When Benchmarks Fail:
The Causes and Consequences of Negative Oil Prices

Erik P. Gilje* Robert Ready† Nick Roussanov‡ Jérôme P. Taillard§

January 21, 2026

Abstract

On April 20, 2020, the crude oil benchmark in North America, the West Texas Intermediate (WTI) futures contract for delivery in Cushing, Oklahoma, settled at a negative price for the first time in history. We combine new empirical evidence with a stylized theoretical model to show that, while local storage constraints created the conditions for negative prices, a key catalyst was unusually large long positions in the expiring contract held by financial traders unable to take physical delivery. These positions distorted the demand signal in the futures market, intensifying pressure on the limited storage capacity and precipitating a sharp price dislocation. We then document that this dislocation significantly influenced oil production decisions through contractual exposure to WTI-based pricing. Even oil producers far from Cushing that were not directly impacted by the storage constraints responded with sharp output curtailments in the face of heightened benchmark risk. Our findings highlight how transitory futures price dislocations due to noise trader demandcan have real economic consequences.

*No affiliation

†Lundquist College of Business, 382 Lillis Hall, 1208 University of Oregon, Eugene, OR 97403. Email: rready@uoregon.edu

‡The Wharton School, University of Pennsylvania, 3620 Locust Walk Suite, 2400, Philadelphia, PA 19104, and National Bureau of Economic Research. Email: nroussan@wharton.upenn.edu

§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 Univeristy, 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 before expiry, the benchmark front-month futures price for West Texas Intermediate (WTI) crude oil delivered to Cushing, Oklahoma fell below zero for the first time in history, settling at $-\$37$ per barrel. We use this event as a laboratory to examine the ability of financial market benchmarks utilized for coordinating real activity to efficiently aggregate information in the presence of uninformed financial investors ("noise traders").

We provide empirical evidence and develop a stylized model to demonstrate that the negative price event was driven, at least in part, by unusually large long positions held by retail traders in the soon-to-expire futures contract. These positions conveyed a misleading signal of strong physical demand at the Cushing storage hub, prompting excess physical supply that exceeded the hub's limited storage capacity. We further show that this dislocation in futures prices affected production decisions of oil producers through their contractual exposure to the WTI benchmark, including those not directly affected by storage constraints at Cushing. Our findings highlight a previously underappreciated dimension through which asset prices and trading activity of financial market participants, including largely uninformed investors ("noise traders"), can influence the real decisions of firms due to the widespread reliance on a financial market-determined benchmark. .

The NYMEX WTI contract, traded on the Chicago Mercantile Exchange, is the most liquid and actively traded crude oil futures contract globally. It specifies a crude oil grade, a specific delivery location at Cushing, Oklahoma, and a delivery month. Both financial and physical traders participate in this market. Any trader holding a position at expiry must make or take physical delivery. For a long position, taking delivery at Cushing requires either a pipeline allocation out the hub or access to a local storage terminal. There is no option to load oil onto ships or trucks, and as such no free disposal. As a result, Cushing operates as a closed system with a fixed, exhaustible volume of storage and pipeline capacity. This landlocked constraint stands in sharp contrast to the other primary global benchmark, Brent crude, which is priced on seaborne oil that can be shipped and stored anywhere in the world.[1]

[1] See, for instance: "What are the differences between ICE Brent and NYMEX WTI futures?," ICE, June 2020, https://www.ice.com/insights/market-pulse/what-are-the-differences-between-ice-brent-

On April 20, 2020, the NYMEX WTI Light Sweet Crude Oil May 2020 (CLK2020) futures contract was the front month contract.[2] Any traders still holding positions at settlement on April 21, 2020 (the next day) would have been required to either receive or deliver physical crude to settle their remaining open futures positions. At the time, storage levels at Cushing had risen sharply to roughly 70% of listed physical capacity, and most of the remaining capacity was "committed" or pre-sold, leaving effectively no additional space for physical delivery.[3] Selling pressure began to rise on the evening of Sunday, April 19, 2020, and intensified the next morning, sending prices into negative territory for the first time around 1PM on April 20, 2020 (see Figure 1). Prices remained below zero, and the contract settled at -$37/bbl at 2:30PM (red dot in figure). This settlement price, regardless of the trading volume during that window, is the reference price used for index and benchmark pricing across the U.S. crude oil complex. Although prices remain below -$10/bbl for only five hours on April 20th, 80 crude grades at locations across the United States transacted that day at an average of -$44/bbl, even at locations with ample storage capacity and waterborne flexible storage options.

The timing of this event, just prior to the expiry of a financial futures contract, underscores the link between financial markets tied to delivery at Cushing and the physical market for storage capacity at the hub. In the first part of the paper, we set out a hypothesis for how an expiring futures contract could trigger a negative price event and present a stylized model to formalize the intuition. Our model relies on a newly documented empirical fact: high open interest in an expiring futures contract is typically followed by an increase in available storage capacity at Cushing shortly after expiry. Therefore, open interest near expiry embeds a signal of physical demand for oil. Consistent with the model's predictions, we show that this pattern reflects a market in which uncertain physical demand from crude oil users (e.g.,

---

and-nymex-wti-futures.

[2]The front month contract is the contract with the nearest expiration date.

[3]See, for instance, "Today in Energy," EIA, April 27th, 2020, https://www.eia.gov/todayinenergy/detail.php?id=43495, and "No vacancy: Main U.S. oil storage in Cushing is all booked," Reuters, April 21st, 2020, https://www.reuters.com/article/world/no-vacancy-main-us-oil-storage-in-cushing-is-all-booked-idUSKCN22332U/). Storage levels ultimately peaked at 85% of listed physical capacity in May 2020 (see Figure 12).

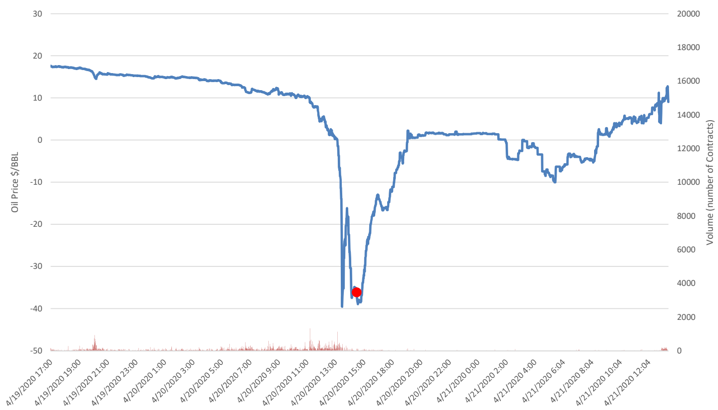

Figure 1: Intraday prices in May 2020 future prior to expiry

refineries) is partially revealed through the aggregate open interest. This demand is met by midstream operators taking the short side in the futures market and delivering physical oil to the hub, effectively pre-selling it. When both sides of the market are primarily physical traders shortly prior to expiry, the balance of supply and demand in the futures market translates into a balance in the subsequent physical spot market.

This balance can be upset, however, if an unusually large share of long positions in the expiring contract is held by financial traders who are unable to take physical delivery (including retail traders). These financial traders act as noise traders that can distort the demand signal in the futures market. The result is a futures market equilibrium in which a large short position from physical traders delivering oil to the hub is matched by a large long position from financial traders (rather than refiners, who demand physical oil). When these financial traders close out their positions prior to expiry so as to avoid taking delivery, the excess physical supply is revealed and prices collapse, since physical demand turns out to be illusory. This mechanism implies that negative prices can arise when unusually high financial open interest coincides with unusually low demand from physical traders. [4] It shows that in the twelve months prior to April 2020, long open interest from financial traders prior to the

[4]This mechanism is consistent with the findings documented in CFTC's Interim Staff Report on negative oil prices published in November of 2020.While aggregate open interest for each contract is available to market participants at the end of each trading day, the report's breakdown of open interest in the expiring contract by trader type was not publicly available prior to its release.

second to last trading day of an expiring contract was relatively steady, averaging roughly 31,000 contracts with a standard deviation of approximately 8,000. Over the same period, long open interest from physical traders averaged roughly 33,000 contracts with a standard deviation of approximately 11,400. For the May 2020 contract, long open interest from financial traders surged by more than seven standard deviations to approximately 96,000 contracts, while long open interest from physical traders was 12,000 contracts, two standard deviations below its mean, and the lowest level observed over the preceding twelve months. The sum corresponds to a total open interest of 108,000 contracts. To clear the market, this unusually large long financial position was met by a large short position from physical traders, producing an unprecedentedly large net physical short position of 43,000 thousand contracts (43 million barrels of oil), a number nearly three standard deviations below the prior twelve-month mean. This physical imbalance would have been revealed only on April 20th, when financial traders began selling in the futures market to close their positions prior to taking physical delivery, triggering a collapse in prices amid a scramble for scarce available storage.

The model also sheds light on another puzzling aspect of April 20th, namely that the collapse appears to have taken market participants by surprise. This is evident from the sharp negative returns on futures and the lack of premiums on near-zero strike put options just prior to the event, despite the fact that the possibility of negative prices had been publicly acknowledged as early as April 8th, when the CME announced a "Clearing Plan" to ensure normal market functioning even if prices fell below zero. In our baseline calibration, given the available storage capacity prior to the April 20th expiry (approximately 16 million barrels), we set the unconditional probability of negative prices at 1 percent; unlikely but not implausible. Conditional on observing 108 thousand contracts of total open interest, however, the high open interest is interpreted as a strong signal of physical demand, which in our calibration drives the perceived probability of an overflow effectively to zero. This same demand signal also induces suppliers to increase deliveries to the hub, creating a physical imbalance and precipitating the price collapse when financial traders close their positions and the lack of physical demand is revealed. Under this calibration, the model also implies an essentially zero probability of negative prices when storage capacity is near its historical average, regardless of financial trader activity. Overall, these results suggest that the conditions for negative prices stemmed from fundamentals, namely concerns about limited storage capacity, but that unusually high financial demand near expiry temporarily dampened those concerns, leading to oversupply at the hub and the subsequent crash.

In our model, financial traders are modeled as exogenous noise traders. In a negative price event, they incur large losses when closing out their long positions without taking delivery, while physical traders remain fully hedged. Although we lack direct evidence on the precise source of the April 2020 spike in financial investment, we provide anecdotal evidence suggesting that it was linked to increased long positions of retail investors.[5] Contemporaneous reporting indicates that retail traders directly trading futures on U.S. platforms (e.g. Interactive Brokers, E-Trade, and TD Ameritrade) were among those who suffered losses on April 20th, and that retail traders with long futures positions through the Bank of China's (BOC) "Crude Oil Treasure" (COT) product also incurred substantial losses.[6] We also document that online search volume for the COT product spiked in late March and early April, suggesting increased interest and potentially higher inflows into this product. Back-of-the-envelope computations further indicate that the reported losses for COT investors are broadly consistent with the magnitude of the open interest spike observed in April 2020.

The Bank of China's decision to allow COT retail investors to hold the front month futures contract until the day before final settlement (futures expiry), and then forcibly roll those positions, was highly unusual for such a financial product.[7] By contrast, most financial traders (such as the United States Oil Fund (USO), the largest oil ETF) close out their front month positions between 15 to five days prior to expiry, rolling into the next closest contract to expire. As a result, open interest in the expiring contract is typically only a small fraction

[5] Ozik et al. (2021) document similar increases of retail investor holdings in equities.

[6] See "Day Traders are a New Wrinkle in the Negative Oil Price Mystery" https://www.bloomberg.com/news/articles/2020-06-08/are-day-traders-a-possible-cause-for-oil-prices-going-negative, as well as the references found in Section 2.1 regarding the Bank of China's "Crude Oil Treasure" financial product.

[7] These contract terms are cited among the causes of the substantial losses alleged by COT retail investors in an ongoing class action lawsuit (S.G. v. Bank of China Ltd (2023)).

of the overall market, so the impact of April 20th was likely negligible for most traders who had already exited the May contract.

Nevertheless, the closing price of the front-month contract remains the benchmark for many crude sales agreements benchmarked to WTI prices at Cushing. These contracts often cover oil produced far from Cushing and of different grades, relying on the deep liquidity of WTI futures for price discovery. This convention, however, exposes oil producers to the risk that benchmark prices may become disconnected from their own local fundamentals, whether due to physical constraints at Cushing, such as limits to storage or deliverability at the hub, or to the trading activity of market participants.

This setting therefore provides a unique opportunity to explore how asset price dislocations driven by uninformed financial traders can affect real decisions through the benchmark channel. While crude oil fundamentals were undoubtedly strained by the COVID-19 pandemic, we present evidence that the WTI price collapse was consistent with short-term frictions, particularly concerns about deliverability and storage issues at the delivery hub, rather than a reflection of broader fundamentals.

Our empirical analysis investigates how oil producers respond to the heightened benchmark risk that followed this unprecedented episode of negative WTI prices. Although prices quickly rebounded, options market data indicate that fears of a recurrence, especially around the expiry of the next WTI futures contract (June 2020 contract expiring May 19th), remained elevated well into May. Identifying the effect of benchmark-specific pricing risk requires disentangling it from concurrent macroeconomic disruptions caused by the COVID-19 pandemic. To this end, we implement two complementary empirical strategies.

Our first set of tests uses high-frequency proxies for daily oil production and leverages the structure of common crude oil purchase agreements between producers and midstream operators. Most physical oil trades in North America use Calendar Month Average (CMA) pricing, under which every barrel sold in a given month receives the average of daily benchmark prices over that month. This pricing arrangement ties daily production to the full month's average benchmark price, making benchmark risk a source of testable empirical predictions around two key dates. Specifically, if producers anticipated a second benchmark dislocation near the

May 19th expiry of the June contract, they would have strong incentives to shut in production at the start of May, when barrels first become exposed to that month's CMA pricing. This behavior would be expected to generate sharp discontinuities in production on May 1st, when the pricing window opens, and May 20, when the uncertainty tied to the contract expiry is resolved.

Consistent with these predictions, we observe a four standard deviation drop in our production proxy between April 30 and May 1, the first day barrels were exposed to the May CMA price. This sharp decline does not reflect a broader trend but rather a discrete discontinuity that aligns precisely with the start of the pricing window tied to the next futures contract expiry. Production rebounds beginning on May 20th, the day after the June 2020 WTI contract expired and the risk of another extreme price dislocation had passed, and continues to recover through month-end. These high-frequency patterns provide compelling evidence that producers responded to benchmark risk not only by reacting to realized prices but also by anticipating future benchmark-driven pricing shocks. The sharp timing and magnitude of these shifts are difficult to reconcile with broader COVID-related uncertainty, which would not be expected to produce such discrete shifts in daily production.

