long as the financial traders place a positive probability on the inventories being constrained, the longer end of the futures curve should respond less than the current price to a temporary shock.

3 Methodology

Our goal is to analyze the market reaction to inventory surprises and using the empirical predictions outlined above, test if the financial traders acquire additional information or simply treat inventories as a sufficient statistic. In this section, we describe identification of oil market surprises, outline how we estimate market response to these surprises, and discuss our data sources and data preparation.

3.1 Institutional background

Weekly estimates of crude oil inventories in the U.S. are provided by the U.S. Energy Information Administration (EIA), a statistical and analytical agency within the U.S. Department of Energy. Any company which carries or stores more than 1000 barrels of oil may be selected into the EIA weekly sample based on a procedure that assures coverage of 90% of the market. Typically, the sample includes gathering and pipeline companies, and storers of crude oil. The selected firms are required to report the end-of-week amount of oil in their storage facilities. On the following Wednesday, a summary report is released in the form of an EIA publication, the *Weekly Petroleum Status Report*.[25] The report becomes available to the public at 10:30am Eastern time and is closely followed by the media.

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

[25]For some weeks which include holidays, releases are delayed by one day and released at 11 am.

costly. However, for more specialized traders, API reports represent quite valuable information, as historically, API and EIA estimates tend to be close to each other. Discrepancies are believed to occur due to different procedures utilized to estimate the remaining 10% of the market.

Figure 1 gives a glimpse of the market reaction to EIA and API announcements. It depicts average (across days) returns over each 5 minute interval of the trading day. The red line takes the average over non-report days which are Monday, Thursday, and Friday. The blue lines correspond to report days: the top panel represents Wednesdays when an EIA report is released, while the bottom panel stands for Tuesdays and API reports. We see that for most part the blue and red line coincide, but at times of report releases (10:30 and 4:30), we observe considerable spikes in absolute returns, meaning that the market reacts quite strongly on average to inventory news. The market reaction to API reports is comparable in magnitude to the market reaction to an EIA release, despite its restricted access and the late time of a release. Hence, neglecting information reported by API may significantly bias the identified surprises and thus distort the estimates of market impact.

Figure 1: Market reaction to oil inventory announcements.

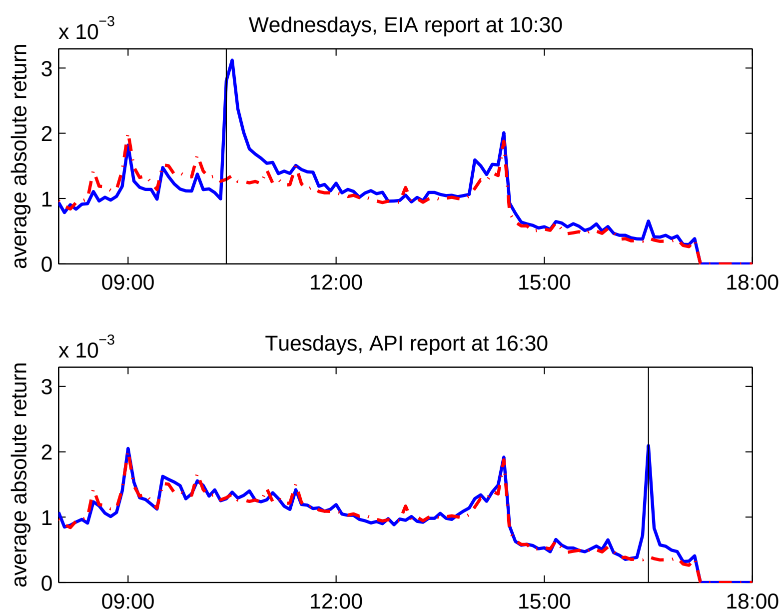

3.2 Definition of market surprises

Let us denote by $\Delta\text{Inv}_\tau^{EIA}$ the change of inventories (normalized by the total oil inventories for the prior week) released by the EIA on Wednesday of week τ. The market surprise is defined as the difference between the realized and expected value at a moment just before the announcement:

$$x_\tau = \Delta\text{Inv}_\tau^{EIA} - E^M[\Delta\text{Inv}_\tau^{EIA}]. \tag{1}$$

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

Unfortunately, we cannot directly observe market expectations and thus have to make assumptions on how the market forms expectations. It has become common in the literature to use surveys of professional forecasters to proxy for market expectations. Given general public interest in oil inventories, various surveys are available that directly ask about agents' expectations of EIA announced changes, including the surveys conducted by Reuters, Bloomberg, and Platt's.[26] The pairwise correlations between the three median forecasts are well above 0.97, which suggests that the information content of these three surveys is the same. Thus, we follow the literature and use the Bloomberg survey; we denote the median *forecast* of the survey of professional analysts conducted by Bloomberg and released on Monday of week τ by $\Delta\text{Inv}_\tau^{BBG}$. In contrast to other studies, we also use the numbers released by the American Petroleum Institute. We denote the reported change in total inventories on Tuesday of week τ by $\Delta\text{Inv}_\tau^{API}$. Finally, we assume that the market forms expectations according to the following linear function:

$$E^M[\Delta\text{Inv}_\tau^{EIA}] = \alpha + \omega\Delta\text{Inv}_\tau^{API} + \gamma\Delta\text{Inv}_\tau^{BBG}. \tag{2}$$

Under the normality assumption, the linear form is optimal and the weights simply reflect precision of the two available signals; otherwise, the linear function can be viewed as an approximation. The weight placed on API information by the market as a whole should depend on the relative quality

[26] Halova Wolfe & Rosenman (2014), Halova et al. (2014), Miao et al. (2018) use the Bloomberg survey. Bu (2014) uses Reuters.

of the API signal, as well as on the overall access of investors to API reports.

We suggest estimating $\{\alpha, \omega, \gamma\}$ jointly with other parameters of the model. Our approach has two main advantages over the standard two step procedure. First, the standard errors of the market response coefficients of interest automatically account for the estimation uncertainty of $\{\alpha, \omega, \gamma\}$. Even more importantly, we impose fewer assumptions on how the market forms expectations. For example, one could estimate the best linear forecast of the actual changes of inventories using a reasonable rolling window, and then use the estimated values of $\{\alpha, \omega, \gamma\}$ to form the surprises. However, one may have doubts that the market correctly estimates the precision of signals and chooses the best linear forecast to form expectations.[27,28] Later we will argue that even though the quality of the API signal was improving over time, the market was slow adjusting the weight placed on the API report accordingly. It should be noted that our approach does not create identification issues, as the variables used to model expectations are predetermined at the time of announcement.

3.3 Econometric framework

To estimate the market response to inventory surprises defined above, we develop an econometric model of joint high frequency dynamics of returns, return volatility and trading activity around announcements.

Let us use the time index τ to denote announcement days (the EIA reports are weekly, so we have about 50 announcements per year). The market surprise in week τ identified according to our procedure outline above is denoted by x_τ.

Our analysis focuses on one hour around EIA announcements, that is, from 10 to 11 am. The sampling is done at a 5-second frequency, which yields 720 data points for each announcement day τ. We use the time index $t = 1 : 720$ to denote 5-second intervals.

Trading inactivity 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; we

[27]There is ample evidence of various behavioral mistakes that the traders tend to make, including underreaction to news and overreaction to stale information.

[28]We could also estimate a simple time series model of weekly inventory changes, and use it to make the forecast, see Roesch & Schmidbauer (2011). However, that would again imply making strong predictions about the way the market forms expectations.

observe substantial periods with no recorded transactions. To illustrate the severity of this issue, we calculate the fraction of 5-second intervals with zero trading volume over the entire sample, separately for each contract. We find that the first four contracts by maturity have 22%, 55%, 82%, and 91% of intervals with no trading, respectively. Clearly, trading activity decreases dramatically with maturity. Thus, we aim to build a model that can handle illiquidity of long maturity contracts by explicitly accounting for trading inactivity. For that purpose, we use data on trading volumes, though we do not model their evolution.

