# When Benchmarks Fail:

# The Causes and Consequences of Negative Oil Prices

Erik P. Gilje<sup>\*</sup> Robert Ready<sup>†</sup> Nick Roussanov<sup>‡</sup> Jérôme P. Taillard<sup>§</sup>

December 15, 2023

#### Abstract

On April 20, 2020 the flagship North American benchmark price for crude oil, West Texas Intermediate (WTI) for delivery in Cushing, OK, fell below \$0 for the first time in its history. We explore the causes and consequences of this event. Using a simple theoretical model, we argue that negative prices were the result of constraints on physical storage, but that these constraints were precipitated by positions of financial traders unable to take physical delivery prior to the expiration of the May 2020 futures contract. We then provide evidence that this event impacted the decisions of oil producers, including those not directly impacted by the storage constraints at Cushing, through their exposure to the WTI benchmark in their purchasing contracts. Our evidence as contracting benchmarks.

<sup>\*</sup>Teocalli Exploration and NBER, 1000 W. Wilshire Blvd Ste 400, Nichols Hills, OK 73116, USA

 $<sup>^{\</sup>dagger}$ Lundquist College of Business, 382 Lillis Hall, 1208 University of Oregon, Eugene, OR 97403. Email: rready@u<br/>oregon.edu

<sup>&</sup>lt;sup>‡</sup>The Wharton School, University of Pennsylvania, 3620 Locust Walk Suite, 2400, Philadelphia, PA 19104. Email: nroussan@wharton.upenn.edu

<sup>&</sup>lt;sup>§</sup>Babson College, 323 Tomasso Hall, Babson Park, MA 02457. Email: jtaillard1@babson.edu

We would like to thank our discussants James Weston, Praveen Kumar, and Dmitry Livdan, as well as seminar and conference participants at the University of Alberta, University of Oregon, MIT, Carnegie Mellon University, The Ohio State University, Babson College, HEC Paris, Georgia Tech, Nazarbayev University, The University of Texas-Austin, Wharton, Cornell University, the Virtual Finance Seminar, the 2021 SFS Cavalcade, the 2022 AFA meetings, the 2022 University of Oklahoma Energy and Climate Finance Research Conference, the 2023 Pacific Northwest Finance Conference, and the 5th Annual JPMCC International Commodities Symposium for helpful comments. This paper subsumes results from a previous draft titled "The day that wti died: Asset prices and firm production decisions".

## 1 Introduction

On April 20, 2020, one day prior to the expiry of the futures contract for May 2020 delivery, the flagship North American benchmark for crude oil, West Texas Intermediate (WTI) for delivery in Cushing, Oklahoma, turned negative, settling at -\$37/bbl. We explore the causes and consequences of this event. In the first portion of our paper, we present a new empirical fact: we show that open interest in the soon-to-expire futures contract provides a strong signal of demand for physical oil at the hub in Cushing. We then incorporate this fact into a novel theoretical model and use it to argue that 1) a negative price event is the result of a real physical constraint (lack of storage), but that this constraint can be at least partly caused by non-physical financial traders who obscure the demand signal, and 2) this event is likely to occur near the expiry of a futures contract, when financial traders are forced to close their positions without taking physical delivery and the true demand is revealed.

In the second portion of the paper we trace out the impact of these negative prices on oil production. Since most crude purchase contracts in the United States are indexed to WTI, millions of barrels of crude in America were effectively sold for -\$37 on April 20, 2020. We present evidence that the fear of a repeat event near the expiry of the June 2020 contract led many producers, even those in areas not directly impacted by the storage capacity constraints in Cushing, to temporarily shut in production in May despite prices having already recovered to their early April levels. Our results provide novel evidence on a previously under-appreciated dimension through which asset prices and the activity of financial market participants can affect real decisions of firms: the benchmark channel.

The NYMEX WTI contract traded on the Chicago Mercantile Exchange is the most liquid and actively traded crude oil contract in the world. This contract specifies a certain crude oil quality, to be delivered at a specific location (Cushing, Oklahoma) at a given point in time. Both financial traders and physical traders can engage in trading on the exchange. On April 20, 2020, the NYMEX WTI Light Sweet Crude Oil May 2020 (CLK2020) futures contract was the front month, and any traders that had not exited their positions by settlement on April 21, 2020 (the next day) would need to either deliver or receive physical crude to settle their



Figure 1: Intraday prices in May 2020 future prior to expiry

remaining open futures positions.<sup>1</sup> As sellers began to outnumber buyers (in dollar terms) on the evening of Sunday, April 19, 2020, the price of crude fell. The selling accelerated the next morning, culminating in prices going negative for the first time around 1PM on April 20, 2020 (see Figure 1). Prices remained negative, so that the daily contract settled at 2:30PM at a price of -\$37/bbl. This settlement price, regardless of the trading volume at the time, is what is used for index and benchmark pricing across the U.S. crude oil complex. Although prices remain below -\$10/bbl for only 5 hours on April 20th, 80 crude grades at locations across the United States transacted at an average of -\$44/bbl for the day, even at locations that had ample storage capacity as well as waterborne flexible storage capacity.

The fact that this event occurred near the expiry of a financial futures contract suggests a link between the financial and physical markets at Cushing. In the first portion of the paper we lay out a hypothesis for how an expiring futures contract can precipitate a negative price event, and present a simple model to formalize the intuition. The primary mechanism in our model relies on the newly-documented fact that high open interest in an expiring futures contract typically predicts more plentiful storage at Cushing shortly after expiry. There is

<sup>&</sup>lt;sup>1</sup>Delivery at Cushing requires either a volume allocation on a pipeline in or out of Cushing or a physical terminal at Cushing, there is no additional storage capacity or pickup available via waterborne delivery or truck, and as such no free disposal. Therefore Cushing is a closed system with a fixed amount of storage and pipeline capacity. While storage levels rose sharply around the negative price event, they peaked at roughly 85% of the physical capacity. However, most of the remaining available storage was likely "committed" or pre-sold so it was not available to take physical delivery (see "Today in Energy", EIA, April 27th, 2020 (https://www.eia.gov/todayinenergy/detail.php?id=43495).

therefore a signal embedded in futures market open interest. If open positions in an expiring future are typically an indicator of demand from physical traders who are willing to take delivery, then high open interest is a signal for midstream pipeline operators to deliver more oil to the hub (delivery point), as there is more demand and hence more available storage capacity. However, if these operators misinterpret the long positions of financial investors as positions of physical traders, then they will deliver excess oil to the hub and potentially outstrip the available storage capacity. When financial investors are forced to close out their positions without taking delivery prior to expiry, the excess supply is revealed and prices crash.

This mechanism means that negative prices can arise when an unusually high amount of financial open interest coincides with an unusually low amount of demand from physical traders. Our hypothesis regarding April 20th is that this financial open interest was the result of retail investors taking long positions in the expiring contract in unusually large numbers.<sup>2</sup> This is supported by anecdotal evidence from news reports suggesting that retail traders directly trading futures on U.S. platforms (e.g. Interactive Brokers, E-Trade, and TD Ameritrade) were among those who lost money on April 20th, and that retail traders with long futures positions through the Bank of China's (BOC) "Crude Oil Treasure" product also incurred substantial losses.<sup>3</sup> To further support this hypothesis, we present additional evidence using online searches in China for the Crude Oil Treasure product. We find that online search volume for this product rose suddenly in late March and early April suggesting increased interest, and possibly increased investment in this product. Back of the envelope computations from these reports are broadly consistent with the open interest spike in April of 2020. This increased demand for futures from retail traders could have then masked the extent of the declining refinery demand due to COVID-19. When the terms of the Crude Oil Treasure product forced its investors to close their positions in the May 2020 contract on

<sup>&</sup>lt;sup>2</sup>Ozik et al. (2021) document similar increases of retail investor holdings in equities.

<sup>&</sup>lt;sup>3</sup>See Traders New "Dav are а Wrinkle inthe Negative Oil Price Mystery" https://www.bloomberg.com/news/articles/2020-06-08/are-day-traders-a-possible-cause-for-oil-pricesgoing-negative, as well as the references found in Section 2.1 regarding the Bank of China's "Crude Oil Treasure" financial product.

April 20th, it is possible that this sudden drop in open interest revealed the lack of physical demand and caused prices to crash, consistent with the mechanism of the model.

