UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

THREE ESSAYS ON ENERGY FINANCE

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

HUIMING ZHANG
Norman, Oklahoma
2023

THREE ESSAYS ON ENERGY FINANCE

A DISSERTATION APPROVED FOR THE
MICHAEL F. PRICE COLLEGE OF BUSINESS

BY THE COMMITTEE CONSISTING OF

Dr. Scott C. Linn, Chair

Dr. Chitru S. Fernando

Dr. Xuhui (Nick) Pan

Dr. Aaron Burt

Dr. Pallab Ghosh

Acknowledgement

Firstly, I would like to express my deepest gratitude to my supervisor, Scott Linn, for his unwavering support and invaluable guidance throughout my PhD research and studies. Dr. Linn has been an enlightening star, illuminating my academic path and inspiring me in all aspects of my life. I sincerely appreciate his selfless and patient guidance.

I am immensely grateful to my committee members, Chitru Fernando, with whom I have enjoyed fruitful co-authoring collaborations resulting in several important papers included in this dissertation. I would also like to acknowledge Nick Pan for his unwavering support and the unique and useful information he shared with me. Aaron Burt's insightful advice has been invaluable throughout my research journey, and Palab Ghosh dedicated his time and efforts to enhancing the quality of my dissertation.

A special mention goes to Louis Ederington, whose influence on my academic career cannot be expressed in words. Though I can no longer express my gratitude to him in person, our co-authored paper holds a special place in this dissertation.

I am indebted to my fellow cohorts, Asif Malik, Bing Xu, Summer Liu, Peter Mueller, and Ping Zhang. Your companionship and mutual support have been instrumental in my achievements thus far. Furthermore, I extend heartfelt thanks to my friends and family. To Zirui Liu and Sean Wang, your unwavering belief in me has given me confidence during moments of self-doubt. To my mom, dad, grandma, grandpa, and my little sister, thank you for always standing by my side and supporting me in every decision I have made.

Abstract

The field of energy finance is currently confronted with new challenges in the ever-evolving global landscape, characterized by dynamic changes in fundamental factors and shifting economic conditions. Understanding the information that influences energy product prices continues to be a promising yet challenging aspect within financial research in the 2020s.

This dissertation addresses key challenges in energy finance studies in the current era. It examines the impact of price uncertainty on oil investments, focusing on decision-making under tail risks. Additionally, it investigates the information spread in natural gas markets, including the motivations and accuracy of forecasters. Furthermore, it explores the shift from pit trading to electronic trading in oil markets, analyzing price discovery in the US and Europe. The research aims to provide valuable insights for investors, decision-makers, and market participants in navigating the complexities of energy finance.

Chapter 1 emphasizes the significance of incorporating non-normality assumptions in analyzing real option investments. It assesses the impact of price change tail risks - skewness, kurtosis, and volatility on the likelihood of exercising real options by scrutinizing oil well operators' investment decisions. The results demonstrate a significant influence of these factors, revealing the substantial implications of tail risks on real investments: a one standard deviation increase in skewness escalates the closing probability by 13%, while a comparable rise in kurtosis augments the drilling probability by 11%. In contrast, volatility diminishes these probabilities, with reductions of 44% and 33% in closing and drilling respectively. These robust findings dismiss the theory of personal preferences as an explanatory factor for these impacts.

Chapter 2 undertakes an investigation into the relations between the decision-making attributes, herding and timing choices, and forecast accuracy in the realm of natural gas storage change forecasts. The empirical evidence reveals a prevalence of anti-herding, despite the negative association between anti-herding and accuracy. Forecasts issued later in each forecasting cycle demonstrate higher accuracy. Decision-making with respect to herding and timing contributes to approximately 50% of the cross-sectional variation in accuracy. We discern an enhancement in accuracy with a moderate anti-herding decline since 2013; however, these developments are contingent upon forecasters' average performance or timing choices. The results suggest that forecasters make trade-offs between forecast accuracy with signaling considerations and early dissemination advantages.

Chapter 3 studies the impact on price discovery when commodity futures move to electronic trading, using data from the West Texas Intermediate (WTI) and Brent oil markets. Electronic trading resulted in explosive increases in trading in both futures markets. After the shift, WTI futures became significantly more important for (indeed, the sole contributor to) price discovery but Brent futures became significantly less important. Our findings show that liquidity increases are not necessarily accompanied by price discovery improvements and suggest that the benefits of electronic trading need to be assessed from the standpoint of liquidity improvements as well as price discovery changes.

Table of Contents

List of Tables and Figures

Chapter 1

Skewness and Kurtosis, Real Options, and Investment Under Uncertainty

1. Introduction

Output price change uncertainty and sunk costs play a central role in modern theories of investment choice and the valuation and exercise of real options (Dixit, 1989 & 1991, and Dixit and Pindyck, 1994). These features can, in theory, produce incentives to delay profitable investments and the abandonment of unprofitable investments. A typical assumption is that output price changes have a Normal distribution, but this is not necessarily always true.

A case in point is the price of oil. In recent years, we have witnessed wide swings in oil prices. Oil price changes exhibit high volatility (Ferderer, 1996) but also non-zero skewness and excess kurtosis, which are inconsistent with price changes described by a Gaussian/Normal distribution. However, the Dixit (Dixit, 1989 & 1991) and Dixit and Pindyck (1994) derivations of the relationship between price uncertainty and investment trigger prices relies on the assumption that price uncertainty can be described entirely by volatility. With many existing studies (Dai et al., 2021) demonstrating that there exists non-zero skewness in oil prices, which I confirm along with the presence of excess kurtosis, this calls for a revisiting the empirical relations between investment and production decisions and the moments of the price change distribution, accounting for the higher order moments. That is, does the valuation and optimal exercise of real options depend on skewness and kurtosis in addition to volatility. Instead, an analysis based upon an extended option pricing model is called for.

In theory, option values can be impacted by the skewness and kurtosis of the price change distribution (Jarrow and Rudd, 1982, Corrado and Su, 1996a & 1996b), and Brown and

Robinson (2002). Oil exploration and production is one sector where there can be substantial uncertainty surrounding future oil prices and significant sunk costs of investment and that price changes exhibit non-Normal behavior. Thus the investment and production behavior of firms operating in this sector provides an ideal setting for studying the impact of expected skewness and kurtosis on those decisions.

The primary question I ask and answer is: Are investment and production choices influenced by the expected skewness and kurtosis of the price change distribution as well as volatility? I accomplish this task through the specification and estimation of discrete choice models utilizing a comprehensive sample of over 53 million well-month observations of the micro-level oil well drilling and operation decisions of operators in five major oil producing states. Specifically, I test whether the oil well production decisions are affected by oil price distribution moments - skewness and kurtosis, in this study. This study develops and empirically estimates, discrete choice models of the decision to drill and complete an oil well, to produce oil, to temporarily cease production (mothball) and to abandon (shut-in) that address the specific questions of how and to what extent skewness and kurtosis of the oil price change distribution, estimated with the model-free option-implied risk-neutral central moments using the methods in Bakshi, Madan, and Kapadia (2003), affect the choices we observe in addition to volatility and sunk costs.

The set of decisions I examine in this study are tailor made for addressing the question of whether skewness and kurtosis of the price change distribution influence real investment and operating decisions. The detailed monthly investment and production records that form the basis for the study, the fact that the product produced is identifiable along with the expected output price and price change distribution moments, and the availability of detailed information about the characteristics of the investments, provide the ability to conduct a micro-level investigation of the choices to invest and produce generally not possible. The utility of such data has not gone unnoticed, having formed the basis of the investigations by Kellogg (2014), Gilje (2019), Anderson et al. (2018), and Décaire et al. (2020).

Oil and gas producers tend to expand drilling and completion activity only when prices are very high, and at the same time do not rapidly close or plug wells when prices are below the marginal cost of production. This research is the first study, to the author's knowledge, to investigate the implications of non-normality of output prices on the actual investment and production decisions of firms.

In addition to a number of controls, I specifically examine whether the main results are impacted by 1) the presence of sunk costs, 2) the influence of the productivity of nearby wells, 3) the regulatory environment of the state in which a well is located, and 4) whether the increased intensity of focus on Environmental, Social and Governance issues has impacted the choices observed and the relations between those choices and skewness and kurtosis. The mood created by the 'Environmental, Social, and Governance' (ESG) movement has had what many believe to be a profound impact on business decisions (see for instance the website maintained by the World Bank, https://datatopics.worldbank.org/esg/). No industry has been more impacted than the energy sector, and within that sector oil and natural gas producers. In this study, I also examine how the increased attention to ESG marked with the Paris Climate Agreement impacted the decisions studied.

The data set I examine allows me to investigate the generality of my findings through a study of oil producers' investment and production decisions in multiple states of the United States and geographically different production areas (basins). I use oil well drilling, shutdown, and production records from California, Pennsylvania, Texas, North Dakota, and Oklahoma to study the effects of price uncertainty on investment and production activity. While a company's quarterly or yearly capital expenditures provide a glimpse of capital spending decisions, a deeper understanding requires disaggregated micro-level data observed at a higher frequency. There are multiple benefits to using these data to explore the impact of uncertainty on investment choices. First, the data are measured monthly and provide information on new investment (drilling), disinvestment (abandonment) choices, and shut-in (producing or mothballing) choices. Second, a primary uncertainty these companies face is price uncertainty reflected in the distributional characteristics of oil prices. Third, the data set examined allows a comparison of choices across different geographical locations and states. Fourth, the data allow an investigation of whether public companies behave differently than private companies. Fifth, the data allow me to examine the impact of an important shift brought on by the increased focus on 'Environmental, Social and Governance' (ESG) issues triggered by the adoption of the Paris Climate Accord in 2015. The data set examined includes upwards of 53 million observations that span ten years and multiple geographic areas and states of the United States.

My analysis involves the estimation of multiple panel choice models. The statistical approaches involve the estimation of Cox proportional hazard models structured to investigate the choice to drill and shut-down (abandon) and dynamic panel probit models to investigate the choice to continue producing or temporarily stop production. Implied moments of the oil price change distribution are central to the investigation and are model-free estimates computed from options on oil futures prices following the methods developed by Bakshi et al. (2003). My analysis reveals that higher-order moments (expected skewness and kurtosis of the price change distribution) are important determinants of oil producers' shut-in, investment, and disinvestment decisions for oil wells in California, North Dakota, Oklahoma, Pennsylvania, and Texas from 2010 to 2019. The results suggest that the option-implied higher-order moments proxy for jump risks or crash risks in output prices (oil futures prices) and leptokurtic risk. These characteristics have a significant impact on the propensity to choose to switch between "producing" or "mothballing". These characteristics also are significant in explaining the choice to shut down a well permanently. The expected price and future price change volatility are significant determinants of each of the decisions evaluated and their impact on the probability of choice is in the direction predicted by models of value maximizing choice. My results show that when option-implied oil price skewness increases by one standard deviation, the propensity to exercise a real option for irreversible investment decisions (reopen a mothballed oil well) increases by 1.75%. When implied kurtosis in oil prices increases by one standard deviation, investing likelihood increases by 5.38%. The propensity to immediately close an oil well increase by 13.10% when skewness increases by one standard deviation and the propensity to drill a new well increases by 11.41% when kurtosis increases by one standard deviation. Finally, the results I present are more consistent with decision makers following value maximizing decision rules than imposing their own person investor preferences for skewness and kurtosis.

In Section 2, I review the literature on real option valuation (ROV) and ROV’s applications in the case of irreversible investment decisions. An economic model of investment choice is presented in Section 3 and develops hypotheses on the relationship between real options exercise and implied skewness and kurtosis. In Section 4 an econometric framework is developed for testing the hypotheses and the data are discussed. Section 5 presents empirical results and discussions. Section 6 presents a summary and conclusions.

2. Related Literature

2.1 Investment Under Uncertainty

There are three important characteristics of many real investment decisions. First, investment is wholly or partially irreversible. That is investment involves sunk costs. Second, investment involves risky payoff streams. Third, new information arrives as time passes which can influence expectations about future risky payoffs. And fourth, most investment opportunities do not necessarily vanish if not taken up immediately. That is the investment decision involves not only whether to invest but also when.

The study of real investment activity and the theory of optimal investment choice by firms has a long and rich history (see Caballero, 1999; Girardi, 2021, and the references therein). The foundations of the optimal decision rule for investment choices by firms stem significantly from the Fisher Separation Theorem (Fisher, 1930) and Hirshleifer (1958, 1970), while the study of such choices under uncertainty was significantly impacted by the work of Tobin and Brainard (1976) and Hayashi (1982). The implications of sunk costs and future uncertainty play an important role in what many would call the modern theory of investment under uncertainty which brings to the fore the value of real options embedded in investment and production opportunities.

2.2 Irreversible Investment: Sunk Costs

Sunk costs play an important role in the modern theory of real investment, especially in the presence of future uncertainty regarding investment payoffs due to future output price change uncertainty. Sunk costs of investment have been explored by Bernanke (1983), Pindyck (1990), Dixit and Pindyck (1994), Carruth et al. (2000), Holland et al. (2000), Kogan (2001), Cooper (2006), and Davis and Cairns (2017). Such costs are not reversible or recoverable. This characteristic is crucial when irreversible investments incur high sunk costs—for instance, mining and capital investments. For example, Hausman (1999) discusses the impact of sunk costs in telecommunication services investments. He finds that ignoring sunk costs and irreversibility leads to underinvestment in the telecommunication market. That regulation on price (the TSLRIC Standard) does not consider the critical fact that telecommunication services investments are irreversible and delay innovation in the industry. The two crucial questions regarding irreversible investment decisions are: first, whether to invest or abandon the option; second, when is the optimal time to invest, i.e., the optimal investment timing. Baldwin (1982) and Pindyck (1991) investigate the former problem and show investment rules for irreversible projects. As for the second question on irreversible investment, McDonald and Siegel (1986), Majd and Pindyck (1987), and Dixit (1991) explores the problem of "when to invest" when projects are irreversible. Due to its nature of irreversibility, choosing whether to invest and the optimal time to invest is a crucial issue.

ROV is a valuable tool in solving these questions. ROV (Myers, 1977) solves the optimal timing problem of irreversible investment by taking the investment option as analogous to a financial option: the choice can be exercised at the expense of the investment cost (option premium) to obtain the cash flows from the investment project (option payoffs). Whether exercising the option depends on real options value and project value that jointly determine optimal exercise times. If investing is not feasible to the end of the investment horizon (option maturity), ROV suggests not investing. ROV explains the irreversible investment decisions by showing that it is optimal to exercise the "investment option" when the option value to wait does not exceed the value of immediate exercise rather than when project NPV exceeds zero (Dixit and Pindyck, 1994). I will discuss the real options (hereafter *RO*) exercises in more detail in the following section.

2.3 Price Change Volatility and Investment

The negative relationship between price change volatility and the exercise of a real option to either invest or divest lies at the heart of the developments in the theory of real options (Dixit and Pindyck, 1994). The choice problem is intimately related to the presence of sunk costs of investment. It is a the essential prediction of ROV compared with NPV when investment is not reversible. While NPV does not implicitly indicate the relationship between uncertainty and investment but the price risk and investment (Samis et al., 2005), ROV suggests that uncertainty leads to postponement of investment (Bloom et al., 2007). The underlying well-known principle is that an option's value increases as expected volatility increases, ceteris paribus. ROV predicts that more significant price uncertainty leads to an increase in real options value and postponement of options exercise (Bernanke, 1983 and Abel, et al., 1996). As the value of the option to wait increases with greater forward asset value volatility if the value of the option exceeds the value gained from immediate exercise the exercise will be postponed. Dixit and Pindyck (1994) conclude that more significant uncertainty increases the "waiting option" value and the propensity to postpone investment. In addition, they show that the threshold at which it becomes optimal to invest deviates from the Marshallian solution and increases with price uncertainty.