Returns and volatility To simplify the notation we drop the index τ when describing the intraday dynamics of returns; we denote the return over the 5-second interval t of announcement day τ by r_t, not $r_{t,\tau}$.

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

$$\mu_t = \mu + \sum_{k=1}^{q_r} \rho_k r_{t-k} + \sum_{k=1}^{q_r^0} \rho_k^0 \mathbb{I}_{\{V_{t-k}=0\}} + \mathbb{I}_{\{t=t^\star\}} R(x_\tau)\sigma_{\tau,S}, \tag{3}$$

where the return response function is given by piecewise linear schedule

$$R(x_\tau) = \mathbb{I}_{\{-\bar{x}\leq x_\tau \leq \bar{x}\}} c_r^0 x_\tau + \mathbb{I}_{\{x_\tau > \bar{x}\}} \left(c_r^+(x_\tau - \bar{x}) + c_r^0 \bar{x}\right) + \mathbb{I}_{\{x_\tau < -\bar{x}\}} \left(c_r^-(x_\tau - (-\bar{x})) + c_r^0(-\bar{x})\right), \tag{4}$$

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

$$\ln \sigma_t^2 = w + \phi \ln \sigma_{\tau,S}^2 + \sum_{k=1}^{p_\sigma} \psi_{1,k} \ln \sigma_{t-k}^2 + \sum_{k=1}^{q_{\sigma,1}} \psi_{2,k}\eta_{t-k} + \sum_{k=1}^{q_{\sigma,2}} \psi_{3,k}|\eta_{t-k}| + \mathbb{I}_{\{t=t^\star\}} \left(c_\sigma^0 + c_\sigma^1 |x_\tau|\right), \tag{5}$$

where $\eta_t = r_t/\sigma_t$ are standardized returns.

In the benchmark case, we classify inventory announcements into large positive (negative) surprises, if $x_\tau > \bar{x}$ ($< \bar{x}$). The main coefficients of interest are $\{c_r^0, c_r^+, c_r^-\}$ that define the return response function $R(x_\tau)$. If the surprise is relatively small, the conditional mean jumps by $c_r^0 x_\tau$.

Large positive surprises increase the conditional mean by $c_\tau^+(x_\tau - \bar{x}) + c_\tau^0 \bar{x}$, and similarly for negative surprises. Thus, we allow for asymmetry and non-linearity in the return response to surprises.[29]

The adjustment of the conditional variance is proportional to the magnitude of the surprise: $c_\sigma^0 + c_\sigma^1 |x_\tau|$. The EGARCH equation has a number of advantages among ARCH models with leverage (see Rodriguez & Ruiz, 2012), one of which is positiveness of conditional variances. We include in the right hand side the daily level of volatility, $\sigma_{\tau,S}^2$, for day τ to account for slow moving changes in volatility. We compute $\sigma_{\tau,S}^2$ as filtered realized volatility for day τ (see Appendix A.4).

We also normalize the reaction of the conditional mean of returns to news by filtered realized volatility, $\sigma_{\tau,S}$. The purpose is to simplify the comparison of the results over time. Without normalization we would not be able to meaningfully interpret the intensification of returns reaction to inventory news. This could be misleading, as the market could have become more sensitive to *all types* of news, not just to inventories. We choose the volatility level as the normalizing factor, because volatility increases when prices respond more strongly to each news arrival.[30,31]

Time varying probability Following Hautsch et al. (2013), we assume that the probability of inactivity may also vary over time. We define the conditional log odds ratio as $h_t = \ln \frac{\pi_t}{1-\pi_t}$, and adopt the following specification for the time varying probability of inactivity:

$$h_t = w_h + \varkappa_\tau + \sum_{k=1}^{p_h} \zeta_k h_{t-k} + \sum_{k=1}^{q_h} \xi_k \mathbb{I}_{\{V_{t-k}>0\}}, \tag{6}$$

where $\varkappa_\tau$ is a daily component meant to pick up changes in the average probability of trading.

Estimation Our model describes the trading dynamics of a single futures contract. The only parameters that are the same across the contracts are $\{\alpha, \omega, \gamma\}$, which drive the formation of expectations. We assume that both trading activity and returns conditional on trading are drawn independently across contracts. This is a reasonable simplification for the following reasons. First,

[29] We have experimented with different forms of the return function. The results remain similar.

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

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

we are working with ultra high frequencies, hence arbitrage conditions are unlikely to hold precisely at each time period. Second, trading inactivity differs significantly even across the first three most liquid futures contracts (fractions of zero trading intervals are 22%, 52% and 82%). Finally, the moment of an announcement, when we expect all three contracts to be traded simultaneously, is captured directly; while the joint informational content is captured via requiring market surprises to coincide across contracts.

We estimate the model jointly for the first three futures contracts. This means that $\{\alpha, \omega, \gamma\}$ are estimated using the observed responses of the three most liquid contracts. To account for possible time variation in trading patterns and thus in parameters, we perform estimation separately for each calendar year.[32] Our approach slightly differs for longer maturity contracts, as will be discussed below.

The model is estimated using the quasi-maximum likelihood approach, with the standard errors constructed using the "sandwich" Bollerslev & Wooldridge (1992) form and computed via numerical derivatives.

3.4 Data sources and preparation

Data We utilize changes in weekly U.S. ending stocks excluding SPR and including lease stocks of crude oil before September 2016 and U.S. ending stocks excluding SPR afterwards as published by the Energy Information Administration.[33] For identification of market surprises we use the estimates of inventory changes published by the American Petroleum Institute, as well as the median consensus forecast from the Bloomberg survey of analysts.

The high frequency data on WTI oil futures traded at NYMEX (CME group) was obtained from TickData. Our sample covers the period from 2010 to 2019 and contains 521 announcement days. We focus on one hour around EIA announcements, from 10 to 11 am. A one hour long interval is long enough to provide a reasonably precise parameter estimates of our dynamic model.

[32]This is necessary due to dramatic changes in trading intensity over time. If we take the second-month futures contract, the empirical probability of no-trading over 5-second intervals decreased from 0.6 in 2014, to 0.45 in 2015, and to 0.4 in 2016.

[33]The EIA terminated publication of the first series in September 2016. Lease stocks have been relatively stable with a range of 30.6 mln barrels to 33.1 mln barrels from January 2014 through June 2016.

The sampling is done at a 5-second frequency.[34]

Figure 2: Weekly U.S. ending stocks of crude oil and capacity utilization.