The fact that the BOC allowed their retail investors to hold the front month futures contract until the day prior to the final settlement date (futures expiry) is highly unusual for a financial product. Most financial traders (for instance the United States Oil Fund, the largest Oil ETF) close out their positions in the "front" month contract between 15 and five days prior to expiry, when they "roll" out of the expiring contract and into the next closest contract to expire, which becomes the "active" month. Therefore, open interest in the expiring contract is a small portion of the overall market, meaning that the impact of the event for most traders was likely negligible as they had already closed their positions in the May contract. However, the closing price of the front month contract is still used as a benchmark by many physical buyers and sellers of oil through crude sales agreements benchmarked to the price of oil in Cushing (WTI prices). Many of these contracts are for oil produced geographically far from Cushing or are for different grades of crude, but these contracts seek to benefit from the price discovery that occurs in the much more liquid WTI futures market for delivery at Cushing, the world's largest crude oil storage hub. However, this convention can also set the stage for crude producers being tethered to pricing that could become disconnected from fundamentals due to potential dislocations in the front month WTI prices as a consequence of physical constraints, such as deliverability/storage, and/or actions from financial traders.

This setting therefore provides a unique opportunity to explore how asset price dislocations caused by uninformed financial traders can affect real decisions via the benchmark channel. Specifically, while crude oil fundamentals were undoubtedly challenging due to the COVID-19 pandemic, we provide evidence that the negative price spike in WTI was consistent with short-term frictions, and in particular concerns regarding deliverability and storage issues at Cushing, and not reflective of broader fundamentals. Therefore, this context provides an attractive setting to estimate how the functioning of asset markets that firms reference for contractual arrangements may affect real production decisions.

Our analysis of production decisions focuses on producers' responses to the risk of a repeat

event in May of 2020. While the negative oil prices quickly resolved after April 20th, option prices suggest that there was substantial fear of a repeat event associated with the futures expiry the following month. Our challenge for this analysis is to disentangle the impact of broader risk from COVID-19 from the risk related to a repeat event in the futures market impacting producers through the benchmark. To do this, we pursue two distinct empirical strategies.

In our first set of tests, we use proxies for daily oil production and exploit the structure of purchasing agreements. If producers decide to shut in wells due to concerns about potential negative price spikes around the contract expiration of crude oil contracts in Cushing, there are predictions that can be put to the test using daily production data. One such prediction arises from fact that most of these agreements rely on calendar month average (CMA) pricing, where every barrel of oil sold in a calendar month receives a price equal to the average of the daily prices across the month. Therefore if, following the April 20th dislocation, producers anticipate the potential for the CMA price to be affected by another significant negative WTI price dislocation in May (specifically when the front month contract expires on May 19th), then one might expect that production would be shut in on May 1st, the first day of a producer's calendar month average (CMA) price for May, as this is the first day of production that is exposed to benchmark pricing risk from the May 19th expiry. Consistent with this hypothesis, we observe a four standard deviation reduction in our production proxies from April 30th to May 1st. One might also expect that, after the June 2020 contract expiration in May, producers may start re-opening wells if a negative price dislocation does not materialize. This is in fact what we observe. The day after contract expiration (May 20th), our production proxies begin to increase, and continue to increase through the end of the month. Put together, these pieces of evidence are consistent with the potential for negative price shocks and contract expiry risk at Cushing (i.e., benchmark risk) having a significant real effect on production decisions.

Our second set of results compare production decisions made for oil wells indexed to WTI to a control set of wells that are not benchmarked to WTI. We focus on two neighboring geographies, North Dakota and Alberta. Both of these areas feature similar geologic formations, and both are unlikely to be directly impacted by the storage constraints at Cushing. The difference is that the wells in North Dakota are primarily benchmarked to the WTI, while the Canadian wells in our sample are primarily benchmarked to Edmonton prices that did not turn negative.

While we do not have daily data for oil production in these areas, we do have monthly data including the number of days in the month a well is producing. Our main outcome measure is whether firms temporarily dial back production by shutting in production from wells in May of 2020. Namely, we search for wells that are producing in April, cut back production for a significant portion of May, and then return to full production shortly thereafter. We find that approximately 8% of the wells in North Dakota have this kind of temporary shut-in in May of 2020, while roughly 2% of wells in Alberta have a similar shut-in. Both the 8% temporary shut-in rate and the 6% difference between North Dakota and Alberta rates are extreme outliers in the 48 months where we have well-level production data for both regions. This fact is particularly striking in the context of prices, which had fully recovered to April 1st levels by May 1st, and continued to improve throughout the month of May so that the average May price for both Alberta and North Dakota were nearly double the April price. Outside of the risk associated with the futures expiry in May, it is not clear what risks would lead to the temporary shut-ins we see in the high frequency and monthly data. We therefore believe that theses two pieces of evidence demonstrate that this event had a substantial impact on production decisions through producers' exposure to the WTI benchmark. Interestingly, in our final tests we find that the shut-ins in North Dakota had little impact on the long-term productivity of the wells, suggesting that these shut-ins may have been an effective response by producers to increased benchmark risk associated with dislocation in the WTI.

Our results link to the literature in finance and economics that focuses on the interaction between financial markets and the real decisions of firms. This literature, going back at least to Hayek (1945) and more recently surveyed by Bond et al. (2012), argues that secondary market prices for assets are important sources of information. For example, equity prices can: (1) help managers decide when to invest and what to invest in (Chen et al. (2007), Foucault and Fresard (2012), Barro (1990)), (2) help managers decide whether to undertake merger decisions (Luo (2005) and Edmans et al. (2012)), (3) inform government regulatory interventions (Bond and Goldstein (2015)), and (4) help customers or employees determine which firm to work for. Due to the rich information content of asset prices provided by liquid financial markets, they are also used as price-setting benchmarks (e.g., Duffie and Stein (2015), and Duffie et al. (2017)). While benchmarking to market prices can alleviate a number of frictions and search costs, it also potentially exposes firms to additional risks if benchmark prices were to deviate from fundamental value. Yet, to date, substantially less attention has been given to analyzing the risks associated with such contractual arrangements, and the potential ramifications in terms of real effects when contract prices are affected by market dislocations and/or limits to arbitrage.<sup>4</sup>

We also contribute to the literature on the financialization of commodity markets. Commodity markets have become an important asset class over recent years. The extant empirical literature has generated mixed results as to the role for financialization on commodity price movements (e.g., Fama and French (1987), Masters (2008), Hong and Yogo (2012), Tang and Xiong (2012), Buyuksahin et al. (2013), Gorton et al. (2013), Buyuksahin and Robe (2014), Cheng and Xiong (2014), and Ready and Ready (2022)). The financialization of oil as an asset class and the events of April 20, 2020 provide an interesting framework to analyze how market price dislocations could have real effects on producers. A related strand of this literature provides theoretical and empirical evidence on how the composition of market participants, namely hedgers versus speculators, may be important for commodity markets (e.g., Hirshleifer (1988), Hirshleifer (1990), Faulkender (2005), Rouwenhorst and Tang (2012), Kang et al. (2020)) and in particular how dynamics may emerge in commodity markets, which drive prices away from fundamental value (Singleton (2012)). Our evidence suggests that if there are indeed price movements in benchmarks driven by financialization then these price movements can have significant real effects on firm production decisions.

<sup>&</sup>lt;sup>4</sup>One such risk is related to the manipulation of benchmark prices. The LIBOR scandal highlighted the strong incentives market participants may have in colluding to set benchmark prices (e.g., see Abrantes-Metz et al. (2012)). It also engendered interest in manipulation-proof benchmarks (e.g., Duffie and Dworczak (2021)).

# 2 A model of an expiring futures contract

In this section we empirically motivate and derive a simple model of an expiring oil futures contract, featuring interactions between financial and physical traders. The primary mechanism in the model involves open interest in the futures market serving as a signal for demand at the hub (delivery point). We therefore first present some findings on the empirical relation between open interest and storage to motivate the mechanism, as well as some evidence on the source of financial open interest in the expiring May contract.

### 2.1 Empirical motivation

As Figure 2 shows, the expiry of the May contract featured a very unusual combination of extremely high open interest in the soon-to-expire May 2020 contract along with a collapse in refinery utilization. While we do not have direct data, there is anecdotal evidence to suggest that this high open interest was due at least in part to long positions of retail traders, and in particular retail traders using the Crude Oil Treasure product offered by the BOC.<sup>5</sup> This product is unusual in that it requires retail investors to roll their position out of an expiring contract and into the next closest contract on their own accord. If they do not roll the contract they are automatically rolled over the day before the final contract settlement (contract expiry). This is in contrast to more typical financial investment products, such as the United States Oil Fund (USO), that roll well prior to expiry (approximately two weeks before expiry in the case of the USO, see Bessembinder et al. (2016)). Reports also suggest that retail investors in this product owed the BOC approximately \$1.4 Billion after the negative price event, which, given the -\$37/barrel price and 1,000 barrel contract size, corresponds to a position of approximately 38,000 contracts, broadly consistent with the open interest spike in April of 2020.