The relationship between price uncertianty and investment choices has been the theme of numerous studies focusing on other (than oil) industries and more general industry sectors. For instance, Drakos and Konstantinou (2013) present the connection of price uncertainty to investment decisions in the manufacturing industry. Doshi, Kumar, and Yerramilli (2018) show that the tendency to follow the real options exercising rules depends on the size of the firms through their examination of capital investment decisions. They find that larger firms are less likely to follow these rules as they mainly hedge output price risk and can avoid the detrimental price movements – a price uncertainty. Other studies in real options include Pindyck (1993) on the effects of cost and technology uncertainty on real options investment, Aizenman and Marion (1993a, 1993b) on the impacts of policy and macroeconomic uncertainty on investments, and Lensink and Morrissey (2000) who examines the investment-uncertainty relationship in the aggregate level; the studies examine in firm-level include Leahy and Whited (1995) as well as Kang et al. (2014) and Gulen and Ion (2016) both examine the influences of policy uncertainty on investment; Jens (2017) who discuss political uncertainty effects.[1]

The studies on the relationship between price uncertainty and oil project investments reveal that higher uncertainty is associated with the postponement of oil drilling decisions. They show that oil price uncertainty is a critical determinant of whether to invest as well as investment timing for oil projects. Using oil well production data from Oklahoma and Texas, Molls (2001) shows that uncertainty is significantly related to the likelihood of oil wells "producing" status. Kellogg (2014) finds a negative association between uncertainty and the propensity to exercise real options in oil drilling. Décaire et al. (2020) reveal the Peer Exercise Behavior – the propensity for oil producers to decide whether to drill a new oil well is influenced by drilling decisions made in neighboring areas, suggesting that information spillover is important and that this information helps reduce information asymmetry when facing investment projects with sunk costs and price uncertainty.

Dixit and Pindyck (1994) also study the relationship between price uncertainty and the propensity to divest and find that in theory greater output price uncertainty can lead to the postponement of investment but also divestment. Moel and Tufano (2002) find that gold

[1] Gulen and Ion (2016) find evidence of a negative relationship between firm capital investment and policy uncertainty captured by a news-based index.

mining’s entry and exit options exercising behaviors are negatively related to uncertainty. Moel and Tufano (2002) extend the real options literature to the application in exits. That is, price uncertainty does not only postpone investment plans but also delays the exercise of divestment. Investors are reluctant to close an operating project when facing greater uncertainty in product output prices.

The “Bad News Principle” proposed by Bernanke (1983) suggests an asymmetry in the association between uncertainty and the probability of exercising a real option. Dixit and Pindyck (1994) show that more significant price uncertainty increases the optimal trigger price at which a real option will be exercised. The “Bad News Principle” suggests that only the downside risk potential of price matters. A mean-preserving increase in price uncertainty leads to higher expected “bad news” and postponement of a real option’s exercise while increasing the value of the option to wait.

2.4 Skewness, Kurtosis and Investment

While the study of skewness and kurtosis in financial security returns has a long history (Conrad et al., 2013), the literature by and large has ignored the implications of skewness and kurtosis of an asset’s price change distribution on real investment activity, choosing to focus exclusively on either the expected price and or expected volatility. This is despite the fact as described earlier that higher order moments can influence the value of an option and consequently when it is optimal to exercise that option. One exception is the study by Schneider and Spalt (2016), which shows the relationship between project return skewness and investment decisions but not in the frame of ROV. Boyarchenko and Levendorskiĭ (2002) and Boyarchenko (2004) derive the optimal exercising rules of RO with jump-diffusion price process and suggest that their model potentially indicates the skewness-kurtosis and investment relationship.[2] Bernanke's theoretical development of his "Bad News Principal" (Bernanke, 1983) indicates that only adverse price shocks matter in the optimal timing of real investment. Thus, changes in skewness can be an indicator of the probability of "bad news" and result in a lower propensity to exercise a real investment option. Another critical higher central moment is kurtosis which I show in Section 4 can be a determinant of an option's value when the price change distribution is not a Normal distribution. I return to a discussion of option pricing adjusting for skewness and kurtosis in Section 3.

To the best of my knowledge, this is the first research exploring the relationship between irreversible investments real options exercise and higher-order central moments of price probability distributions. This study extends the study on the relationship between uncertainty and investment to higher-order moments of price.

3. Skewness, Kurtosis, and Real Options Investment

Dixit and Pindyck (1994) explains the basic framework for solving the optimal price trigger problem under the Real Options Valuation method, in which investment projects with

---

[2] A separate literature has focused on the implications of skewness and kurtosis of financial security returns. One strand of the literature reveals that the relationship between equity investment and probability distribution skewness and kurtosis depends on forms of investor utilities (Brunnermeier and Parker, 2004). For example, Kraus and Litzenberger (1976) show investors' preference for positive skewness in the equity market, suggesting a non-quadratic form of investor utility. Another strand of studies explains the influences of skewness and kurtosis on investor preferences. For instance, Harvey and Siddique (2000) find that investors require a positive premium for systematic skewness, and Boyer et al. (2010) find a negative return premium for idiosyncratic skewness. While the previously mentioned studies use central return moments, others focus on the risk-neutral moments of returns. For instance, Conrad et al. (2013) reveal that more positive risk-neutral skewness is strongly related to lower subsequent returns. Bali and Murray (2013) show that risk-neutral skewness is negatively related to the equity portfolio returns. Both studies suggest investors' tendency to hold assets with positive skewness. Furthermore, Bali and Murray (2013) show that, although both sides of skewness influence asset returns, left-hand side skewness is remarkably priced into skewness assets by investors. Bali and Murray (2013) point out that exposure to skewness can be asymmetric by decomposing the left- and right-hand sides of skewness risk. Studies on implied higher moments and security price include Datta et al. (2017). Kurtosis is strongly related to investors' preferences and returns premiums. Related studies include Dittmar (2002) focus on the investor's preference on kurtosis, and other studies like Arouri and Nguyen (2010), Diavatopoulos et al. (2012), and Bachmann and Bayer (2014) explore the relationship of investment and both skewness and kurtosis.

large sunk costs are considered to be embedded with the "delay" option to postpone investment decisions to next decision period. This "delay" option enables the postponement of investment decisions that allows better information on price distribution to be collected in the next decision period, possibly achieving greater project outcomes. However, an important underlying assumption in that analysis is that relative price changes are only a function of an expected price change and a normally distributed error, specifically Geometric Brownian Motion:

$$dP/P = \alpha dt + \sigma dz$$

Where α is the drift parameter and σ is the variance parameter, which fully describe the distribution of the price change process. A key result of the analysis is the identification of a 'price' trigger for investment or divestment that is a function of volatility. The Brownian Motion assumption does not deal with the case in which price changes exhibit skewness and kurtosis. Although studies including Boyarchenko and Levendorskiĭ (2002) discussed the solution for trigger investment prices and optimal exercising times with Levy process output prices and suggested the effects of skewness and kurtosis on the solution outcomes, no studies so far, to my best knowledge, considers the explicit relationship between skewness, kurtosis, and real options solutions. That is, the relationship between the higher-order moments of a non-normally distributed output prices and the optimal investment triggering prices has not been revealed. This study aims at solving this problem and filling this research gap.

An unique but rich dataset becomes adopted by more recent studies in the Real Options Valuation. The oil well drilling and closing activities and crude oil prices. The company's capital investment data has a relatively lower observation frequency – most capital investments are reported on a quarterly basis and some are even reported on an annual basis. This low frequency of reporting limits the usage of company's capital investments data in achieving the goal of examining the effects of prices of output products on investment choices. The company-level data however is unable to obtain a direct observation of its output prices. The output product's prices, as well as its risk-neutral central moments, cannot be observed directly in financial markets. Instead, I use a relatively new but useful dataset in the commodity markets – the oil well drilling, production, and closing records and the WTI (West Texas Intermediate) crude oil futures prices and risk-neutral central moments of the price returns. I focus on several major oil-producing states but located in different geometric places over the U.S. – California, North Dakota, Texas, Oklahoma, and Pennsylvania, to analyze the effects of price moments on investment choices. The WTI crude oil is the most actively traded oil contract in the world. Its daily settlement prices and options can be collected over the past four decades. By using this unique dataset, I can observe the investment choices on a monthly basis for about 600,000 individual oil wells from the time span January 2010 to September 2019 with the monthly-average of the daily prices and moments of oil price returns.

Investment under uncertainty is an ongoing debate in financial literature. The Real Options Valuation method has been used in many existing studies to explore the investment under uncertainty problem. In particular, the Real Options Valuation has a unique advantage in its power to explain the choices in investments considering sunk costs. Myers (1977) indicates that Real Options Valuation explains the "conservative" leverages of firms with market investment opportunities real options and whose value is negatively related to risky debt. Aguerrevere (2009) studies product market competition and finds that the relationship explains variations in firm returns among demand, concentration, and growth options. Aguerrevere (2009) suggests the importance of real options in influencing investment decisions under different market conditions. There are other studies on real options, including Décaire et al. (2020), who show the effect of peer exercises on real options, Berger et al. (1996), who explore the exit value of real options, and Mayers (1998) who analyzes convertible bond value with ROV. The Real Options Valuation suggests a negative relationship between price uncertainty and investment tendency – the higher price uncertainty leads to greater value of the "delay" option, leading to lower likelihood to invest. Other studies on the relationship between uncertainty and oil investment include Mohn and Misund (2009) in the aggregate level and Yang et al. (2008) in the micro and firm-level. The studies focusing on the impact of price uncertainty include Elder and Serletis (2010), Yoon and Ratti (2011), Ahmadi et al. (2019), Phan et al. (2019), Cao et al. (2020), Maghyereh and Abdoh (2020), and Doshi et al. (2018). However, some recent studies do not support the negative uncertainty-investment relationship. For instance, Miao and Wang (2007) show that uncertainty is positively associated with investment when the market is incomplete and risk cannot be perfectly hedged. Lambrecht (2017) provides a survey of the literature.

Kellogg (2011) and Covert (2015) examine learning and productivity in well drilling and fracking, respectively. Molls (2001) studies Oklahoma oil well production and finds that sunk costs and historic price change volatility help to explain oil well drilling decisions. Kellogg (2014) studies oil well drilling activity in Texas and concludes that oil well drilling activity is strongly affected by price uncertainty measured as price change volatility. He finds that greater volatility is associated with a larger probability of postponement of drilling activity. Using similar oil wells' drilling data, Décaire et al. (2020) using data similar to Kellogg find that an information spillover effect influences drilling activity. Specifically, drilling in neighboring oil fields encourages oil well drilling in adjacent areas. Those authors however do not control for price uncertainty. Doshi et al. (2018) study quarterly investment expenditures by energy companies and also document an inverse relation between oil price volatility and aggregate quarterly expenditures among smaller firms. Boomhower (2019) and Muehlenbachs (2015) study firms' decisions to either abandon or environmentally remediate wells that are no longer productive. Anderson, Kellogg and Salant (2018) study production and drilling activity in Texas and find that while the expected price influences drilling activity it does not influence production activity. Chen and Linn (2017) study rig activity in the U.S. and find that rig activity in developed economies is positively related to oil futures prices.

Given that option values can be influenced by the skewness and excess kurtosis of the price change distribution of the underlying asset and that oil futures price changes exhibit such characteristics, I reexamine drilling activity in Texas as well as four additional major oil producing states accounting for both volatility as in Kellogg (2014), and the influence of peer activity as in Décaire et al. (2020), but importantly accounting for sunk costs as well as a comprehensive list of additional controls. The drilling choice does not involve a previous choice but does involve a sunk unrecoverable cost. On the other hand, the production choice, that is the selection between producing and temporarily shutting down a well, or permanently shutting down a well, do reflect a previous choice. The state of the well's production status coming into a time *t* will matter, do to both a cost of switching, as well as a change in per period operating cost. I will explore these issues.

Optimal value-maximizing behavior conditioned on the expected price change distribution serves as a benchmark for assessing actual choices. If implicitly or explicitly, decision makers account for higher orders moments of the price change distribution, that is, they recognize the implications of these moments for the value of the real options they possess, then I should expect to see a relation between those moments and investment choices. This benchmark therefore serves as a null hypothesis and allows me to test whether that hypothesis is supported by the actual choice data. In order to set the stage, I begin by formulating the decision maker's problem and its solution conditional on the assumption of value maximization. I then translate that solution into an empirical model of choice that can be estimated with the data I have available. I focus on micro-level decisions at the level of the individual well and on the choices to drill, to produce, to temporarily shut down (mothball) and to permanently shut down. Its output prices are easy to measure with financialized crude oil commodities, oil futures and options, and oil well investment costs tractable with well depth and age. I use oil well production data of Pennsylvania, North Dakota, California, Oklahoma, and Texas and the BKM option-implied crude oil risk-neutral skewness and kurtosis to examine the impacts of changes in higher-order uncertainty on irreversible investments. So far, there is little research shedding light on the tail risks of price uncertainty on irreversible investment in the framework of real options. I present pieces of evidence that there is a strong association between skewness, kurtosis, and investment decisions, regardless of whether investors intend to apply the Real Options Valuation or not.

I set up several different empirical models to examine the impacts: the production status changing options, drilling options, and shutdown options in the oil wells' production history profile. I find significant effects of tail risks uncertainty on real options investment exercise in these models. I show that skewness and kurtosis have substantial power in explaining oil wells' production status and shutdown. The majority of the coefficients are statistically significant at a 1% significance level. When oil wells are not shut down and plugged in, I split producing status into two classes: "producing" and "mothballing." A more negative skewness (and a lower kurtosis) leads to more extraordinary expected "bad news" to increase the in-the-money "delay option" value, resulting in a greater tendency to postpone investments. A higher skewness (and lower kurtosis) results in lower expected "bad news" for in-the-money "delay" put option value, leading to a higher propensity to postpone "shut-in" plan. When focusing on the drilling and shutdown options, I run the regressions of the exercise dummy on skewness and kurtosis to find that the coefficients are significant and consistent with the Real Options Valuation predictions for shutdown options. The results suggest a research gap in the impacts of tail risks uncertainty on real options exercise (irreversible investments). Investors may implicitly apply the Real Options Valuation in making irreversible investments, considering the strong association of uncertainty on decision making.

Another significant contribution of this paper to related literature is the examination of the "Bad News Principle" on real options. If the "Bad News Principle" holds for real options investment, will its conclusion hold for higher-order uncertainty of price? I show that the "Bad News Principle" is valid for the higher-order central moments, the more severe the expected "bad news" (downside potential of price risk) correlates positively with the optimal exercise time and the probability of postponing investment. This finding formally addresses the asymmetry in the effect of price changes on real options' value. I build the relation between expected "bad news" and skewness and kurtosis. I reveal that the relationship between expected "bad news" and optimal timing reflects in the relationship between oil wells' production status and BKM implied skewness and kurtosis, so as the relationship between drilling and shutdown and the central moments.

4. Model: The Choices to Invest, Produce, Temporarily Shut-in and Shut down

4.1 Option Values and Higher Order moments

Theoretically, higher-order expected moments of the price change distribution of the underlying asset can influence the value of options written on that asset. Jarrow and Rudd (1982), Corrado and Su (1996a, 1996b), and Brown and Robinson (2002) develop models of option valuation in which the log price change distribution of the underlying asset may exhibit higher order moments, that is skewness and kurtosis which deviate from the parameters of the Normal distribution. Jarrow and Rudd employ the Edgeworth expansion of an arbitrary probability distribution and Corrado and Su use the Gram-Charlier Type A series expansion in their developments. As these methods essentially lead to equivalent results aside from the organization of terms, I will highlight only the Corrado and Su model.