Our second empirical strategy compares production decisions across two adjacent regions with similar geology and market access but different benchmark exposures: North Dakota and Alberta. Both regions are geographically remote from Cushing and were not subject to localized storage constraints during this period. However, oil from North Dakota is priced relative to the WTI benchmark, whereas Alberta's light crude is priced off the Edmonton Par benchmark. Although the two benchmarks are typically highly correlated, Edmonton prices neither turned negative nor experienced a large dislocation on April 20.

This setting yields clear predictions: if producers are responding to WTI-specific benchmark risk, we would expect a larger and more pronounced production response in North Dakota than in Alberta, despite broadly similar macroeconomic conditions. Using monthly well-level production data that report the number of days a well is active, we define temporary shut-ins as wells that were producing in April, reducing activity for a significant portion of May, and then returning to near-full production shortly thereafter.

We find that approximately eight percent of wells in North Dakota meet this shut-in definition in May 2020, compared to just two percent of wells in Alberta. Both the level of shut-ins in North Dakota and the six percentage point difference between the regions represent sharp discontinuities relative to the surrounding 48-month period. This divergence is particularly striking given the price environment: by May 1, prices had already recovered to early-April levels and continued to rise throughout the month. The average May price in both Alberta and North Dakota was nearly double the April average. Absent WTI-specific benchmark risk, it is difficult to reconcile the magnitude and timing of these shut-ins with improving fundamentals. Finally, we show that the shut-ins in North Dakota had little impact on the long-term productivity of the affected wells. This finding suggests that producers were able to manage benchmark risk effectively, using temporary shut-ins as a flexible response to the WTI dislocation without incurring lasting production losses.

Together with our high-frequency results, our empirical evidence suggests that benchmark-linked CMA pricing transmitted the WTI dislocation directly into production decisions in areas without local storage constraints. Producers facing WTI-based pricing shut in wells despite improving fundamentals, while those benchmarked to alternative prices did not. The stark contrast across otherwise similar geographies underscores the real effects of benchmark risk in commodity markets.

Our paper contributes to the literature on the role of uninformed traders in financial markets While early literature on noisy rational expectations, such as Radner (1979), Grossman (1976), and Grossman and Stiglitz (1980) relied on the presence of non-informational traders to "break" the classic "no trade" result of Milgrom and Stokey (1982), if such "noise" traders generate substantial aggregate asset price volatility it can also contribute to systematic mispricing, which is the central insight of De Long et al. (1990). In our model noise trader risk impacts prices not because it deters arbitrageurs (who are risk-neutral) from equating futures prices with their (rational) expectations, e.g. as in De Long et al. (1990), but simply because their presence obscures the "fundamental" signal, thus distorting the optimal response of the other market participants (e.g., Black (1986)).

Our results contribute to the broader literature on the interaction between financial markets and firms' real economic decisions. This literature, dating back at least to Hayek (1945), and more recently surveyed by Bond et al. (2012), emphasizes the role of secondary market prices as valuable sources of information for firms. Prior work shows that feedback from asset prices can (1) help managers decide when and where to invest (Chen et al. (2007), Foucault and Fresard (2012), Barro (1990)), (2) influence corporate decisions such as mergers and acquisitions (Luo (2005) and Edmans et al. (2012)), (3) inform government regulatory interventions (Bond and Goldstein (2015)), and (4) help employees determine where to seek employment Gao et al. (2021). Because of the rich information embedded in liquid financial instruments, their prices can serve as benchmarks in contractual pricing arrangements (e.g., Duffie and Stein (2015), and Duffie et al. (2017)). While benchmarking can reduce frictions and search costs, it also exposes firms to the risk that contract prices deviate from fundamentals. Yet, to date, far less attention has been given to these benchmark-related risks and their potential real effects when prices are distorted by market dislocations or limits to arbitrage.[8]

We also contribute to the literature on the financialization of commodity markets. Commodities have become an important asset class and theory predicts that uninformed financial traders can influence commodity prices (e.g. Hamilton and Wu (2014), Sockin and Xiong (2015), Basak and Pavlova (2016), Baker (2021), and Goldstein and Yang (2022)). Empirical work, however, has yielded mixed findings on the price impact of financialization (e.g., Fama and French (1987), Masters (2008), Sanders and Irwin (2010), Irwin and Sanders (2011), Tang and Xiong (2012), Cheng and Xiong (2014), Hamilton and Wu (2015), and Basak and Pavlova (2016)) and on its effects on real production decisions (e.g., Brogaard et al. (2019) and Bohl et al. (2023)). Much of this literature focuses on longer time series of prices and investment, and generally finds modest price impact of financial investors, reflecting the depth and liquidity of major commodity futures markets, which enables them to accommodate even large investor flows (Ready and Ready (2022)). By contrast, we provide an in-depth analy-

[8]One such risk is the potential manipulation of benchmark prices. The LIBOR scandal illustrates the strong incentives market participants may have to collude in setting benchmark rates (e.g., see Abrantes-Metz et al. (2012)), and has spurred interest in designing manipulation-proof benchmarks (e.g., Duffie and Dworczak (2021)).

sis of an episode in which retail investor flows into an illiquid front-month futures contract precipitated a sharp price dislocation. We then trace how this financial market disruption transmitted through benchmark pricing arrangements and document significant real effects on production decisions.[9]

2 Causes of negative oil prices

In this section, we empirically motivate and derive a simple model of an expiring oil futures contract that captures the interactions between financial and physical traders. The central mechanism in the model is that open interest in the futures market serves as a signal for demand at the storage hub. To motivate this mechanism, we first present empirical evidence on the relationship between open interest and storage, followed by a detailed examination of the source of financial open interest in the expiring May contract.

2.1 Empirical motivation

We begin our empirical investigation with a notable fact: prior to the negative price event on April 20th, 2020, the May 2020 WTI futures contract exhibited an unusually high level of open interest at the close of trading on April 17th, just two trading days before its expiration on April 21st. Panel A of Figure 2 plots open interest two days prior to expiry for contracts maturing over the three years preceding the negative price event. The red bar highlights the stark increase of approximately 40,000 contracts (corresponding to 40 million barrels of oil) for the May 2020 contract. These data are published by the exchange for each contract maturity at the end of each trading day and are therefore observable by all market participants. What is not observable in real time, however, is the composition of this open interest by trader type. The CFTC's "Disaggregated Positions of Traders" reports provide such a breakdown,

[9] A related strand of this literature provides theoretical and empirical evidence on how the composition of market participants, namely hedgers versus speculators, can be important for commodity markets (e.g., Hirshleifer (1988), Hirshleifer (1990), Faulkender (2005), Hong and Yogo (2012), Rouwenhorst and Tang (2012), Acharya et al. (2013), Gorton et al. (2013), Buyuksahin and Robe (2014), Kang et al. (2020)) and, in particular, how certain dynamics may arise, which drive prices away from fundamental values (Singleton (2012)).

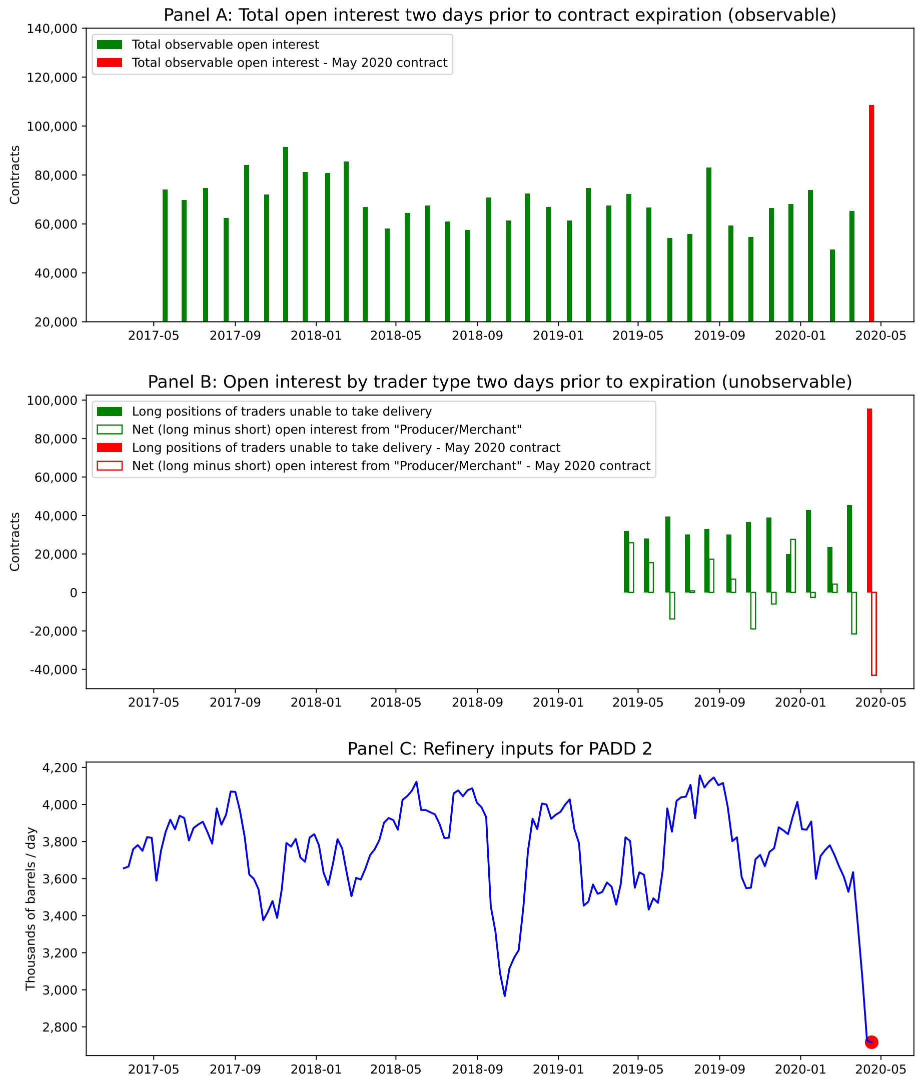

Figure 2: Open interest in expiring futures contracts and refinery utilization

Panel A plots open interest (number of contracts) in the expiring futures contract measured two trading days prior to its monthly expiry. For example, for the May 2020 contract we plot open interest at the close of trading on Friday, April 17th 2020, which was two trading days prior to April 21st, 2020, the final settlement day (expiry) of the contract. These data are available immediately after market close on the CME website. Panel B shows open interest two trading days prior to contract expiry, broken out by trader type. Long positions of traders unable to take delivery are calculated as total long open interest minus the long positions of traders classified as Producer/Merchant by the CFTC. The figure also reports net positions (long minus short) of Producer/Merchant traders. These data are drawn from Figure 14 of the CFTC Interim Staff Report, published on November 23, 2020, which only included the previous 12 months of observations, and were not publicly available before that date. Panel C plots weekly refinery inputs for PADD 2, which covers the Midwest, including Oklahoma. Refinery utilization for the week ending April 17th is highlighted (dot), as reported by the EIA on April 22, 2020.

but these data are aggregated across all active futures contract for a given product and are published weekly and with a lag.

Following the negative price event, the CFTC released an Interim Staff Report on November 23rd, 2020. This report offered a more detailed breakdown, including a figure (Figure 14) that discloses open interest by trader type two days prior to expiry for the May 2020 contract and the 12 previous contract expiries. Panel B of Figure 2 reproduces data from that figure. As shown, the elevated open interest for the May 2020 contract was driven almost entirely by an unusual increase in long positions of financial traders who could not take delivery and were therefore required to close out their positions prior to contract expiry.[10] Panel B also reports the net position (long minus short) of physical traders (classified as Producer/Merchant). The figure also indicates that the May 2020 contract exhibited a highly unusual imbalance among these traders, with shorts positions exceeding long positions by approximately 40,000 contracts. Panel C provides further context, showing that weak buying interest was likely driven in part by a sharp decline in refinery demand. While some portion of this drop would have been public knowledge at the time, the full extent remained uncertain because refinery utilization data are only released by the EIA a week later (e.g., April 17th data were released on April 22nd). In particular, if the surge in long financial positions in the futures market came from an atypical source, it could have been interpreted as physical demand, and thus met with offsetting short positions from physical traders planning to deliver oil to the hub. In this scenario, the conditions for an oversupply and binding storage constraints would have been in place by April 17th, but the imbalance would only have been revealed on April 20th, when financial traders attempted to liquidate their positions without taking physical delivery.

In our model, an *unexpectedly* high financial open interest from unsophisticated (noise) traders can precipitate a price collapse near a futures expiry, with these noise traders ulti-

[10] According to CME Group, traders must close or roll their physical-delivery futures positions prior to the final trading day to avoid being obliged to make or take delivery. In addition, Trading-at-Settlement (TAS) orders for expiring contracts must be entered by the end of the business day immediately preceding the final trading day, as TAS is not available on the contract's last trading day. Financial trader long positions are calculated as the long positions of all trader types excluding the category Producer/Merchant, as this is the only category with remaining long positions on the final trading day of the contract (see Figure 16 of the Interim Staff Report).

mately bearing substantial losses. While we lack detailed disaggregated data for the unusually large long financial positions shown in Panel B, strong anecdotal evidence suggests that a portion of this surge in financial open interest near contract expiry was driven by retail traders, particularly those using the *Crude Oil Treasure* (COT) product offered by the Bank of China (BOC).[11] The COT product is unusual in that retail investors are responsible for rolling their expiring contract into the next nearest contract themselves. If they fail to roll voluntarily, the rollover occurs automatically the day before the final contract settlement (contract expiry). This structure stands in sharp contrast to more conventional financial investment products, such as the United States Oil Fund (USO), which roll positions well before expiry (approximately two weeks before expiry in the case of the USO, see Bessembinder et al. (2016)). Reports also indicate that retail investors in the COT product owed the BOC approximately \$1.4 billion after the negative price event. At a price of -37 dollars per barrel and a contract size of 1,000 barrels, this loss corresponds to a position of approximately 38,000 contracts; a figure broadly consistent with the spike in open interest observed in April, 2020.

Figure 3 presents further evidence suggesting that at least a portion of the large spike in financial open interest in April of 2020 might have been driven by the presence of Chinese retail investors. 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

[11]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.