In our model it is *unexpectedly* high financial open interest that can cause a price crash

<sup>&</sup>lt;sup>5</sup>See "Explainer: How China's retail investor army got burned by the shock oil collapse", Emily Chow and Cheng Leng, Reuters, April 24, 2020. (https://www.reuters.com/article/us-global-oil-china-investors-explainer/explainer-how-chinas-retail-investor-army-got-burned-by-the-shock-oil-collapse-idUSKCN2261MH) and "China's 'Crude Oil Treasure' Promised Riches. Now Investors Owe the Bank," Alexandra Stevenson and Cao Li, https://www.nytimes.com/2020/05/21/business/china-oil-investors.html.



Figure 2: Open interest in expiring futures and refining utilization

The first panel of the figure plots open interest (number of contracts) in the expiring futures contract two trading days prior to its monthly expiry. Friday, April 17th 2020 was the trading day prior to April 20th, 2020, the day with negative prices and April 21st 2020 was the final settlement day (expiry) of the contract. The second panel plots weekly refinery inputs for PADD 2, which covers the Midwest, including Oklahoma. Refinery utilization for the week ending April 17th is shown, as reported by the EIA on April 22nd.

#### Panel A: Crude Oil Treasure advertisement



Panel B: Crude Oil Treasure search volumes 1/1/2020 - 4/17/2020







Figure 3: Interest in Crude Oil Treasure product

Panel A of the figure shows an advertisement for the Bank of China's crude oil treasure product in its original version as well as translated via Google translate. Panel B shows online search volume data for "Crude Oil Treasure" from QiHoo 360 for January 1st, 2020 to April 17th, 2020. The vertical line denotes the first day in which new investors would have been invested in the May 2020 delivery contract. Panel C repeats the search volume plot for January 1, 2020 through June 1, 2020. The spike in the plot in Panel C occurs on April 20th when prices go negative, the much larger scale of these searches relating to negative prices renders the earlier search volume from Panel B invisible.

11

near a futures expiry. We therefore explore whether the open interest in this product, which was first introduced in 2018, was unusually high in April of 2020. While we do not have direct data, Figure 3 presents two pieces of evidence to suggest this may have been the case. The first is an advertisement for the product from this period. Panel A shows the original version of this advertisement along with a translation created by Google Translate. The advertisement uses the slogan "Crude oil is cheaper than water", referring to the low prices in the latter half of March 2020 and first part of April 2020, when WTI was trading at approximately \$20 a barrel, down from nearly \$50 a barrel on March 1st. If retail investors misunderstood the structure of the product, they may have assumed that it would allow them to benefit from a subsequent rise in prices. Data from search volumes suggests this advertising may have been effective. Panel B of Figure 3 shows search volumes collected from the Chinese search engine QiHoo 360 from January 1st to April 17th of 2020. As the plot shows, interest in this product (as proxied by search volumes) rose drastically in the latter half of March and first half of April. It is notable that this period, from March 19th to April 17th, is precisely when new investors would have been entering into positions in the May 2020 contract that ended up going negative on April 20th. Panel C shows that the search volume prior to April 20th was dwarfed by the broader interest relating to the losses associated with the product during the negative price event.

The presence of an unusually high financial open interest late in the contract's life may have created an issue in the spot market if market participants viewed this open interest as a signal of demand from physical traders willing to take delivery. To provide evidence that futures open interest may serve as a signal for the spot market, Figure 4 plots total open interest in the nearest term future two days before its expiry against changes in storage levels at the Cushing hub reported in the next Weekly Petroleum Status Report released by the Energy Information Agency (EIA). The figure shows that in the three years prior to April 2020, the open interest prior to the contract expiry does provide a predictive signal for the physical demand at the hub. As the plot shows, in the 36 prior months, high open interest typically signaled decreased levels of storage, and thus more available storage capacity, in the near term. Again, the May 2020 contract is a notable outlier. Though there was an



Figure 4: Open interest prior to contract expiry and subsequent storage changes at Cushing: May 2017 to May 2020

Figure plots the open interest before the second to last trading day prior to expiry for a futures contract on the X-axis and the reported weekly change in storage at Cushing in the subsequent week on the Y-axis. The blue dots represent the 36 futures contracts with delivery from May 2017 to April of 2020, and the red square is the May 2020 future. The regression line is fitted to the data excluding the May 2020 future.

extremely high level of open interest, this coincided with a very large increase in storage, consistent with producers mistakenly inferring that the high open interest reflected a high number of physical traders able to take delivery.

Table 1 shows that this relation between open interest in the expiring contract and storage holds for longer periods dating back to the financial crisis.<sup>6</sup> Panel A shows regressions of subsequent storage changes at Cushing on open interest in the expiring contract, controlling for the lag change in the futures basis, which typically drives levels of storage via the cost-of-carry arbitrage relation. The four columns show the negative relation between open interest and changes in storage is significant in expanding windows back to the financial crisis. Panel B repeats these regressions using changes in storage outside of Cushing, and here we see no significant relation between futures open interest and storage. To our knowledge this relation between open interest and storage at Cushing is a new finding, and provides evidence that unusual levels of financial open interest could have provided a misleading signal to the physical

<sup>&</sup>lt;sup>6</sup>Storage data at Cushing are available back to 2004, and including this full period the relation is still negative but no longer statistically significant due to highly volatile observations during the financial crisis.

	Change in storage in next EIA report following expiry						
	Prior 3 years	Prior 5 years	Prior 10 years	Post Crisis			
	(1)	(2)	(3)	(4)			
Open Interest	$-0.056^{***}$ (0.020)	$-0.034^{**}$ (0.015)	$-0.027^{***}$ (0.009)	$-0.023^{**}$ (0.009)			
Basis Change	$126, 332.600^{***}$ (41, 928.700)	$55,995.500^{***}$ $(15,997.360)$	$41,712.630^{***}$ (10,511.200)	$15,691.250^{***}$ (5,226.696)			
Constant	$3,480.581^{**}$ (1,396.412)	$2,013.826^{**}$ (980.938)	$1,658.373^{***}$ (595.518)	$1,367.082^{**}$ (570.212)			
$\begin{array}{c} \text{Observations} \\ \text{R}^2 \end{array}$	36 0.371	$\begin{array}{c} 60\\ 0.244\end{array}$	$\begin{array}{c} 120 \\ 0.169 \end{array}$	$\begin{array}{c} 135\\ 0.098\end{array}$			

Panel A: Regressions of changes in storage at Cushing on open interest two days prior to expiry

Panel B: Regressions of changes in storage outside of Cushing on open interest two days prior to expiry

	Change in storage in next EIA report following expiry							
	Prior 3 years	Prior 5 years	Prior 10 years	Post Crisis				
	(1)	(2)	(3)	(4)				
Open Interest	0.018	0.024	-0.043	-0.055				
	(0.081)	(0.060)	(0.036)	(0.034)				
Basis Change	216,468.700	-40,023.370	12,796.870	14,810.160				
	(169, 210.400)	(65, 760.750)	(41, 452.780)	(19, 863.490)				
Constant	-2,579.498	-2,171.073	3,318.653	$4,100.639^{*}$				
	(5, 635.460)	(4, 032.364)	(2, 348.529)	(2, 167.028)				
Observations	36	60	120	135				
$\mathbb{R}^2$	0.047	0.010	0.012	0.022				
Note:		*p<0.1; **p	o<0.05; *** p<0.01					

Table 1: Regressions of storage changes on open interest prior to contract expiry

The table shows the results from regressions of storage changes (in 1000s of barrels) on the opening interest in the front month WTI future at the close of trading two days prior to contract expiry. The storage changes are reported weekly by the EIA, and we therefore use the storage change for the week that encompasses the day that is two days prior to contract expiry. The regressions also include the change in the futures basis from the previous month, where the basis is the log-difference between the next month contract and the expiring contract, again measured two days prior to expiry. Panel A shows this regression for storage at the hub in Cushing, Oklahoma, and Panel B shows the regression using all U.S. storage excluding Cushing and the Strategic Petroleum Reserve. Heteroskedasticity robust standard errors are in parentheses.

market for oil in Cushing.

To summarize, our hypothesis is therefore that an unexpectedly large amount of financial open interest near the expiry of the May 2020 contract, possibly caused by retail investors, may have led pipeline operators to infer that there was ample refinery capacity, and therefore to send too much oil to the storage hub. This was revealed when the contract neared expiry and the financial traders attempted to roll out of their positions without taking delivery, causing the prices to suddenly plunge. We next present a model to formalize this intuition.