Corrado and Su (1996a) develop the skewness- and kurtosis-adjusted Black-Scholes option pricing formula based upon the Gram-Charlier Type A series expansion. Subsequently, Brown and Robinson (2002) correct an error in the Corrado and Su model involving a sign in the expansion. The call option price when the log price change distribution of the underlying asset exhibits skewness and kurtosis equals (Brown and Robinson, 2002):

$$C = C_{BS} + \mu_3 \cdot Q_3 + (\mu_4 - 3) \cdot Q_4$$

where μ_3 and μ_4 are the skewness and kurtosis coefficients of the log price change distribution, respectively. The parameters Q_3, and Q_4 are linear functions of the option's moneyness. If skewness equals 0 and kurtosis equals 3, and volatility is positive, the model reduces to the Black-Scholes price. Therefore, for a non-normal price change distribution, skewness and kurtosis in theory will influence an option's value. This result presents the motivation for considering that real investment choices, the exercise of real options may be influenced by not only expected future volatility but also expected future skewness and kurtosis. [3]

I do not directly explore the issue of the underlying process that may give rise to skewness and kurtosis. However, empirical simulations show that a general model that incorporates mean reversion and random jumps, in addition to Brownian motion, produces data that exhibit skewness and kurtosis for reasonable parameters. Such models have been found to fit oil futures price change data.[4]

As the choices I study involve decisions to exercise, or not exercise, an option, it is instructive to consider the interplay of skewness and kurtosis on the optimal exercise strategy within the context of the model described above. Figure 1 presents a graph in which skewness and kurtosis are allowed to vary along two axes and the trigger price at which it is optimal to exercise the option is identified on the third axis. The computations are based upon the following parameters: Current price = 92.00, constant annual risk-free rate of 2%, annual standard deviation in underlying security prices 20%, time to expiration one month, dividend yield 6%, and strike price 92.00. The figure shows that the optimal exercise price of the call option declines as skewness increases. The optimal exercise price does not experience a significant change as kurtosis increases.

[3] See also Borland and Bouchaud (2004) and Potters et al. (1998). The option pricing model with non-zero skewness and excess kurtosis has been shown to perform better than the standard Black-Scholes option pricing model using market-observed option prices (Corrado and Su, 1996b).

[4] In the double exponential jump-diffusion process suggested by Kou (2002), the intensity of jumps on either side changes the higher-order distribution of prices. A greater intensity of jumps in the lower side decreases skewness and any non-zero jump intensity increases the excess kurtosis of distribution. A more recent study by Boyarchenko and Levendorskiĭ (2000) derives the theory for the optimal exercising rule under the Lévy processes of prices which they suggest could solve the real options problems with skewed prices and outliers, as well as Asmussen, Avram, and Pistorius (2004) who derive the put options with Lévy processes. See also Merton (1976), Bates (1997), Zhang (1997), Arnold and Crack (2000), Gukhal (2001), Lewis (2001), Kou and Wang (2004), Levendorskiĭ (2005), and Sepp (2008).

The value of an option as described above under a non-Normal price change distribution serves as my base assumption for the valuation and exercise of real options when the underlying exhibits skewness and kurtosis. To establish the baseline framework for the evaluation of actual observed choices I next present a series of choice models based upon the assumption of value maximization in the presence of real options whose values should in theory be related to the skewness and excess kurtosis of the underlying price change distribution. Optimal choice behavior by producers should therefore be a function of those two parameters in addition to others including volatility. However, whether actual choices are related to skewness and kurtosis is an empirical matter which I evaluate using empirical choice models taking as the basis for their specifications the value maximizing theoretical frameworks.

4.2 Models of Value Maximizing Choices

Oil producers face several decisions which include the decision to drill and complete a well versus delaying investment, whether to abandon (shut-in) a well versus continuing to produce, and whether to mothball (temporarily shut-in) or continue to produce.[5] These choices represent options for the producer. Choices based upon a value maximizing criteria serve as the benchmark for the evaluation of the actual choices selected.

4.3 The Production and Temporary Shut-in (Mothball) Choices

I begin with a discussion of the choices regarding the producing status of a well, under the assumption that the status is chosen by selecting those decisions that are value maximizing at each future date over the expected life of the well. Assume the well has been drilled and completed and from that point on could either be producing or not producing (I will take up the

[5] I abstract from the decision to enter into a lease arrangement that then provides an opportunity to drill, taking the opportunity as a given.

decision to drill later). Assume a risk-neutral expected-payoff maximizing representative producer who must decide whether to make an irreversible investment when the producer's oil well's current status is "mothballing", i.e., production is temporarily shut-in. Define the state of a well as *prod=p* for "producing" and *prod=m* for "mothballing", respectively. The producer may immediately re-open the oil well and re-start production incurring an irreversible (sunk) investment cost, but will then receive the net payoffs from produced crude oil based upon a random future stream of output prices P,operating costs C^p, and production volumes Q. The cost to switch from mothballed to producing is assumed to equal $\mathbb{C}^{m,p}$ and is not recoverable. However, the producer may choose to defer the decision – use the "delay option" to delay the decision to a future date. By waiting the producer can observe an updated price level resolving uncertainty from *t* to to $t+1$.

Similarly, the producer can choose to temporarily shut in an operating well, that is, to change to "mothballing" when current status is "producing". The revenues from production (as discussed previously, a function of output prices P, operating costs C^p, and production volumes Q) will be sacrificed and the representative producer maintain the oil well by costing maintenance cost C^m. The producer may choose to defer the decision with the "delay option" to postpone until next date. Price uncertainty is resolved by using the "delay option". To summarize, I assume that if a well is producing an operating cost is incurred and if the well is mothballed a maintenance cost is incurred each period. The operating cost of the producing well is C^p and the maintaining cost of the mothballed well is C^m. The operating cost is incurred only when *prod=p,* and the maintaining cost is incurred each period only when *prod=m*. The value of the immediate choice to switch equals the discounted expected payoff stream less the switching cost $\mathbb{C}^{p,m}$ (the net present value of the immediate choice). Assume that for a given well, C, $\mathbb{C}$, and θ are static parameters for each well where θ includes observed and unobserved other variables affecting investment value. Assume that $\mathbb{C} = \mathbb{C}^{m,p} = \mathbb{C}^{p,m}$.

Define the value function $V(P, Q, C, \theta, O(\mathbb{C}))$, where value depends upon future oil prices P and the distribution of those prices, production quantity Q, operating cost C^p, maintaining cost C^m, other well characteristics θ, and the options to "mothball" and to "produce" which are exercised at an investment cost $\mathbb{C}$. If the state of the well is mothballed then the option to switch to production is in force and if the state of the well is producing the option to switch to mothballing in is force. The options have values denoted as $O^m(P, Q, C^m, \mathbb{C}, \theta)$ and $O^p(P, Q, C^p, \mathbb{C}, \theta)$ which arise from the options to switch from "producing" to "mothballing" and from "mothballing" to "producing", respectively. So the value function consists of three components: the profit that depends on P_t and Q_t, the costs C^p or C^m, and the real options O^m or O^p. The producer faces a multiperiod problem of selecting production as well as the optimal policy for exercising the options. The producer is assumed to follow a policy that maximizes an expected payoff function by choosing the optimal investment time *t* to exercise the option to switch, conditional on the current state of the well and an optimal policy concerning production.

Assume for illustration that the current state is *prod* = *m* (mothballing). The optimal choice to switch or continue with the current state of the well (that is delay switching) will define an optimal investment trigger price $P*$. If the actual price exceeds the trigger price then the option would be exercised. That trigger price will be a function of the expected future distributional characteristics of the price change distribution. Define the implicit objective function γ where the optimal policy defines the selection of the state of the well at each date as:

$$\gamma = \max_{prod} V(P, Q, C, \mathbb{C}, \theta, O)$$

where P represents the expected price distribution at t, Q is the set of quantities produced each period over the production horizon of the well, C represents the operating cost ($denoted\ C^p$) stream of each period, including both fixed and variable cost if $prod = p$, and the cost of maintaining the well if $prod = m$ ($denoted\ C^m$). $\mathbb{C}$ is the switching investment cost. The vector θ contains both observed and unobserved well characteristics that affect production decisions, and O is the set of options the producer faces in each period (for instance, $O = O^p\ if\ prod = m$). The behavior and distribution of P is assumed to be exogenously determined and is a function of time *t* and depends both on the expected price change, and the anticipated volatility, skewness, and kurtosis reflected in the price distribution.

$$P = P(\mu, \sigma, sk, k|t)$$

where μ is the expected change, σ is volatility, sk is skewness, and k is the kurtosis of the price change distribution. Oil prices are determined in a world market and so this assumption conforms to both the actual market environment and the data (Kaminski, 2012, U.S. Energy Information Administration https://www.eia.gov/energyexplained/oil-and-petroleum-products/prices-and-outlook.php).[6]

Define the multi-period dynamic problem of the producer as:

$$\gamma = \max_{prod} E_t\big(V(P, Q, C, \mathbb{C}, \theta, O)\big)$$

$$= \max_{prod}\left(\pi(P_t, Q_t, C, \mathbb{C}, \theta|\mathbb{I}_t) + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t)\right)$$

where $\frac{1}{(1+r)}$ is the discount factor, $\pi(P_t, Q_t, C, \mathbb{C}, \theta|\mathbb{I}_t)$ is the profit function at *t* given the current information set and $E_t(V_{t+1}|prod_t)$ is the expectation of the continuing value function at *t+1*

[6] It is worth pointing out that a process for oil prices that exhibits mean reversion and random jumps as well as a continuous stochastic error, as has been found in the literature (Al-harthy, 2007) produces a price change distribution that exhibits both non-zero skewness and excess kurtosis for reasonable parameter values.

given the producing status choice at t, $prod_t$ assuming an optimal policy subsequent to *t+1*, where I have assumed that the expectation exists and that the price process is exogenous to the firm.[7] The value function at *t* can be rearranged to contain two maximized terms: the first term is the profit function at *t*, and the second is the discounted expected value function at *t+1* given producing status choice at *t* and an optimal policy following *t*. The production state *prod* is selected at *t* to maximize the sum of the two terms. The investment choice at *t* therefore depends upon the production state at *t-1*. If the representative oil producer chooses $prod_t = p$ when $prod_{t-1} = p$, the above equation becomes:

$$V_t(prod_t = p, prod_{t-1} = p) = \pi(P_t, Q_t, C, \mathbb{C}, \theta|\mathbb{I}_t) + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = p)$$

$$= \breve{\pi}(P_t, Q_t, \theta) - C^p + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = p)$$

$\breve{\pi}(P_t, Q_t, \theta) - C^p = \pi(P_t, Q_t, C, \mathbb{C}, \theta|\mathbb{I}_t)$ for $prod_t = p$ when $prod_{t-1} = p$.

I henceforth drop the notation for the information set $\mathbb{I}_t$ as it is implicit.

When $prod_{t-1} = m$ and the representative producer chooses $prod_t = p$, the value function becomes:

$$V_t(prod_t = p, prod_{t-1} = m) = \breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = p)$$

$\breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} = \pi(P_t, Q_t, C, \mathbb{C}, \theta|\mathbb{I}_t)$ for $prod_t = p$ when $prod_{t-1} = m$. The only difference is that transitioning from "mothballing" to "producing" incurs a one-time sunk cost $\mathbb{C}$. Similarly if the representative producer chooses $prod_t = m$ when $prod_{t-1} = m$, the value equation is:

$$V_t(prod_t = m, prod_{t-1} = m) = -C^m + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = m)$$

[7] The reader will recognize this as a Bellman equation. Molls (2001) presents a similar specification.

If the representative producer chooses $prod_t = m$ when $prod_{t-1} = p$:

$$V_t(prod_t = m, prod_{t-1} = p) = -C^m - \mathbb{C} + \frac{1}{(1+r)} E_t(V_{t+1} | prod_t = m)$$

Where C^m is the maintenance cost in "mothballing" periods, and the only difference between the two equations is the term $-\mathbb{C}$, the one-time transitioning cost from "producing" to "mothballing".

The well operator is assumed to select the choice that maximizes value conditional on the current state of the well (m or p). Therefore, if maximizing behavior is assumed, the choice observed implies that choice has the largest current value, this of course also implies the choice carries with it an optimal policy going forward. Mothballing incurs a per period cost while a complete shut down does not, hence it would never be optimal to switch from complete shutdown to the mothball state.

When $prod_{t-1} = m$, the difference in the value equation between choosing $prod_t = p$ and choosing $prod_t = m$ will equal:

$$\begin{aligned}
&\Delta V_t(prod_{t-1} = m) \\
&\qquad = \breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} + \frac{1}{(1+r)} E_t(V_{t+1} | prod_t = p) \\
&\qquad - \left(-C^m + \frac{1}{(1+r)} E_t(\mathbb{V}_{t+1} | prod_t = m) \right) \\
&\qquad = \breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} + C^m \\
&\qquad + \frac{1}{(1+r)} \big(E_t(V_{t+1} | prod_t = p) - E_t(V_{t+1} | prod_t = m) \big)
\end{aligned}$$

When $prod_{t-1} = p$, the difference in the value equation between choosing $prod_t = p$ and choosing $prod_t = m$ equals:

$$\Delta V_t(prod_{t-1} = p)$$

$$= \breve{\pi}(P_t, Q_t, \theta) - C^p + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = p)$$

$$- \left(-C^m - \mathbb{C} + \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = m)\right)$$

$$= \breve{\pi}(P_t, Q_t, \theta) - C^p + \mathbb{C} + C^m$$

$$+ \frac{1}{(1+r)} \big(E_t(V_{t+1}|prod_t = p) - E_t(V_{t+1}|prod_t = m)\big)$$

The difference is twice the one-time transitioning cost from "mothballing" to "producing" (or from "producing" to "mothballing"), $2\mathbb{C}$.[8]

If there were no sunk costs associated with switching between production and mothballing then switching would be costless and instantaneous. In that case the future continuation value vanishes because the lagged state, that is the state coming into time *t*, will have no influence on future decisions. In the presence of sunk costs however, the lagged production state matters for the current choice of production state. If sunk switching costs are zero, the firm will produce if current profit is positive and will shut down if profit is negative. Mothballing would never occur since mothballing incurs a per period cost. Conversely the larger are the non-recoverable switching sunk costs, the more important is the lagged production status for the current period production choice.

The value effect from the optimal choice to switch from production to mothballing, or vice versa therefore implicitly defines a decision rule. The decision to switch will depend upon the distributional properties of future price changes and the remaining parameters of the model including the cost of switching. The profit function is determined by the oil price, quantity

[8] I assume that the transitioning cost from "mothballing" to "producing" is the same as the cost from "producing" to "mothballing", as a consequence, the difference between these two equations is twice the cost. If the transitioning costs are different ($\mathbb{C}^{mp}$ and $\mathbb{C}^{pm}$), the difference between these two equations become $\mathbb{C}^{mp}$+$\mathbb{C}^{pm}$.

produced, and well characteristics, $\breve{\pi}(P_t, Q_t, \theta)$. Given that the price is a random variable it can be characterized by the parameters of its probability distribution which I define as the expected future price, volatility, skewness, and kurtosis at *t*, conditional on the information set at time *t.* In addition, the difference in value function $(-C^p + \mathbb{C} + C^m)$ is related to the operating cost, maintaining cost, as well as the transition cost. I re-write the difference in the value function ΔV_t as reduced form equations:

$$\Delta V_t(prod_{t-1} = m) = \gamma_0^m + \gamma_P^m \cdot P_t + \gamma_Q^m \cdot Q_t + \gamma_C^m \cdot C + \gamma_\theta^m \cdot \theta + \gamma_{C^m}^m \cdot C^m + \gamma_{\mathbb{C}}^m \cdot \mathbb{C} + \varepsilon_t^m$$

and:

$$\Delta V_t(prod_{t-1} = p) = \gamma_0^p + \gamma_P^p \cdot P_t + \gamma_Q^p \cdot Q_t + \gamma_C^p \cdot C + \gamma_\theta^p \cdot \theta + \gamma_{C^m}^p \cdot C^m + \gamma_{\mathbb{C}}^p \cdot \mathbb{C} + \varepsilon_t^p$$

4.3.1 The Empirical Model of the Choice to Produce or Mothball

I specify a dynamic panel binary choice model (i.e., panel probit) to examine the relationship between the producers' propensity to choose the "producing" status rather than the "mothballing" status and the characteristics (implied volatility, skewness, and kurtosis) of the oil price change distribution.