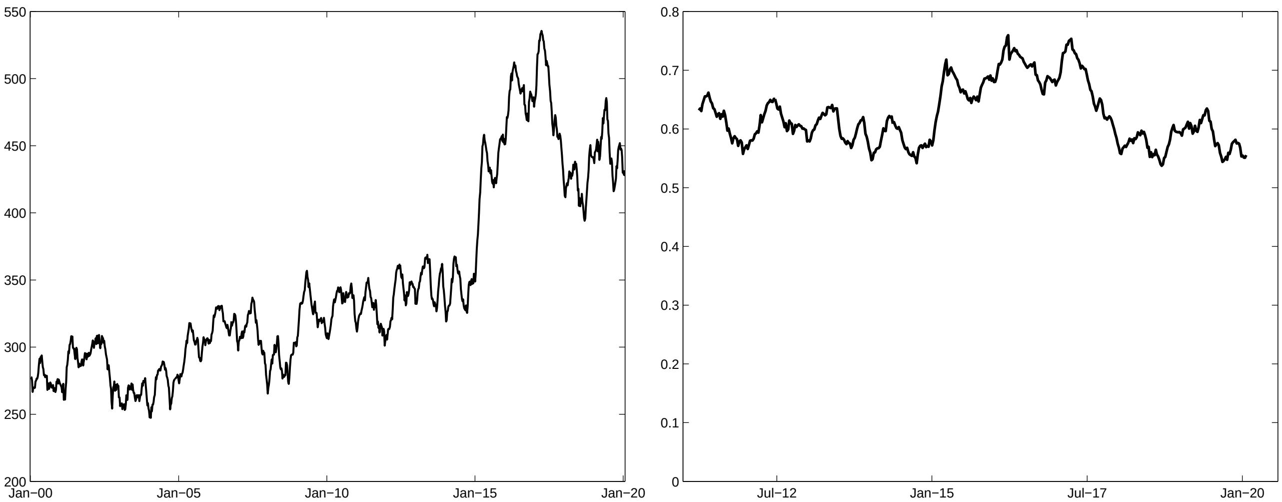

Notes: Left panel: weekly U.S. ending stocks of crude oil excluding SPR (mln barrels). Right panel: Capacity Utilization = U.S. ending stocks of crude oil excluding SPR / U.S. commercial crude oil stocks and storage capacity (available biannually since March 2011).

Identifying times of binding inventories For the benchmark exercise, we identify the period of binding inventories using the level of total inventories in the U.S. Weekly estimates of crude oil inventories in the U.S. are shown in Figure 2. Until about the end of 2014, oil inventories had been fluctuating around 350 mln barrels. At the beginning of 2015, inventories skyrocketed and reached unprecedented levels. Over the next two years the oil stocks remained high, at about 500 mln barrels, and only gradually decreased in 2017. Figure 2 also displays capacity utilization. We see elevated levels in 2015-2016, and perhaps over a first few months of 2017. Finally, we also analyze the stocks at Cushing, Oklahoma, the world's largest crude oil storage facility and the main deliver hub for the oil futures contracts. By April 2015, the oil stock reached 60 mln barrels, getting close to the maximal 70 mln barrels of reported capacity, and remained high at this level for the next two years.[35] Thus, for our purposes we will consider the period from 2015 to 2016 as a period of constrained inventories, and the year 2017 as a transition year. Finally, in section 4.3 we identify

[34]The results remain qualitatively unchanged if we use 5-minute frequency and are available upon request.

[35]See the EIA working storage capacity data as of March 31, 2015, which are available from https://www.eia.gov/petroleum/storagecapacity/.

the moment and the speed of the transitions from unconstrained to constrained regimes and back.

Of course, the spare capacity was still available even in 2015, but at much higher prices. For example, storing oil in tankers is significantly more expensive than in on shore facilities, but it can definitely be used as a facility of last resort. In 2015, storing oil even in crude oil tankers became profitable. E.A. Gibson Shipbrokers Ltd and Frontline Ltd document that up to 20 VLCCs were used for temporary storage by the end of January, 2015, which is equivalent to 30-40 mln barrels of oil in storage. The tanker rates spiked, the cost of renting a VLCC for 1 year increased from \$33,000/day at the beginning of 2014 to \$65,000/day by mid January 2015.[36] Therefore, our assumption on the exhaustion of spare capacity in 2015-2016 seems reasonable.

Choice of futures contracts and data preparation We need to deal with a number of issues specific to the futures market. One such feature is expiration. We follow a standard approach in the literature and rely on a rolling procedure to create continuous futures contracts (the details can be found in Appendix A.3).[37]

Another issue is illiquidity. The estimation of market responses for long-maturity futures contracts requires a modification of our approach. Even though our model is specifically designed to handle a certain level of illiquidity, once we move beyond the first three futures contracts, liquidity drops considerably and precludes any meaningful analysis. Fortunately, we find that certain contracts maintain reasonable liquidity throughout the year: the December and June contracts. Disproportional interest in these contracts perhaps reflects the convenience that mid year and end of the year expiry brings to reporting of hedging procedures; it may also reflect certain market coordination over time. However, if we plan to utilize contracts with fixed maturity dates rather than continuous contracts, we have to revisit the stationarity issue. As the expiration date approaches, the contracts are used for different trading strategies, but we do not allow the model parameters to vary over time.

We could restrict the sample to only a few months, e.g., to consider only three months from January to March of 2015, and thus limit the variation in maturity.[38] However, we would not have

[36]Hellenic Shipping News from January 19, 2015 and March 16, 2015.

[37]See, for example, Halova et al. (2014) or Gorton et al. (2012).

[38]The resulting maturities of each contract at the start of each quarter are displayed in Table A.1 in Appendix A.3.

enough weekly inventory announcements to perform meaningful estimation. To compensate for that, when working with longer maturity contracts, we combine the same quarters of different years into one sample.[39] Given the evolution of inventories, it is reasonable to combine years 2010-2012, 2013-2014, 2015-2016, and 2018-2019. We exclude 2017 as a year of gradual transition from one regime to another (see Section 4.3).

We also simplify the return response function when dealing with longer maturity contracts. In particular, we do not distinguish large or small surprises, and do not separate positive and negative surprises, so that $R(x_\tau) = c_r x_\tau$. This is due to the insufficient number of large surprises over some quarters. In this exercise, we also consider the values of $\{\alpha, \omega, \gamma\}$ as given and equal to the estimated values for the first three most liquid contracts.

Other details In the benchmark exercise, we fix $\bar{x} = 0.006$. This value was chosen as optimal for the front month contract for the vast majority of years, and also generates a reasonable number of large surprise as will be discussed below. We decide to fix $\bar{x}$, rather than allow it to vary across years, to facilitate the comparison of the magnitude of the market reaction over years and also over maturities.[40]

The number of lags is chosen using BIC information criterion for the front month futures contract. This result in $q_r = 7$, $q_r^0 = 0$, $p_\sigma = 2$, $q_{\sigma,1} = 2$, $q_{\sigma,2} = 2$, $p_h = 8$, and $q_h = 7$.[41]

4 Results

To facilitate the comparison, we present our results separately for the periods of unconstrained and constrained inventories. However, because the predictions made by the theory of storage are

[39]For example, we combine the first quarters of 2015 and 2016 into one sample and call it 'Q1,15-16'. Over this period of time, we analyze the trading activity of the so called 'current December contract', which is constructed by combining the data on the CLZ 15 contract during 2015Q1 and on the CLZ 16 contract during 2016Q1. Similarly, the 'next year June contract' is constructed by combining the data on CLM 16 for 2015Q1 and CLM 17 for 2016Q1. The 'next year December contract' similarly combines CLZ 16 and CLZ 17. We use the standard NYMEX terminology, where CLZ T denotes the crude light futures contract that expires in December of year T, and CLM T is the contract that expires in June of year T.

[40]The results remain unchanged when we allow $\bar{x}$ to be chosen optimally for the first three futures contracts, or when we fix it at other reasonable levels. The results are available upon request.

[41]We do not do complete search over all possible combinations of the number of lags of each kind; instead, we check that a change in one number at a time is not improving the objective function.

Figure 3: Estimated weight of API signal and correlation of realized inventories.