Panel A: Crude Oil Treasure advertisement

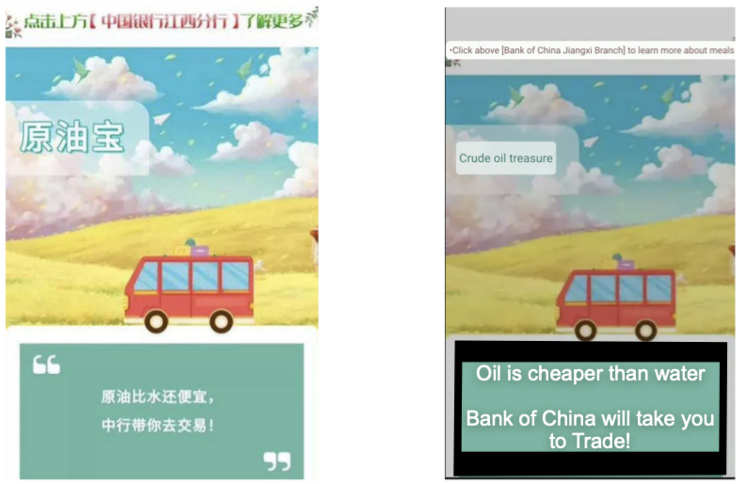

Panel B: Crude Oil Treasure search volumes 1/1/2020 - 4/17/2020

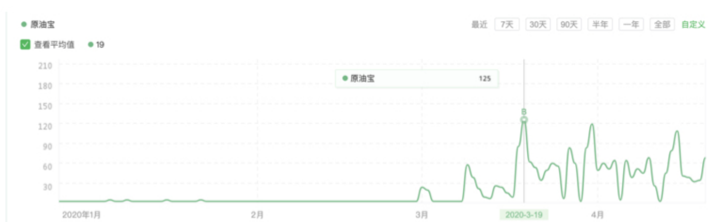

Panel C: Crude Oil Treasure search volumes 1/1/2020 - 6/1/2020

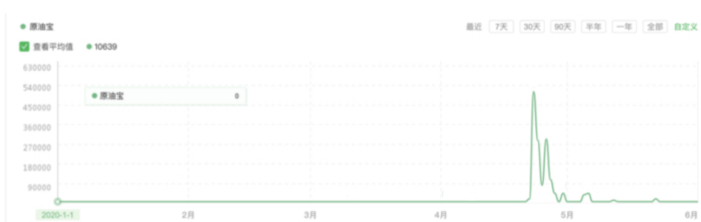

Figure 3: Interest in Crude Oil Treasure product

Panel A displays an advertisement for the Bank of China's "Crude Oil Treasure" product in its original form alongside a Google Translate version. Panel B plots online search volume data for "Crude Oil Treasure" from QiHoo 360 over the period January 1, 2020 to April 17, 2020. The vertical line marks the first day on which new investors would have been invested in the May 2020 delivery contract. Panel C repeats the search volume plot for the period January 1, 2020 through June 1, 2020. The spike on April 20, when prices turned negative, is much larger in scale, rendering the earlier search activity from Panel B invisible.

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.

An unusually high financial open interest late in the contract's life can impact the spot market if participants interpret this open interest as evidence of demand from physical traders willing to take delivery. Figure 4 provides evidence that futures open interest can serve as such a signal. The figure presents three plots with observable total open interest two days prior to contract expiry on the X-axis each time. Panels A and B draw on data from the CFTC Interim Staff Report to show that open interest is typically a reliable indicator of physical demand. Panel A plots the long open interest of physical traders prior to expiry, illustrating that over the previous twelve months, this unobservable physical open interest was highly correlated with observable total open interest. The May 2020 contract, highlighted in red, is a stark outlier with elevated open interest driven largely by financial rather than physical traders. Panel B shows that high observable open interest generally corresponds to excess physical demand, reflected in a more positive net imbalance from physical traders, with May 2020 again standing out as an exception. Finally, Panel C examines the predictive relation between total open interest two days before contract expiry and changes in storage levels at the Cushing hub, as reported in the following week's *Weekly Petroleum Status Report* from the Energy Information Administration (EIA).[12]

The plot reveals a strong relationship between open interest in the expiring contract and subsequent changes in storage over the three-year period preceding April 2020. The evidence across these panels provides clear empirical support for the conjecture that high open interest

[12]These reports are released on Wednesdays at 10:30 AM, and provide storage data as of the end of the previous week (typically Friday). Accordingly, this information is not available to the market participants in real time. We measure the weekly change in storage between the first reported value that occurs after the second day prior to expiry for a contract, and the value reported the following week. For example, for the May 2020 contract, the second day prior to expiry was Friday, April 17th. We use open interest from that day and the weekly change in storage between April 17th (reported on Wednesday April 22nd) and April 24th (reported on Wednesday April 29th).

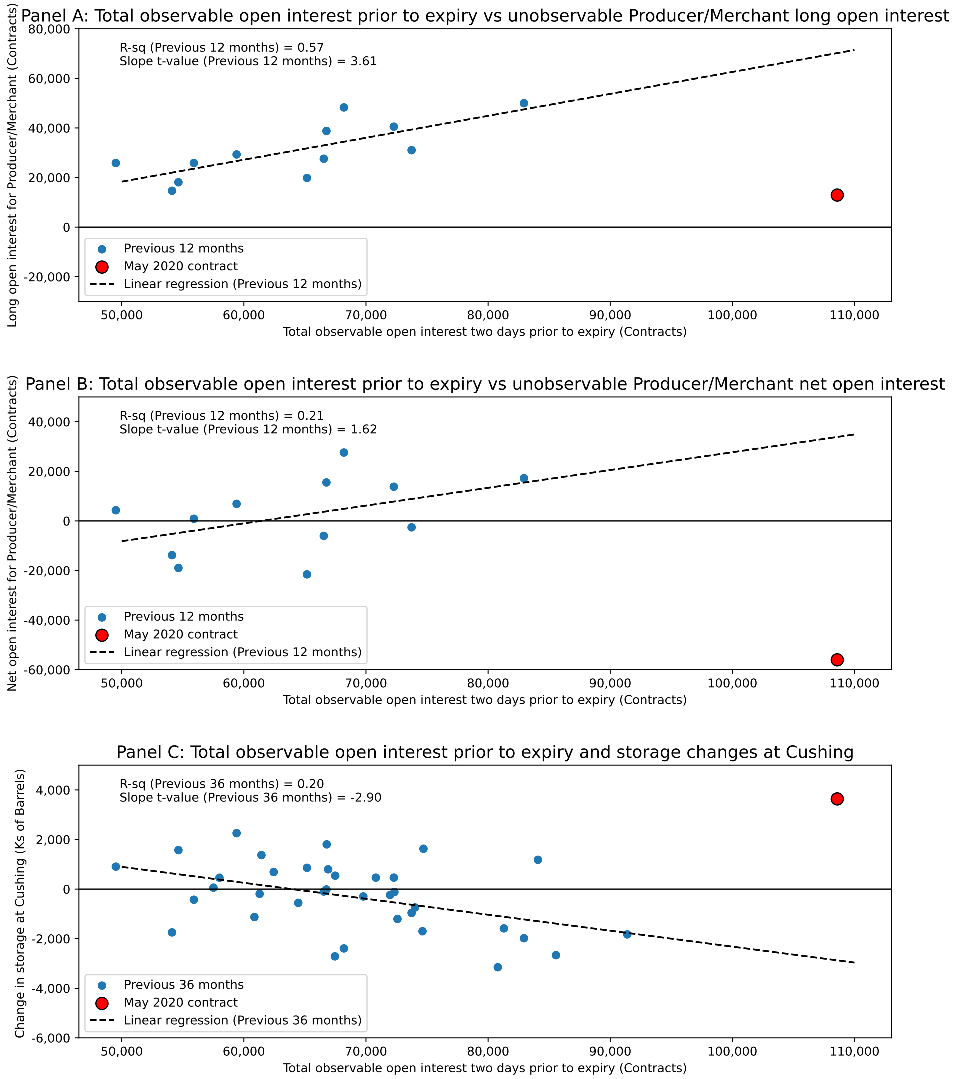

Figure 4: Signals of physical demand from open interest prior to futures contract expiry

Panels A and B plot total publicly observable open interest two days prior to contract expiry (X-axis) against, respectively, the long and net (long minus short) positions for traders classified as Producer/Merchant. Data on Producer/Merchant positions are drawn from Figure 14 of the CFTC Interim Staff Report, published on November 23, 2020, which only included the previous 12 months of observations. Panel C plots publicly observable open interest two days prior to expiry, reported at the close of trading by the CME, against the weekly change in storage at Cushing reported in the EIA Weekly Petroleum Status Report. Storage changes are reported weekly by the EIA; thus, for each contract, we use the storage change for the week containing the day two days prior to contract expiry. The blue dots represent futures contracts prior to May 2020 (12 months for Panels A and B, 36 months for Panel C). The red dots represent the May 2020 contract. The regression lines are fitted excluding the May 2020 contract.

	Change in storage at Cushing			
	Prior 3 years	Prior 5 years	Prior 10 years	Post Crisis
	(1)	(2)	(3)	(4)
Open interest	-0.056^{***}	-0.034^{**}	-0.027^{***}	-0.023^{**}
	(0.020)	(0.015)	(0.009)	(0.009)
Futures basis change	$126,332.600^{***}$	$55,995.500^{***}$	$41,712.630^{***}$	$15,691.250^{***}$
	(41,928.700)	(15,997.360)	(10,511.200)	(5,226.696)
Constant	$3,480.581^{**}$	$2,013.826^{**}$	$1,658.373^{***}$	$1,367.082^{**}$
	(1,396.412)	(980.938)	(595.518)	(570.212)
Observations	36	60	120	135
R^2	0.371	0.244	0.169	0.098

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

This table shows results from regressions of weekly storage changes at Cushing (in thousands of barrels) on the open interest in the front-month WTI futures contract, measured at the close of trading two days prior to contract expiry. Storage changes are reported weekly by the EIA; thus, for each contract, we use the storage change for the week containing the day two days prior to contract expiry. The regressions also control for 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, measured two days prior to expiry. Heteroskedasticity robust standard errors are in parentheses.

signals strong physical demand, and thus greater available storage. In sharp contrast, the May 2020 contract stands out as an outlier, with elevated open interest coinciding with a substantial buildup in storage (i.e., a *decrease* in storage capacity).

Table 1 extends this analysis to a longer historical window, reaching back to the 08-09 financial crisis. The table reports 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 influences storage levels through the cost-of-carry arbitrage relation (e.g., Ederington et al. (2021)). Across the four specifications, using expanding windows back to the financial crisis, the negative relation between open interest and changes in storage remains statistically significant. To our knowledge this relation between open interest and storage at Cushing has not been documented previously, and it again supports the view that unusually high financial open interest could have conveyed a misleading signal to the physical oil market in Cushing.

In summary, our hypothesis is that the unexpectedly large amount of financial open interest near the expiry of the May 2020 contract, driven in part by retail investors, contributed to a misalignment between physical demand and supply. Physical suppliers, interpreting the high open interest as indicative of delivery-ready buyers, pre-sold in the futures market to financial investors unable to take delivery. As the contract approached expiry, these financial traders attempted to unwind their positions, triggering a sudden price collapse as excess supply met severely limited storage capacity. In the next section, we present a model that formalizes this mechanism.

2.2 Model

Our model builds on classic models of noise trading (e.g., Grossman and Stiglitz (1980)), and relates to recent models of commodity financialization, including Hamilton and Wu (2014), Sockin and Xiong (2015), Basak and Pavlova (2016), Goldstein and Yang (2022), and Ge et al. (2022), but it differs along several key dimensions. With the exception of Ge et al. (2022), these models aim to explain long-term patterns in commodity prices rather than price behavior around specific contract expiry dates, and none focus on the potential signaling role of open interest that we see in the data. While Ge et al. (2022) also examine the April 20, 2020 negative price event, their focus is on the microstructure effects of Trade-at-Settlement (TAS) contracts and the closing pressure they exert on settlement prices. Their model takes market fundamentals as given and examines how TAS trading can exacerbate price declines. In contrast, our model centers on the interplay between financial investment and physical storage constraints. Many contemporaneous accounts of the event emphasized the sudden realization among market participants that storage was nearly exhausted. For instance, one storage broker remarked on April 21: "I have never been contacted by as many hedge funds as I did yesterday looking for storage."[13]

[13]See "Remaining Oil in Storage is Already Booked", Pipeline & Gas Journal, 4/21/2020 (https://pgjonline.com/news/2020/04-april/remaining-oil-storage-in-cushing-ok-is-already-booked-traders). It is unclear from this account whether the hedge funds in question already held WTI futures prior to April 20 and, unwilling to liquidate at negative prices, were seeking storage capacity to take delivery, or

Additionally, our model provides novel insight into how negative prices in April 2020 could be both “expected” and “unexpected.” They were “expected” in the sense that market participants recognized the possibility, as evidence by the CME’s April 8th announcement on April of a “Clearing Plan” confirming that futures and options trading would proceed normally even if prices turned negative. Yet they were also “unexpected” in the sense that prices turned negative with extraordinary speed, producing steep losses for long futures positions just before the April 21st expiry, while options markets had assigned near zero probability to such an outcome only days earlier.

Our model provides an explanation for this apparent inconsistency. Negative prices are unconditionally more likely if storage capacity is scarce, as in April of 2020. However, their conditional probability is highly sensitive to total open interest prior to expiry. High open interest , such as that observed in the week before April 21, signals ample demand, leading market participants to believe that supply is unlikely to exceed even the limited available storage. This perception, driven by financial traders’ long positions, reduces the perceived probability of a storage overflow, prompting midstream operators to maintain or even increase supply. When financial traders then reveal their unwillingness to take delivery by closing positions en masse, the resulting excess supply triggers a sharp price collapse. We view the model of Ge et al. (2022) as complementary to ours. It is plausible that the mechanism we describe initiated the price drop, which was subsequently amplified by the microstructure dynamics emphasized in their analysis.[14]

---

whether they were attempting to capitalize on negative prices by going long and taking delivery contingent on securing storage in Cushing. However, what is clear from this and other contemporaneous accounts (see footnote 3), is that virtually all storage requests made on April 20 were denied.

[14]We also note that while prices began to recover after the April 20 settlement, they remained negative through the night, with approximately 4,000 contracts transacting at negative prices before moving above zero on the morning of April 21. This pattern is consistent with a gradual easing of the storage constraint as operators found ways to free up previously contracted capacity, and is less consistent with a pure “closing pressure” explanation. After the settlement on April 20th 2020, there were no more TAS trades in the contract, and yet negative prices persisted.