### 2.2 Model

The model presented here is related to recent models of commodity financialization, most notably Sockin and Xiong (2015), Goldstein and Yang (2022), and Ge et al. (2022), but differs along several key dimensions. For instance, the model of Sockin and Xiong (2015) features similar uncertainty regarding the positions of financial traders that distort signals of aggregate demand. However, this distortion is not revealed due to continued supply shocks, and therefore the price effect persists in the spot market. While the model of Goldstein and Yang (2022), examines commodity markets with different levels of financial investment, this investment is treated as a known parameter, and therefore this model does not capture the uncertainty regarding the amount of financial investment that we focus on here. Another difference is that both of these models are attempting to explain long-term patterns in commodity prices, not the behavior around a specific expiry. We also note that neither of these models highlight the potential signaling role of open interest that we see in the data.

More recent work by Ge et al. (2022) focuses on the April 20th, 2020 negative price event. However the focus of that paper is the microstructure impact related to Trade-at-Settlement (TAS) contracts. The model in that paper takes the fundamentals of the market as given and examines how TAS trading can contribute to a further deterioration of prices. We view this model as highly complementary to our model here. We think it is quite plausible that the fundamental effect we lay out here initially caused the negative prices, which were subsequently exacerbated by the structure of trading in the market.

15

#### 2.2.1 Structure of model

We specify a two-period model with a futures market at time 0, and a spot market with physical delivery at time 1. The model is meant to represent a futures market very near expiry of the contract, so we think of time 0 as corresponding to the period just prior to expiry (e.g. Friday, April 17th, 2020), and time 1 corresponding to the period when financial traders must close their positions without taking delivery (e.g. Monday, April 20th, 2020). We also assume that there is a fundamental price  $\hat{P}$  at which stored oil or refined oil can be sold at after time 1. Deviations from this price therefore represent short-term dislocations from fundamentals caused by conditions at the storage hub.

At time 0, risk-neutral profit maximizing producers make a decision to send an amount of oil Z to the hub. Not thate we dub these agents "producers" to emphasize that they control the supply of oil to the hub, but they more closely correspond to the midstream pipeline operators rather than operators of the actual oil wells. The total refining capacity ( $\tilde{R}$ ) is random and unobserved at time 0. This capacity is controlled by a measure of competitive risk-averse agents, each with access to a unit of refining capacity. We assume refining is costless, so each refiner makes a simple yes or no decision to operate their refinery.

At time 0 there is a futures market in a future with expiry at time 1. We assume refiners are risk averse with mean-variance utility, and therefore hedge their purchases by buying futures (taking a long position) in this market. Financial traders who are unable to take delivery also trade in this market. We assume for clarity that they only take long positions (though this is not necessary), and this amount of open interest is given by  $OI_F$ . The other side of the futures market (the short positions) are taken by competitive, risk-neutral arbitrageurs. The total open interest,  $OI_{Total} = \tilde{R} + OI_F$ , is observable to all agents in the model.

Note that we are assuming that refiners are risk-averse and hedge purchases in the futures market, but that the producers are risk-neutral and therefore do not. This assumption is made for tractability, as it greatly simplifies the solution for total open interest, but does not impact the qualitative implications of the model. However, it is important that all of the

Futures Market	Physical Market					
t = 0	t = 1					

- Risk-averse refiners hedge their oil purchases (*R̃*) in futures market
- Price-insensitive financial investors purchase long position in oil futures (OI<sub>F</sub>)
- Risk-neutral arbitrageurs take other side for positions of refiners and financial investors (total open interest  $\tilde{R} + O\tilde{I}_F$ )
- Risk-neutral producers observe open interest and choose amount of oil to send to hub (Z barrels) to maximize profits

- Oil (Z) is delivered to the hub
- Refiners take delivery of their oil allocation  $(\tilde{R})$
- Financial investors close futures positions but do not take delivery
- $Z \tilde{R}$  barrels of oil are added to (or removed from) storage
- Spot price clears competitive storage market

#### Figure 5: Model Timing

physical uncertainty in the model is related to the level of refinery capacity, and therefore reflects demand uncertainty. If the uncertainty at time 1 was primarily supply uncertainty (for instance uncertainty regarding the cost parameter for producers), high open interest would reflect high supply, and therefore coincide with a higher level of storage, and this would be inconsistent with the empirically negative relation between open interest and storage.

At time 1, the oil is delivered to the hub and used by refiners. We also assume that there is storage available, and again for tractability we assume this storage is supplied by competitive, risk-neutral, agents. Current storage is  $S_0$  and we assume there is a capacity of costless storage available  $\bar{S}$  with  $\bar{S} > S_0$ . We also assume for simplicity that  $S_0$  is greater than the maximum possible refinery usage  $\bar{R}$ , so that there is no possibility that storage goes to zero. Storage at time one is therefore  $S_1 = S_0 + Z - \tilde{R}$ . If  $S_1$  is greater than the capacity of costless storage  $(\bar{S})$  then costly storage is competitively supplied with an exponential cost of  $\frac{\tau}{v}(S_1 - \bar{S})^v$ . The timing of the model is summarized in Figure 5.

The model is very stark, and does not include many features of this market (notably convenience yields, refining costs, or fundamental price uncertainty). However, as we show, it can generate dynamics that at are similar to some of the empirical observations around the negative price event in April of 2020.

### 2.2.2 Time 1: Spot Market

To focus on the uncertainty regarding financial investment and the short term nature of the market near expiry, we assume that both stored oil and refined oil may be sold at some guaranteed fundamental price  $\hat{P}$  in the future. Differences between  $P_1$  and  $\hat{P}$  therefore represent a short-term dislocation of prices from fundamentals. We assume that there is no discounting between time 0 and time 1. Since refining is zero marginal cost and refiners are competitive, they will pay any price up to  $\hat{P}$  in the spot market. Storage agents choose  $S_1$  to maximize their profit which is given by:

$$\Pi_{S} = (\hat{P} - P_{1})(S_{1} - S_{0}) \qquad \text{if } S_{1} \le \bar{S}$$
  
$$\Pi_{S} = (\hat{P} - P_{1})(S_{1} - S_{0}) - \frac{\tau}{v}(S_{1} - \bar{S})^{v} \qquad \text{if } S_{1} > \bar{S}$$

Since  $S_1 = Z - \tilde{R} + S_0$ , storage agents' F.O.C.s imply that the spot market price  $P_1$ , given an amount of production Z and a refinery capacity  $\tilde{R}$ , is:

$$P_1 = \hat{P} \qquad \qquad \text{if } (Z - \tilde{R} + S_0) \le \bar{S} \qquad (1)$$

$$P_1 = \hat{P} - \tau \left( (Z - \tilde{R} + S_0) - \bar{S} \right)^{\nu - 1} \qquad \text{if } (Z - \tilde{R} + S_0) > \bar{S} \qquad (2)$$

Note that in regions incurring costly storage, the price falls below the longer-term fundamental, and prices can be negative if excess production  $(Z - \tilde{R})$  is high enough relative to available costless storage capacity  $(\bar{S} - S_0)$ .  $P_1$  is therefore a function of the endogenous choice of production Z and the random realization of refinery capacity  $\tilde{R}$ .

### 2.2.3 Time 0: Futures market and producers' decision to ship oil to the hub

Competitive producers choose to send Z to the hub at time 0. Producers are risk-neutral and maximize profits. Producers' profits (realized at time 1) are

$$\Pi_Z = P_1 Z - \frac{\phi}{2} Z^2$$

Where  $\phi$  determines the marginal cost sending an additional unit of oil to the hub. Producer's optimization implies:

$$Z = \frac{E[P_1|OI_{Total}]}{\phi} \tag{3}$$

This equation highlights the connection between the financial market and the real production decision. Any change in open interest that impacts the expected future spot price will lead to a change in the amount of oil sent to the hub.

We assume that producers know the probability distribution of  $\tilde{R} \in [0, \bar{R}]$  that falls between zero and an upper capacity constraint  $\bar{R}$ . We also assume that the fraction of this capacity available follows a beta distribution  $\frac{\tilde{R}}{\bar{R}} \sim B(\alpha, \beta)$ . When we solve the model we parameterize this distribution so that refinery capacity is typically near its maximum, but there is a small probability of significant shortfalls.<sup>7</sup>

We make the stark assumption here that producers have no information at time 0 about  $\tilde{R}$  beyond its unconditional distribution. While this is likely unrealistic, it is a useful simplification to illustrate how a futures market can provide information about the level of refinery capacity. However we first consider the model without a futures market.

<sup>&</sup>lt;sup>7</sup>Solving for equilibrium in the model for a given amount of refinery capacity and financial open interest involves choosing an endogenous amount of production so that Equation (3) holds using the definition of  $P_1$ from Equations (1) and (2). The probabilities for the expectation are the conditional probabilities of  $\tilde{R}$  given the observed level  $OI_{Total}$  which can be calculated using the distributions of  $\tilde{R}$  and  $OI_F$ . Since the price  $P_1$ is either flat or decreasing in production for all values of  $\tilde{R}$  and  $Z \ge 0$ , the equilibrium is unique.