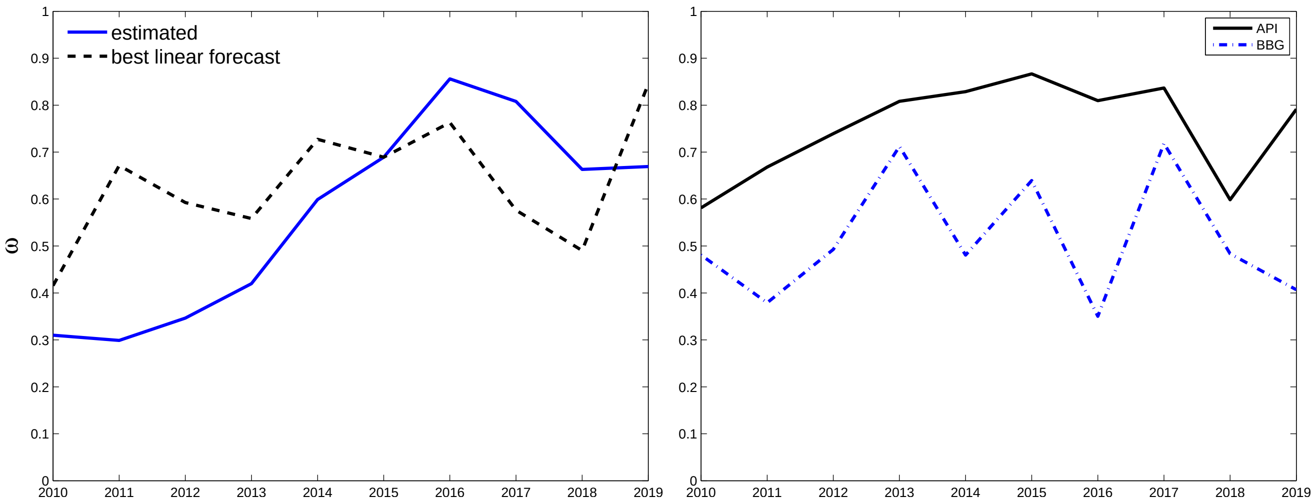

Notes: Left panel: Estimated weight ω of the API signal. The solid blue line corresponds to ω estimated jointly for the first three most liquid contracts. The black dotted line displays the value of ω representing the best linear forecast. Right panel: Correlation of EIA realized inventories changes and API estimates (solid line); of EIA and Bloomberg median forecast (dashed line).

all conditional on a shock, a key element of our analysis is the identification of oil shocks. Thus, we start by describing the evolution of market perception of API information and characterize the distribution of realized market surprises.

4.1 Formation of expectations and market surprises

Market perception of API information

The left panel of Figure 3 displays the estimated value of ω, the weight that the market places on API reports for the first month futures contract. We can see that the weight increases from 0.3 in 2010 to 0.85 in 2016, but falls below 0.7 again in 2018.

To analyze the actual accuracy of the two signals over time, the right panel of Figure 3 shows yearly pairwise correlations of announced and predicted values. The precision of the Bloomberg signal fluctuates quite a lot over time, but no apparent trend is visible. In contrast, the precision of the API signal seems to have been improving over time; the correlation increases from 0.6 in 2010 to over 0.8 in 2016. However, we also can see a sizable decline in accuracy in 2018.

To provide a more formal argument and facilitate the comparison, we calculate the weight to

Table 1: Distribution of realized market surprises, x.

Number of surprises				Magnitude of large surprises (absolute values)							
				Conditional on I^+				Conditional on I^-			
year	I^+	I^0	I^-	mean	$q1$	median	$q3$	mean	$q1$	median	$q3$
2010	12	34	6	0.010	0.007	0.008	0.011	0.012	0.010	0.011	0.013
2011	11	28	13	0.009	0.008	0.009	0.011	0.012	0.007	0.013	0.016
2012	12	33	7	0.011	0.009	0.011	0.013	0.009	0.009	0.010	0.010
2013	7	35	10	0.009	0.007	0.007	0.010	0.010	0.007	0.009	0.012
2014	9	36	8	0.010	0.008	0.009	0.011	0.007	0.006	0.007	0.008
2015	8	37	7	0.010	0.007	0.008	0.012	0.007	0.007	0.007	0.008
2016	11	35	6	0.009	0.007	0.007	0.009	0.009	0.007	0.009	0.011
2017	5	40	7	0.010	0.007	0.008	0.013	0.009	0.007	0.008	0.011
2018	17	27	8	0.013	0.008	0.011	0.017	0.010	0.007	0.008	0.012
2019	15	26	11	0.010	0.008	0.009	0.012	0.011	0.008	0.010	0.014
Total	107	331	83	0.010	0.007	0.009	0.012	0.010	0.007	0.008	0.012

Notes: The table provides descriptive statistics on the identified market surprise. $I^+(I^-)$ corresponds to large positive (negative) inventory surprises $x > \bar{x}$ ($< -\bar{x}$), with $\bar{x} = 0.006$, all other smaller surprises both negative and positive are denoted by I^0. The inventory surprise, x, is defined as the difference between the realized and the expected value of a change in oil inventories (normalized by the total oil inventories for the prior week). We take absolute values of negative surprises to facilitate exposition.

be placed of API reports that would consistent with the best linear projection or the best linear forecast in year t. We plot it on the same graph with the estimated weight. Although both lines display a clear upward trend, there are also noticeable differences. The market used to significantly underweight API signals in the early years. One reason for this can be the relatively low accuracy of API signals in early years. As the API subscription is costly, the benefits of additional noisy signal may had not outweigh the cost. However, as the accuracy improved, potentially larger fraction of people became willing to pay for API reports.[42] The subsequent decrease in 2018 is roughly consistent with the decreased accuracy of API signals. It should also be noted that the estimated weight was changing over time gradually, which may be consistent with the behavior of market participants who only infrequently reevaluate the precision of available signals.

[42]If markets are efficient, API information has to be fully revealed in prices. In reality, information percolation may be slow, especially given that API reports come long after main trading hours end. The only problem with this explanation is the fact that the highlights of the API report, at least in terms of total crude oil stocks, tend to be published by major independent news providers, and thus can be freely accessed 5 minutes after the release of the API report. An alternative explanation may be related to heterogeneity in sophistication. Some traders, especially those who do not specialize in oil trading, may be unaware of the existence of this additional private source of information.

The results of this subsection have two important implications. First, it is crucial to account for API information when modeling market expectations. The weight place on API is large, especially in recent years. Second, it is also important to account for its time-variability. Sharpening of identification of market surprises is one of the methodological contributions of our paper.

Distribution of surprises

Panel A in Table 1 reports the number of market surprises of each type realized in each calendar year. We do not see any significant changes over time; the number of surprises of both types stays roughly constant, about 20 per year. Panel B provides further information about the magnitude of realized surprises by depicting the mean, median, and two additional quantiles for each type. We do not see that one type of surprises (positive or negative) is systematically larger than the other. We will rely on this result when we analyze asymmetry of market responses. Our results also do not indicate any particular trend in the magnitude of surprises over time, which allows us to compare market responses over time. Overall, we do not find any evidence of systematic mistakes made by the market participants.

4.2 Market response to inventory shocks

Before we proceed to our main results, let us formally establish the importance of the EIA announcements. For each year in our sample we perform a Wald test for the joint significance of the coefficients associated with the arrival of EIA reports, which corresponds to the 15 exclusion restrictions in total (5 coefficients, $\{c_r^+, c_r^-, c_r^0, c_\sigma^0, c_\sigma^1\}$, for each of the first three contracts). The null hypothesis is rejected for all years with p-values less than 0.01%.

We first outline our results separately for the periods of unconstrained and constrained inventories, and then interpret our findings.