2.2.1 Model structure

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 designed to capture conditions in a futures market very close to the contract expiry. We interpret time 0 as the period just prior to expiry (e.g. Friday, April 17th, 2020), and time 1 as the period when financial traders must close their positions to avoid taking delivery (e.g. Monday, April 20th, 2020). We assume the existence of a fundamental price $\hat{P}$ at which stored oil or refined oil can be sold after time 1. Deviations from this price therefore represent short-term dislocations from fundamentals, driven by storage constraints at the delivery hub.

At time 0, risk-averse profit maximizing midstream operators with mean-variance preferences decide how much oil Z to deliver to the hub. This decision is irreversible, reflecting the constraints of the pipeline-based transportation network around Cushing. Physical demand for oil comes from refineries, and this physical demand ($\tilde{R}$) is random and unobserved at time 0. This demand is provided by a measure of competitive risk-averse refiners with mean-variance preferences, each operating a single unit of refining capacity. We assume refining is costless, so each refiner makes a binary decision at time 0 of whether or not to commit to operate at time 1 based on a private signal about its production capacity $I^i = \{0, 1\}$; a measure $\tilde{R}$ of refiners receive a positive signal $I^i = 1$; and may choose to produce one unit of oil (if doing so yields positive profits), while the rest remain idle. Crucially, these decisions are not observable by other market participants and are therefore revealed only at time 1.[15]

At time 0, a futures market opens for a contract maturing at time 1. Both risk-averse refiners and producers hedge their exposure by buying or selling futures in this market. Financial traders who cannot take delivery also participate in the futures market. They are comprised of two groups: noise traders (“financial investors”), whose demand is exogenous and thus perfectly inelastic, and “rational arbitrageurs,” whose demand is perfectly elastic. For expositional clarity, we assume that inelastic financial investors take only long positions

[15]This reduced-form specification is meant to capture both private information about local demand (similar to Sockin and Xiong (2015)), and shocks to the refineries’ ability to produce, such as a Covid-related outbreak among workers at the refinery.

(though this is not essential for our results), and the exogenous random quantity of such inelastic financial demand is denoted by $\tilde{OI}_F$.

We assume that all physical uncertainty in the model pertains to demand. If, instead, the main uncertainty at time 1 were on the supply side (for instance uncertainty regarding the cost parameter for midstream operators), then high open interest would signal high expected supply and be associated with higher storage levels. This outcome would run counter to the empirical finding from the previous section that high open interest is associated with lower storage levels.

At time 1, oil is delivered to the hub and used by refiners. We also assume that there is storage capacity available that is supplied by competitive, risk-neutral agents.[16] We assume that the total available (unused) costless storage capacity at time 0 is C_0. At time 1, oil is delivered to the hub and sent to refiners so that the total available capacity at time 1 is $C_1 = C_0 - Z + \tilde{R}$. When $C_1 < 0$, additional storage is supplied competitively with a quadratic cost of $\frac{\tau}{2}(C_1)^2$. This specification implies that marginal storage costs rise rapidly once a certain level of storage is reached, reflecting the exhaustible nature of storage capacity at Cushing discussed earlier in the paper.[17] The timing of the model is summarized in Figure 5.

The model is intentionally stylized, and omits several features of the crude oil market, including convenience yields, refining costs, and fundamental price uncertainty. Nevertheless, as shown below, it generates dynamics consistent with the empirical patterns observed around the negative price event of April 2020.

[16]The risk-neutrality of the storage agents implies that they do not hedge in the futures market at time 0. This assumption greatly simplifies the open interest dynamics in the futures market. In contrast to midstream suppliers and refiners, these agents are not exposed to significant losses at time 1, as their physical decisions are made at time 1 rather than committed at time 0.

[17]By focusing on available capacity rather than total storage levels, we implicitly assume that the probability of a complete depletion of oil inventories (a "stockout") is zero. Including a potential stockout condition does not change the implications of the model.

Futures Market

$t = 0$

- Risk-averse refiners hedge their oil purchases ($\tilde{R}$) in futures market
- Price-insensitive financial investors purchase long position in oil futures ($\tilde{OI}_F$)
- Risk-averse midstream operators choose amount of oil to send to hub (Z) to maximize profits and hedge in futures markets
- Risk-neutral arbitrageurs take the other side of futures positions
- Aggregate futures open interest is observable to all participants

Physical Market

$t = 1$

- 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

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 can be sold in the future at some guaranteed fundamental price $\hat{P}$. Differences between P_1 and $\hat{P}$ therefore represent short-term price dislocations from fundamentals due to storage constraints. We assume no discounting between time 0 and time 1. Since refining has zero marginal cost and refiners are competitive, they will pay any spot price up to $\hat{P}$. Storage agents choose period-1 storage by selecting remaining available capacity C_1 to maximize their profit given by:

$$\begin{aligned} \Pi_S =& (\hat{P} - P_1)(C_0 - C_1) && \text{if } C_1 \geq 0 \quad \text{(Normal Conditions)} \\ \Pi_S =& (\hat{P} - P_1)(C_0 - C_1) - \frac{\tau}{2}(C_1)^2 && \text{if } C_1 < 0 \quad \text{(Over Capacity)} \end{aligned}$$

Since $C_1 = C_0 - Z + \tilde{R}$, the first-order conditions for storage agents imply that the spot market price P_1 given production Z, refinery demand $\tilde{R}$, and initial capacity C_0 is:

$$P_1 = \hat{P} \qquad \text{if } C_0 - Z + \tilde{R} \geq 0 \quad \text{(Normal Conditions)} \tag{1}$$

$$P_1 = \hat{P} - \tau\left(C_0 - Z + \tilde{R}\right) \qquad \text{if } C_0 - Z + \tilde{R} < 0 \quad \text{(Over Capacity)} \tag{2}$$

In the costly-storage region, the price falls below the longer-term fundamental $\hat{P}$, and can be negative if excess production $(Z - \tilde{R})$ is sufficiently large relative to available costless storage capacity (C_0). P_1 is therefore a function of supply to the hub Z (determined at time 0) and the random realization of refinery capacity $\tilde{R}$ (revealed at time 1).

2.2.3 Time 0: Futures market and midstream operators' supply decision

Competitive midstream operators choose quantity Z to send to the hub at time 0. These operators have mean-variance preferences over profits and thereforemaximize

$$E[\Pi_Z] - \lambda_Z Var(\Pi_z)$$

Midstream operators' profits (realized at time 1) are given by:

$$\Pi_Z = P_1 Z - (\phi_0 Z + \frac{\phi_1}{2} Z^2) + w_z(P_1 - F)$$

We model the midstream operators' cost function as having both a linear component determined by ϕ_0, and a quadratic component determined by ϕ_1.We include the two separate terms so that we can calibrate the model to match both the average level of oil sent to the hub (determined by both ϕ_0 and ϕ_1) and the responsiveness of midstream operators to changes in futures prices (determined primarily by ϕ_1). Midstream operators also have a risk aversion coefficient $\lambda_Z \geq 0$; F denotes the futures price at time 0; and w_z is the futures position held by midstream operators at time 0. The refiners solve an analogous problem, choosing at time 0 whether or not to operate their refinery at time 1 and how much to hedge in the futures market, leading them to maximize:

$$E[\Pi_R] - \lambda_R Var(\Pi_R)$$

Since the spot price cannot exceed the fundamental price ($\hat{P}$), refiners who received a positive production signal will always commit to operate at full capacity and their profit Π_R^i at time 1 is:

$$\Pi_R^i = \left(\hat{P} - P_1\right) + w_R^i(P_1 - F)$$

Risk-neutral arbitrageurs who speculate in the futures market and maximize the expected value of future profits given by:

$$\Pi_A = w_A(P_1 - F)$$

Finally, there are financial investors whose demand for long futures position $\tilde{OI}_F$ is exogenous and completely inelastic. All market participants can observe total open interest. To calculate total open interest, we first make the assumption that no trader simultaneously holds both a long and a short position (such positions would be netted out by the exchange and not appear in the final open interest tally), and we assume that financial traders take the minimum position required to clear the market.

The definition of equilibrium is standard: all agents maximize their utility and futures markets clear at time 0, while spot markets clear at time 1. The first-order condition of the risk-neutral arbitrageur implies that the time-0 futures price equals the expected time-1 spot price, conditional on the observed open interest $F = E[P_1|OI_{Total}]$. This specification differs from models that study the sources of expected returns in commodity futures markets (e.g. Acharya et al. (2013), Hamilton and Wu (2014), and Goldstein and Yang (2022)), which feature hedging demand from physical traders combined with limits to arbitrage in the form of risk-averse financial arbitrageurs or speculators. In those settings, non-zero expected returns on futures contracts creates a motive for speculation. In contrast, in the present framework, there is no speculative motive: the expected payoff from a futures position is zero. Producers therefore fully hedge their production by setting $w_z = -Z$, which eliminates all variance in their profits. Likewise, each refiner hedges its unit of production by setting $w_R^i = 1$, making total open interest from refiners equal to $\tilde{R}$.[18] Total observable open interest is thus the combined long positions of financial investors and refiners (who demand long positions) or the total quantity of oil supplied by the midstream operators (who take short positions), with any mismatch (either short or long) being met by the arbitrageurs, which implies that $w_A = Z^* - (\tilde{OI_f} + \tilde{R})$ and $OI_{Total} = max(Z^*, \tilde{OI_F} + \tilde{R})$. In our calibration, the combined long demand from financial traders and physical traders almost always exceeds Z^*, so that $OI_{Total} = \tilde{OI_F} + \tilde{R}$ in most cases. Midstream operators maximize their expected profit by supplying the optimal level of oil to the hub:

$$Z^* = \frac{E[P_1|OI_{Total}] - \phi_0}{\phi_1} = \frac{F - \phi_0}{\phi_1} \tag{3}$$

[18]The first order condition for the producers' futures position is: $0 = -\lambda_z(w_z + Z)Var(P_1)$, and for refiners: $0 = -\lambda_R(w_R^i - 1)Var(P_1)$.

This relationship highlights the link between the financial and physical markets: any change in open interest that alters the expected future spot price will lead to a change in the futures price, and therefore a change in the quantity of oil supplied to the hub.

2.2.4 Parameterizing and solving the model

We solve the model numerically, which requires specifying distributions for the exogenous realizations of physical demand and financial open interest, as well as choosing values for the parameters $(\tau, \phi_0, \phi_1, C_0, \hat{P})$. The model is intentionally stylized and not intended to provide a fully quantitative description of the data. Rather,, matching certain salient features of observed market behavior helps to demonstrate that the conditions leading to negative prices in April 2020 were highly unusual, yet the outcome was consistent with the model's predictions .

For the distributions of the two exogenous random variables, we draw on data from the 12 months preceding April 2020, as reported in the CFTC Interim Report. Although these data were not directly observable to market participants in real time, we treat them as a reasonable proxy for the longer-term expectations of market participants who have observed patterns in this market over many years. As our proxy for physical demand we use the long positions of open interest from the Producer/Merchant category for the expiring contract (vertical-axis in Panel A of Figure 4). Financial open interest is calculated as the remaining long open interest, which corresponds to traders who are unable to take physical delivery.

We fit three alternative distributions to each series: 1) A truncated normal distribution, 2) a log-normal distribution, and 3) a non-parametric distribution obtained via Kernel-Density Estimation (KDE). Figure 6 presents the results. Panels A and B show that both physical and financial long open interest have a mean of roughly 30 thousand contracts. However, the variance of physical long open interest is approximately 50% greater than that of financial long open interest. Panel C presents the posterior distribution of physical long open interest conditional on observing the actual total long open interest of 108.6 thousand contracts at the close of trading on April 17, 2020, the final trading day before the negative price event of April 20. A market participant who knew the underlying unconditional distributions would

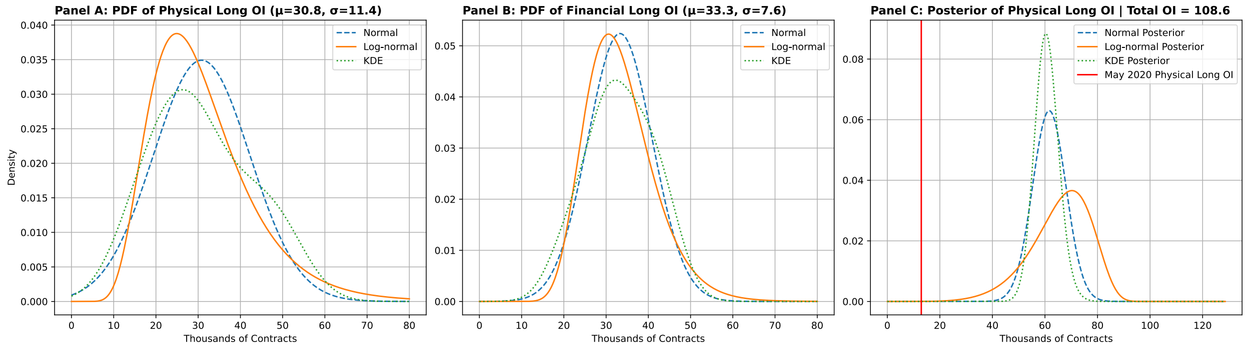

Figure 6: Fitted Distributions for Physical and Financial Open Interest of Expiring Contracts The figure presents distributions fitted to open interest data for the 12 expiring futures contracts prior to the May 2020 contract, obtained from the CFTC Interim Report. These data were not publicly available in real time. Panel A fits three distributions: (1) truncated normal, (2) log-normal, and (3) truncated kernel density estimation (KDE), to the realized long open interest for the Producer/Merchant category, which is able to take physical delivery. Panel B fits the same distributions to the realized long open interest for all other traders (total open interest less Producer/Merchant long open interest). Panel C shows the posterior distribution for Producer/Merchant long open interest, conditional on the observed total reported open interest prior to the May 2020 expiry (108.6 thousand contracts) and assuming independence between the realizations of physical and financial long positions. The red vertical line in Panel C denotes the actual, unobservable Producer/Merchant long open interest (12.9 thousand contracts).

have inferred a conditional mean physical long open interest of approximately 60 thousand contracts, assuming that the unusually high total open interest was driven primarily by higher total physical open interest. In reality, the true physical long open interest was only 12.9 thousand contracts, just under eight standard deviations below the posterior mean given a normal fit, and even further away under the more flexible KDE. This extreme deviation is the "surprise" at the heart of the negative price event. The drop in April 2020 was not driven by new information regarding broad fundamentals. Instead, the market learned that the large outstanding long positions in the expiring futures contract were not linked to physical demand from producers capable of taking physical delivery. To formalize this intuition, we take the empirical distributions fitted to the historical data, and embed them directly into our model as the distributions for the two exogenous variables $\tilde{R}$ and $O\tilde{I}_F$.