By estimating the above set of equations and testing whether coefficients on the lagged state dummy variables are significant and positive, I can test whether sunk costs matter in the presence of expected price uncertainty, but importantly test how expected prices and the expected future moments of the price change distribution influence actual choices. The above equations specify the problem of choosing producing status for each period *t* within the horizon 1 to *T*. Note that the implied value changes depend upon the lagged production status of the well. Combining the two expressions and introducing a lagged indicator for $prod_{t-1}$ yields an integrated reduced form model. The binary choice problem depends upon the change in value from optimal selection of the production status choice going forward. Therefore, the implicit decision functions can be summarized as

$$I_{i,t} = \theta(\alpha + \beta \cdot prod_{t-1} + \lambda \cdot P_{i,t} + \omega \cdot Q_{i,t} + \gamma \cdot C_{i,t} + \beta \cdot \mathbb{C}_{i,t} + \tau \cdot \theta_{i,t} + \varepsilon_{i,t})$$

where $P_{i,t}$ is a vector including the expected (output) price level and risk-neutral central moments (volatility, skewness, and kurtosis) and λ is a vector of coefficients. $Q_{i,t}$ is the produced quantity of crude oil in month t for well i. Notice that the oil well production rate declines over time with well age. $C_{i,t}$ contains all costs for $prod = m, p$ including operating fixed and variable costs C_i^p and mothballing maintaining costs C_i^m. Well depth is positively correlated with the operating costs. As noted previously the profit function and the value function depend on the current producing status, so the price distribution moments ($P_{i,t-3}^{18}, Vol_{i,t-3}^{18}, Skew_{i,t-3}^{18}, Kurt_{i,t-3}^{18}$) should interact with lagged producing status $prod_{i,t-1}$. $\mathbb{C}_{i,t}$ is the investment cost to switch producing status ($\mathbb{C}_{i,t}$=$\mathbb{C}_{i,t}^{pm}$ for a producing well changing to mothballing and $\mathbb{C}_{i,t}$=$\mathbb{C}_{i,t}^{mp}$ for a mothballing well to re-open). $\theta_{i,t}$ contains other observed determinants of well production choices, including drilling types, basin location, etc.

I define the following variables $horizontal\ drilling_i, directional\ drilling_i, \overline{basin}_i$, $productivity_i$. "vertical drilling" is omitted so as to avoid the dummy variable trap. The error $\varepsilon_{i,t}$ is assumed to capture any additional unobserved variables affecting oil well production decisions. The complete binary model can be estimated with the Maximum Likelihood Estimator. Here Φ is the standard normal distribution function for the probit model.

$$
\begin{aligned}
prod_{i,t} = \Phi(\alpha_0 + \beta_0 \cdot prod_{i,t-1} + \ \beta_1 \cdot P_{i,t}^{18} + \beta_2 \cdot \{prod_{i,t-1} = 1\} \cdot Vol_{i,t}^{18} \\
+\beta_3 \cdot \{prod_{i,t-1} = 0\} \cdot Vol_{i,t}^{18} + \beta_4 \cdot \{prod_{i,t-1} = 1\} \cdot Skew_{i,t}^{18} \\
+\beta_5 \cdot \{prod_{i,t-1} = 0\} \cdot Skew_{i,t}^{18} + \beta_6\{prod_{i,t-1} = 1\} \cdot Kurt_{i,t}^{18} \\
+\beta_7\{prod_{i,t-1} = 0\} \cdot Kurt_{i,t}^{18} + \nabla_1 \cdot \overline{basin}_i + \delta_1 \\
\cdot horizontal\ drilling_i \\
+\delta_2 \cdot directional\ drilling_i + \delta_3 \cdot undetermined\ drilling_i + \gamma_1 \\
\cdot age_{i,t} \\
+\gamma_2 \cdot depth_i + \gamma_3 \cdot productivity_i + \varepsilon_{i,t}) \qquad (1)
\end{aligned}
$$

4.4 The Choice to Drill Now or Delay

Similarly, new drillings are real options for oil producers. Kellogg (2014) shows that the drilling choice is positively related to price but negatively to price volatility. Drilling are the call options to producers. I predict that skewness and kurtosis are significant influences for the real options decisions according to the “Bad News Principle” and I use the following econometric model to examine the impacts of moments:

Similar to the exercise of the “producing” or “mothballing” option, producers will exercise the “drilling” option if the “waiting option” value is less than the immediate exercise value of drilling. The choice problem is one in which drilling has not occurred but can at a fixed cost, generating a stream of net cash flows in the future. Conversely, the choice to not drill at date *t* but to wait, will be governed by the value of waiting with the option to exercise (invest) at a future date. In summary form the choice is to select to drill to maximize a value function that includes the option to drill

$$\gamma = \max_{drill} V(P, Q, C, \mathbb{C}, \theta, O^d)$$

where $P = P(\mu, \sigma, sk, k)$ and O^d is the drilling option. Conversely, the choice not to drill implies that the value of the option to delay exceeds the immediate value of drilling.

4.4.1 The Empirical Model of the Choice to Drill or Delay

The empirical reduced form model is specified as, where *h* is the decision function

$$I_{i,t}^{drilling} = \theta(\alpha + \beta \cdot prod_{t-1} + \lambda \cdot P_{i,t} + \omega \cdot Q_{i,t} + \gamma \cdot C_{i,t} + \beta \cdot \mathbb{C}_{i,t}^{drilling} + \tau \cdot \theta_{i,t} + \varepsilon_{i,t})$$

Unlike the choice to continue production or mothball which is a dynamic decision-making process, the exercise of the drilling option is a one-time event. The purpose of the Cox Proportional Hazard model is to evaluate simultaneously the effect of several factors on survival. In the case at hand the initial state is that the well has not been drilled. The well survives in this state until drilling begins. In other words, it allows me to examine how specified factors influence the rate of a particular event happening (e.g., drill) at a particular point in time. The Cox proportional hazard model asserts that the rate of an event occurring at a point in time is related to a set of explanatory variables and a baseline rate. Within the context evaluated here, the model allows me to examine how specified economic factors or characteristics influence the rate of a particular event happening at a particular point in time. The rate is commonly referred to as the hazard rate. The Cox proportional hazard model considers that the probability of exercising the option (the event occurs) increases as time passes; it is rather than a static process in which the likelihood to exercise an option does not change over time. For example, the probability that a new oil well will be drilled in 10-th month is higher than that in the second month after the drilling option becomes available. The Cox model will fit this element when estimating the coefficients. So, the Cox proportional hazard model is more suitable for exploring drilling options' exercises.

Define the variable $exer_{i,t}$ for well i at date t. The variable $exer_{i,t}$ which takes the value 1 if the choice to drill at t occurs and 0 otherwise. Specifically, drilling is defined as those months which occur on and after the spud dates of oil wells, and the dates on and before well completion. Once drilling has begun, the costs incurred are not reversible. Hence the choice to drill is an irreversible investment decision. That is, once drilling begins the cost is sunk. As emphasized earlier, the model allows a test of several hypotheses. First, that expected prices, volatility, skewness and kurtosis, influence the choice to drill. Second, a test of the hypothesis that sunk costs matter. The estimated model is specified to include the variables $P_{i,t-3}^{18}$, $Vol_{i,t-3}^{18}$, $Skew_{i,t-3}^{18}$, and $Kurt_{i,t-3}^{18}$ as defined earlier and in Appendix A. I define dummy variables for the primary drilling operations: horizontal drilling, directional drill and undetermined, with vertical drilling excluded from the set of dummies to avoid the dummy variable trap. The model also includes two variables that reflect costs $depth_i$ is the well depth in tens of thousands of feet t, and $total\ costs_i$ is equals the sum of drilling, tie-in, and completion costs. I include completion costs as these would be the costs required to implement actual production. Expected productivity, $productivity_i$, is a measure of expected reserves based upon first test oil volume in bbl. The model is expressed as the following where $h_{0,i}$ is the hazard rate when all explanatory variables equal 0. The hazard rate represents the probability of drilling, given that no drilling has occurred up to a specific time.

$$
\begin{aligned}
h_i(t) = {} & h_{0,i}(t)\exp(\beta_1 \cdot P_{i,t-3}^{18} + \beta_2 \cdot Vol_{i,t-3}^{18} + \beta_3 \cdot Skew_{i,t-3}^{18} + \beta_4 \cdot Kurt_{i,t-3}^{18} \\
& + \nabla_1 \cdot \overline{basin}_i + \delta_1 \cdot horizontal\ drilling_i + \delta_2 \cdot directional\ drilling_i \\
& + \delta_3 \cdot undetermined\ drilling_i + \gamma_1 \cdot depth_i + \gamma_2 \cdot productivity_i \qquad (2)
\end{aligned}
$$

$+\gamma_3 \cdot total\ costs_i)$

4.5 The Choice to Permanently Shut Down or Continue

Shutting down an oil well is a put option to producers. I derive the payoff functions of the producing and shutdown statues in the following equations: The producer will choose "shutdown" over "producing" when $-\mathbb{C}^{shutdown}(\theta) > V_t(prod_t = p)$[9]. Where $\mathbb{C}^{shutdown}$ is the one-time sunk costs to switch from the producing status to shut down, which include but are not limited to the costs of the well hole plug-in cost, the cost to dispose of well pollutes, and transportation cost. Also, I have $P = P(\mu, \sigma, sk, k)$. Intuitively, it is optimal to switch to shutdown status from producing if $-\mathbb{C}^{ps} > V_t(prod_t = p)$ (Notice that $V_t(shutdown) = 0$). When maintaining producing status it is not economical for oil wells, i.e., when $V_t(prod_t = p)$ becomes negative as operating costs surge or price declines, closing a well may be optimal. Similar to the derivation for the "defer" option, I can derive the difference in value functions (when $prod = p$ vs. when $prod = s$):

$\Delta V_t(prod_{t-1} = p)$

$$= \breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} + \frac{1}{(1+r)} E_t(V_{t+1} | prod_t = p)$$

$$- \frac{1}{(1+r)} E_t(V_{t+1} | prod_t = s)$$

$$= \breve{\pi}(P_t, Q_t, \theta) - C^p - \mathbb{C} + \frac{1}{(1+r)} \big(E_t(V_{t+1} | prod_t = p) - E_t(V_{t+1} | prod_t = s)\big)$$

and:

[9] I ignore the case when a shutdown well is re-opened as this is a very rare case due to significant costs. I also ignore the case when switching from mothballing to shutdown as the data do not distinguish this case from direct shutdown.

$$\Delta V_t(prod_{t-1} = s) = \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = s) - \frac{1}{(1+r)} E_t(V_{t+1}|prod_t = s)$$

$$= \frac{1}{(1+r)} \left(E_t(V_{t+1}|prod_t = s) - E_t(V_{t+1}|prod_t = s)\right)$$

I use a similar model to examine the relationship between skewness (kurtosis) and shutdown real option exercise:

$$I_{i,t}^{shutdown} = \theta(\alpha + \beta \cdot prod_{t-1} + \lambda \cdot P_{i,t} + \omega \cdot Q_{i,t} + \gamma \cdot C_{i,t} + \beta \cdot \mathbb{C}_{i,t}^{shutdown} + \tau \cdot \theta_{i,t} + \varepsilon_{i,t})$$

$$exer_{i,t} = h(\alpha_0 + \beta_1 \cdot P_{i,t-3}^{18} + \beta_2 \cdot Vol_{i,t-3}^{18} + \beta_3 \cdot Skew_{i,t-3}^{18} + \beta_4 \cdot Kurt_{i,t-3}^{18} + \nabla_1 \cdot \overline{basin}_i$$
$$+\delta_1 \cdot horizontal\ drilling_i + \delta_2 \cdot directional\ drilling_i + \delta_3$$
$$\cdot\ undetermined\ drilling_i$$
$$+\gamma_1 \cdot depth_i + \gamma_2 \cdot productivity_i + \gamma_3 \cdot total\ costs_i + \varepsilon_{i,t})$$

It is optimal to exercise the shutdown real option when the output price falls below the investment threshold, which is determined by the relative size of the “defer” to shut down option and the immediate shutdown value.

Exits of producers are also real options. By exercising the option to shut down, producers avoid the costs associated with operating the properties when continuing production is no longer economical. For instance, declines in oil price and rise in operating costs. Dixit and Pindyck (1994) show discuss the option to abandon, whose threshold decreases with the upside potential of price. Moel and Tufano (2002) find that the propensity to exercise exit options for gold mines is negatively associated with uncertainty after controlling for price. The “Bad News Principle” applies to the “mothball” option and exit option. I predict the following: the probability of shutting down a well falls when expected good news becomes a more likely event; wells are more prone to be closed when upside potential price changes are smaller.

4.5.1 The Empirical Model of the Choice to Shut Down or Continue

The shutdown model is similar to the drilling model but with the dependent variable as the dummy "exer"=1 for well closing months, i.e., the month when the oil well was closed. The initial state is therefore that the well is producing. The choice variable "exer" equals 1 for the closing month and equals 0 for those months before the closing month. The explanatory variables are the same as those in the drilling model in except for $productivity_i$, which is replaced with a variable labeled "last12" which equals the production rate of the well for the last 12 months'. A more productive producing oil well (higher "last12") is less likely to be closed. The structure of the empirical model is:

$$
\begin{aligned}
h_i(t) = \; & h_{0,i}(t)\exp(\beta_1 \cdot P_{i,t-3}^{18} + \beta_2 \cdot Vol_{i,t-3}^{18} + \beta_3 \cdot Skew_{i,t-3}^{18} + \beta_4 \cdot Kurt_{i,t-3}^{18} \\
& +\nabla_1 \cdot \overline{basin}_i + \delta_1 \cdot horizontal\ drilling_i + \delta_2 \cdot directional\ drilling_i \\
& +\delta_3 \cdot undetermined\ drilling_i + \gamma_1 \cdot depth_i + \gamma_2 \cdot productivity_i \\
& +\gamma_3 \cdot total\ costs_i) \qquad (3)
\end{aligned}
$$

4.6 Predictions

Based upon the generic model described in section 4.1 in which call option values are a function of both volatility, skewness and kurtosis and the developments above on the optimal choice problems, Table 1 presents the predicted relations between the probability of each choice studied and expected volatility, skewness, kurtosis and prices. In addition, the table lists well characteristics discussed in the following section, and the predicted influence of those characteristics on the choices studied.

5. Data and Methods

5.1 Oil Well Lifecycle - Drilling, Producing, Mothballing, and Shutdown

The lifecycle of an oil well goes through several stages, from testing to plugging in. Firstly, a seismic survey of an oil well usually includes testing for the estimated quantity of expected total production and measured depth of the oil reserve. It takes about three months on average to complete the drilling process. The porous rocks are broken by a rotary drilling rig, which is usually rented and thus incurs rig-rental cost. The method also involves removing the mud from the drilling - a transportation cost. The drilling also uses materials to stabilize the rock with cement and steel - a casing process that incurs completion costs. Water may be pumped into the reservoir to increase pressure and help pump out the oil. While the well is producing, additional horizontal wells (infill wells) may be drilled to increase oil production from the same reservoir – these are the infill wells.[10]

After drilling ends and the well is completed, production may begin, and the oil well moves to its second stage, "producing." At this stage, producers earn revenues from produced crude oil and which is positively related to the produced quantity of oil product and market oil prices, as well as is negatively related to the operating costs, including the administration cost and equipment depreciation, thus net revenue from operating a producing oil well comes from two parts – the generated income from selling the produced crude oil at market prices and the cost to maintain the production. At this stage, oil producers consider the expected oil price (and

[10] For a detailed description of the drilling process see https://production-technology.org/.

distribution moments) in the futures and how much the cost to maintain the production to decide on production status.

Mothballing is a temporary suspension of production. When facing bad news for production, such as price shocks or a surge in operating costs, producers have the option to mothball the producing oil well. The difference between mothballing and shutdown (plugging in) is that a mothballed well can be quickly re-opened to resume production, and the cost of doing so is lower than reopening a plugged well, where in fact in might be impossible to reopen a plugged well and a plugged in well is permanently unusable. Mothballing an oil well incurs less cost than completely closing a well and re-opening the mothballed well costs less than drilling a new one. Mothballing status is an intermediate status between "producing" and "shut down" but it requires a medium level of maintaining cost to keep this status. The benefit of maintaining the "mothballing" status is that oil producers still have the option to resume production.

The last stage of an oil well is "shutdown" – when producing or maintaining the well is no longer economical or when the oil wells are exhausted. As I discussed before, closing an oil well takes far more cost than mothballing a well – it takes plug materials to permanently close a well to prevent pollution or expansion of the abandoned wells. Most used plug material is cement. Re-opening a plugged oil well is generally not possible and where possible is extremely costly.