Period of unconstrained inventories

The 2010-2014 period and the period after 2017 are characterized by unconstrained inventories, when oil can be easily moved in or out of storage in response to a temporary shock. Our findings

Table 2: Returns reaction to inventory surprises.

year	c_r^+	c_r^0	c_r^-	$\overline{\sigma_{\tau,S}}$	P-value
2010	-0.10	-0.32	-0.09	0.016	0.95
	(0.07)	(0.05)	(0.06)		
2011	-0.07	-0.21	-0.10	0.018	0.73
	(0.11)	(0.04)	(0.03)		
2012	-0.18	-0.26	-0.10	0.013	0.32
	(0.05)	(0.04)	(0.08)		
2013	-0.10	-0.19	-0.11	0.011	0.91
	(0.10)	(0.03)	(0.06)		
2014	-0.11	-0.15	-0.22	0.014	0.59
	(0.05)	(0.04)	(0.19)		
2015	-0.04	-0.30	-0.89	0.026	<0.01
	(0.06)	(0.05)	(0.22)		
2016	-0.12	-0.57	-0.78	0.024	<0.01
	(0.11)	(0.04)	(0.15)		
2017	-0.17	-0.42	-0.08	0.015	0.73
	(0.11)	(0.08)	(0.24)		
2018	-0.12	-0.41	-0.14	0.016	0.71
	(0.03)	(0.03)	(0.03)		
2019	-0.27	-0.29	-0.16	0.017	0.16
	(0.10)	(0.07)	(0.05)		

Notes: The table displays the estimated values of $\{c_r^+, c_r^0, c_r^-\}$ for each year for the front month futures contract. The return response to x_τ, the inventory surprise in week τ, is given by the function $R(x_\tau) = \mathbb{I}_{\{-\bar{x}\le x_\tau\le\bar{x}\}} c_r^0 x_\tau + \mathbb{I}_{\{x_\tau>\bar{x}\}}\left(c_r^+(x_\tau-\bar{x})+c_r^0\bar{x}\right)+\mathbb{I}_{\{x_\tau<-\bar{x}\}}\left(c_r^-(x_\tau-(-\bar{x}))+c_r^0(-\bar{x})\right)$ where $\bar{x}=0.006$. The average effect on returns is given by $R(x_\tau)\overline{\sigma_{\tau,S}}$; $R(x_\tau)$ is also plotted on Figure 4. P-value is a probability value for the null $H_0 : c_r^+ = -c_r^-$ against $H_A : c_r^+ \neq -c_r^-$. Bollerslev–Wooldridge standard errors are in parenthesis.

for these periods can be summarized as follows:

Summary I. *The results for the unconstrained period:*

1. *Negative and significant relation between inventories and returns.*
2. *Market responses are symmetric.*
3. *No effects of inventory news on the term premium on the short end of the curve and only a weak effect for longer maturity contracts.*

1. Effect of inventory surprises on returns

Table 2 displays the estimated coefficients on inventory surprises in the return equation for the front month contract, $\{c_r^+, c_r^0, c_r^-\}$; Figure 4 plots the resulting response function, $R(x_\tau)$, for $x_\tau = \pm 0.01$.

Figure 4: Return reaction to inventory surprises for the first three futures contracts by maturity.

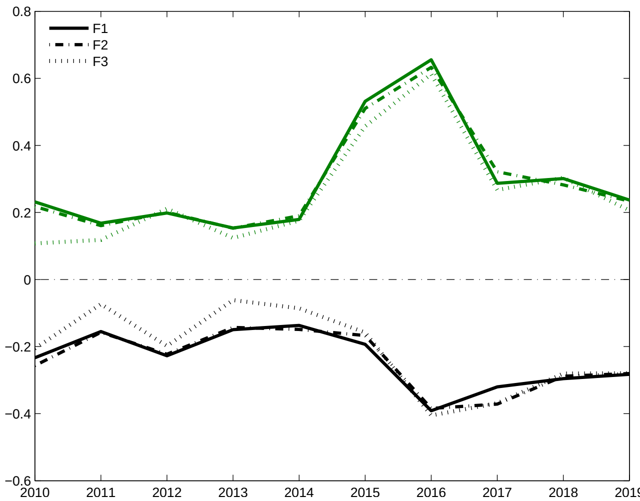

Notes: The solid line corresponds to the front month contract and depicts $R(x_\tau)$ for $x_\tau = 0.01$ in black (1% positive inventory surprise) and $x_\tau = -0.01$ in green (1% negative inventory surprise). The average effect on returns is given by $R(x_\tau)\overline{\sigma_{\tau,S}}$, where $\overline{\sigma_{\tau,S}}$ is given in Table 2. For example, an unexpected decline of inventories by 1% in 2010 would cause the oil price to immediately increase by $(-0.32 \cdot (-0.006) - 0.09 \cdot (-0.01 - (-0.006))) \cdot 0.016 \cdot 100 = 0.4\%$.

The results indicate a strong negative link between oil inventory surprises and returns during this period. The coefficients corresponding to large inventory surprises are all statistically significantly different from zero. As for the magnitude of the effects, we find that an unexpected decline of inventories by 1% in 2010 would cause the oil price to immediately increase by 0.4%.

The negative link between oil futures returns and inventory surprises is observed not just on average. If we consider large positive or negative announcements, the negative comovement of returns and surprises is observed in 137 weeks out of 146 (94%).

2. Asymmetry of market responses

Even though on average the market is equally likely to underestimate or overestimate oil inventories (as can be seen from Table 1), the market reaction to these surprises can be different. To investigate if that is indeed the case, the last column in Table 2 presents the p-values for testing the equality of coefficients on negative and positive surprises. For all years from 2010 to 2014 and from 2018 to 2019 we see that the null hypothesis cannot be rejected. Thus, our results indicate similar reaction to similar shocks at times of normal inventories.

3. Effect of inventory surprises of term premium

Our estimation approach differs slightly for short and long maturity contracts, thus we split our results into two parts.

Short maturities Figure 4 displays the estimated reaction of returns to 1% inventory surprises, both positive and negative, for the first three futures contracts by maturity, for each calendar year. We are interested in the part of the curve corresponding to the periods before 2014 and after 2018. The solid and dashed lines depict the estimates for the first two contracts. We can see that these lines are almost indistinguishable, implying that the prices of the first two contracts adjust by exactly the same amount in response to inventory news. The dotted line shows the estimates for the third month contract. It slightly diverges from the other two lines at times, but clearly follows the same adjustments. Hence, we can say that the term premium remains constant at the short end of the term structure curve.

Long maturities Next we analyze the adjustments of the contracts with longer maturities. Panel A in Table 3 compares the return reaction of the first month futures contract with reaction of the most liquid long maturity contracts; the results are presented for each quarter of each of the three unconstrained periods: 2010-2012, 2013-2014, and 2018-2019.

We observe a similar magnitude of price adjustments in response to inventory surprises. The prices of the front month and the current December contracts, in particular, adjust to news by the same amount. The difference in the estimated announcement returns for these contracts is only visible in the first quarters when the liquidity of the December contract is still relatively low. The reaction of the two longest maturity contracts, which are the next June and next December contracts, is also similar to the front month contract. Moreover, in some quarters the relationship is even reversed and we find that the long maturity contracts respond stronger to news than the front month contract, for example in the second quarter of 2012-2013 period.

Overall, the futures curve shifts in parallel in response to news, without any term structure adjustments. The reaction is especially uniform in the 2013-2014 period.

Table 3: Returns reaction to inventory surprises for longer maturity contracts.