Figure 6: Model results: No futures market

The figure plots various model outcomes as a function of refinery capacity  $\hat{R}$  in a model without a futures market. The maximum refinery capacity is set as  $\bar{R} = 4$ , and the ratio of realized capacity to this maximum capacity has a Beta Distribution with  $\alpha = 5$  and  $\beta = 1$ . Other parameters are producer cost  $\phi = .25$ ; storage costs  $\tau = 0.2, \nu = 5$ ; storage level and costless capacity  $S_0 = 4$ ,  $\bar{S} = 5$ , and long-term price  $\hat{P} = 1$ .

#### 2.2.4 Model solution without a futures market

In the case without a futures market, the unconditional expectation of spot prices determines the level of production. If  $\tilde{R}$  is extremely low relative to this expectation, there can be a very low, or even negative spot price.

Figure 6 plots the results of the model without a futures market. Price at time 0, and therefore the level of production, are constant, and producers choose production so that the expected spot price is slightly lower than  $\hat{P} = 1$ , to reflect the small probability that there is a shortfall of refinery capacity. In cases where the shortfall is severe enough, near-zero or negative prices are possible in the spot market, with the severity depending on the parameters of the model. Note that the price change occurs at time 1 when uncertainty is revealed.

### 2.2.5 Model with a futures market: Without financial traders

We now add a futures market. The presence of the risk-neutral arbitrageur implies that the time-0 futures price equals the expected time-1 spot price, conditional on the observed open

20

interest  $F = E[P_1|OI_{Total}]$ . Since there is no motive for speculation, refiners will simply hedge their full amount of refinery capacity. To see this, first suppose that each refiner buys a long position with a notional of w. The Profit  $\Pi_R$  at time 1 will then be:

$$\Pi_R = \hat{P} - P_1 + w(P_1 - F)$$

Since  $F = E[P_1|OI_{Total}]$ , we have that  $E[\Pi_R] = \hat{P} - P_1$  and  $VAR[\Pi_R] = (1-w)^2 VAR[P_1]$ , so setting w = 1 is the optimal strategy. We assume that even if there is no uncertainty, refiners will still choose to hedge the full amount of their capacity. Given this assumption, in the absence of financial traders, total open interest will be  $OI_{Total} = \tilde{R}$ , and will fully reveal the amount of refinery capacity and thus remove any uncertainty about the spot price at time 0. We note that the assumption of full hedging in the absence of uncertainty is only necessary for the illustrative case with no financial traders, alternatively we could assume a very small mass of financial traders to achieve a similar result.

Since there is no uncertainty, producers will never over-deliver, and will send an amount to the hub that satisfies  $\frac{P_1}{\phi} = Z$ . In this case the futures market removes the uncertainty that could lead to a surprise drop in prices. Figure 7 illustrates the results of the model in this case. Producers' production is always more than  $\tilde{R}$  as they are willing to incur some small amount of high-cost storage with this parameterization. However, spot prices are never negative, and returns in the futures market are always zero since there is no uncertainty to be resolved in the spot market.

#### 2.2.6 Model with a futures market: With financial traders

We assume financial traders inelastically demand a long position of a given size at any price in the futures market, and denote this amount as  $OI_F$ . We assume that  $OI_F$  is strictly positive, is independent with respect to  $\tilde{R}$  and is log-normally distributed, with  $\log(OI_F) \sim \mathcal{N}(\mu_F, \sigma_F)$ .

Since refiners always hedge exactly their full amount of capacity by the same argument as in the previous section, we have that



Figure 7: Model results: Futures market without financial traders

The figure plots various model outcomes as a function of refinery capacity. The maximum refinery capacity is set as  $\bar{R} = 4$ , and the ratio of realized capacity to this maximum capacity has a Beta Distribution with  $\alpha = 5$  and  $\beta = 1$ . Other parameters are producer cost  $\phi = .25$ ; storage costs  $\tau = 0.2, \nu = 5$ ; storage level and costless capacity  $S_0 = 4, \bar{S} = 5$ ; and long-term price  $\hat{P} = 1$ .

$$OI_{Total} = \tilde{R} + \tilde{OI}_F$$

Therefore, total open interest  $OI_{Total}$  is now a noisy signal about refinery capacity R. If there is a high level of financial open interest, producers may rationally, but mistakenly, interpret the large total open interest as an indication of high refinery capacity.

Figure 8 plots the outcomes for various levels of refinery capacity in two cases, one with high financial open interest, and the other with low. In the high open financial open interest case, the producer infers a higher expected value of refinery capacity due to the high open interest, and sends relatively more oil to the hub. This in turn leads to lower prices in the spot market, and in extreme cases can lead to negative spot prices.

To help understand when negative prices may occur, Figure 9 shows the joint probability distribution function of refinery capacity and financial open interest in the first graph, and spot prices as a function of refinery capacity and financial open interest in the second graph. Negative spot prices occur rarely (in this calibration approximately 0.2% (0.002) of the time),



Figure 8: Model results: Futures market with financial traders

The figure plots various model outcomes as a function of refinery capacity. The maximum refinery capacity is set as  $\bar{R} = 4$ , and the ratio of realized capacity to this maximum capacity has a Beta Distribution with  $\alpha = 5$  and  $\beta = 1$ . Other parameters are producer cost  $\phi = .25$ ; storage costs  $\tau = 0.2, \nu = 5$ ; storage level and costless capacity  $S_0 = 4, \bar{S} = 5$ ; long-term price  $\hat{P} = 1$ . Financial trader open interest is log-normally distributed with  $\mu_F = 0$  and  $\sigma_F = 0.75$ . The blue line represents a realization of financial open interest in the 5th percentile of the distribution, and the red line represents a realization of financial open interest in the 99th percentile of the distibution to capture a large positive surprise to financial open interest.



Figure 9: Model results: Joint distribution of refinery capacity and financial open interest, and equilibrium spot prices

The left plot shows the joint probability distribution function of refinery capacity and financial open interest. The right plot shows spot prices as a function of refinery capacity and financial open interest. The black contour line in the right hand graph defines the region where spot prices transition from positive to negative. Negative spot prices occur with very low probability, in regions where there is jointly high financial open interest and low refinery capacity.

and in cases with unexpectedly high financial open interest coinciding with unexpectedly low refinery capacity.

To understand how these negative prices relate to overall open interest and storage, we simulate 10,000 draws of the model, and record the spot price and storage utilization at time 1 along with the open interest at time 0. Figure 10 shows the results. As the figure shows, high open interest typically signals a large amount of refinery capacity, and so levels of storage fall. However, in some cases with high financial open interest, you have low levels of refinery capacity and over-production from producers. This necessitates a high level of storage, and therefore high marginal cost, and in extreme cases spot prices go negative. Note



Figure 10: Model simulations of open interest in time 0 vs. change in storage at time 1 The figure shows values of time 0 open interest and time 1 change in storage for 10,000 simulations of the model. Red stars denote simulations where negative spot prices occurred.

that this pattern is qualitatively similar to the data shown in Figure 4.

To summarize the model, our hypothesis is that an unusually high level of financial open interest just prior to the expiry of the May 2020 futures contract caused oil producers to incorrectly infer the presence of a large amount of physical traders able to take delivery. This in turn led to an excess amount of oil being sent to the hub, where there was in fact an extremely low levels of refiner demand and available storage capacity. The situation was revealed on April 20th when the financial traders attempted to close the positions without being able to take physical delivery, causing a sudden collapse in spot prices as the marginal cost of storing the oil rose to extremely high levels. Our simple two-period model captures this intuition, and is able to generate results consistent with empirical patterns in open interest and storage observed in the WTI market before and during the negative price event. We now turn to the impact of this event on producer decisions.

## **3** Benchmark risk and oil production decisions

In this section we explore how oil well production decisions were impacted by the benchmark risk associated with expiring futures contracts. Our analysis of the impact on producers focuses on the uncertainty regarding the June 2020 futures occurring in May. As Figure 11 shows, prices and volumes of put options with strike prices very close to zero spiked immediately after the negative prices in late April and stayed high well into May. These prices and volumes were unprecedented in the data, and, as the figure also shows, did not occur in corresponding futures for the global waterborne benchmark Brent Crude contracts. This is despite the two contracts trading at similar prices levels over this period.

The challenge here is that we need to disentangle the risk associated with futures expiry and the benchmark with overall macroeconomic risk during the COVID-19 pandemic. To accomplish this, we pursue two distinct empirical strategies. In the first, we will exploit the calendar moving average (CMA) nature of crude oil purchasing contracts along with proxies for daily oil production. In the second, we compare production (shut-in) decisions in North Dakota, to a control area with less exposure to WTI, namely Alberta, Canada. To motivate these methods, we first provide some institutional background.