5.2 Oil Well Investment and Production Data

5.2.1 Records of Oil Wells' Drilling, Production, and Closing Dates

Oil well production rates and dates and dates production begins and the dates production ends are obtained from Enverus.[11] Each state of the United States typically requires producers to file well drilling and production records with a state agency on a monthly basis. Enverus gathers these data records and assembles them into a common format data base. Data obtained include a well identifier code (an unique well identification 14-digit number: American Petroleum Institute, short for "API"), the well's operator company's names, the well's monthly oil production rates, gas production rates, and water production rates, well's measured depth in tens of thousands of feet, a well lease identifier (and lease names), the basin where the well is located, well's field name, and location's county and state names, total drilling costs, the date drilling begins (the spud dates), the dates when the well begins production (the completion dates), well's peak production rates, well's first 6-months' and first 12-months' production rates, and age of the well measured as the number of months from the first production date, and the date a well is closed off so that it stops producing (the "last production" dates).

Records on over 600,000 oil wells are obtained for the states of Texas, North Dakota, Oklahoma, California, and Pennsylvania. The total number of well-month observations equal roughly 53,000,000. According to their monthly oil production rates and their first and last production months, I construct the production histories/status for each of the 600,000 oil well's on a monthly basis. The earliest date is January 2010 and the last date is September 2019. The spud date identifies when drilling begins and investment costs are incurred. The production status histories, allow recovery of the investment decisions (choices between producing and mothballing) of oil producers through every oil well's entire production history during the ten-year period studied. The histories allow identification of the investment date (which I classify as

[11] Enverus.com. I am grateful to Enverus for providing me with access to their well production and cost datasets and especially to Ms. Annie Shen and Mr. Jason Eleson for their very generous help.

the spud date) and the decisions to temporarily suspend production and to abandon (shut in) a well. The production histories therefore provide information on the state of the well at any date. Each oil wells' drilling and shutdown profiles are constructed from the well spud dates and their last production dates.

5.2.2 Data on Expected Revenues and Costs

Age, Depth, and Productivity: The likelihood of resuming, mothballing, drilling or closing a well depends on potential productivity. Each oil well's total reserve is used to proxy for expected productivity. The accumulated oil production rate is used as the proxy for operating and mothballing oil wells; oil wells' last 12-months' production rate is used as the proxy for the productivity of oil wells facing closing down decisions. Age controls for the decline in production rates over time as oil reserves are exhausted. An oil well's oil production rates changes over time and can be predicted by the Hubbert curve with the prediction function $Q(t) = \frac{Q_{max}}{1+e^{a(b-t)}}$ [12] which decreases with age (the *t* in the function). Well age, depth, reserve, last 12-months' production rate, and cumulative oil production rate are rescaled to fit the magnitudes of other variables in my model and are measured as natural logs. Age equals the log of the number of months since production began. Well depth (in natural logs of tens of thousands of feet) is included to control for differential operating costs as those costs vary with depth.[13]

Drilling and Shutdown Costs: With the drilling cost dataset from Enverus, I can measure individual wells' total drilling costs (total drilling costs consist of the drilling, completion, and tie-in costs), which are positively related to oil wells' closing costs.[14] So, I use the drilling costs

[12] See refence: Claerbout, J., & Muir, F. (2008). Hubbert math. *Sepstanford. edu.*
[13] OilGasEquity.com. (n.d.). *Drilling & Completion Facts.* OilGasEquity.com. Retrieved September 6, 2022, from https://www.oilgasequity.com/resources/drilling-completion-facts/#:~:text=The%20study%20found%20that%20drilling,%241.6m%20for%20drilling%20alone.
[14] This is confirmed by a conversation with industry professional.

data to measure the total costs to drill a new oil well in a lease field and as a proxy for the cost to plug in a closed oil well, however plugging an existing well incur on average less costs than drilling a new well.[15][16] The drilling costs are deflated[17] in order to control for the difference in dollar value across time. Natural logs of the data are used to fit to the scales of other variables.

Drilling Types and Basins: Directional or horizontal drilling allows producers more flexibility and precision in reaching and extracting oil/gas compared to vertical drilling and directional drilling generally yields greater oil extraction. By some estimates however vertical drilling costs only about 25-50% of what it costs for a horizontally drilled well.[18] Basin is a depression, or dip, which may be filled with oil. Geographic characteristics of the basin in which the well is located can influence the amount of oil produced. I control for well head location by identifying the basin in which the well is located. Controlling for the basin location accounts for differential geographic issues such as crude oil quality, gravity[19], sulfur, and rock porosity.

5.2.3. Construction of the Dataset

I assemble the following data for each well in the set of wells with spud dates between March 2010 and September 2019. I access the age (in months) and depth (in feet) of oil wells, oil well's accumulated production rates (monthly in bbl), total reserve (in bbl), drilling types

---

[15] Raimi, D., Krupnick, A. J., Shah, J. S., & Thompson, A. (2021). Decommissioning orphaned and abandoned oil and gas wells: New estimates and cost drivers. *Environmental Science & Technology*, 55(15), 10224-10230.
[16] HydroCarbonWell.com (n.d.). *How Much Does It Cost to Plug an Oil and Gas Well?* Hydrocarbon Well Services. Retrieved September 6, 2022, from https://hydrocarbonwell.com/how-much-does-it-cost-to-plug-an-oil-and-gas-well/
[17] Both price and costs are deflated by inflation rate, source of CPI: https://fred.stlouisfed.org/.
[18] Rig Worker. (2022, September 5). *Horizontal well costs - oil well drilling.* Rig Worker. Retrieved September 6, 2022, from https://www.rigworker.com/oil-well/horizontal-well-costs.html
[19]API gravity, is a measure of how heavy or light a petroleum liquid is compared to water. High API gravity referred to as 'light'. Oil with a high sulfur content is considered "sour" and low-sulfur oil is "sweet". Light-sweet oil typically commands a higher price.

(horizontal, directional[20], vertical, or undetermined), oil basin (for instance, Fort Worth, Palo Duro, and east Texas coastal), last 12-month's production rates (in bbl), total drilling costs (including drilling cost, completion cost, and tie-in cost) (in millions of USD), oil lease names, and state in which the well is located. I examine activity in Texas, North Dakota, Oklahoma, California, and Pennsylvania, the five states with the most complete records as well as large numbers of wells. I use these variables to control investment costs, operating costs, expected oil production rate, expected lifespan of the oil wells, etc. I define total drilling costs to include both completion cost as well as tie-in cost (the cost of tying into a pipeline for transportation) as these costs must be incurred before oil can physically be produced and transported for sale.

I exclude observations with any missing variables. Table 3 presents descriptive statistics for the sample production data.[21] This table summarizes the variables in the production dataset (except for the *exer(drilling)* and *exer(shutdown)* variables which are only available in the drilling dataset and the shutdown dataset, respectively). On average, half of the well-month observations are from operating oil wells – they are producing crude oil, but the other half of the observations are from mothballed wells. This mean in "*prod*" suggests that mothballed oil wells are not rare cases – actually, a shut-in oil well is very common. Notice that the variables *age*, *cum. oil production*, *reserve*, *depth*, and *last12* are taken natural logarithm to fit to the scale of other variables. So, on average, the oil wells in my dataset is about 126 months' old – they have been producing for about 10 years. This is consistent with the fact that most oil wells can produce up to about 20 years. On average, a decision to drill an infill well happens in the middle of the oil lease's lifespan – the mean of *exer(drilling)* is 0.3281. It suggests that from the drilling

[20] In the Enverus well dataset, directional drilling is distinguished from horizontal drilling; but (horizontal) directional drilling is defined as a specific type of well drilling method and generally belongs to the directional drilling class. Horizontal drilling is different from directional drilling in that it is generally used to boost production rates of an existing reservoir by drilling several additional wells along the field.
[21] Sample sizes used in estimation may differ due to missing observations for some explanatory variables.

of the first oil well in a lease, it takes about half of the total production life of a lease to drill an infill well to boost production. *Exer(shutdown)* is much smaller than *exer(drilling)*. Notice that a shutdown oil well will not remain in the dataset for a long time. Among all the oil wells the majority are vertical drilling oil wells. Horizontal and directional drilling wells consist of about less than 20% of the sample. Texas occupies a large proportion of oil wells in the sample – about 58.33% are Texas oil wells. North Dakota only consists of a small portion of the sample (2.34%).

Drilling: In my examination models of oil drilling decisions, I used only the first "infill" wells within a lease – that is, the first wells being drilled after the spuds of the first oil wells within a lease. Using infill well drilling decisions represents a clean way of defining optionable choices to drill as they are made only after the initial stage of the drilling on the leased area. Kellogg (2014) and Décaire et al. (2020) follow a similar approach in their analyses. The "infill" wells are the oil wells drilled to increase production rates of existing oil well fields when there already exists a producing oil well in the leased field.

Shutdown and Production: In my examination of well production status changes and shutdown choices, I used all oil wells ever producing within the months March 2010 – September 2019, which is the complete sample of oil wells from Enverus in the states California, Oklahoma, North Dakota, Texas, and Pennsylvania. Instead of using just the first infill wells as in the oil well drilling sample, I used all producing wells from 2010 to 2019 as all producing oil wells become a put option for oil producers – they could choose to shut in the oil wells to reduce operating cost but maintain the option to re-open it. In the case of extremely low expected oil prices and price change uncertainty, they might permanently close the existing oil wells to reduce potential financial loss, since a shut-in oil well still requires a maintaining cost to enable future re-opening but a closed oil well does not incur further costs but a one-time plugging in cost.

As mentioned previously, drilling an oil well is an example of exercising an irreversible investment, and closing an oil well of abandoning a previous investment. I recover the beginning dates of each oil well's investment horizon from their oil lease information. It is the spud month of the first oil well in production from the same oil lease field when drilling an additional oil well becomes an exercisable option within the lease.[22] The option remains exercisable until the last oil well in the lease ceases to produce. I recover the closing dates of each oil well from their last production dates on record. The beginning of the oil well's potential disinvestment horizon is the oil well's first production date. To summarize, I take the oil well lease after the spuds of first wells as the real options to drill (thus call options to open a new oil well) until the spuds of the first infill wells which are the dates options are closed, and the producing oil wells as the put real options to shut in or close wells; in addition, a mothballed and temporarily shut-in oil well is a call option to re-open productions.

5.3 Oil Futures Prices and Implied Moments

I use BKM's (Bakshi, Kapadia, and Madan, 2003) risk-neutral model-free option-implied central moments to estimate volatility, skewness, and kurtosis of the oil price change distribution. These central moments are constructed from futures and futures' options prices.

Daily crude oil (symbol: CL) futures prices and options on futures prices (symbol: LO) are obtained from the Chicago Mercantile Exchange, CME, Group Datasets (End of Day Complete database). The data are used to compute a proxy for the expected oil price and option-implied volatility, skewness, and kurtosis at future time horizons. Only option prices with

[22] Also see definitions of "infill wells" in this section. This "infill" wells follow the set up in Kellogg (2014) and Décaire et al. (2020).

maturities between 10 and 180 days, the most liquid options contracts, are utilized, as contracts expiring or far from maturity are traded thinly and whose settlements are more noisy. Options and futures prices that are clearly out of line relative to averages are assumed to be recording errors and are excluded. Risk-free rates are obtained from the OptionMetrics files through WRDS. Risk-free rates are extrapolated and interpolated to ensure sufficient data for maturities spanning 10-180 days.

I follow Chang, Christoffersen, and Jacobs (2013) and Ruan and Zhang (2018) in computing the risk-neutral central moments of the price change distribution. First the data are filtered by a) dropping option prices lower than 3/8 (minimum tick), b) dropping deep-in-the-money options (put options with an exercise price higher than 103% of futures price and call options with an exercise price lower than 97% of futures price), c) dropping prices on days with fewer than 2 put or call prices, and/or options prices violating arbitrage conditions.[23] Second, to expand the option prices set from the CME option settlement prices, I expand the range of moneyness within a date and maturity using existing observations' moneyness values. The expansion uses the cubic spline interpolation method. I expand moneyness for each date and maturity to the range between 0.0001 and 3 and use the implied volatility (from CME's modified Black-Scholes option pricing model) to interpolate and extrapolate implied volatilities in that range. Third, I use the implied volatilities to infer option prices using the Black-Scholes option on futures pricing model to obtain a smooth option price function in the moneyness 0.0001 and 3 range for each date and maturity. Fourth, futures prices and in the moneynesses [0.0001,3] option prices are then used to calculate option implied volatility, skewness, and kurtosis following the algorithm proposed by Bakshi, Kapadia, and Madan (2003) (BKM). The trapezoidal integration

[23] Arbitrage conditions: excluding observations of option prices when a call option's price is greater than or equal to present value of futures prices or when a call options' price is lower than or equal to the present value of futures prices minus present value of strike prices.

method utilized by Chang, Christoffersen, and Jacobs (2013) and Ruan and Zhang (2018) is employed in the calculations.[24] Fifth, I compute 18-month BKM risk-neutral central moments. Kellogg (2014) points out that "18 months" is a typical horizon for oil operators to observe prices. Décaire et al. (2020) use similar measures for price and volatility. Also, Slade (2001) mentions that "Most firms use a long-run commodity price" for decision purposes. Those producers usually consult forecasted long-run output prices for investment decisions according to survey results with managers of copper producers. Specifically using futures prices I estimate the term structure of oil return realized central moments, and then compute 10-180 days central moments to interpolate the one-month maturity central moments.[25] The one-month central moments and term structure are then used to extrapolate 18-month maturity central moments. However, for the production examination model, I use the one-month price and central moments, as transitioning between producing and mothballing can happen immediately.

I assume the physical timing to drill or shutdown is not immediate but the time between the formal evaluation and actual implementation of the choice occurs with a lag.[26] Three-month lags of the estimated central moments are employed in the analyses presented later for the drilling and shutdown examination models.[27] 18-month forward central moments are calculate by taking averages of the 18-month maturity futures settlement prices. Both the 18-month and the one-month oil price are deflated and scaled by 100 to fit to the scale of the risk-neutral moments. The oil producers can observe monthly average forward price and distribution moments to determine if to continue producing or shut in it. My sample is on a monthly-basis for each individual well and I can only observe the changes in oil well producing states in the months.

[24] For more detail on the trapezoidal integration of central moments, please see p.587-588 of Ruan and Zhang(2018).
[25] The method follows Kellogg (2014, Appendix A).
[26] This follows Kellogg (2014) – it usually takes three months to commence drilling after decision.
[27] Computation details are available upon request.

6. Empirical Results

6.1 Higher Order Central Moments and Oil Production Choices

This sub-section explores the impacts of price distribution moments (volatility, skewness, and kurtosis) on irreversible investment decisions. I explore the problem of the exercise of the delay option, that is the choice to continue production versus temporarily shutting down, after controlling for a complete set of controls, including oil well productivity, existing producing status, and drilling type. I examine the impacts of the distribution moments in two steps. Firstly, in section 6.2, I look at the effect of volatility on oil production investment decisions, i.e., the relationship between volatility and the probability of choosing "producing" status. The results show that increases in expected future price volatility are inversely related to increased real options value, and the increase in real options value postpones irreversible investments. My conclusions are consistent with the existing literature, increases in price volatility lead to a higher probability of delayed investment as the delay option value increases. Secondly, in section 6.3, I examine the impacts of the higher-order moments on oil production investment decisions. I examine the relationship between skewness and kurtosis and the probability of choosing "producing" status to explore the effects of the non-normality of the price change distribution on the option to delay and the likelihood of investing immediately. In this part, I explore the relation between the higher-order moments of the price change distribution on the choice to delay decisions. Based upon the modified option pricing model in which skewness and kurtosis impact option value. I expect to find that skewness and kurtosis affect the choice to produce or mothball. More positive skewness should be associated with a greater probability of mothballing an operating oil well. As for kurtosis, if resuming production (and mothballing a well) is a not-deep-in-the-money option, I expect to find a positive relationship between kurtosis and the likelihood of re-opening a mothballed well (to shut in a producing well).