Panel A: The periods of unconstrained inventories

	2010-2011-2012					2013-2014					2018-2019			
	F1	C Dec	N Jun	N Dec		F1	C Dec	N Jun	N Dec		F1	C Dec	N Jun	N Dec
Q1	−0.20	−0.11	–	–	Q1	−0.17	−0.14	–	−0.10	Q1	−0.23	−0.13	−0.08	−0.10
	(0.02)	(0.03)				(0.03)	(0.05)		(0.04)		(0.02)	(0.03)	(0.04)	(0.01)
Q2	−0.23	−0.18	–	–	Q2	−0.23	−0.21	−0.25	−0.26	Q2	−0.28	−0.25	−0.24	−0.19
	(0.04)	(0.02)				(0.03)	(0.02)	(0.05)	(0.06)		(0.02)	(0.03)	(0.04)	(0.03)
Q3	−0.21	−0.20	–	−0.17	Q3	−0.14	−0.13	−0.21	−0.16	Q3	−0.29	−0.29	−0.16	−0.21
	(0.02)	(0.01)		(0.03)		(0.03)	(0.03)	(0.02)	(0.03)		(0.03)	(0.04)	(0.05)	(0.06)
Q4	−0.16	–	–	−0.05	Q4	−0.12	–	−0.08	−0.09	Q4	−0.24	–	−0.18	−0.12
	(0.02)			(0.04)		(0.01)		(0.02)	(0.02)		(0.02)		(0.03)	(0.04)

Panel B: The period of constrained inventories

	2015-2016			
	F1	C Dec	N Jun	N Dec
Q1	−0.30	−0.15	−0.12	−0.08
	(0.06)	(0.07)	(0.01)	(0.05)
Q2	−0.46	−0.38	−0.29	−0.27
	(0.05)	(0.06)	(0.05)	(0.02)
Q3	−0.39	−0.31	−0.26	−0.20
	(0.10)	(0.05)	(0.10)	(0.06)
Q4	−0.46	–	−0.39	−0.34
	(0.06)		(0.08)	(0.06)

Notes: The tables show the estimated return response function, $R(x_\tau)$ evaluated at $x_\tau = 0.01$. In this exercise, we do not distinguish large or small surprises, and do not separate positive and negative surprises: $R(x_\tau) = c_r x_\tau$. To calculate the magnitude of the return reaction, one needs to multiply the value in the table by the average level of filtered realized volatility. 'C' stands for "Current", 'N' stands for "Next". Bollerslev–Wooldridge standard errors are in parentheses. In 'Q4' we do not consider the current year December contract due to its approaching expiration. In 2010-2012 we do not consider the next year June contract due to its insufficient liquidity. In the first half of 2010-2012 we do not consider the next year December contract due to its insufficient liquidity.

4. Effect on volatility

Finally, we briefly discuss volatility results (the tables are not included, but available on request).

We find that volatility jumps on announcement; c_σ^0 is positive and significant for all years and all contracts. The effect is quite large, on average, $|\ln\sigma^2|$ increases on announcement by about 20%. At the same time, we find that the announcement jump in volatility is not proportional to the size of the announcement surprise; c_σ^1 is insignificant for almost all years and all contracts.

Period of constrained inventories

The 2015-2016 period is characterized by constrained inventories, when positive supply or negative demand shocks cannot be easily smoothed out by putting oil in storage. Our findings for this period can be summarized as follows:

Summary II. *The results for the constrained period:*

1. *Negative and significant relation between inventories and returns. The market responds more strongly to inventory news when inventories are constrained.*
2. *Market responses are asymmetric, but the response is stronger to negative shocks.*
3. *No effects of inventory news on the term premium on the short end of the curve. However, the prices of longer maturity contracts are less reactive to news.*

1. Effect of inventory surprises on returns

We have already established a negative link between market surprises and returns reaction irrespective of the status of inventories, see Figure 4. It is worth emphasizing, however, that the negative comovement of inventory surprises and returns is observed in every single one of the 32 large inventory surprises.

Figure 4 also clearly shows intensification of the market reaction to inventory news in 2015 and 2016 relative to earlier and later years. In absolute terms, the difference in magnitudes is striking: the returns reaction to negative shocks is 3 times larger in 2016 than in 2014, despite that the average magnitude of the shocks remains roughly the same over time (see Table 1). Moreover, the normalization of returns reaction by volatility allows us to attribute this increase to the greater sensitivity to inventory information, and not just higher overall market sensitivity to news.[43]

2. Asymmetry of market responses

Figure 4 also reveals asymmetry. The market reacts stronger when a negative surprise is realized in 2015 and 2016. To formally test this for asymmetry, the last column in Table 2 presents the p-values for testing the equality of coefficients on negative and positive surprises. The null hypothesis of symmetry is rejected at the 0.01% level in both 2015 and 2016.

3. Effect of inventory surprises of term premium

Short maturities Now we are interested in the part of the curve on Figure 4 that corresponds to the period from 2015 to 2016. Our results indicate absolutely no difference in the reaction of the first three futures contracts. The formal tests cannot reject the null hypothesis.

Long maturities Panel B in Table 3 displays estimation results for longer maturity contracts.[44] We can see that the period of constrained inventories is characterized by muted reaction of the long maturity contracts relative to the front month contract. Moreover, the results indicate a robust decline in the magnitude of reaction with maturity. It could be natural to argue that muted reaction is observed because the long maturity contracts become unresponsive to news when inventories become constrained. We would like to emphasize that this is not the case. Quite the opposite, in absolute terms, the reaction of all long maturity contrast actually becomes stronger relative to unconstrained years (which can be seen by comparing the results in Panels A vs B for each contract and each quarter).

[43]To illustrate that fluctuations in market sensitivity to news may be substantial, Figure A.1 in Appendix A.4 shows realized variance estimated using high-frequency returns on the front month futures contracts. We can see a distinct change in the volatility regime at the end of 2014, when volatility increased dramatically and remained high for a long period of time.

[44]The results remain the same when we consider negative and positive surprises separately.

4. Effect on volatility

The volatility results are similar relative to the period of unconstrained inventories. We find that volatility jumps on announcement; c_σ^0 is positive and significant for all years and all contracts; however the size of the surprise has no significant effect on the volatility response.

We also find intensification of volatility reaction when inventories become constrained. The volatility responses becomes about 20% larger in 2015 relative to 2014, and decreases back to original levels in 2018.

Interpretation of results

So what do our findings tell us about beliefs of financial traders?

Identifying the effective date of the shocks (news shocks vs realized shocks) Our results indicate a strong negative link between inventory surprises and returns. This result is robust and confirms previous findings of Bu (2014), Halova Wolfe & Rosenman (2014), Halova et al. (2014), and Miao et al. (2018). This negative relationship is observed not only on average; we document negative comovement in 94% of weeks with large positive or negative surprises. The negative sign of the relationship implies that the market views inventory changes as reflecting *already realized demand or supply shocks*, rather than shocks to expectations of future oil market conditions. Not distinguishing between realized shocks and news shocks does not necessarily reflect irrationality. As shown recently by Mackowiak et al. (2018), rationally inattentive agents may optimally decide to economize on information processing costs needed to distinguish current changes in fundamentals from future changes in fundamentals.

The negative sign of the relationship could also be explained by the absence of insignificance of news shocks. However, there is substantial empirical evidence against this assumption. Arezki et al. (2017) use giant oil discoveries as directly observed news shocks and find a significant anticipation effect on the current accounts of small open economies through saving and investment channels. Kilian & Murphy (2014) show that shifts in expectations play a sizable role in the monthly fluctuations of the real price of oil. The news of coronavirus outbreak outside mainland

China in 2020 decreased the price of oil to below 30 $/bbl as the market started to expect a massive collapse in demand. Interestingly, in a different setting, Crouzet & Oh (2016) do not find a negative comovement of sales and finished-goods inventories, and interpret this as evidence of the insignificance of news shocks in business cycle fluctuations. However, we suggest an alternative explanation for this result.

Identifying the state of inventories and the persistence of shocks Our second main finding is a lack of any effect of inventory surprises on the short end of the curve. When inventories are unconstrained, this result is consistent with the behavior of the real agents who smooth out all temporary shocks by moving oil into or out of storage. In addition to a weak effect on the term premium on longer maturity contracts (or no effect at all in some years), our results are consistent with *uniform revision of expectations of current and future oil prices*. Our results can be viewed as yet another evidence that in the presence of inventories, all oil price movements become permanent and unpredictable as conjectured by Hamilton (2009). The symmetry in market responses at times of unconstrained inventories is also consistent with theoretical predictions.