We solve the model under all three distributional assumptions, selecting a single set of remaining parameters to match specific empirical targets in the data. We set $C_0 = 16.3$ to reflect that storage at Cushing, reported for April 17, 2020, was approximately 59.7 million barrels, compared to a total working storage capacity of 76 million barrels given by the EIA. We normalize the long-term fundamental $\hat{P} = 1$ and then choose the remaining three parameters so that, when solving the model using the non-parametrically estimated distributions, we obtain (1) an unconditional average midstream supply (determined primarily by ϕ_0) equal to the empirically observed average physical demand of approximately 30,000 contracts (or 30 M barrels); (2) an unconditional probability of negative prices (determined primarily by ϕ_1) equal to 1%; and (3) a price of oil given the empirical realizations of financial and physical demand for the May 2020 expiry equal to -4 (determined primarily by τ to reflect the negative price realization on April 20 of approximately four times the magnitude of the previous day's price). The 1% probability of negative prices is somewhat arbitrary, chosen to represent a low but non-trivial likelihood of such an event; alternative choices for this value yield only minor changes in the model's implications. This calibration yields parameter values for midstream operators of $\phi_0 = -0.525$, $\phi_1 = 0.0495$, and a storage cost parameter of $\tau = 1.98$.

Solving for equilibrium in the model, given a realization of refinery capacity and financial open interest, involves determining the endogenous amount of production so that Equation (3) holds using the definition of P_1 from Equations (1) and (2). The probabilities used in the expectation are derived from the conditional probability distribution of $\tilde{R}$ given the observed level OI_{Total}, which can be computed from the distributions of $\tilde{R}$ and $\tilde{OI_F}$. Since the price P_1 is either flat or decreasing in production for all values of $\tilde{R}$ and $Z \geq 0$, the equilibrium is unique. Furthermore, since optimal supply (Z^*) is a monotonically increasing function of expected price, observing either supply or total futures demand ($\tilde{OI_F} + \tilde{R}$) is sufficient to determine the equilibrium. Therefore, observing total open interest at time zero fully reveals the sum of physical and financial demand. Observing the futures price at time zero would also suffice to achieve equilibrium, as in Grossman and Stiglitz (1980).

2.2.5 Model results

Figure 7 plots various equilibrium outcomes in the model as a function of the relevant state variable, namely total demand from financial and physical traders ($\tilde{OI}_F + \tilde{R}$). Panel A shows the expected physical demand as a function of total demand, highlighting how financial demand can provide a misleading signal of physical demand. Panel B shows the futures price, equal to the expected time-1 spot price, at time 0, which directly determines midstream operator supply shown in Panel C. As total demand increases, the probability of exceeding available storage capacity declines, and midstream operators increase their supply to meet this demand, effectively responding to the perceived signal. Panel D shows that observable open interest is a monotonically increasing function of total demand, with a kink at the point where midstream operator supply equals the sum of financial and physical demand. Below this point, open interest is equal to midstream supply and above it equal to the sum of financial and physical demand. In this calibration the latter case is far more likely.

Panel E displays the expected change in storage as a function of physical demand, directly relating to our empirical finding that open interest in the expiring contract predicts changes in storage at the hub. Comparing Panels D and E, we observe that open interest and the expected change in storage are approximately linear in total demand, making open interest an excellent empirical predictor of expected storage change. By contrast, while the futures price is also monotonically related to expected changes in storage, the convex nature of the price function produces a highly non-linear relationship as shown in Panel B. In particular, large regions of the state space feature essentially flat futures prices. This fact provides a potential explanation for our empirical finding that open interest predicts storage changes even when futures prices are included as a control.

Lastly, Panel F shows the probability of negative prices conditional on total demand, illustrating how such negative prices can arrive as a sudden surprise even when their possibility is recognized ahead of time. In this calibration, negative prices occur with a one percent unconditional probability, but are conditionally likely only in regions with low total open interest, which typically correspond to low demand. Conditional on high open interest, such as the 108,000 contracts observed for the May 2020 contract, the probability of negative prices is essentially zero. This highlights the extremely low ex ante likelihood of the seven standard deviation outlier in financial open interest observed for the May 2020 contract.

Figure 8 provides further context for the occurrence of negative prices in the model, and highlights the role of financial open interest in creating the sudden surprise on April 20th, 2020. Panel A shows the joint probability density function (PDF) of financial and physical demand obtained from the KDE, along with the region of this joint distribution corresponding to the calibrated one percent unconditional probability of negative prices. As Panel A illustrates, negative prices arise only in regions of low physical demand, with high financial open interest increasing the likelihood of such outcomes. The red star marks the realized financial and physical demand for the May 2020 contract, which lies within the negative price region. While the model is stylized and should therefore be interpreted cautiously in quantitative terms, Panel A suggests that a more typical level of financial trader demand would likely have avoided negative prices, as it would have supported a lower futures price and thus reduced midstream operator supply. In this sense, the unexpected large financial demand appears to be a contributing driver of the negative price event.

Looking at Panel A however, it is clear that high financial open interest is not a necessary condition for negative prices, as they can also arise with low financial demand if physical demand is sufficiently low. Nevertheless, as Panel B demonstrates, instances of negative prices driven solely by low physical demand are unlikely to produce the type of "surprise" negative prices observed in April 2020. In Panel B, we outline the region in which negative prices occur despite the time 0 futures price being very close to the long-run fundamental (defined here as $F > 0.99$), indicating a very low probability of negative prices conditional on the observed open interest.[19] The figure shows that unusually high financial open interest is indeed a necessary condition for such an outcome. The red star, marking the observed financial and physical long demand observed prior to the May 2020 contract expiry, lies within this region, underscoring that the highly unusual combination of extremely high financial open interest and low physical demand created precisely the conditions for a sudden collapse

[19]Recall: $F = E[P_1|OI_{Total}]$, so $F > 0.99$ implies a very low conditional probability of negative prices.

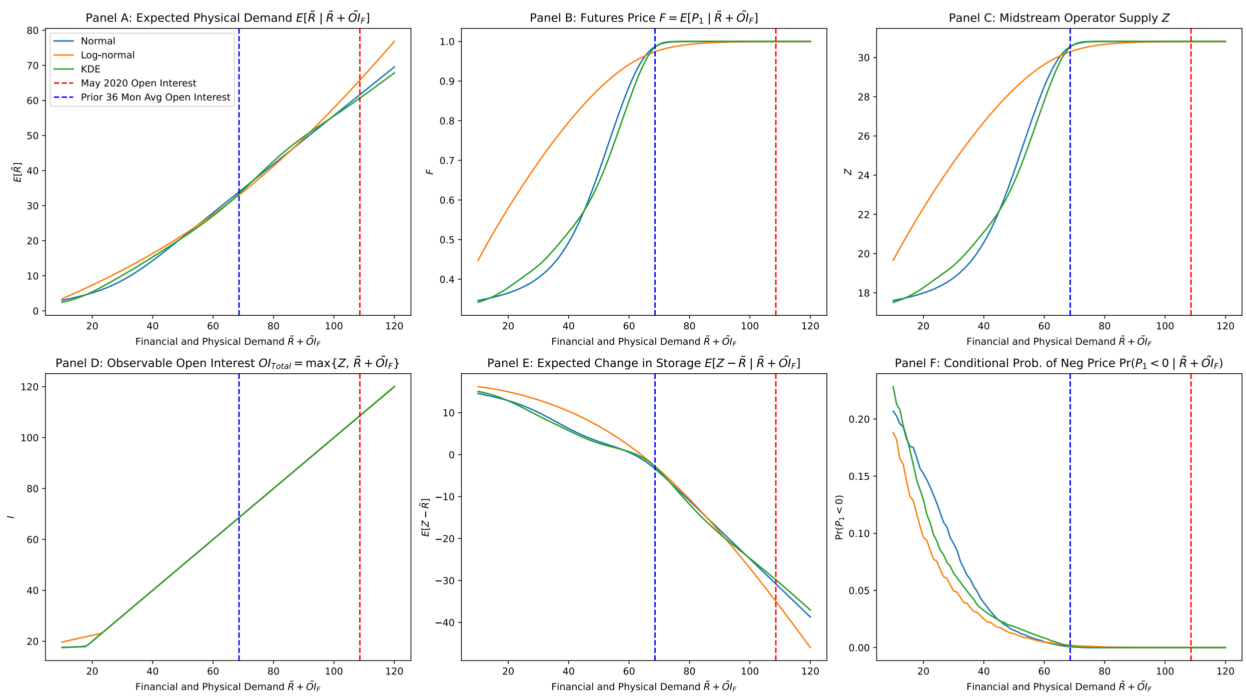

Figure 7: Model Outcomes

The figure plots various model outcomes as a function of the total demand for futures contracts, defined as the sum of financial and physical demand, $\widetilde{OI}_F+\tilde{R}$, measured in thousands of contracts (equivalently, millions of barrels). The model is solved under each of the three distributional assumptions for financial and physical demand described in Figure 6. Available storage capacity is set at C_0 =16.3 million barrels to match the storage capacity situation prior to the May 2020 contract expiry. The remaining parameters are $\phi_0 = -0.525$, $\phi_1 = 0.0495$, and $\tau = 1.98$. Panel A plots the expected futures price as a function of total demand, Panel B plots the midstream operator supply, Panel C plots total open interest, Panel D plots expected physical demand conditional on the total demand, Panel E plots the expected change in storage conditional on total demand, and Panel F plots the conditional probabilities of negative prices conditional on total demand. The unconditional probability of negative prices is calibrated to be 1%. The vertical blue line denotes the average observed open interest for expiring contracts over the prior 36 months (68 thousand contracts), while the red line denotes the observed open interest in the May 2020 contract.

in prices.

Finally Panel C illustrates how the unconditional probability of negative prices depends on available capacity. To construct this plot, we hold all model parameters fixed and vary the available storage capacity at time 0 (C_0) and compute the corresponding unconditional probability of a negative price event. The results show that negative prices are only possible when available capacity is scarce and completely disappear at more typical capacity levels. This again helps to rationalize the events of April 2020. As capacity diminished, the prospect of negative prices became more salient, prompting the CME to clarify how the exchange would handle such outcomes. However, the high observed open interest in the May 2020 contract led the market to believe there was sufficient physical demand to avert this scenario. The true lack of physical demand only became apparent when financial investors sought to close their positions before having to take delivery, triggering the sharp price collapse.

To provide a final perspective on the model's predictions, we simulate 100,000 draws of financial and physical long open interest and record the spot price and storage utilization at time 1, along with the futures price and open interest at time 0. Figure 9 presents the results. Panel A plots total observed open interest in the futures market against changes in storage in the spot market. Since high open interest typically signals a large amount of physical demand, it generally leads to a subsequent drop in storage levels, consistent with the empirical pattern documented in Panel C of Figure 4. In nearly all simulations (99% of the time in our calibration), negative prices do not occur (blue dots). When refinery demand is low, oil inventories accumulate, creating the potential for negative prices. In most cases, lower total open interest signals weak demand, market participants anticipate the risk of storage capacity being exceeded, and futures prices fall below fundamentals (the yellow dots). By contrast, when weak demand coincides with high long open interest from financial traders, the signal from elevated total open interest implies a low conditional probability of exceeding capacity and futures prices remain near the fundamental value. This, in turn, induces greater supply from midstream operators than would arise in the absence of financial demand. This dynamic can unexpectedly push prices into negative territory (the red dots). The scarcity of red dots highlights how rare this outcome is. Panel B shows how these different scenarios

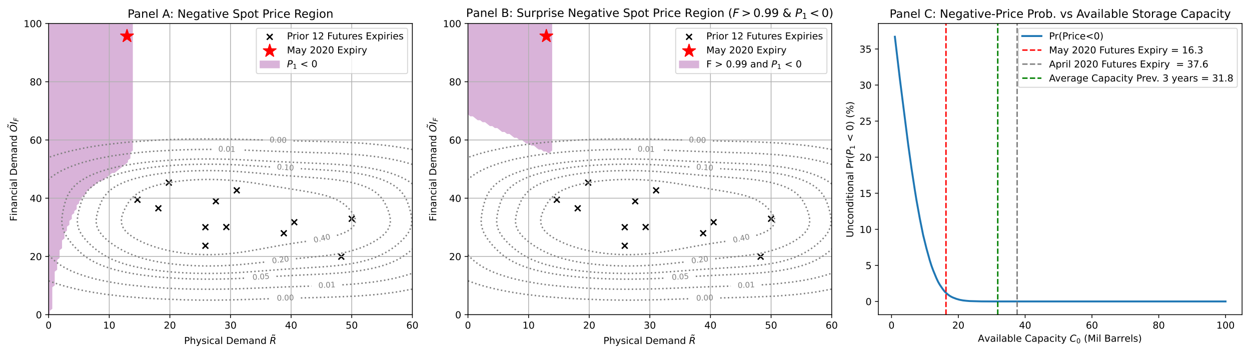

Figure 8: Negative Price Regions in the Model

The figure highlights the parameters and realizations of exogenous variables that lead to negative prices in the model, solved using Kernel Density Estimator (KDE) distributions for the exogenous physical and financial demand as described in Figure 6. Panel A assumes available storage capacity of 16.3 M barrels (the approximate available capacity prior to the May 2020 contract expiry), and highlights the region in which the joint realization of the two exogenous variables, Physical Demand ($\tilde{R}$) and Financial Demand ($\tilde{OI}_F$) generate negative prices under the various distributional assumptions for the two variables. The plot also marks the historical realizations for Financial and Physical Demand for the 12 months preceding May 2020, as well as the realization for the May 2020 expiry itself (see Figure 6for descriptions), and overlays the fitted joint PDF from the KDE. The model is calibrated so that the unconditional probability of negative prices (i.e. the unconditional probability of a joint realization in the negative-price region) under the KDE distributional assumption equals 1% at the given available capacity of 16.3 M barrels. Panel B replicates Panel A but identifies the region corresponding to "surprise" negative prices, defined as negative spot prices in period 1 following a futures price at time zero that is greater than 0.99. Panel C plots the unconditional probability of negative prices changes as a function of available storage capacity. To create the plot, all model parameters are held fixed, while initial available capacity C_0 is varied and the unconditional probability of negative prices is calculated for each value. The vertical lines shows the capacity levels prior to the May 2020 expiry (red line), the April 2020 expiry (grey line) and the average available capacity over the previous 12 months (green line).