6.2 The Choice to Produce or Mothball: Volatility and Sunk Costs

More significant uncertainty (increases in price volatility) increases the severity of "bad news" and in theory leads to a postponement of investment. As price volatility increases, the delay option value increases for call and put options. When the delay option value exceeds the value of "immediate investment" producers are more likely to choose to postpone investment (choosing the delay option) rather than invest immediately (choosing the immediate investment). This choice leads to the relationship between price volatility and the probability of changing the producing status of a wells. Changing from "mothballing" to "producing "is an investment (which incurs costs and brings cash flows from selling crude oil from production). If the current state of a well is 'mothballed' to resume production requires a sunk cost investment. That is the value of the choice depends upon the state of the well at the time of the decision. The predicted coefficient for the interaction of "mothballing" lagged status and volatility is negative, higher uncertainty reduces the probability of making the investment to switch immediately when the current state is mothballed as the delay option value exceeds the value of resuming production.

Similarly, the interaction of "producing" lagged status and volatility is positive, as more significant uncertainty reduces the probability of exercising the option to "temporarily shut-in" the well. The value of remaining open and producing exceeds the value of shutting in the well. The dependent variable *prod=1* for "producing" indicates that the oil well's status is unchanged when oil price uncertainty increases. I illustrate the predicted coefficients for the interactions for oil price volatility in section 4 in detail.

Column (1) of Table 4, reports panel probit estimation results in which the choice variable is defined as follows. The dependent variable here is the oil well's production status, and $prod = 1$ for "producing" and $prod = 0$ for "mothballing." The production choice prod=1 if production is chosen and equals 0 if mothballing is chosen. The first part of the table displays the estimated coefficients of the models while the bottom section of the table displays the average marginal effects of each of each of a set of variables related to the price change distribution and the lagged state of the well. The estimated coefficient on *Price* is positive and significantly different from zero (0.266 and significant at 1%) as is lagged production status, *l.prod* (0.412 significant at 1%). Note that the value of "immediate investment" and "delay option" increase with price. However, increases in output price have a more substantial impact on "immediate investment" than on the "delay option." The conclusion is intuitive. The "delayed" cash flows are discounted, and price increases should have a smaller impact on the more discounted cash flows than on the less discounted ones. So, a mothballed oil well is more likely to be reopened by oil producers when oil price increases. For a producing well, the value of the put option of "mothballing" the oil well decreases with increases in output price, making it less attractive to shut in the wells, thus leading to the positive coefficient of price. NPV analysis leads to a similar conclusion that "shutting in" the wells will reduce the value of the wells, so oil producers will not choose to shut in the oil wells. Thus, both methods predict that a producing well is less likely to be shut in if oil price increases, so the *price* coefficient should be positive.

When the lagged production status is "producing," high expected price volatility is associated with a higher probability of choosing "producing" as the production status. That is, oil producers decide not to temporarily shut in an oil well immediately since the shut-in put option value increases, thus, it is optimal to wait (to not exercise the option) rather than to shut in immediately. Increases in price volatility increase the put option value. The fact that lagged production status as well as volatility have an impact on the choice is consistent with producers recognizing the presence of sunk costs and future uncertainty.

Price is not the only relevant variable, characteristics related to expected future production as well as costs and technologies that impact costs should play a role. I account for those characteristics. An older oil well is less likely to be in production or to produce as much in the future relative to a newer oil well. An older oil well's oil production rate exponentially declines after it spuds.[28] It is also more expensive to maintain an older oil well than a newer one – both its operating cost and maintenance cost increase with age, so it is less profitable to operate a more aged oil well. The coefficient estimate on the variable age is negative and statistically significant, as is the marginal effect. Therefore, the probability selecting the production state is lower for older wells. The conclusion is unanimous for whether the lagged production status is "producing" or "mothballing."

Selection of the status producing is positively related to cumulative well production (*cumulative oil production*). More productive wells are more likely to be selected to remain in production than less productive wells, as the expected production rates are higher.

Well productivity will generally depend upon the basin where the well is located. Many of the *basin* location dummies included in the model have significant coefficients ($p<0.10$), for instance, the productive Texas *Permian* basin. This finding suggests that it is crucial to control *basin* which potentially captures the net effect of geographic differences, including rock porosity and the age of the future potential of the location..

6.2.1 Skewness and Kurtosis and the Oil Production Choice

[28] Other than if the producing oil well is re-worked which itself is costly.

As shown earlier, the "skewness- and kurtosis-adjusted option pricing formula" predicts that increases in skewness lead to a reduction in the value of the delay option. As the delay option value declines, the immediate investment value increases, thus, oil producers are more likely to choose immediate investment. In other words, the likelihood of immediate shut down increases. Oil producers may decide to shut in wells immediately or postpone the shut-in plan. The former is determined by the "immediate investment" value (the NPV of shutting in oil wells), and the "delay option" value decides the latter. Larger skewness values lead to declines in the delay option's value increasing the probability of immediate shut down. The impact of skewness on investment decisions is thus opposite of the effect of volatility – while an increase in volatility increases option value, a similar rise in skewness reduces option value. While higher volatility postpones investment and keeps oil wells in current status, higher skewness works in the opposite direction.

When price skewness is high, oil wells are more likely to change status than in a month when skewness is low. Thus, if value maximization prevails I predict a negative coefficient for the interaction of "producing" lagged status and skewness, and I expect a positive coefficient for that of "mothballing" lagged production status and skewness.

In column (2), the interaction of "producing" lagged status and skewness has a negative coefficient (-0.007). That coefficient on "mothballing" lagged status and skewness is positive (0.093) and significant (with a p-value < 0.001). Both results are consistent with the predictions under value maximization in the presence of the real options to produce or mothball. Higher skewness leads to acceleration of immediate investment and a stronger propensity to reopen (a mothballed oil well) or shut-in (a producing oil well). Although skewness is not statistically significant when the lagged status is "producing," the coefficient sign is consistent with the prediction. This finding indicates that the skewness- and kurtosis-adjusted option pricing formula is a helpful tool to predict the impact of skewness (or crash risks and jumps in prices) in price distribution. These results indicate that operators make the choice to produce or mothball consistent with them considering the value of the real options involved.

Kurtosis is a measure of the price distribution's leptokurtic nature and increases with the possibility of outliers. Higher kurtosis is indicative of a distribution with fatter tails. When an option is deep-in-the-money, a higher kurtosis may be related to a higher probability of higher payoff as the result of an outlier price change. The modified option pricing formula predicts that when the option is deep-in-the-money, an increase in kurtosis positively impacts the option's value. When an option is not deep-in-the-money, however, an increase in kurtosis (an increase in the probability of outliers) increases the likelihood of price "jumps" out of the "in-the-money" domain, which makes the option lose value. So, when the option is not deep-in-the-money, increased kurtosis harms option value and accelerates immediate investment. Thus, I predict that kurtosis has the same coefficients as skewness - the interaction between "producing" status and kurtosis should have a negative sign, and the interaction between "mothballing" status and kurtosis should have a positive sign if value maximizing behavior in the presence of these real options. When kurtosis increases, oil well status is more likely to be changed as the delay option value declines, and the immediate investment value is more likely to exceed the delay option value.

As predicted, the interaction of "producing" lagged status and kurtosis has a negative and significant coefficient (-0.192 at 1%), and the coefficient of the interaction of "mothballing" lagged status and kurtosis is positive and significant (0.173 at 1%). Increases in kurtosis accelerate investment, suggesting that an increase in kurtosis reduces the delay option value and makes it less optimal to postpone investment. The results indicate that asymmetry and dispersion have significant impacts on investment propensities. I conclude the results support the predictions that within the context of producing versus mothballing, the operators in the sample behave as if they account for skewness and kurtosis when making these decisions in a dynamic setting. These results are in addition to the influence of volatility in general.

The marginal effects for these variables are also positive and significantly different from zero.[29] The positive coefficient for $price$ (and the positive marginal effect) indicate the probability of selecting to produce is positively related to expected prices. Likewise for the skewness and kurtosis interactions with lagged status, indicating that higher-order moments affect investment choices.

Consistent with the results in panel (1) of Table 4, the coefficient on *age* is negative and significant (-0.307 and p-value <0.001), indicating that older wells are less profitable, consistent with older wells being less productive and more costly to operate. The results on *age* are thus unanimous for the models that include and exclude higher-order moments of the price distribution. I find that the coefficient for *depth* is negative and significant too (-0.427 and p-value <0.001). The estimated coefficient on the variable cumulative oil production *cumulative oil production* is positive and statistically significant. Interestingly however the estimated coefficient on the reserve variable is not significantly different from zero. However, in unreported results excluding basin, drilling types, and depth in the empirical model, reserve becomes significant and it is positive.

The estimated coefficient on lagged production status, *l.prod*, has a positive and significant coefficient. This finding suggests that lagged production status can be an essential determinant of production status. When the lagged production status is "producing," the positive

[29] Marginal effects are a function of the estimated model parameters (see Greene, 2008, ch. 23).

coefficient indicates that the oil well is more likely to stay in the "producing" status. Similarly, when the lagged production status is "mothballing," the coefficient indicates that the oil well is more likely to remain in the "mothballing" status. The switching cost for changing the oil well's production status or the sunk cost that the producers have paid to establish its current status drives this result. The higher the cost to switch from "producing" to "mothballing" (or vice versa), the less likely producers choose to change. Also, the higher the sunk cost of establishing the current production status, the less likely they will be willing to switch.

6.3 The Choice to Drill or Defer

Drilling an oil well costs substantially more than reopening a mothballed well. Permanently closing a well also incurs costs, but it eliminates operating and maintaining costs.I begin with an examination of the choice to drill and in the following section of the paper present an examination of the shut down decision. The examination of the drilling decision is in a spirit similar to Kellogg (2014) and Décaire et al. (2020) except that here I account for the influence of skewness and kurtosis of the price change distribution in addition to volatility. I examine the relationship between drilling an "infill oil well" and the magnitude of price distribution moments to explore the effects of price uncertainty on irreversible investment decisions, as the uncertainty affects investment options value.[30]

I examine the impact of the expected price and expected skewness and kurtosis on the decision to drill. Following Kellogg (2014) and Décaire et al. (2020), I examine infill wells only and use the Cox proportional hazard model. Examining infill wells allows me to focus on a set of

[30] Schlumberger describes infill drilling as "The addition of wells in a field that decreases average well spacing. This practice both accelerates expected recovery and increases estimated ultimate recovery in heterogeneous reservoirs by improving the continuity between injectors and producers. As well spacing is decreased, the shifting well patterns alter the formation-fluid flow paths and increase sweep to areas where greater hydrocarbon saturations exist." https://glossary.slb.com/en/terms/i/infill_drilling

decisions that are as reflective of new decisions as possible. To identify infill wells, I exclude the first wells in a field (the exploratory wells) and keep only the first drilled well whose spud dates are between Mar 2010 and Sep 2019 when oil price data are available in my dataset. I drop wells with spud dates before or in Mar 2010 or which are shut down after Sep 2019. Undrilled fields are identified as "unexercised options" with *drilling*=0 through date range for all months up to the spud date at which I assign *drilling* = 1. There are 32,817 options among which 12,729 are exercised.

The analysis of the option to drill when considering infill drilling has the feature that as new wells are drilled the availability of new drilling options falls. As Kellogg (2014) points out a probit specification (or for that matter the linear probability model) is therefore not the appropriate structure to account for this dynamic and Decaire et al. (2020) agree in their analyses. However, the Cox proportional hazard model does fit the setting described.[31]

I construct the dataset of the unexercised "drilling" options by assigning *drilling=0* for the months between the beginning of a lease's production and the last month before drilling. The dependent variable *drilling=1* for the spud month. So, the months with *drilling=0* are when the wells remain undrilled, and the months with *drilling=1* are when drilling wells. Each well enters data when it becomes available for drilling (when the first well in the lease spuds so oil producers can prepare an infill oil well to boost oil production). This oil well has *drilling=0* for the months after it enters the sample up to the month of drilling, but after it no longer is an opportunity as it has been drilled. Hence at any date up to drilling there is a probability that it

[31] The Cox proportional-hazards model (Cox, 1972) is a model used for investigating the association between the survival time of an entity, person, etc., and one or more predictor variables. Basically, it is formulated to model how a set of specified factors influence the rate of a particular event happening (e.g., drilling a well) at a particular point in time. This rate is commonly referred as the hazard rate. See Cameron and Trivedi (2005, Ch. 17) for a review of the Cox Proportional Hazard Model.

can be drilled conditional on the economic factors prevailing at the time. Hence, this dynamic fits the Cox model nicely.

I truncate the months after the spud dates.[32] I report the estimation results in column (4) of Table 5. The results indicate that price is positively related to the probability of drilling and that volatility is negatively associated with the likelihood. The conclusions are consistent with existing literature on real options investment. The result that price (volatility) is positively (negatively) related to the probability of drilling is consistent with the results presented in Moel and Tufano (2002), Kellogg (2014), and Décaire et al. (2020). Anderson, Kellogg and Salant (2018) find that oil well production does not repond to price changes.

I find evidence that higher-order moments influence the choice to drill. Column (4) of Table 5 shows that when kurtosis increases by 1.00, the probability of drilling a new oil well increases by 44.34% (exponential of 0.367). It suggests that the tail risk in the price distribution of output influences the decisions to invest. Although skewness is not significant, its coefficient (0.039) suggests that it is consistent with the prediction of the Real Options Valuation and modified option pricing models.

I also find that a well's oil reserve is an essential determinant for drilling decisions – *reserve* has a positive and significant coefficient (0.142 and p-value <0.001). The higher the potential of being a productive oil well increases the probability of being drilled. Oil producers expect a higher production rate for wells with high reserves, and the high reserve oil wells are more likely to bring higher expected profits. The expected profit affects both the delay option value and the immediate investment value. However, the "immediate investment" value increases more than the "delay option," causing a higher likelihood of drilling a high reserve oil

[32] The hazard model will not use the months after the event happens (when the dependent variable is 1). Also, the hazard model will not use any unexercised options for estimation. These observations (after exercised and never exercised options) contain no useful information for estimating the hazard model.

well. The deflated drilling cost has a positive sign, but it is not statistically significant. In an unreported result, I exclude *basin*, *depth*, and *drill*, and deflated drilling cost *cost_def* has a negative and significant coefficient. This finding suggests that the controls correlate with *cost_def*. I expect the drilling cost to have a negative coefficient as the cost decreases project value. An increase in drilling cost is analogous to the rise in option strike price, thus reducing the delay option value. However, the immediate investment value will decrease more than the delay option value.

Next, I exclude skewness and kurtosis from the econometric model to verify the existing literature showing that volatility is negatively related to the possibility of "immediate investment." I present the results in column (3). *Volatility* has a negative and significant coefficient (-5.453 and p-value<0.01). When controlling for *price*, *depth*, *drill*, *basin*, *reserve*, and *cost_def*, increases in price uncertainty lead to declines in the propensity to invest immediately. Higher price uncertainty is positively related to a delay option value increase, so postponing investment is more likely to be the optimal choice for investors. I show that *price* remains a positive and significant coefficient and whose effect rarely changes after removing *skewness* and *kurtosis* from the model (1.267 and p-value<0.01). Price remains the direct impactor for investment choices – a higher price is directly related to a higher expected payoff from investing and delay option value; however, the increase in value of the "immediate investment" exceeds that in value of the option.

6.4 The Choice to Permanently Shut-down (plug) or Continue Producing

I define "shutdown" (*shutdown=1*) as the last month of production for the oil well. *Shutdown=0* up until the final month of production, beginning with the in first production date and ending the previous month before closing. *Shutdown=0* when the well begins production.

The well has *shutdown=1* in its last production month. I exclude the months when wells stop producing in the middle of their lives but then resume as these constitute temporary shut-ins (mothballing). The cost of reopening a permanently shut-down well is prohibitive and likewise as mentioned earlier switching to mothball status would cause a shift from no periodic maintenance cost to a per period cost, which would make no economic sense.