However, when inventories are constrained, we would expect to see term structure adjustments, especially at the shorter end of the futures curve. As our model illustrates, the near-term prices should respond stronger to shocks, thus diverging from the more distant end of the curve. The lack of such adjustments is a striking result and contradicts conventional wisdom. The steep term structure curve is often observed at times of high inventories; see Figures 2 and 5. The slope is especially pronounced at the shorter end of the curve for maturities below four months. In particular, we can see a huge buildup of inventories in 2015. In 2008/09, the increase in inventories was less dramatic, but still quite sizable. In both cases, a steep term structure curve was attributed to the realization of a large temporary shock. In 2008, a negative demand shock was believed to have created an abundance of oil and depressed spot oil prices. In 2015, the market was believed to have hit by a positive supply shock due to rising shale oil production in the U.S. and by a negative demand shock due to slowdown of the Chinese economy. The association between steep term structure curve and high inventories turns out to be a more general phenomenon, not limited to the oil market as documented by Gorton et al. (2012).

The reaction of longer maturity contracts may seem to be more consistent with the theory. Indeed, we do observe muted reaction of longer maturity contracts relative to the front month contract in the period of constrained inventories. However, to be completely consistent with conventional theory, this muted reaction should be the result of the *unresponsiveness* of longer maturity contracts to news when inventories become constrained. Whereas in reality, the reaction of all long maturity contrast actually intensifies when inventories become constrained. This is puzzling, because if inventories can no longer smooth out temporary shocks, the longer end of the curve should become less sensitive to news, because only truly persistent shocks should be able to affect it, but that is not what we observe.

The lack of term structure adjustments may reflect an unawareness of the financial sectors of the constrained state of inventories. However, we believe that this is unlikely to be the case. The oil market reaction to inventory surprises abruptly intensified in February 2015 (we identify the exact moment of transition in the next section), exactly when inventories spiked. Moreover, the media had been following the evolution of inventories quite extensively.

We argue that the lack of term premium adjustments is observed because financial agents do not acquire enough information to precisely identify the persistence of shocks. Intuitively, agents are not accustomed to doing this. During normal times, inventories help to smooth out all temporary shocks (basically the real agents transform all shocks into permanent ones). No matter what kind of shock triggers the change in inventories, no term structure adjustments are needed, and the futures curve simply moves up or down as we see in 2013-2014, for example. As a result, traders may rationally develop a habit of disregarding the persistence, and trading as if all shocks were permanent. When suddenly inventories become constrained, traders fail to adapt. Such history dependence is a known feature of dynamic information acquisition, discussed by Mackowiak et al. (2018). In a different setting, Mackowiak & Wiederholt (2018) show that when a rare event occurs, rationally unprepared agents tend to take suboptimal actions, as they prefer to prepare for contingencies that are more likely to occur. Therefore, when inventories suddenly become constrained, traders may not be prepared to distinguish the shocks by their persistence.

Finally, our results also show negative asymmetry when inventories are high. The market

Figure 5: The slope of the term structure at the shorter end.

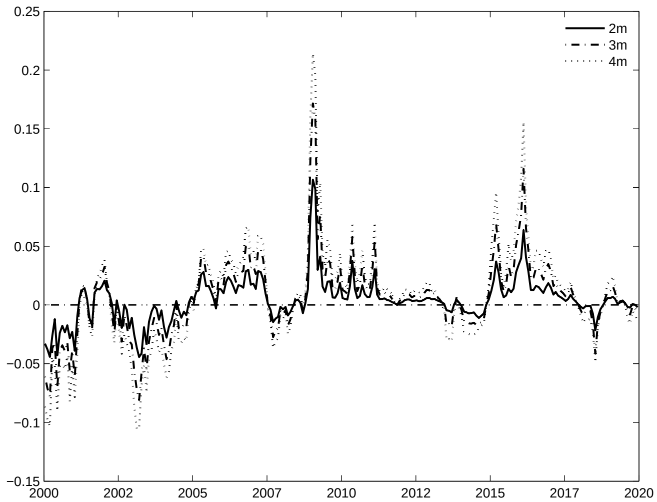

Notes: The figure displays the monthly series of the price difference between the three futures contracts with the shortest maturities and the front month futures contract, normalized by the price of the front month contract, $(F_{n,t} - F_{1,t})/F_{1,t}$, where $n = 2, 3, 4$ months.

responds more strongly to negative inventory surprises, though the theory predicts the opposite. This result may suggest that financial traders do not fully comprehend what the state of constrained inventories implies for equilibrium prices.

Trading activity on announcements Our results indicate that the actual size of an inventory surprise is irrelevant for the volatility response to news. Volatility spikes even following uninformative announcements.[45] The finding that trading activity intensifies even following uninformative releases is not new and has been documented numerous times in the earlier literature. Conventional interpretation of volatility and/or volume spikes on announcements is that both reflect idiosyncratic information processing and convergence of beliefs through trading. In a pioneer paper, Kandel & Pearson (1995) develop a theoretical difference-in-opinion model. In the model, the agents disagree about the interpretation of the public signal at the time of announcement, and thus trade occurs until all individual posterior beliefs converge. Recently the model has been extended and tested using high frequency data by Bollerslev et al. (2018). An alternative explanation is proposed by Crego (2020), who argues that a public signal may endogenously alter the composition

[45]The earlier version of the model also included a dynamic model of trading volumes. Similarly to volatility, trading volume also increases even following uninformative announcements.

of traders in the market. To provide empirical evidence, Crego (2020) studies the effect of EIA oil inventory announcements on the stock market, and shows a lack of the effect of unexpected component of news on volume and spreads of oil-related firms relative to less dependent on oil firms (even though volumes and spreads spike at the moment of announcement) consistent with the proposed model, and in line with our results.

Overall, our results may suggest the presence of significant disagreement among market participants about the meaning of inventory reports. The intensification of volatility response in 2015-2017 period is also consistent with the disagreement interpretation, as the oil market uncertainty clearly spiked during this time period.

Alternative identifying assumptions and discussion The lack of term structure adjustments at the shorter end of the curve is puzzling. It should be noted, that allowing for a *time varying risk premium* does not resolve this puzzle. If we assume that the expectations of the more distant future oil prices are not revised (or revised significantly less with maturity), than the change in the risk premium must account for the documented change in prices. Hence, we necessarily have to assume that the risk premium *decreases* when the price of oil falls (and the more so the larger the maturity); however, this assumption is unlikely to be true.

The second concern might be that in our exercise we do not distinguish supply shocks from the demand shocks. However, Kilian (2009) shows that the distinction matters: both the real price of oil and the macroeconomy react differently to these shocks. However, our approach makes this labeling less important. Importantly, we do not assume that all the shocks are the same; quite the opposite, the oil supply and demand shocks are likely to differ by their size and persistence, and thus should trigger different responses of inventories on average. However, what we study is the market reaction to inventory changes, the elasticity, and thus the source of the shock becomes irrelevant.

Another potential criticism of our approach is that the narrow window around the EIA announcement cannot capture the full market reaction to inventory news. Perhaps it takes longer for the market to fully process information and react. To provide further support for our findings, we perform additional estimation using daily data on futures returns and following a standard approach, outlined, for example, in Miao et al. (2018). The results are similar and confirm our main findings (see Appendix A.6).

4.3 Transition from unconstrained to constrained inventories and back

In our analysis so far we have split the sample at the end of 2014 to distinguish the periods of constrained and unconstrained inventories. We also have considered 2017 as a transition year from high inventories back to normal. In this last exercise, we aim to identify the transition moments more precisely.

To model the evolution of the parameters, we use a threshold autoregression (TAR) and smooth transition autoregression (STAR), with time as the threshold/transition variable. We carefully control for 'background' changes in parameters associated with the proliferation of new trading strategies. Full description of our approach and results can be found in Appendix A.5, here we only briefly comment on the main findings.