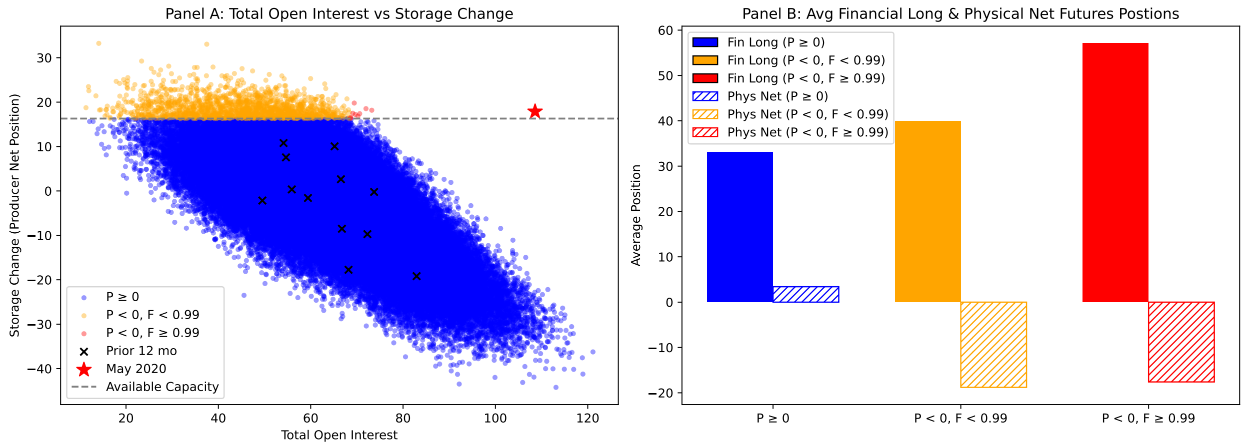

Figure 9: Simulated model outcomes

The figure presents results from 100,000 simulations of the model. Panel A plots open interest in the futures market at time 0 against subsequent storage changes at time 1 $(\tilde{R} - Z)$, segmented into three different regimes. Blue dots represent simulations with positive spot prices, red dots indicate "surprise" negative prices, defined as a negative spot price following a time 0 futures price greater than 0.99. Yellow dots indicate simulations with negative time 1 spot prices following a futures price lower than 0.99 at time 0, reflecting that negative prices at time 1 were considered more likely as of time 0 given observed open interest. Panel A also marks the storage outcome predicted by the model given the realized financial and producer open interest at the May 2020 contract expiry, as well as for the previous 12 expiring futures contract reported in the CFTC Interim Report (see Figure 6 for details). Panel B reports the median values of financial long open interest $\tilde{OI}_F$ and the net position of physical traders $(\tilde{R} - Z)$ in the futures market for each of the three outcome regimes shown in Panel A.

(blue, orange, and red) are reflected in the unobservable long positions of financial traders and the unobservable net positions of physical traders $(\tilde{R} - Z)$. In simulations that produce a surprise negative price (in red), the financial open interest is high enough to mask weak refinery demand, resulting in large negative net physical positions corresponding to excess supply. This pattern closely mirrors the data in the CFTC Interim Staff Report, as shown in Panel B of Figure 2.

In sum, the stylized model above suggests that low levels of available capacity create initial concerns about storage overflows and negative prices. However, the unusually high level of financial open interest just prior to the expiry of the May 2020 futures contract led the market to mistakenly infer substantial demand from physical traders, implying a reduced probability of overflows and a higher futures price. This misperception prompted midstream operators to send more oil to the hub, having pre-sold significant volumes in the futures market to financial traders who were unable to take delivery. The reality became clear on April 20th, when these financial traders sought to close their positions as they could not take physical delivery, triggering a sudden collapse in spot prices as the marginal cost of storing oil surged to extremely high levels. Despite its simplicity, our two-period model captures this mechanism and generates results consistent with the observed empirical patterns in open interest and storage in the WTI market before and during the negative price event. We now turn to the impact of this event on producer decisions.

3 Consequences of negative oil prices

In this section, we examine how oil well production decisions were influenced by the benchmark risk associated with expiring futures contracts. While our theoretical model outlines a basic mechanism through which negative prices can arise, it likely understates the broader uncertainty generated by the unprecedented events of April 20, 2020. Market participants were forced to interpret a pricing outcome that had never occurred before in U.S. oil markets, introducing substantial ambiguity into forward-looking decisions. Figure 10 illustrates this uncertainty: prices and trading volumes of put options with strike prices near zero spiked immediately after the negative WTI settlement and remained elevated well into May. These levels were unprecedented, and, notably, were not observed for Brent Crude options, despite WTI and Brent trading at similar spot prices during this period. The divergence highlights the benchmark risk associated with WTI following the April shock. Our empirical analysis of production responses focuses on this benchmark-based uncertainty, particularly surrounding the expiry of the June 2020 WTI contract in May 2020.

To cleanly identify the effects of benchmark risk as distinct from broader macroeconomic uncertainty associated with the COVID-19 pandemic, we implement two complementary empirical strategies. First, we exploit how benchmark risk interacts with the Calendar Month Average (CMA) structure of crude oil purchase contracts, along with high-frequency proxies

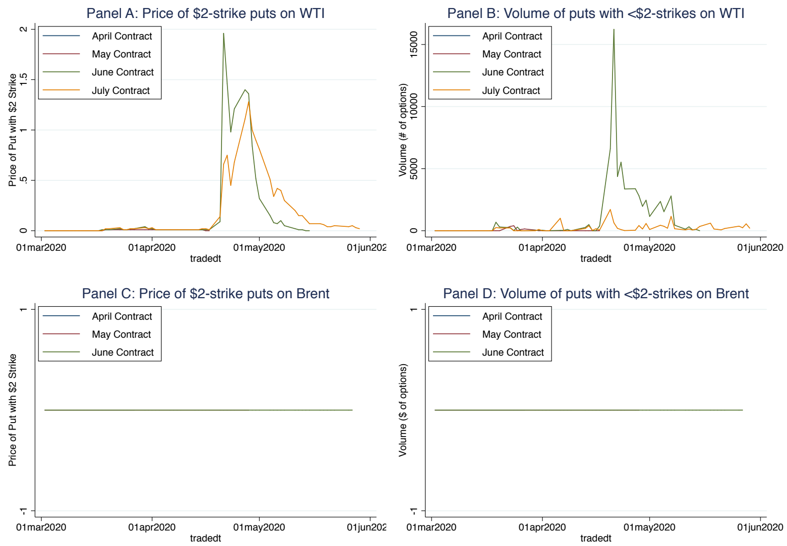

Figure 10: Near-zero-strike put option prices and volumes in April and May 2020

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.

for daily production. Second, we compare production decisions in North Dakota, a region exposed to WTI pricing, to those in Alberta, Canada, where producers are not benchmarked to WTI. To motivate our empirical design, we begin by providing relevant institutional background.

3.1 Institutional background

Several features of the physical crude oil market are essential for understanding our empirical strategy. We outline these features in detail below.

3.1.1 Benchmark prices for different regions

Firms known as midstream operators pay crude oil producers to purchase oil at the wellhead. A typical transaction involves the purchaser driving a company-owned truck to the well site, filling it from on-site storage tanks, and then transporting the oil to a pipeline or terminal.

From there, the oil enters the broader distribution network and is ultimately refined. In practice, crude oil is transacted at tens of thousands of locations each day across North America. This decentralized system raises a core pricing challenge: how to set the transacting price at the wellhead. As Duffie and Stein (2015) discuss, in the absence of a widely accepted benchmark, buyers and sellers may disagree on the fair market value of crude at the point of sale. To address this issue, market participants typically rely on pricing mechanisms tied to traded benchmark indices. According to Duffie and Stein, the use of a benchmark provides significant informational advantages, including "lower search costs, higher market participation, better matching efficiency, and lower moral hazard in delegated execution." To obtain these benefits, firms often substitute their best-fit-for-purpose transaction with a trade tied to a liquid benchmark.

The primary benchmark for crude oil in North America is the West Texas Intermediate (WTI) futures contract for delivery in Cushing, OK. As a result, most daily posted prices for crude in other U.S. regions are not based on actual local "spot" transactions, but rather on the WTI closing price, adjusted by a location-specific and infrequently updated basis differential to account for differences in quality or transportation cost. To illustrate this mechanism, Figure 11 reports the calendar month postings for crude grades purchased by the refiner Phillips 66. Notably, 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 20, 2020. A closer look reveals why: the reference prices at these locations were derived mechanically from the WTI benchmark, applying a fixed differential (e.g., –\$1.25 in the case of LLS). These fixed basis formulas transmitted the WTI benchmark collapse directly across all Central U.S.[20]

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

This type of benchmarking implies that oil sold in other parts of the country is exposed to fluctuations in WTI prices, even when those price movements are driven by events specific to

[20]In response to this episode, two major price reporting agencies (S&P Platts and Argus) have since introduced new regional benchmark prices that do not rely on landlocked WTI pricing. See, for instance: https://www.reuters.com/article/us-usa-oil-prices-idUSKBN23W3CS/.

Phillips 66 Crude Oil Prices for Apr-2020

in dollars ($)

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 Adjustment		A	A	B	H	A	A	C	C	C	A	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
2020-083	4/30/2020	15.46	15.46	12.65	15.02	14.96	14.96	15.26	15.46	15.46	14.21	12.55

Figure 11: Daily posted prices for different crude benchmarks in April 2020

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

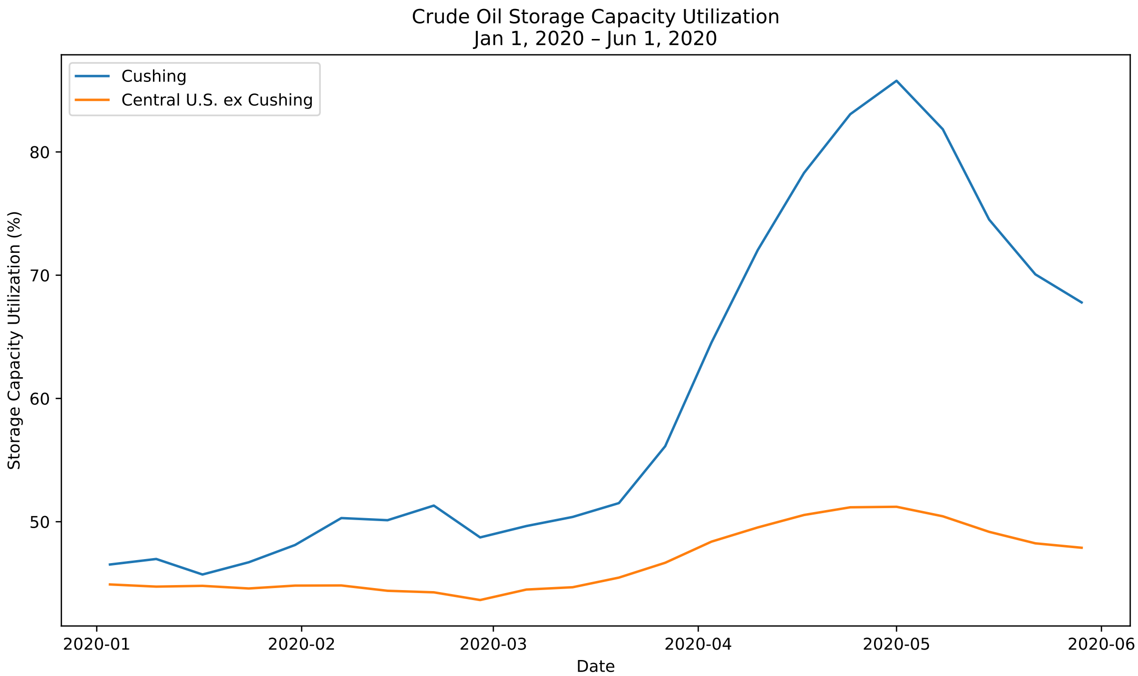

Figure 12: Storage capacity utilization across U.S. regions in 2020
This figure shows storage capacity utilization for Cushing (OK), as well as for the Central U.S. (PADD 2 and PADD3) excluding Cushing for the period from January 1, 2020 to June 1, 2020. Data for stocks and working capacity are from the Energy Information Administration (EIA).

the Cushing delivery point in Oklahoma. One might argue that the observed price behavior simply reflected broader market fundamentals, with storage constraints occurring throughout the Central U.S., not just in Cushing. However, the data do not support this interpretation. Figure 12 shows storage capacity utilization in Cushing alongside the rest of the Central U.S. As the figure indicates, the spike in storage utilization was specific to Cushing, while storage availability remained ample elsewhere in the Central U.S. Despite this ample storage, producers in areas like North Dakota effectively received negative prices, a direct result of the WTI-based benchmarks embedded in their contracts.

To better understand the potential impact of benchmark pricing risk, we analyze production decisions in a region that experienced negative prices despite having no shortage of storage. As shown above, several areas in the Midwest suffered negative pricing despite being geographically distant from Cushing, Oklahoma. We focus our well-level analysis on North

Dakota (ND) for a key reason: it allows for a meaningful comparison to a nearby control region that shares similar fundamentals. The neighboring province of Alberta, Canada, which borders North Dakota to the north serves as an ideal counterfactual. Alberta and North Dakota crude oil share the same end market, predominantly refineries in the upper Midwest, and both feed into the same pipeline network. By excluding heavy oil production from the Alberta sample (i.e, oil sands), we ensure that the wells in both regions are comparable in terms of extraction and product characteristics. However, the two regions differ in a critical respect: Alberta crude is priced off a distinct benchmark tied to prices in Edmonton, Alberta. Figure 13 plots crude oil price movements during April and May of 2020 across these geographies: Bakken prices for North Dakota, Edmonton Par prices for Alberta light oil, and the two major benchmarks for light crude: WTI (North America) and Brent (global). While prices generally move in parallel, the figure shows that the negative price spike was confined to the two U.S. benchmarks, WTI and Bakken. In our empirical design, Alberta production serves as a proxy for regions with low exposure to WTI prices, while North Dakota proxies for regions with high WTI exposure. Again, this approach relies on the implicit assumption that, aside from differences in WTI exposure, technologies and underlying fundamentals are broadly similar across these two geographies (see Gilje and Taillard (2017)for a detailed justification).