The choice is therefore to shut-down or to defer shutting down. The delay option value depends on the expected oil price, price volatility, skewness, and kurtosis. As the modified option pricing formula presented in Section 4 shows, when skewness increases and the call option is out-of-the-money, a call option's value increases. By put-call parity, when the put option is in the money and skewness increases, the call option value increases, so the put option value decreases. A similar analysis can apply to the impact of kurtosis increase on put option value. The prediction is that when kurtosis increases, a deep-in-the-money put option increases in value. Thus, I expect the put option value to decrease when skewness increases, and I anticipate that the put option value increases when kurtosis increases. If the delay option value increases and the immediate investment value remains unchanged, producers are more likely to postpone investment, as the delay option's value exceeds the value of the decision to immediately shut-in the well. Suppose the delay option's value decreases and the "immediate investment" value remains unchanged. In that case, maximizing producers make the immediate decision to close the well. So, when skewness increases, producers are more likely to select to close. When kurtosis increases and the option is deep-in-the-money, producers are more likely to postpone the shutdown choice. Thus, I posit that the probability of the choice to shut down is increasing in skewness. Likewise, the probability of closing is negatively related to kurtosis if the shutdown option is deep-in-the-money. I have made such predictions in the hypothesis development section.

The empirical results are unsurprisingly consistent with these predictions. I report these results in column (6) of Table 6. The results indicate that price is negatively related to the probability of closing an oil well. I showed that price is positively related to the likelihood of drilling, increasing the expected payoffs from drilling an oil well. Such value increases are more for "immediate investment" than "delay option," as the latter is discounted more by time. Similarly, when price decreases, oil producers are more likely to choose to close down the oil wells, concerning that the declines in oil price will make an operating oil well unprofitable. So, when the opposite happens, oil producers are less likely to close down the oil wells. The oil producers have two choices – the first is to shut down the oil wells immediately, and the second is to postpone the shutdown plan. Both become less attractive as oil price increases; however, the first choice is even less appealing to producers as the increases in oil price affect the expected payoffs from the first choice more than the second one. So, when oil price increases, oil producers are more likely to postpone shutdown plans. I also find that the coefficient of *volatility* is negative, indicating increased expected volatility leads to an increased value for the delay option and has a lower probability of selecting to shut-in. When price volatility is high, there is a better chance that, in the future, a price increase will make the shutdown plan unattractive. So, if increases in price volatility increase the delay option value, producers tend to choose to delay shutting down wells. I should obtain a negative coefficient for volatility, and the empirical results suggest that *volatility* is negatively related to the probability of shutdown, verifying ROV predictions. To summarize, the empirical results by far with price and volatility are consistent with existing literature and verify the validity of Real Options Valuation in solving irreversible investment problems.

More importantly, I find that skewness and kurtosis are positively and negatively related to the probability of oil well shutdowns, respectively, which conclusion is consistent with the prediction of Real Options Valuation. In column (6), the coefficient of skewness is significant at the 1% level. It suggests that, when increases in price skewness, i.e., increases in the probability of positive jumps in prices, producers are more likely to shut down oil wells immediately. This finding seems counter-intuitive at first; however, the put option value decreases, suggesting that it is not worth waiting to proceed with the shutting down plan. The expected payoff from the immediate shutdown does not change with skewness but is related to the expected prices. The result is consistent with the immediate shutdown value being perceived as greater than the value of waiting. The coefficient of kurtosis is not statistically significant; however, its sign (-0.083) is consistent with the prediction that kurtosis increases the delay put option value. Increases in the probability of price outliers reduces the value of the immediate decision. To summarize, I find that extreme positive shocks accelerate a close-down plan, and outliers on both sides decrease the value from making such decisions immediately.

I also find that (oil well) depth is positively related to the probability of shutdown. Deeper wells have higher operating costs than shallower wells; thus, it is more likely to become optimal to shut down. I include *basin* in the econometric model, showing that the dummies are significant explanatory variables for investment decisions. I use *last12* as a proxy for the productivity of the well. If the oil well remains productive, its expected payoffs from future production are higher than a less effective oil well. I show that *last12*, i.e., the well’s last 12-month’s production rate, is negatively associated with the propensity to close down the well. An oil well's production rate decays with time, and its decline is exponential. Thus, even for the same well, its profitability changes with time. A more productive and newer well is more profitable to keep in the production state than a less effective and older well. Similar to what I did before, I ran a similar regression excluding skewness and kurtosis to explore the validity of ROV in explaining irreversible investments in the existing literature in column (5). I show that both price and volatility are negatively associated with the probability of shutting down wells (as reported in column (5)), which is consistent with the full model's results. The coefficients of other variables largely remain unchanged.

6.5 Preference Theory

If managers (decision makers) slant the investment decisions I examine to suit their own personal preferences then the relationship between skewness, kurtosis, and real options investments could be driven by personal investment preferences for skewness and kurtosis. If that was the case the empirical results should have been different from what I find. If managers have a personal preference for positive skewness and that colors the decisions I study, then I should have observed that oil wells are more likely to be drilled or operations resumed when skewness is more positive and that they are more likely to be closed or mothballed when skewness is more negative. If, however, investors have a preference over negative skewness, oil wells are more likely to be drilled or resumed when skewness is more negative and they are less likely to be closed or mothballed when skewness becomes more negative. My empirical results show the relations between skewness, kurtosis, and real investments of oil wells is consistent with the predictions of the Real Options Valuation and the modified Black-Scholes option pricing model. This finding suggests that choices studied are consistent largely with value maximizing behavior as laid out in Section 4.

6.6 Summarizing the Results

The empirical results presented in Tables 4, 5 and 6 are found to generally be consistent with the hypothesis that production versus temporary shut down decisions, drilling versus waiting to drill and permanent shutdown versus continuing to produced choices are consistent with value maximizing behavior in the presence of real options (price uncertainty and sunk costs). When comparing the full model (that includes skewness and kurtosis) and the volatility model (everything the same but excluding skewness and kurtosis), I present evidence that the model that includes skewness and kurtosis, in addition to volatility, better explains the changes between the "producing" and the "mothballing" status of oil wells. The choice to drill or delay drilling, while determined in part by expected price change volatility, is influenced only by kurtosis not skewness. Likewise, the choice to continue production or permanently shutting down a well is determined by skewness and volatility but kurtosis does not play a role. The results are found to be robust to alternative specifications of the models to which I now turn.

6.7 Robustness of the Results

6.7.1 Unobserved Heterogeneity in the Production/Temporary Shut-in (Mothball) Model

I investigate whether unobserved heterogeneity across wells affects the main results presented in Tables 4-6. There may exist substantial unobserved heterogeneous differences among oil wells that determine production statuses. Determinants other than depth, age, drilling type, or basin may affect the profitability of the oil wells, such as operating costs. Using a robust model controlling for unobserved heterogeneity, I implement the simple initial condition solution proposed by Wooldridge (2005) to solve the unobserved heterogeneity issue in random-effects panel probit model. In addition to the initial model's specification, I include add the lagged values (initial conditions) of the dependent variable, cumulative oil production, and age's initial values, and averages of cumulative oil production and age to control unobserved heterogeneity among wells. I find that the empirical results rarely change. I obtain positive and significant coefficients for lagged production status, price, the interaction of volatility and “producing” lagged status, the interaction term of skewness and “mothballing” lagged production status, the interaction term of kurtosis and “mothballing” lagged status, and cumulative oil production. Also, I continue to obtain negative and significant coefficients for the interaction term of kurtosis and “producing” lagged status, age, and depth. This result suggests that my original model in section 6.2 does not suffer substantially from the unobserved heterogeneity issue.

6.7.2 The Influence of Nearby Drilling Activity

Décaire et al. (2020) find that the propensity for oil producers to decide whether to drill a new oil well is influenced by the drilling activity of neighbor producers, suggesting that an information spillover effect influences producers decisions to drill. I construct measures of proximity in the following manner: using the arrangement of sections within a lease,[33] I can identify the “neighbors” of an oil well with the section numbers. A section can have up to eight neighbors. For instance, section 16 has neighbor sections 8, 9, 10, 15, 17, 20, 21, and 22. However, some sections have fewer neighbors within the lease. For instance, section 1 only has neighbor sections 2, 11, and 12. A well lease can be identified by its range and township values which, along with section numbers, are available from the Enverus’ database. After identifying each well’s neighbors, I count the number of drilled options among these neighbors. These are the number of exercised options around the well. This number serves as a source of spillover information, as it indicates the productivity of the oil wells and reserve underground. I am able to split between the oil well’s firm’s own drilled neighboring wells and drilled wells from other

[33] An example of the map of the arrangements of section within a well lease can be accessed: https://www.geomore.com/locating-wells/.

firms. Both numbers are included in the empirical model to examine the effect of proximity options exercises on oil well drilling decisions.

Including these measures of proximity in the base drilling model, I find that my basic results regarding the relation between the choice to drill, volatility, skewness and kurtosis are unchanged. However, I do find that proximity does matter in my sample, that the probability of drilling is positively related to proximity.

6.7.3 The Regulatory Environment of the State in Which a Well is Located

While the activities I have been studying are all potentially affected by U.S. Federal oversight and regulation, the regulatory environments may differ across the states included in the sample (Texas, Oklahoma, California, North Dakota and Pennsylvania). The models evaluated to this point control for the basin location of wells, that only indirectly controls for the state regulatory environment. Evaluation of these states indicates that the regulatory environment is significantly more restrictive in California relative to the other states represented in the sample. I divided the sample into two subsamples, one which includes only activity in California and the second which includes all the remaining states. The variable *reserve* is replaced by the first 12-months' production rate (*first12*) as reserve is not available for California oil wells. I repeat the estimation of the drilling model for each subsample. The results with only California oil wells and excluding California oil wells do not differ.

6.7.4 Did the Focus on Environmental, Social And Governance Issues (ESG) Impact the Production, Drilling and Shut-down Decisions?

I end by separating the sample into two time periods. Up to the end of 2015 and after. The choice of point at which to divide the sample period is based upon the Paris Climate Agreement that was adopted by 196 Parties at COP 21 in Paris, on 12 December 2015 and entered into force on 4 November 2016. [34] Inspection of the Google Trend Index based upon the phrase 'ESG' shows there was a sharp escalation in attention to ESG issues following 2015 and so seems to be a natural point of demarcation. Oil and natural gas producers were directly or indirectly impacted by the adoption of this accord and may as a consequence have significantly altered how they make decisions. I find that the ESG shift has no impacts on the effects of higher-order moments (skewness and kurtosis) of price distribution of oil after 2015 – the coefficients on skewness and kurtosis are still consistent with the Real Options Valuation and modified option pricing models' predictions.

7. Summary and Conclusions

In this study, I show that implied skewness and implied kurtosis significantly affect irreversible investment decisions by finding the association between oil price implied skewness and kurtosis and oil well production. By focusing on production histories for all oil wells domiciled in California, North Dakota, Pennsylvania, Texas, and Oklahoma from 2010 to 2019 and relating the changes in producing status to oil price moments (volatility, skewness, and kurtosis), I show that oil producers account for higher order oil price moments consistent with the predictions of value maximizing behavior in the presence of real options (price change uncertainty and sunk costs). The choices to produce or temporarily shut down,are determined by expected skewness, kurtosis and volatility of the price change distribution in addition to other well-specific characteristics. The choice to continue producing or permanently shut down are determined by expected skewness and volatility. The choice to drill or wait to drill is determined however only by expected kurtosis and volatility. The evidence overall indicates that sunk costs matter in the presence of future price uncertainty. The choices of the oil producers in the sample

[34] See https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement

are consistent with value maximizing behavior in the presence of real options, however not all choices are influenced by skewness and kurtosis despite the fact that in theory the value of the real options involved are influenced by those parameters of the price change distribution.

Chapter 2

Forecaster Accuracy, Herding, Anti-herding, and Timing Over Time

Coauthored with Louis H. Ederington and Scott C. Linn

"What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention…" (Herbert Simon, 'Designing Organizations for an Information-Rich World' in Martin Greenberger (ed.) *Computers, Communications, and the Public Interest* (1971))

"Information quantity is not equal to information quality or information value, but the second requires the first." (Hilbert, "How to Measure "How Much Information"? Theoretical, Methodological, and Statistical Challenges for the Social Sciences", *International Journal of Communication* 6 (2012), 1042–1055

"Can everyone just stop whining about information overload? I mean, in the knowledge economy, information is our most valuable commodity*….[but],*
The flood of information that swamps me daily seems to produce more pain than gain…. Current research suggests that the surging volume of available information—and its interruption of people's work—can adversely affect not only personal well-being but also decision making, innovation, and productivity." (Paul Hemp*, Harvard Business Review*, September 2009)

1. Introduction

We have witnessed an explosion of accessible financial and commodity market related information over the last two decades. By two accounts roughly half the U.S. adult population had broadband access in the year 2000, while the number exceeds 90% today (World Economic Forum, 2020; Pew Research Center, 2021). A leading expert on information theory has stressed that the digital age essentially began around 2002 and has grown exponentially since, emphasizing however that information quantity does not necessarily imply information quality or value (Hilbert, 2012). Such access has vastly expanded the information set available to amongst others, professional forecasters. A critical element in the assessment of a forecast's value is, of course, its accuracy. In this study we examine how forecaster behavior and forecast accuracy changed over the period 2003 through 2021. We investigate the relations between forecast accuracy, the propensity and intensity of forecaster herding, and the timing of forecast issuance by forecasters, and how these relations have changed over the roughly 20-year period we study. Our investigation sheds new light on these relations over a period of increasing information availability.

Analysts forecasting economic variables, whether of a micro or macro nature, face several decisions, each of which could affect the accuracy of their forecasts. First is the decision whether to form and release their forecast before others, thereby satisfying a demand by users for early but potentially less accurate forecasts. Intertwined with this choice is the companion decision of whether to release a less accurate versus a more accurate forecast. Conversely a forecaster may choose to delay releasing a forecast until close to the time when the actual data being forecasted are disclosed, thereby accumulating more information to reflect in the forecast and possibly improving the forecast's accuracy. Finally, the analyst must decide whether to herd—that is, whether to base the forecast predominantly or wholly on the prior forecasts of others rather than relying on private information. We examine the relation between forecast timing release, forecaster herding behavior and accuracy and whether these relations have changed over time.

We view five characteristics as being important for the design of our study. First, the actual values of the forecasted variable should be assembled and released by an independent organization that has no operating relation in the market being studied. Second, the focus should be on a forecast whose accuracy has been shown to have economic relevance, such as through deviations of the forecast from the actual value having an impact on the market price of the asset associated with the forecast. Third, to ensure that the forecasts are not slanted by a prior relationship between the analyst and the user of the forecast, the forecasts should not be related either to estimates of the value or performance of the forecast users. Fourth, the forecasts should not have the potential to influence the actual value of the economic variable being forecasted. Fifth, forecasts should be available over an extended period associated with increasing information availability about the market being studied.

The setting we investigate meets the conditions just described. First, we investigate professional analyst forecasts of changes in the weekly natural gas storage figures reported in the Weekly Natural Gas Storage Report (WNGSR) compiled and released by the United States Energy Information Administration (EIA). The EIA is an arm of the United States Department of Energy and focuses solely on assembling and reporting energy market data. Second, numerous studies have shown a significant natural gas futures price reaction at the time the EIA report is issued to deviations between forecasts of the change in the actual amount of natural gas in storage (for instance, Ederington, Lin, Linn, and Yang, 2019; Gu and Kurov, 2018; Gay, Simkins, Turac, 2009; Linn and Zhu, 2004). Third, unlike the case with earnings forecasts, the variable being forecasted is an aggregate economy-wide variable and is not directly tied to the performance of a particular firm (Lin and McNichols, 1998; Dechow, Hutton and Sloan, 2000; Richardson et al., 2004; Ke and Yu, 2006; Kothari, So, and Verdi, 2016). This eliminates the incentive for analysts to issue biased forecasts that cater to specific clients, as can potentially be the case with earnings forecasts. Fourth, the EIA issues the WNGSR each Thursday morning, but the data being reported are measured as of the prior Friday. Forecasts of the amount of natural gas in storage are issued between the date the actual amount in storage is measured and the date the report of the actual amount is issued by the EIA. Thus, forecast users cannot take actions to manipulate the actual storage numbers and thereby influence the accuracy of the forecasts. Fifth, a large number of forecasters follow and make forecasts for this single variable, ostensibly competing against one another. Sixth, forecasts are issued weekly, which provides both relatively high frequency and a large sample size focused on a single economic variable. Finally, forecasts are available for the period 2003-2021, a period of increasing information availability.