### 3.1 Institutional background

There are several features of the market for physical crude oil that are important context for our empirical strategy. Here we cover these features in detail.

#### 3.1.1 Benchmark prices for different regions

Firms pay crude oil purchasers to purchase oil from the wellhead. Crude purchasers typically drive up to a well with a truck owned by the purchaser and fills the truck up from tanks on the well location. The purchaser then trucks this oil to a pipeline or terminal and then the crude oil is brought to market and processed at a refinery. In essence crude is transacted at tens of thousands of locations every day. The problem of setting the transacting price is laid out by Duffie and Stein (2015), without an independent benchmark price to settle the transaction,

26



Figure 11: Near-zero-strike put option prices and volumes in April and May 2020

Panel A of the plot shows the price of put options with a strike price of \$2 for various future contracts around the events of May 2020. Panels A and B show prices and volumes in WTI futures contracts, and Panels C and D show similar plots for Brent futures.

the buyer and seller could easily disagree on what the fair market price is for crude at the well site, hence there are a number of factors that drive a pricing mechanism based on a tradable index. As Duffie and Stein (2015) argue, market participants that use a benchmark in their transaction will "reap the information benefits of a benchmark...including lower search costs, higher market participation, better matching efficiency, and lower moral hazard in delegated execution." In order to obtain these benefits, market participants or their agents will often choose to substitute their "best-fit-for purpose" trade with a benchmark trade. The flagship, most liquid benchmark for crude oil in North America is given by the West Texas Intermediate (WTI) for delivery in Cushing, OK. Therefore most of the daily published prices for other geographies in the U.S. may not in fact be representations of specific recent transactions in the area, they are instead the daily WTI price with an infrequently updated differential (basis) applied to capture differences in quality or transportation costs. To illustrate this, Figure 12 reports the Calendar month postings for crude grades purchased by the refiner Phillips 66. As can be seen, crude grades at locations far from Cushing, including Texas (WT Inter),

Phillips 66 Crude Oil Prices for Apr-2020												
				in do	llars (	<b>(S)</b>						
Bulletin Number	Effective Dates	WT Inter	NM Inter	WT Sour	NM Sour	TX Pan All Fields	OK Pan All Fields	Central OK Swt	W Central TX Inter	NTX Inter	LLS Onshore	Central MT
Gravity A	Adjustment	Α	Α	В	H	Α	Α	С	С	С	Α	G
2020-063	4/1/2020	16.93	16.93	14.12	16.49	16.43	16.43	16.73	16.93	16.93	15.68	14.02
2020-064	4/2/2020	21.94	21.94	19.13	21.50	21.44	21.44	21.74	21.94	21.94	20.69	19.03
2020-065	4/3/2020	24.96	24.96	22.15	24.52	24.46	24.46	24.76	24.96	24.96	23.71	22.05
	4/4/2020	24.96	24.96	22.15	24.52	24.46	24.46	24.76	24.96	24.96	23.71	22.05
	4/5/2020	24.96	24.96	22.15	24.52	24.46	24.46	24.76	24.96	24.96	23.71	22.05
2020-066	4/6/2020	22.70	22.70	19.89	22.26	22.20	22.20	22.50	22.70	22.70	21.45	19.79
2020-067	4/7/2020	20.25	20.25	17.44	19.81	19.75	19.75	20.05	20.25	20.25	19.00	17.34
2020-068	4/8/2020	21.71	21.71	18.90	21.27	21.21	21.21	21.51	21.71	21.71	20.46	18.80
2020-069	4/9/2020	19.38	19.38	16.57	18.94	18.88	18.88	19.18	19.38	19.38	18.13	16.47
	4/10/2020	19.38	19.38	16.57	18.94	18.88	18.88	19.18	19.38	19.38	18.13	16.47
	4/11/2020	19.38	19.38	16.57	18.94	18.88	18.88	19.18	19.38	19.38	18.13	16.47
	4/12/2020	19.38	19.38	16.57	18.94	18.88	18.88	19.18	19.38	19.38	18.13	16.47
2020-070	4/13/2020	19.03	19.03	16.22	18.59	18.53	18.53	18.83	19.03	19.03	17.78	16.12
2020-071	4/14/2020	16.73	16.73	13.92	16.29	16.23	16.23	16.53	16.73	16.73	15.48	13.82
2020-072	4/15/2020	16.49	16.49	13.68	16.05	15.99	15.99	16.29	16.49	16.49	15.24	13.58
2020-073	4/16/2020	16.49	16.49	13.68	16.05	15.99	15.99	16.29	16.49	16.49	15.24	13.58
2020-074	4/17/2020	14.89	14.89	12.08	14.45	14.39	14.39	14.69	14.89	14.89	13.64	11.98
	4/18/2020	14.89	14.89	12.08	14.45	14.39	14.39	14.69	14.89	14.89	13.64	11.98
	4/19/2020	14.89	14.89	12.08	14.45	14.39	14.39	14.69	14.89	14.89	13.64	11.98
2020-075	4/20/2020	-41.01	-41.01	-43.82	-41.45	-41.51	-41.51	-41.21	-41.01	-41.01	-42.26	-43.92
2020-076	4/21/2020	6.63	6.63	3.82	6.19	6.13	6.13	6.43	6.63	6.63	5.38	3.72
2020-077	4/22/2020	10.40	10.40	7.59	9.96	9.90	9.90	10.20	10.40	10.40	9.15	7.49
2020-078	4/23/2020	13.12	13.12	10.31	12.68	12.62	12.62	12.92	13.12	13.12	11.87	10.21
2020-079	4/24/2020	13.56	13.56	10.75	13.12	13.06	13.06	13.36	13.56	13.56	12.31	10.65
	4/25/2020	13.56	13.56	10.75	13.12	13.06	13.06	13.36	13.56	13.56	12.31	10.65
	4/26/2020	13.56	13.56	10.75	13.12	13.06	13.06	13.36	13.56	13.56	12.31	10.65
2020-080	4/27/2020	9.40	9.40	6.59	8.96	8.90	8.90	9.20	9.40	9.40	8.15	6.49
2020-081	4/28/2020	8.96	8.96	6.15	8.52	8.46	8.46	8.76	8.96	8.96	7.71	6.05
2020-082	4/29/2020	11.68	11.68	8.87	11.24	11.18	11.18	11.48	11.68	11.68	10.43	8.77
2020 002	4/20/2020	15.46	15.46	12.65	15.02	14.96	1/ 96	15.26	15.46	15.46	14.21	12.55

Figure 12: Daily posted prices for different crude benchmarks in April 2020

This figure provides data on the crude oil posted prices for Phillips 66, these prices were used to compute the calendar month average for the months of April for different geographies. Note that for each of these prices, the price is equal to WTI less a constant differential (for instance \$2.90 for Central MT in the last column).

New Mexico (NM Inter), and Louisiana (LLS Onshore), were marked at negative values on April 20th. Taking a closer look at each column, the reason for this becomes clear, namely that the reference prices at different locations are simply the WTI benchmark daily close less a constant differential (for instance \$1.25 in the case of LLS). Of note, since the event, two major price reporting agencies (S&P Platts and Argus) have developed new regional benchmark prices that do not rely on the landlocked WTI pricing.<sup>8</sup>

<sup>&</sup>lt;sup>8</sup>See, for instance: https://www.reuters.com/article/us-usa-oil-prices-idUSKBN23W3CS/.



Figure 13: The geography of U.S. oil prices on April 20, 2020

### 3.1.2 Price and fundamentals across North America on April 20th, 2020

This type of benchmarking means that oil sold in other parts of the country will be exposed to WTI price changes even if they are caused by events specific to the facility at Cushing. Figure 13 illustrates this. On April 20th, prices in Louisiana, Michigan, North Dakota, and Colorado all settled negative. Notably, prices in California did not, as they use a different pricing mechanism.

One possibility of course is that the storage constraints in Cushing were in fact a common phenomenon across all of the regions that saw negative prices, and that California's separate network of pipeline and storage facilities did not see similar constraints. However, the data do not bear this out. Figure 14 shows storage capacity utilization in California, Cushing, and the rest of the Central U.S. As the plot shows, the spike in storage utilization was a unique feature of Cushing, and in particular there was ample storage in the rest of the Central U.S. Despite this ample storage, producers in places such as North Dakota effectively received negative prices due to the benchmarking in their contracts.