Another similarity across both regions lies in the regulatory response to the sharp decline in oil prices at the onset of the COVID-19 pandemic in 2020. Regulators across North America issued emergency orders permitting well operators to shut in production without facing penalties such as the loss of leasing rights. While the timing of these orders varied across states and provinces, a key commonality is that they took effect before May 1st.[21] This timing is crucial: it ensures that both Canadian and U.S. producers were allowed to reduce output, enabling us to observe the *unconstrained behavior* of both groups in response to the significant increase in benchmark price risk.

[21] North Dakota issued an announcement on March 24, 2020 (link), Oklahoma on April 17, 2020 (link), Texas on May 5, 2020, (link) and Alberta on March 17, 2020 (link).

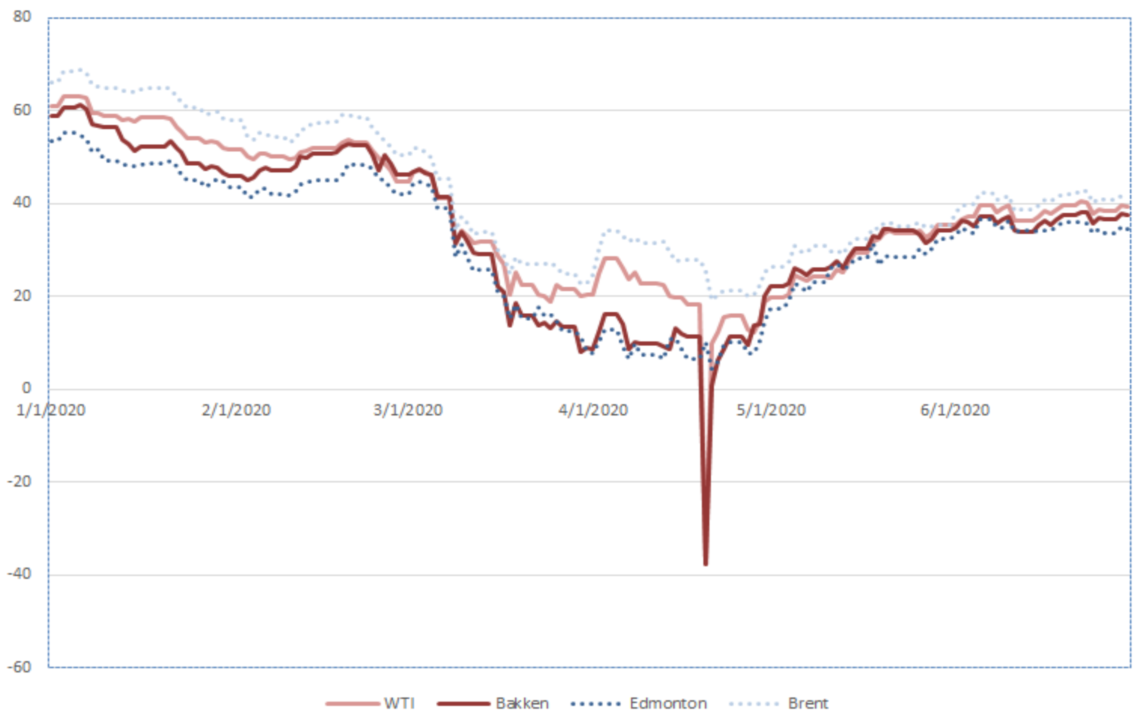

Figure 13: International and domestic crude oil prices around April 2020

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

3.1.3 Calendar month average purchase agreements

Crude oil purchase agreements are typically not based on the spot price of oil on a given day, but rather, on the Calendar Month Average (CMA) of reference prices. As illustrated in Figure 14, this feature implies that oil sold early in the month is still subject to price fluctuations that occur later in the month. Crucially, if crude benchmarks settle at a negative value on a given day, that negative number enters into the monthly average, effectively causing crude sold on that day to be transacted at the negative settlement price. The implications of this pricing structure were emphasized by Harold Hamm, CEO of Continental Resources, in a letter to the CFTC, where he stated that the events of April 20th "materially impact the Calendar Month Average (CMA) pricing of physical crude" (see Figure 14). Further evidence is shown in Figure 11, which reports the CMA postings for crude grades purchased by the refiner Phillips 66. Notably, 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 how producers responded to the perceived risk of a repeat event associated with the June 2020 expiry on May 19th. Under the CMA price structure, the first day of production exposed to this benchmark risk is May 1st, a key feature we exploit in our first empirical test.

3.2 Evidence from a daily proxy for oil production

Our first test of production responses to benchmark risk relies on the structure of the CMA contract. Specifically, we test whether producers adjusted their production decisions in response to the risk that negative prices could occur again around the expiry of the June 2020 WTI futures contract on May 19th. Since any negative price on May 19 would be included in that month's CMA, barrels produced on the first day of May would be exposed to this risk, whereas barrels produced on the last day of April would receive the prior month's CMA and remain unaffected by events occurring on May 19. As such, we hypothesize that producers would begin reducing output on May 1 to avoid this heightened benchmark risk. Conversely, once the May 19 expiry passes, and assuming no major negative price event, benchmark risk

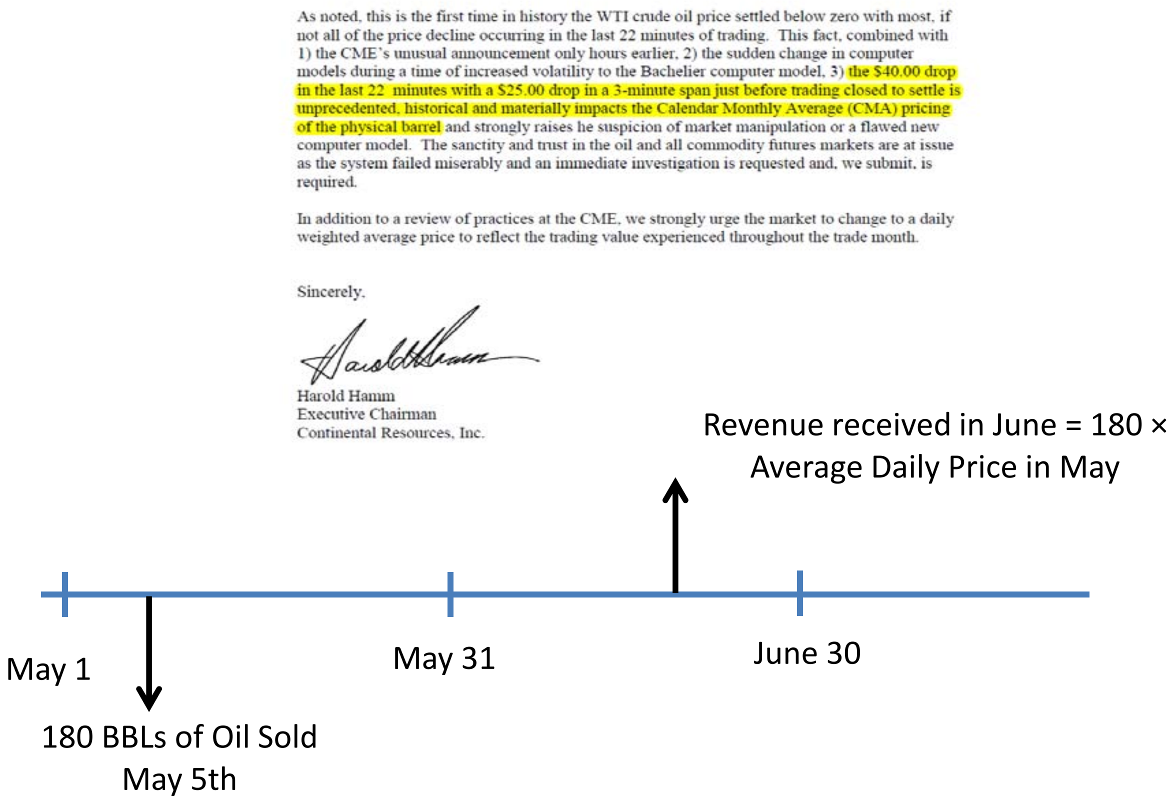

As noted, this is the first time in history the WTI crude oil price settled below zero with most, if not all of the price decline occurring in the last 22 minutes of trading. This fact, combined with 1) the CME's unusual announcement only hours earlier, 2) the sudden change in computer models during a time of increased volatility to the Bachelier computer model, 3) the $40.00 drop in the last 22 minutes with a $25.00 drop in a 3-minute span just before trading closed to settle is unprecedented, historical and materially impacts the Calendar Monthly Average (CMA) pricing of the physical barrel and strongly raises he suspicion of market manipulation or a flawed new computer model. The sanctity and trust in the oil and all commodity futures markets are at issue as the system failed miserably and an immediate investigation is requested and, we submit, is required.

In addition to a review of practices at the CME, we strongly urge the market to change to a daily weighted average price to reflect the trading value experienced throughout the trade month.

Sincerely,

Harold Hamm
Executive Chairman
Continental Resources, Inc.

Figure 14: 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.

would be lessened, potentially prompting a subsequent increase in production under the null.

To test these predictions, we require a high-frequency proxy for oil production, as daily oil output data are not publicly available. We address this challenge by using pipeline flow data for natural gas from Bloomberg, which can serve as a useful proxy because a significant share of U.S. natural gas is produced essentially as a by-product of oil production as it is extracted from the same wells that were drilled and fracked primarily for the purpose of oil production (so called "associated" gas). In regions with commingled production, a drop in natural gas flow is a direct indicator of reduced oil production. To isolate the production response specifically tied to commingled oil and gas wells, we compare the "associated" gas production, that is, natural gas from oil-producing wells in states like North Dakota, Oklahoma, and Texas, to the so-called "dry" gas production, that is, natural gas from non-oil-producing wells in states such as Pennsylvania, West Virginia, and Kentucky.[22]

We plot daily associated gas and dry gas production in Figure 15. Associated gas production drops by approximately 5% from April 30th to May 1st, a four standard deviation day-to-day change. In contrast, dry gas production remains flat over this 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 turn negative on this contract expiry, and we observe that associated gas production begins to rise the following day, consistent with more oil wells being brought back online after prior shut-ins. The sharp production decline on May 1st is especially striking given that crude oil prices changed very little that day, and no similar decline is observed in dry gas states. Absent exposure to benchmark risk embedded in the June 2020 CMA contract, it is difficult to explain such a sharp drop in daily production on that day. We interpret this pattern as compelling evidence that benchmark risk plays a significant role in shaping production decisions. Motivated by this high-frequency evidence, we now turn to our second set of tests using monthly well-level data.

[22]For example, gas production in Pennsylvania from the Marcellus and Utica shale formations is predominantly dry gas, while Texas produces significant associated gas from oil-rich regions like the Permian basin (see: https://www.eia.gov/todayinenergy/detail.php?id=63704).

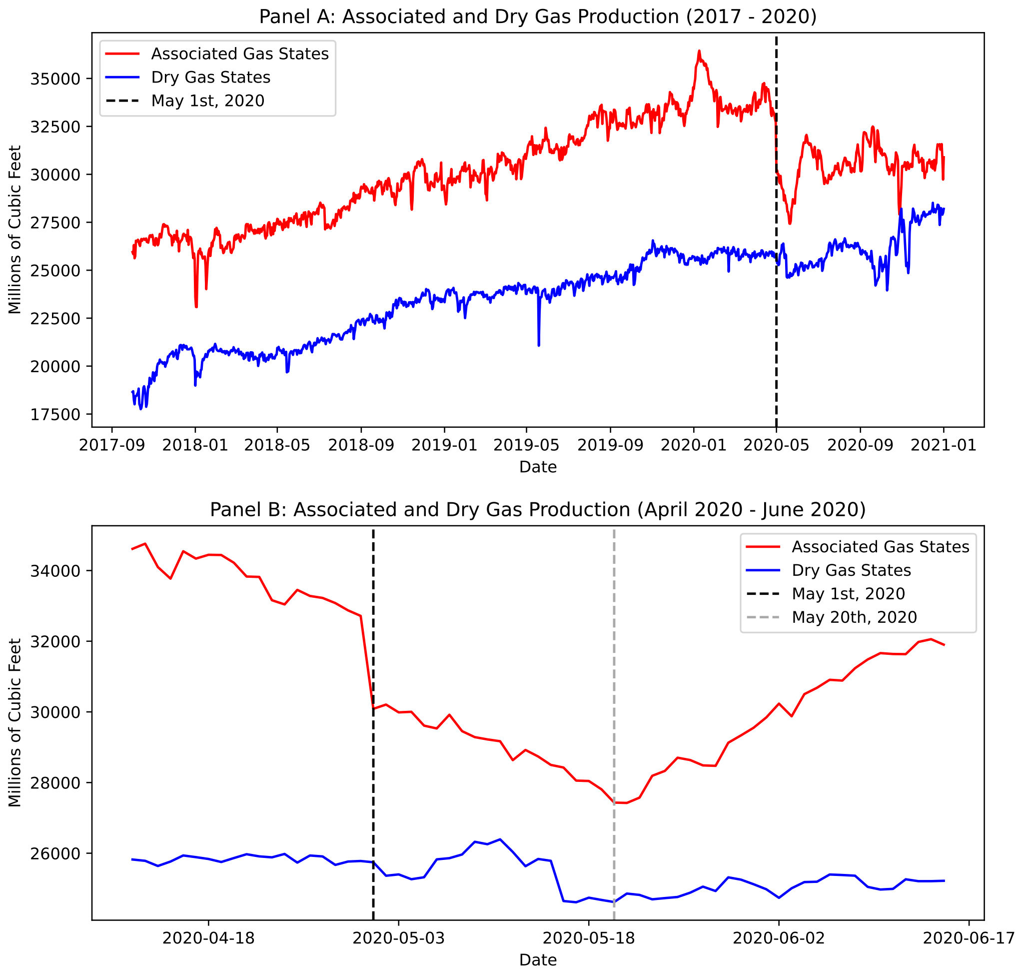

Figure 15: Daily natural gas production for associated and dry gas states

Figure plots daily production from “associated” gas states (ND, TX, and OK) and “dry gas” states (PA, WV, and KY). Associated gas 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 day after the expiry of the June 2020 future.