Our results suggest that despite the dramatic fall of oil prices and steepening of the term structure curve since November 2014, the market responded to inventory shocks in 2014 in *exactly the same* way as before. The break in the market response to inventory news only occurred in the last week of February 2015, precisely when the term structure curve spiked (see Figure A.2), potentially indicating that inventories reached a certain critical level.

Our transition results further reinforce the term structure puzzle. The spike in the term structure curve coincides with an abrupt intensification of the market response to inventory news. However, when oil inventory announcements come, traders do not revise expectations accordingly and do not adjust the term premium.

In contrast, the transition back to normal regime occurred gradually from March to September of 2017 and was consistent with the dynamics of the total oil inventories and the term spreads. The gradual transition may reflect the heterogeneity of traders' beliefs and overall uncertainty regardless the state of inventories. As more traders update their beliefs and place higher probability on capacity no longer being maxed out, we observe more and more muted reaction to news.

5 Conclusion

We provide a framework that integrates inventories data and ultra high frequency futures prices data into an analysis of the formation of beliefs of financial traders. Our results imply that financial traders fail to acquire additional information and treat inventories as a sufficient statistic. As a result, they fail to distinguish already realized shocks and news shocks, and they also treat all shocks as persistent.

The most surprising result is the lack of any effect of inventory news on the term premium at the shorter end of the futures curve when inventories are constrained. This result contradicts conventional wisdom, as 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. However, surprisingly, we see that when inventory news announcements come, traders do not revise expectations accordingly.

The evolution of coronavirus fears outside mainland China provides an excellent example of a news shock. Consider a short period of time from the end of February until March 13, right before the national state of emergency was finally declared in the U.S. Although the virus had already affected China, the worst was yet to come for the U.S. and most European countries. Early in March general hopes to contain the virus outside the US were gradually replaced with expectations of soon-to-be-imposed harsh social distancing measures.[46] Even before the U.S. government imposed a lockdown, the public voluntary decided to follow social distancing measures. Indeed, a voluntary drop in mobility as measured by the GPS locations of US cellphones was documented by Simonov et al. (2020); similarly, an early drop in restaurants reservations was documented by an online restaurant-reservation service company OpenTable. Thus, although some changes in economic activity already started to occur, they were relatively mild compared to what were to happen next. Not surprisingly, the short term forecasts began to deteriorate quickly as the governments were expected to issue harsh measures to curtail the spread of the virus. Therefore, coronavirus can be seen as a shock with a sizable news component.

The response of the oil market was largely consistent with the narrative above. Since mid-

[46]Despite efforts of some controversial news outlets such as Fox News to convince the public otherwise; see Simonov et al. (2020).

January the oil price had been falling reflecting drop in demand from China. However, it was not until March 9 when the price of oil plunged by historical 25% in a single day, potentially reflecting a sudden reassessment of the severity of pandemic coming to the U.S, and partially the announcement by Saudi Arabia on March 6 to keep the production unchanged. Somewhat puzzling, however, was again the behavior of the oil futures term structure curve. Given pessimistic projections of economic activity, one could expect to see backwardation in March, as the worst was clearly yet to come. The economic activity was expected to plummet in the next few months. In reality, however, the oil market moved into a deep contango (see Figure A.4)[47].

Analyzing the oil market response to coronavirus news is a fruitful topic for future research. Our findings suggest that one could need to consider alternative drivers of the term structure curve.[48]

References

Alquist, R. & Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4), 539–573.

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

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

Arezki, R., Ramey, V. A., & Sheng, L. (2017). News shocks in open economies: Evidence from giant oil discoveries. *Quarterly Journal of Economics*, 132(1), 103–155.

[47] It should be noted that the hectic fall of the front month futures price right before expiration in April 2020 have a completely different nature. We are more interested of what moved the market in contango already in March.

[48] Alternative theory of the steep term structure curve was developed by Mou (2010), Hamilton & Wu (2014) and Selezneva (2015), and emphasizes the role of the dramatic inflow of investment in ETFs tracking oil prices. When such funds become large enough relative to the open position of the rest of the market, a significant effect of their operations on the term premium can be expected.

In two months since the end of January 2020 the United States Oil Fund, which is the largest oil ETF, grew from \$1.2 bln to \$3.2 bln in assets under management. As of April 9, 2020, the fund holds 67,844 of May20 contracts and 56,420 of Jun20 contracts which represent 17% and 14% of the open interest in these contracts, respectively.

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

Bianchi, F., Kung, H., & Kind, T. (2019). Threats to central bank independence: High-frequency identification with twitter. NBER Working Paper No. w26308.

Bollerslev, T., Li, J., & Xue, Y. (2018). Volume, volatility and public news announcements. *Review of Economic Studies*, 85(4), 2005–2041.

Bollerslev, T. & Wooldridge, J. M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. *Econometric Reviews*, 11, 143–72.

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

Coibion, O. & Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1), 116–159.

Crego, J. A. (2020). Why does public news augment information asymmetries? *Journal of Financial Economics*, 137(1), 72–89.

Crouzet, N. & Oh, H. (2016). What do inventories tell us about news-driven business cycles? *Journal of Monetary Economics*, 79, 49–66.

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

Ederington, L. H., Fernando, C. S., Holland, K. V., Lee, T. K., & Linn, S. C. (2020). Dynamics of arbitrage. *Journal of Financial and Quantitative Analysis*. forthcoming.

Fama, E. F. & French, K. R. (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1), 55–73.

Froot, K. A. (1989). New hope for the expectations hypothesis of the term structure of interest rates. *Journal of Finance*, 44(2), 283–305.

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

Goldstein, I. & Yang, L. (2019). Commodity financialization and information transmission. Working Paper.

Gorton, G. B., Hayashi, F., & Rouwenhorst, G. (2012). The fundamentals of commodity futures returns. *Review of Finance*, 17(1), 35–105.

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

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

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

Hamilton, J. D. (2009). Understanding crude oil prices. *Energy Journal*, 30(2), 179–206.

Hamilton, J. D. & Wu, J. C. (2014). Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42, 9–37.

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

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

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3), 1053–69.

Kilian, L. & Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29, 454–478.

Li, J. (2019). Rational inattention and price underreaction. Manuscript, Stanford GSB.

Mackowiak, B., Matejka, F., & Wiederholt, M. (2018). Dynamic rational inattention: Analytical results. *Journal of Economic Theory*, 176, 650–692.

Mackowiak, B. & Wiederholt, M. (2018). Lack of preparation for rare events. *Journal of Monetary Economics*, 100, 35–47.

Mankiw, N. G., Reis, R., & Wolfers, J. (2013). Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 18, 209–248.

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

Mou, Y. (2010). Limits to arbitrage and commodity index investment: Front-running the goldman roll. Working Paper.

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

Ng, V. K. & Pirrong, S. C. (1994). Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. *Journal of Business*, 67(2), 203–230.

Nielsen, M. J. & Schwartz, E. S. (2004). Theory of storage and the pricing of commodity claims. *Review of Derivatives Research*, 7(1), 5–24.

Pindyck, R. S. (1990). Inventories and the short-run dynamics of commodity prices. NBER Working Paper 3295.

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

Roesch, A. & Schmidbauer, H. (2011). Crude oil spot prices and the market's perception of inventory news. Conference: International Symposium on Forecasting (ISF), Prague.

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

Simonov, A., Sacher, S., Dube, J., & Biswas, S. (2020). The persuasive effect of fox news: non-compliance with social distancing during the covid-19 pandemic. NBER Working Paper No. w27237.

Wen, Y. (2005). Understanding the inventory cycle. *Journal of Monetary Economics*, 52(8), 1533–1555.

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