The fact that North Dakota did not experience storage shortages, but did experience



Figure 14: Storage capacity utilization across U.S. regions in 2020

negative prices, makes North Dakota oil wells an interesting laboratory to study the risk posed by the benchmark. However, to control for changes in fundamentals, we need a control group. For these purposes, we will use oil production in Alberta, which has several attractive features. Specifically, Alberta crude oil has the same end market as North Dakota, predominantly refineries in the upper Midwest, and their oil feeds into the same pipeline network as North Dakota crude. However, Alberta crude has a distinct pricing mechanism related to crude oil prices set in Edmonton, Alberta. Excluding oil sands (heavy oil), the wells across Alberta and North Dakota share similar characteristics. Figure 15 plots the crude oil price movements in April and May of 2020 across these geographies as well as the WTI and Brent (Global waterborne) crude oil prices. While the prices generally all track each other, the figure shows that the negative price spike was confined to the two U.S. benchmarks. In effect, Alberta production will serve as a proxy for "low WTI" exposure regions, while North Dakota serves as a proxy for "high WTI" exposure regions, with the implicit assumptions that technologies and fundamentals are otherwise similar across these two geographies.



Figure 15: International and domestic crude oil prices around April 2020

Figure plots daily closing prices for the North Dakota (Bakken), Cushing (WTI), Global waterborne index (Brent), and Alberta (Edmonton) around April 2020.

### 3.1.3 Calendar month average purchase agreements

Crude oil purchase agreements are not based on the price of crude on a given day, but rather, on the calendar month average (CMA) of reference prices. As reported in Figure 16, this means that oil sold early in the month will be subject to price fluctuations that occur later in the month. Moreover, this means that if crude benchmarks were to settle at a negative value on a given day, this negative number is incorporated in the calendar month average, effectively rendering all crude sold on that given day as being transacted at the negative price. The nature of this contract arrangement is supported by the letter Harold Hamm, CEO of Continental Resources wrote to the CFTC stating the events of April 20th "materially impacts the Calendar Month Average (CMA) pricing of physical crude," as presented in Figure 16. Figure 12 reports the Calendar month postings for crude grades purchased by the refiner Phillips 66. As can be seen, crude grades at locations far from Cushing, including Texas (WT Inter), New Mexico (NM Inter), and Louisiana (LLS Onshore), were marked at negative values on April 20th. We focus next our analysis on the behavior of producers faced with the risk of a repeat event associated with the June 2020 expiry on May 19th. The structure of the contract means that the first day of production that will face this benchmark risk occurs on May 1st, which is a feature we will exploit in our first empirical test.

### 3.2 Evidence from a daily proxy for oil production

Our first test for production responses to benchmark risk relies on the structure of the CMA contract. As mentioned previously, we are interested in evidence that producers changed their production decision in response to the risk associated with a repeat of negative prices around the expiry of the June 2020 WTI future on May 20th. Because any potential negative prices on this day would enter the CMA for the entire month of May, a barrel produced on May 1st will be exposed to this risk, while a barrel produced on April 30th will receive the April CMA and therefore not be exposed. Our hypothesis is that, to avoid such risk, producers may reduce production starting on May 1st. Conversely, on May 20th, after the expiry occurs, the price risk for barrels produced in May is now resolved, and without a



Figure 16: CMA purchase contracts

This figure documents the Calendar Monthly Average (CMA) purchase price mechanism crude producers use to sell their oil. The bottom part of the figure documents the computation, the top reports the excerpt from a letter that an oil company CEO wrote to the CFTC stating the effect the events of April 20th had on the price his firm received under this pricing mechanism.

major negative price shock, production may begin to increase.

To test these predictions we need a proxy for daily oil production, as oil production data are not publicly available at this frequency. We therefore proceed by using data on pipeline flows for natural gas. Natural gas flow data can be a good indicator of daily oil production changes because a portion of natural gas in the United States comes from commingled production with oil. That is, if natural gas goes down from a region that has commingled production, it is a direct indicator that oil production has also gone down. We compare how this "associated" gas production compares to "dry" gas or unassociated production by looking at production from states that have predominantly associated gas (from oil wells), namely North Dakota, Oklahoma, and Texas, versus states that have predominantly dry gas (from natural gas wells), namely Pennsylvania, West Virginia, and Kentucky.<sup>9</sup>

We plot associated gas and dry gas in Figure 17, and find that associated gas production drops roughly 5% from April 30th to May 1st, a four standard deviation day-to-day change. As can be seen dry gas remains flat during this time period. Our second empirical prediction concerns changes in associated gas production around the May contract expiry on May 19th. As it turned out, prices did not go negative on this contract expiry, and what we observe is that the day after this contract roll, associated natural gas begins to increase in production, consistent with more oil wells being reopened after being shut-in. The production drop on May 1st is particularly striking, because this occurred on a day with very little price change for crude oil, and there is no associated drop from the dry gas states. Outside of the changing exposure to the benchmark risk associated with the June 2020 expiry embedded in the CMA contract, it is unclear why this historic drop in production would take place. We therefore believe this result is strong evidence for the importance of benchmark risk in determining production decisions. Motivated by this high frequency data, we now move to our second set of tests using monthly well data.

<sup>&</sup>lt;sup>9</sup>As an example, the gas production in Pennsylvania from the Marcellus and Utica shale formations is predominantly dry gas, not crude oil. In contrast, Texas has significant "associated gas" from oil wells from places such as the Permian basin.



Figure 17: Daily natural gas production for associated and dry gas states

Figure plots daily production from "associated" gas states (ND, TX, and OK) and "dry gas" states (PA, WV, and KY). Associated states are states where natural gas comes primarily from oil producing wells, and dry gas states are states where natural gas comes primarily from wells that do not produce oil. Panel A plots the daily time series from 2017 to 2021, and Panel B plots the data from April to June of 2020. May 1st, 2020 is the first day that oil production is exposed to the expiry of the June 2020 futures contract via calendar moving average purchasing contracts, and May 20th is the expiry of the June 2020 future.

### 3.3 Evidence from monthly data on individual well production

In this section we estimate how oil well shut-in decisions relate to having crude oil purchase agreements indexed to WTI. We will compare production decisions in North Dakota to those made in Alberta, so that Alberta will serve as a "control" group in what is effectively a difference-in-difference approach. Motivated by our high frequency results above, we search for wells that are producing in April, cut back production for at least part of May, and then return to full production shortly thereafter. Formally we define this "temporary shut-in" as a well that is producing more than three fourths of available days in the previous month, less than one third of available days this month, and returns to more than three fourths of available days within two months (for instance, a well decreasing its production to less than one third of days in May 2020 would need to rise to more than three fourths of days by July 2020).<sup>10</sup> We calculate the percentage of wells in both Alberta and North Dakota that meet this criteria for each month that we have data, from the beginning of 2019 to the end of 2022.

Figure 18 shows the results of this analysis. Panel A plots the calendar monthly average price for both the Edmonton Mixed Sweet (EMS) benchmark in Alberta, and the WTI. Panel B plots the percentage of wells with partial shut-ins in both North Dakota and Alberta, and Panel C plots the difference between the two series in Panel B. Panel A shows that the benchmark prices for both regions track quite closely across the period, except for the negative price spike only witnessed for the WTI benchmark on April 20th, 2020. Both benchmarks reach their lowest level in April 2020. Panel B shows the temporary shut-ins prior to April 2020 tend to be a steady-but-small portion of wells (approximately 0.5% of wells each month), most likely reflecting routine service of the wells. The trough of the price series in April 2020 sees a roughly equivalent increase to approximately 2% of wells being temporarily shut-in across both regions, potentially as wells are taken out for service in an opportunistic period after prices were low through the beginning of April. While these partial shut-ins remain at a similar level in Alberta for May 2020, there an extreme spike in this behavior observed in

<sup>&</sup>lt;sup>10</sup>We use one third of days as the cutoff to include wells that begin producing shortly after the futures roll on May 19th, 2020, and we use the two-month window since production is still increasing into the second half of June in Figure 17. The qualitative results are largely unchanged when we use a one-month window or different cutoffs to define the temporary shut-in.

North Dakota (ND), with nearly 8% of all ND wells shut in. Panel C plots the difference, and as the plot shows, May 2020 is again an extreme outlier (here approximately 20 standard deviations). What is again striking is that this extreme behavior is happening in May 2020, where prices were substantially higher than April 2020. The behavior is also only present in North Dakota, an area similar to Alberta in nearly every respect except for the benchmarking of the purchase contracts. This analysis provides further support to our hypothesis that real production reacted to the benchmark uncertainty associated with the futures expiry in May 2020.

### 3.4 Long-run productivity of shut-in wells

As Figures 17 and 18 show, there seems to be a substantial impact on production in May in response to the increased risk associated with contract expiry. One natural question is: How costly this behavior was for producers? One could imagine producers responding to this risk with a financial hedge, such as an option, which may be more efficient if shutting in the wells creates long-term damage to their future productivity. Another possibility is that shutting in the wells simply moved production from the current uncertain time to a time in the future if this temporary shut-in did not impact the long-term viability of the wells.