3.3 Evidence from well shut-ins: North Dakota versus Alberta

In this section, we estimate how oil well shut-in decisions relate to having crude oil purchase agreements indexed to WTI. Building on the rationale provided in Section 3.1.2, we compare production decisions in North Dakota to those made in Alberta. In this context, wells in North Dakota serve as the treatment group (i.e., high WTI exposure) and those in Alberta serve as a control group (i.e., low WTI exposure), in what is effectively a difference-in-difference approach. Motivated by the high frequency results presented above, we screen for wells that are producing in April, cut back production for at least part of May, and then returned to full production shortly thereafter.[23] More specifically, we define a well as experiencing a "temporary shut-in" in a given month if it meets the following criteria: it produces more than three-fourths of available days in the previous month, fewer than one-third of available days in the given month, and returns to producing on more than three-fourths of available days within two months. For instance, a well that reduced output in May 2020 to fewer than one-third of available days would need to return to above three-fourths of days by July 2020 to qualify as temporarily shut-in.[24] We calculate the share (percentage) of wells meeting this temporary shut-in definition in both Alberta and North Dakota for each month from the beginning of 2019 to the end of 2022.

Figure 16 presents the results of this analysis. Panel A plots daily oil prices for both the Edmonton Mixed Sweet (EMS) benchmark in Alberta and the WTI benchmark. Panel B shows the percentage of wells with temporary shut-ins in both North Dakota and Alberta, while Panel C displays the difference between the two series in Panel B. Panel A shows that benchmark prices for both regions track quite closely over the sample period, with the notable exception of the extreme price dislocation seen only in the WTI benchmark on April 20, 2020. Both benchmarks reach their lowest level in April 2020. Panel B shows that, prior to April 2020, temporary shut-ins were relatively stable and infrequent in nature (affecting

[23]As discussed in Section 3.1, both North Dakota and Alberta producers received regulatory approval to shut in production prior to May 1st.

[24]We use the one-third threshold to include wells that resumed production shortly after the May 19th futures contract roll, and the two-month window is chosen because, as shown in Figure 15, production continues to rise into the second half of June. Our results are quantitatively similar when using variations in the shut-in definition, such as using a one-month window or different cutoffs.

approximately 0.5% of wells each month), likely reflecting routine maintenance activity. As prices bottomed out in April 2020, shut-ins rose to approximately 2% in both regions, possibly indicating opportunistic maintenance-related downtime during a weak price environment. However, a stark divergence emerges in May 2020. While Alberta maintains a similar rate of shut-ins, North Dakota experiences a sharp spike, with nearly 8% of wells temporarily shut in that month. Panel C highlights this divergence, with May 2020 standing out as an extreme outlier approximately 20 standard deviations from the mean difference computed over the other months of the sample period (i.e. an OLS t-stat of 20 for a non-zero event-month effect with a p-value approaching zero). What makes this particularly striking is that this behavior occurs in May, when prices were rebounding and were significantly higher than in April. Moreover, the spike is confined to North Dakota, a region otherwise comparable to Alberta in most respects except for the pricing benchmark used in crude oil contracts. This analysis provides further support for our hypothesis that production decisions were influenced by benchmark risk associated with the June 2020 futures contract, which expired in late May. The results suggest that in the absence of regulatory or infrastructure constraints, U.S. operators curtailed production preemptively, in contrast to Canadian operators. This behavior is consistent with a risk management response to the uncertainty introduced by calendar month average (CMA) pricing and the possibility of another negative price event at the next contract expiry.

3.4 Long-run productivity of shut-in wells

As Figures 15 and 16 illustrate, production in May 2020 was substantially affected by the heightened benchmark risk associated with the futures contract expiry. One natural follow-up question is: How costly was this behavior for producers? One might expect producers to hedge this risk using financial instruments, e.g., by purchasing put options, particularly if shutting in wells carries long-term consequences for productivity. Alternatively, if temporary shut-ins do not harm future output, then the decision may represent an intertemporal shift in production from a period of heightened uncertainty to a more stable future and as such,

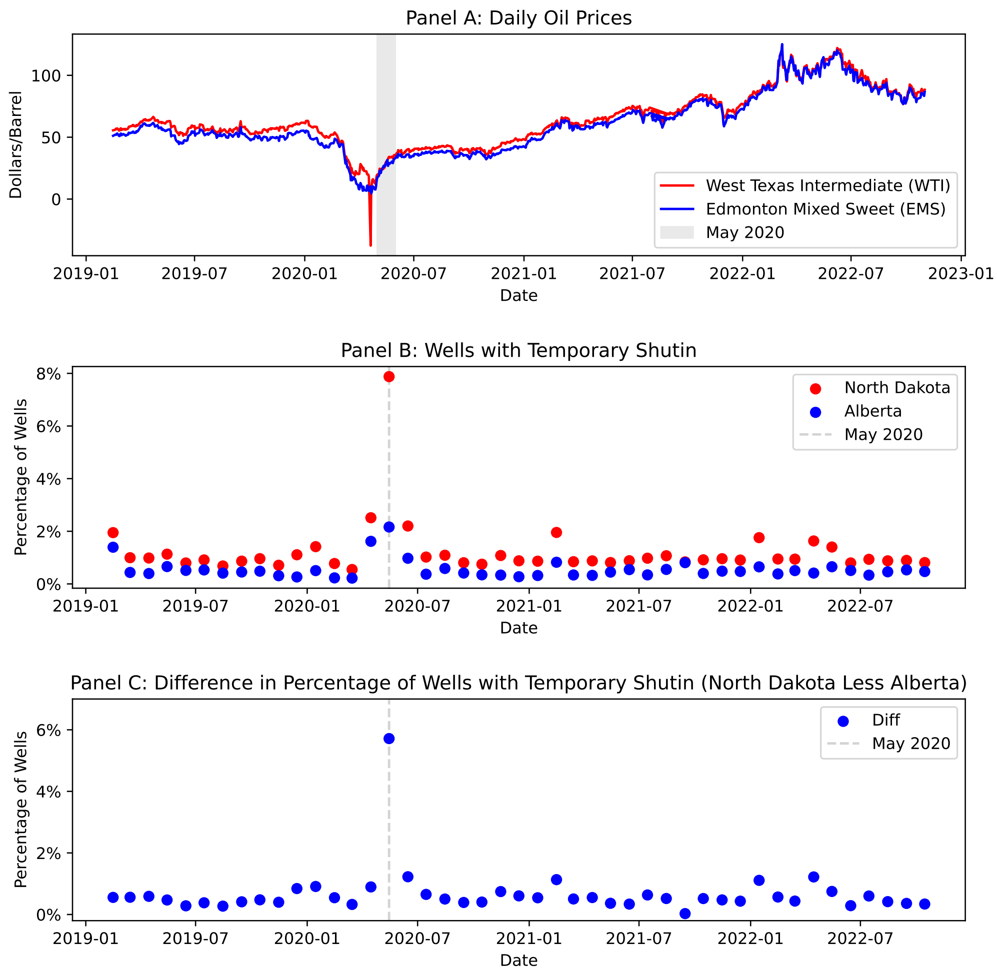

Figure 16: Temporary shut-ins in Alberta and North Dakota (2019 - 2022)

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.

serve as an efficient operational hedge against extreme price volatility.

To explore this question, we examine the long-term productivity of wells in North Dakota that experienced a temporary shut-in in May 2020. We compare these wells to a matched sample of other North Dakota wells that were not shut in during that month, pairing each shut-in well with the one having the closest average monthly production in April 2020. Figure 17 presents the results. By construction, the two groups exhibit nearly identical productivity levels in April 2020, as shown by the overlapping blue and red dots in the figure, which represent the shut-in wells and their matched counterparts, respectively. Shut-in wells exhibit a sharp drop in average production in May, consistent with the shut-in definition, and a partial rebound in June, as some (but not all) return to service. By July 2020, however, the shut-in wells fully resume production, and their average productivity not only recovers to pre-shut-in levels but exceeds them. Over the next two years, the shut-in wells continue to outperform the matched sample, whose output follows a more typical decline curve. By the end of the sample period, the productivity of the two groups converges, and their cumulative production over the full horizon is nearly identical. These findings suggest that the temporary shut-ins did not impose significant long-term productivity costs.[25] Rather, they allowed producers to shift output forward in time as a form of operational hedging in response to the heightened price risk. In sum, while financial distortions around futures expiry clearly influenced short-term behavior, producers were able to respond flexibly, thus mitigating risk without materially sacrificing long-term well performance.

4 Conclusion

This paper explores the causes and consequences of commodity futures market dislocations, focusing on the episode of negative crude oil prices that occurred on April 20th, 2020. We present a simple theoretical model and novel empirical evidence showing that this event was

[25]While we do not observe maintenance or other costs associated with shut-ins, the fact that the average productivity of shut-in wells exceeds their pre-shut-in levels (as of April 2020) suggests that producers may have used the shut-in period to perform maintenance or other well-improving activities. Accordingly, our analysis reflects a comparison of production outcomes, without a full accounting of associated costs.

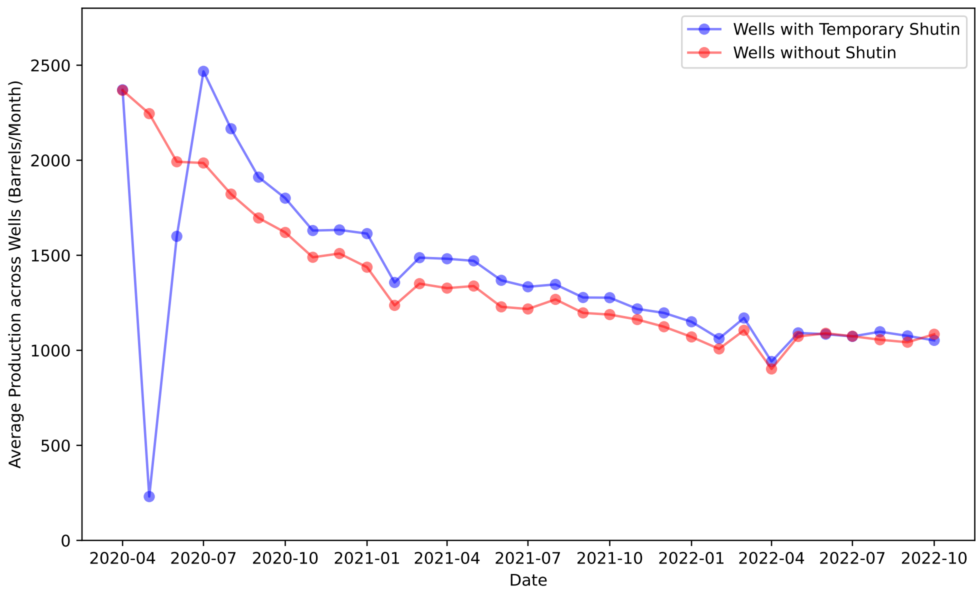

Figure 17: Long-run production from shut-in and open wells in North Dakota

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 of North Dakota wells that were not shut in (producing more than three fourths of days) from April to June of 2020. The matched sample is constructed by selecting, with replacement, oil wells in North Dakota that were not temporarily shut in with monthly productivity that are closest to each shut-in wells in April of 2020.

driven by physical storage constraints in Cushing that were, at least in part, obscured by the large open positions held by financial traders in the expiring futures contract. As these traders closed their positions on April 20, the underlying scarcity of storage was fully revealed, triggering a temporary breakdown in the pricing mechanism. This event led to a situation in which a major oil price benchmark experienced a dislocation from the broader fundamentals of the U.S. oil market. We document that this benchmark dislocation had real effects: oil producers, anticipating a possible repeat of the negative pricing episode in May, preemptively shut in a portion of their wells, even in regions where storage capacity remained ample. Once the benchmark risk subsided, firms resumed normal operations. Overall, our study highlights how dislocations in financial benchmarks can propagate into the real economy.

References

Abrantes-Metz, R., Kraten, M., Metz, A., Seow, G.S., 2012. Libor manipulation? 36, 136–150.

Acharya, V.V., Lochstoer, L.A., Ramadorai, T., 2013. Limits to arbitrage and hedging: Evidence from commodity markets. Journal of Financial Economics 109, 441–465.

Baker, S.D., 2021. The financialization of storable commodities. Management Science 67, 471–499.

Barro, R., 1990. The stock market and investment. Review of Financial Studies 3, 115–31.

Basak, S., Pavlova, A., 2016. A model of financialization of commodities. The Journal of Finance 71, 1511–1556.

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.

Black, F., 1986. Noise. The journal of finance 41, 528–543.

Bohl, M.T., Irwin, S.H., Pütz, A., Sulewski, C., 2023. The impact of financialization on the efficiency of commodity futures markets. Journal of Commodity Markets , Article 100330.

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.

Brogaard, J., Ringgenberg, M.C., Sovich, D., 2019. The economic impact of index investing. The Review of Financial Studies 32, 3461–3499.

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.

De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of political Economy 98, 703–738.

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.

Ederington, L.H., Fernando, C.S., Holland, K.V., Lee, T.K., Linn, S.C., 2021. Dynamics of arbitrage. Journal of Financial and Quantitative Analysis 56, 1350–1380.

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.

Gao, Z., Sockin, M., Xiong, W., 2021. Learning about the neighborhood. The Review of Financial Studies 34, 4323–4372.

Ge, Y., Kang, W., Tang, K., Yan, H., Yang, L., 2022. The road to negative futures prices is paved with financialization. Available at SSRN 4080137 .

Gilje, E.P., Taillard, J.P., 2017. Does hedging affect firm value? evidence from a natural experiment. The Review of Financial Studies 30, 4083–4132.

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.

Grossman, S., 1976. On the efficiency of competitive stock markets where trades have diverse information. The Journal of finance 31, 573–585.

Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. The American economic review 70, 393–408.

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

Hamilton, J.D., Wu, J.C., 2015. Effects of index-fund investing on commodity futures prices. International economic review 56, 187–205.

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.

Irwin, S.H., Sanders, D.R., 2011. Index funds, financialization, and commodity futures markets. Applied Economic Perspectives and Policy 33, 1–31.

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`.

Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. Journal of economic theory 26, 17–27.

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.

Radner, R., 1979. Rational expectations equilibrium: Generic existence and the information revealed by prices. Econometrica: Journal of the Econometric Society , 655–678.

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.

Sanders, D.R., Irwin, S.H., 2010. A speculative bubble in commodity futures prices? cross-sectional evidence. Agricultural Economics 41, 25–32.

S.G. v. Bank of China Ltd, 2023. Civil action no. 23-cv-2866. Class action complaint and jury demand. URL: `https://storage.courtlistener.com/recap/gov.uscourts.nysd.596779/gov.uscourts.nysd.596779.1.0.pdf`. 23-cv-2866 U.S. 483 (2023).

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, 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 focus on North Dakota producers. Our Canadian data is downloaded from Petrinex; it is also at the well-level, by producer, and monthly. Since Canada has substantial production from oil sands, which is a distinct production technology, 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 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 on capacity and stocks for different geographies are collected from the Energy Information Administration (EIA).

A.4 Search volume data

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