To answer this question we examine the average production (barrels/month) of wells in North Dakota that experienced a temporary shut-in in May of 2020. We compare these wells to a matched sample of wells that were not shut in in May of 2020, where the matched sample is created using the well in North Dakota with the closest level of total production (barrels/month) in April of 2020. Figure 19 shows the results. The sample is constructed so that the wells have nearly identical average productivity in April of 2020. The shut-in wells drop their productivity precipitously in May as they are shut in, and since some remain shut-in in June average productivity rebounds but stays relatively low. However, in June of 2020, when the wells are fully re-opened, productivity of the shut-in wells actually increases back above the April 2020 level, and is then well above the matched sample of open wells which experience a more typical decline curve. The shut-in wells' productivity remains above



Figure 18: Temporary shut-ins in Alberta and North Dakota (2019 - 2022)

Panel A of the figure plots the daily price of oil in Alberta (Edmonton Mixed Sweet) and the daily price in Cushing, OK (West Texas Intermediate). Panel B plots the percentage of oil wells in each month experiencing a temporary shut-in. This is defined as a well that was producing more than three fourths of available days in the previous month, less than one third of available days this month, and returns to more than three fourths of available days within two months. Panel C plots the difference between the two series in Panel B.



Figure 19: Long-run production from shut-in and open wells in North Dakota

The plot shows the average monthly production (in barrels/month) for North Dakota wells that were temporarily shut in in May of 2020 relative to a matched sample (on productivity) of North Dakota wells that were not shut in (producing more than three fourths of days) from April to June of 2020.

the open wells for a period of approximately two years, before converging by the end of the sample. Interestingly, the cumulative productivity for the two sets of wells is very similar over this period, suggesting that there was little long-term cost in terms of productivity from the shut-in. Overall this analysis suggests that, while the financial distortions of oil prices impacted producers' behavior, the nature of the wells allowed producers to effectively hedge the rise in uncertainty by shifting production forward rather than forgoing it altogether.

# 4 Conclusion

In this paper, we document an under-appreciated channel through which asset prices can affect the real economy. We present a simple theoretical model and argue that the negative crude oil prices of April 20th, 2020 were due to physical storage constraints in Cushing, but were exacerbated by positions of financial traders in an expiring futures contract. This event created a situation where a major oil price benchmark experienced a dislocation from the fundamentals affecting oil prices throughout most of the U.S., and we find that this channel led to real effects. Specifically, oil producing firms react to the elevated risk of a repeat event in May of 2020 by preemptively shutting a portion of their wells in anticipation of future dislocations of the benchmark, despite fundamentals outside of the benchmark delivery location being supported by ample storage availability. Once this risk has receded, firms resume normal operations. Ultimately, our study highlights the importance of asset prices for firms that adopt benchmark pricing in their purchase and sales agreements.

# References

- Abrantes-Metz, R., Kraten, M., Metz, A., Seow, G.S., 2012. Libor manipulation? 36, 136–150.
- Barro, R., 1990. The stock market and investment. Review of Financial Studies 3, 115–31.
- Bessembinder, H., Carrion, A., Tuttle, L., Venkataraman, K., 2016. Liquidity, resiliency, and market quality around predictable trades: Theory and evidence. Journal of Financial Economics 121, 142–66.
- Bond, P., Edmans, A., Goldstein, I., 2012. The real effects of financial markets. Annual Review of Financial Economics 4, 339–360.
- Bond, P., Goldstein, I., 2015. Government intervention and information aggregation by prices. The Journal of Finance 70, 2777–2812.
- Buyuksahin, B., Lee, T.K., Moser, J.T., Robe, M.A., 2013. Physical markets, paper markets and the WTI-Brent spread. The Energy Journal 34, 129–151.
- Buyuksahin, B., Robe, M.A., 2014. Speculators, commodities and cross-market linkages. Journal of International Money and Finance 42, 38–70.
- Chen, Q., Goldstein, I., Jiang, W., 2007. Price informativeness and investment sensitivity to stock price. Review of Financial Studies 20, 619–50.
- Cheng, I., Xiong, W., 2014. Financialization of commodity markets. Annual Review of Financial Economics 6, 419–441.
- Duffie, D., Dworczak, P., 2021. Robust benchmark design. Journal of Financial Economics 142, 775–802.
- Duffie, D., Dworczak, P., Zhu, H., 2017. Benchmarks in search markets. Journal of Finance 72, 1983–2044.

- Duffie, D., Stein, J.C., 2015. Reforming libor and other financial market benchmarks. Journal of Economics Perspectives 29, 191–212.
- Edmans, A., Goldstein, I., Jiang, W., 2012. The real effects of financial markets: The impact of prices on takeovers. Journal of Finance 67, 933–971.
- Fama, E., French, K., 1987. Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage. Journal of Business 60, 55–73.
- Faulkender, M., 2005. Heding or market timing? selecting the interest rate exposure of corporate debt. Journal of Finance 60, 931–962.
- Foucault, T., Fresard, L., 2012. Cross-listing, investment sensitivity to stock price, and the learning hypothesis. Review of Financial Studies 25, 3305–3350.
- Ge, Y., Kang, W., Tang, K., Yang, L., 2022. Closing pressure, predatory trading, and the negative price of oil. Predatory Trading, and the Negative Price of Oil (April 4, 2022).
- Goldstein, I., Yang, L., 2022. Commodity financialization and information transmission. The Journal of Finance 77, 2613–2667.
- Gorton, G., Hayashi, F., Rouwenhorst, G., 2013. The fundamentals of commodity futures returns. Review of Finance 17, 35–105.
- Hayek, F., 1945. The use of knowledge in society. American Economic Review 35, 519–30.
- Hirshleifer, D., 1988. Residual risk, trading costs, and commodity futures risk premia. Review of Financial Studies 1, 173–93.
- Hirshleifer, D., 1990. Hedging pressure and futures price movements in a general equilibrium model. Econometrica 58, 411–28.
- Hong, H., Yogo, M., 2012. What does futures market interest tell us about the macroeconomy and asset prices? Journal of Financial Economics 105, 473–90.

- Kang, W., Rouwenhorst, K.G., Tang, K., 2020. A tale of two premiums: The role of hedgers and speculators in commodity futures markets. Journal of Finance 75, 377–417.
- Luo, Y., 2005. Do insiders learn from outsiders? evidence from merger and acquisitions. Journal of Finance 60, 1951–82.
- Masters, M., 2008. Testimony before the Committee on Homeland Security and Governmental Affairs. Technical Report. United States Senate. URL: https://www.hsgac.senate.gov/imo/media/doc/052008Masters.pdf.
- Ozik, G., Sadka, R., Shen, S., 2021. Flattening the illiquidity curve: Retail trading during the covid-19 lockdown. Journal of Financial and Quantitative Analysis 56, 2356–2388.
- Ready, M.J., Ready, R.C., 2022. Order flows and financial investor impacts in commodity futures markets. The Review of Financial Studies 35, 4712–4755.
- Rouwenhorst, G.K., Tang, K., 2012. Commodity investing. Annual Review of Financial Economics , 447–67.
- Singleton, K., 2012. Investor flows and the 2008 boom/bust in oil prices. Management Science 60, 300–318.
- Sockin, M., Xiong, W., 2015. Informational frictions and commodity markets. Journal of Finance 70, 2063–2098.
- Tang, K., Xiong, W., 2012. Index investment and financialization of commodities. Financial Analysts Journal 68, 54–74.

## A Data sources

### A.1 Well data

We collect production data on individual wells for the United States and Canada for 2019 to 2022. We rely on two different data sources. For all United States data we rely on DrillingInfo,

43

which provides detailed well level data by month, by producer, with detailed geographic location data for most jurisdictions in the United States. In our study we predominantly focus on North Dakota producers. Our Canadian data is downloaded from a website maintained by the province of Alberta, and is also at the well-level, by producer, and monthly. We focus on wells that are arguably similar to one another, in particular Canada has substantial production from oil sands, which is a distinct production technology, therefore we limit our Canadian data to non-oil sands wells that produce a crude grade similar to WTI.

### A.2 Price, futures, and options data

Data on daily benchmark prices is obtained from Bloomberg, and intraday price and volume data for oil futures (including TAS contracts) are obtained from the CME. Daily option price data are from the CME and ICE. We also hand collect posted prices from crude purchasers off of their websites.

### A.3 Storage data

Storage data for different geographies are collected from the Energy Information Association.

### A.4 Search volume data

Data on searches are from QiHoo 360 obtained at trends.so.com.