Our results show that, first, some forecasters are consistently more accurate than others and their accuracy persists across time. The mean absolute forecast error for the most accurate forecasters is half that of the least accurate, and the null that forecast accuracy varies randomly across forecasters is rejected at the .0001 level. Moreover, forecast accuracy is highly persistent. Forecasters who have been most accurate in the past tend to be the most accurate in the future, and those who have been the least accurate in the past tend to be the least accurate in the future. Relatively inaccurate forecasters rarely improve their forecast accuracy substantially. We do find however that overall forecast accuracy improved over time. Second, forecasts released shortly before the EIA announcement tend to be considerably more accurate than those released just a day or two earlier, indicating that later forecasters have more and better information and this is true over the two subperiods we study. This suggests that information accumulates over time, enabling later forecasters to issue more accurate forecasts. Thus, our results imply that timeliness (early release) and forecast accuracy tend to be inversely related. Third, based upon three alternative measures of herding (Bernhardt, Campello, and Kutsoati, 2006, hereafter BCK, and Chen and Jiang, 2006, hereafter CJ) we find strong evidence that forecasters either do not herd or engage in an anti-herding strategy and this behavior differs sharply across forecasters and is strongly persistent across time, like forecast accuracy. The null that differences in herding strategies differ randomly across forecasters is rejected at the .0001 level, and forecasters generally stick to the same herding strategy over time. Fourth, the most accurate forecasters neither herd nor anti-herd, while the less accurate tend to anti-herd. This gives rise to an association between greater anti-herding behavior and greater inaccuracy. This relation does not change over time. Fifth, together with forecast timing, differences in herding strategy explain on average over 50% of the variation in forecast accuracy. Studies of forecasts of other economic variables have focused on determinants of forecaster accuracy, such as forecaster skill, experience, and available resources We find that in the setting we study, we find that over 50% of the cross-sectional variation in forecaster accuracy is due to strategic choices made by the forecaster—specifically, how much to herd and how far in advance to release forecasts.

We provide additional evidence on the behavior of anti-herding by the forecasters in our sample using individual forecaster panel models in which we control for firm fixed effects and separately forecaster fixed effects. These results also support the conclusion that anti-herding behavior is the norm. We find that anti-herding on average declined between the two subperiods, and that this can be explained by the following. First, the top 20 forecasters in terms of accuracy are on average associated with less anti-herding while the bottom 20 forecasters on average exhibit greater anti-herding and these findings did not change between the two subperiods studied. Finally, within a subperiod, anti-herding is most intense 3 days prior to the release of the EIA's report, declines 2 days out and increases the day before the EIA release. The finding does not change for the 2-day out cases between the subperiods, but anti-herding declines for the 1 day out cases during the second subperiod.

The information obtained from the source, Bloomberg, provides the order in which the forecasts are released for each date for a subset of the forecast dates studied. Based upon this subset of data, which represent within- day releases, we find that our base results on herding are further confirmed. Several additional robustness tests also confirm our basic findings.

Section 2 provides a review of predictions regarding the relations between accuracy, information, herding and anti-herding and forecast release timing. Section 3 provides a discussion of related literature and how this study contributes. We provide a first look and discussion of the sample forecast errors in Section 4. Section 5 provides an examination of the herding behavior of the forecasters in the sample. Section 6 presents and discusses results on the joint impact of herding and timing on forecast accuracy through the estimation of between-firm models as well as panel models focused on individual analysts. In section 7 we explore the behavior of anti-herding in more detail. Section 8 provides a discussion of several robustness tests we performed. Section 9 presents our conclusions.

2. Information, Forecaster Accuracy, Herding, Anti-herding and Timing

2.1 Information

An increasing inventory of information can be beneficial or alternatively might cause information overload and be a detriment. We test whether forecast accuracy, herding, anti-herding and forecast timing changed between the two subperiods we study.

Professional forecasts of economic variables reflect analysts' private information, but also the views of others (public information), both of which are a function of total available information and how forecasters learn. Theory suggests that an expanded information set from which users learn could lead to improvements in accuracy or at worst no improvements, but not degradation. Such would be the case if forecasters are rational users of information and do not face capacity constraints or costs in their ability to process information, unhindered leading to a rational expectations (RE) equilibrium. Foundational examples focusing on convergence to a RE equilibrium include Bray (1982), Bray and Savin (1986), Lucas (1986), Marcet and Sargent (1989a), and Woodford (1990). The implication being that such forecasters would always select the optimal use of public and private information. Conversely, an expanded information set could be associated with information overload. Information overload could include an overflow of information, information that is ambiguous and the presence of limitations on the ability to absorb and understand information especially in a timely manner (Eppler and Mengis, 2004). Such a situation could thereby result in more reliance on the beliefs of others, and potentially a drift away from the optimal use of public and private information (Bernales, Valenzuela, and Zer, 2023). This would especially be the case in the face of limited information processing ability or limited attention on the part of individual forecasters (Kahneman, 1973; Johnston and Pashler, 1998; Lang, 2000; Sims, 2003).

2.2 Herding and Anti-herding

A forecaster that concludes a forecast based on the forecaster's private information would differ substantially from the publicly available forecasts of earlier forecasters might ignore or underweight his private information and adjust his forecast toward those released earlier, even though it might otherwise be optimal to not do so. That is, the forecaster might herd by overly tilting his forecast towards the prior public consensus (BCK, CJ). Conversely, if a forecaster tilts her forecast more towards her private information and ignores the publicly available information (i.e. the consensus) she would be anti-herding. The remaining case is when the forecaster strikes the optimal balance between public and private information and neither herds nor anti-herds relative to the optimal weighting scheme.

The relation between forecast accuracy and herding or anti-herding rests upon how forecasters weight private versus public information when composing their forecasts. Given an optimal weighting scheme between private and public (consensus) information, the prediction is that forecasts based upon those weights would have a lower mean squared error relative to the actual value of the variable being forecasted than forecasts that tilt away from those weights in either direction. The expectation is therefore that accuracy will be worse for both anti-herders and herders relative to forecasters who neither herd nor anti-herd (those who form their forecasts optimally).

2.3 The Timing of the Forecast Release

The timing of a forecast's release relative to when the actual is reported involves a trade-off between accuracy and factors that may influence the supply or demand for forecasts other than accuracy. If fundamental information accumulates, allowing better forecasts to emerge as the EIA report date approaches, we should see that later forecasts are more accurate than earlier forecasts. This of course raises the question then of why forecasts would be released early. Rosen (1981) and Scharfstein and Stein (1990) have argued that a "superstar" (forecaster) effect may manifest if the best forecasters are rewarded monetarily but also with reputation which may include the reputation of being first or breaking away by being early. Ehrbeck and Waldmann (1996) define what they label "rational cheating" as a description of the situation in which forecasters, "…balance their aims of minimizing forecast errors and looking good before the outcome is observed" (p. 22) giving rise to biased but early forecasts. Laster, Bennett, and Geoum (1999) develop a model of competition amongst forecasters who are compensated for both the accuracy of their forecasts and the attraction (through publicity) that their forecasts bring to the firms that employ them, essentially a reputation effect. The upshot is that each forecaster faces a tradeoff between accuracy and publicity and will make the choice depending on how their employer weights, either explicitly or implicitly, the two factors in setting the forecaster's compensation. What the actual relation is between forecast timing and accuracy is of course an empirical question. We test whether later forecasts are more accurate than earlier forecasts and whether the relation has changed over time.

2.4 Herding, Anti-herding and Timing

The impact of timing and herding on forecast accuracy may depend upon the confluence of the two effects. That is, there may be an independent effect of each but also an interaction effect. For instance, as the EIA release date approaches and forecasters' information becomes more consistent and similar, the marginal benefit to the analyst whose forecast signals unique private information may be larger than for forecasts issued at earlier dates. However, the marginal cost to the analyst may be higher for issuing a forecast that deviates from the consensus. Finally, there may be no interaction effect. Which conjecture dominates is an empirical question. We account for these effects by incorporating an interaction variable that is the product of the herding and timing measures.

3. Related Literature

The studies by Gay, Simkins, and Turac (2009) and Gu and Kurov (2018) explore the accuracy of forecaster predictions about the amount of natural gas in storage and find significant differences among natural gas forecasters' accuracy. However, these authors do not explore the reasons for forecast accuracy differences nor whether there have been changes over time. We take up these two issues in this study along with the question of whether the impacts of accuracy determinants have changed over time.

The primary studies that have focused on the determinants of forecaster accuracy have concentrated on the accuracy of earnings forecasts produced by professional analysts (see the surveys: Kothari, So, and Verdi, 2016; Ramnath, Rock, and Shane, 2008). Although some early studies found no significant differences in forecast accuracy among earnings forecasters, Stickel (1992) and Sinha, Brown, and Das (1997) find systematic and persistent differences in forecaster accuracy. Clement (1999) reports that more experienced analysts and analysts employed by large firms tend to be more accurate. He also finds that forecaster accuracy is negatively correlated with the number of firms followed by the analyst as does Hirst, Hopkins, and Wahlen (2004). Both Brown (2001) and Michaely et al. (2023) find that earnings analyst accuracy is persistent. Whereas Clement (1999) reports that forecast accuracy increases with forecaster experience for earnings forecasters, Lamont (2002) finds that as forecasters of macroeconomic variables became older and more established, they produce more radical and less accurate forecasts. Therefore, one of the significant differences between our study and the extant literature is an examination of forecast accuracy over time, and more specifically the relation between accuracy, herding and the timing of forecast releases over time.

Studies of forecaster accuracy do not examine how aggregate information accumulation has impacted accuracy, including for example, Coval and Moskowitz (2001), Ivkovic and Weisbenner (2005), Malloy (2005), Engelberg and Parsons (2011), and Miller and Shanthikumar

(2012). We address this question by studying accuracy over two distinct time periods during which information availability has expanded exponentially.

Our study also contributes to the literature on forecasters' herding behavior by providing empirical evidence on whether herding has changed over time, and importantly whether the relation between herding and forecast accuracy has changed over time.[35] Several authors have presented evidence that at least some corporate earnings forecasters and stock analysts herd, including, Trueman (1994), Hong, Kubik, and Solomon (2000), Welch (2000), Clement, and Tse (2005), Jegadeesh and Kim (2010), Clement, Hales, and Xue (2011), and Huang, Krishnan, Shon and Zhou (2017), and Altınkılıç, Balashov and Hansen (2019). Zitzewitz (2003), Bernhardt, Campello, and Kutsoati (2006) (BCK) and Chen and Jiang (2006) (CJ) argue based upon the evidence they present that most common stock analysts anti-herd, meaning that they tend to overweight their private information relative to what would be optimal behavior. [36] These findings suggest that the question of whether forecasters herd, anti-herd or neither (that is, act optimally) has not been fully resolved. We present evidence on this question through an examination of a sample of forecasters that avoids several issues that could influence earnings forecasters, but yet the laboratory for analysis we select is associated with forecasts of an variable that are closely watched by traders in an important energy market, the market for natural gas. Importantly, studies of past herding behavior do not investigate whether such behavior has changed over time. Our results provide new insight into how herding behavior has changed during a period of increased information availability.

[35] For reviews of the theoretical literature on herding including informational herding and the area of informational cascades see SHirshleifer and Teoh (2003, 2009) Hirshleifer (2020), Bikhchandani, S., Hirshleifer, Tam, and Welch (2021).
[36] See Ramnath, Rock, and Shane (2008) and Spyrou (2013) for other reviews.

With several exceptions, such as Lamont (2002) who studies GDP growth estimates, few studies document or explain differences in herding behavior between forecasters of macro-level variables, which is another contribution we make to the literature. Gallo, Granger, and Jeon (2002) find evidence consistent with herding in GDP forecasts while Lamont (2002) finds that in forecasting GDP growth, older forecasters tend to issue forecasts further from the consensus than younger forecasters. Pierdzioch, Rulke, and Stadtmann present evidence of anti-herding among oil *price* forecasters (2010) and metals price forecasters (2013). Bewley and Fiebig (2002) find evidence of herding in interest rate forecasts, but Ehrbeck and Waldmann (1996) and Pierdzioch and Rulke (2013) find that interest rate forecasters tend to anti-herd. Broughton and Lobo (2018) find that forecasters of the Purchasing Managers' Index (PMI) issued by the Institute of Supply Management tend to anti-herd. We present both cross-sectional evidence as well as evidence across time of the herding and anti-herding behavior of forecasters of natural gas in storage, a macro variable that is widely followed by energy traders and the greater energy industry.

In a recent study, Fernandez-Perez, Garel, and Indriawan (2020) also present tests of whether natural gas storage forecasters engage in herding. However, their analysis uses data that would not have been available to forecasters before the EIA report release date. As explained below, the herding measures we compute involve comparing a forecaster's individual forecast with the consensus of *prior* forecasts when the individual's forecast is made. In addition, we present results for multiple herding measures, establishing the robustness of our conclusions. Another distinguishing characteristic of our study is the examination of whether herding and anti-herding have changed over time as well as whether the relations between these measures and forecast accuracy have changed as well.

Clement (1999) and Sinha, Brown, and Das (1997) find that late earnings forecasts tend to be more accurate than early forecasts. Cooper, Day, and Lewis (2001) explore the importance of forecast release timing for earnings estimates and argue that earnings analysts have strong incentives to be among the first to release their earnings estimates. They find that stock prices respond significantly more strongly to earnings forecasts by forecasters who are generally among the first to release their forecasts than to forecasts by historically late forecasters. In contrast, they do not find a significant difference in the price response to forecasts by the historically most accurate forecasters versus the less accurate. This suggests that investors value earlier forecasts. We specifically provide evidence on the timing of forecast release and its impact on accuracy and whether this relation has changed over time.

4. Forecast Errors Pre and Post Year-End 2012

Each Thursday at 10:30 AM ET, the EIA releases its survey of natural gas in storage in the United States measured as of the preceding Friday.[37] Although the report presents storage levels by U.S. geographic regions, the headline figure is the total of working gas in storage for the lower 48 U.S. states in billions of cubic feet (Bcf).[38] Bloomberg collects analyst forecasts of the change in total gas in storage by email or phone beginning several days prior to the EIA announcement. Bloomberg reports these individual forecasts as they are received, along with the average, median, standard deviation, and high and low values. Bloomberg publishes the date, but not the time, that it receives each individual forecast. We collect these individual forecasts along with the final summary measures for each week during the period 3/4/2003 through 12/9/2021 (978 weeks). The data set includes 19,677 individual storage forecasts, or an average of roughly

[37] Details about the compilation of the Weekly Natural Gas Storage Report by the U.S. EIA are available at http://ir.eia.gov/ngs/ngs.html.

[38] Natural gas is generally stored in depleted oil or natural gas fields or in salt caverns. To maintain pressure, these facilities require the presence of a cushion, or an amount of "base" gas. Working gas in storage is the amount stored in addition to the base gas. https://www.eia.gov/tools/glossary/index.php?id=W#work_gas

20 forecasts per week by 102 different firms.[39] We divide the sample into two time periods: Subperiod 1 spans 3/4/2003 through 12/30/2012 and Subperiod 2 spans 1/1/2013 through 12/9/2021. "Forecast Error" (FEr) is measured as the difference between the individual forecast and the actual change subsequently announced by the EIA and the absolute forecast error, AFE, is the absolute value of FEr.

The data indicate that absolute forecast errors are highest in January, decline monotonically through May, are low through the summer and fall, and then rise in November and December. We control for seasonality by computing what we refer to as the relative absolute forecast error (RAFE) for a forecaster. The RAFE equals the forecaster's absolute forecast error divided by the mean AFE for the month in which the forecast is issued. In addition to controlling for seasonality of the absolute forecast error, the Brown-Forsythe Modified Levene test does not reject the null hypothesis of equality of variances of RAFE across the months-of-the-year (p-value .101). Therefore, our adjustment controls for both mean and variance effects due to seasonality.

Descriptive statistics are reported in Table 7 for the full sample period along with the two subperiods. In Panel A (full sample results) we find that the mean forecast error (FEr) is negative (-0.368) and significantly different from zero (p-value for test that mean equals 0 is .000), but the median FEr is zero. The results also indicate there is considerable cross-sectional variation in the forecast errors. The average absolute forecast error equals a positive 8.542 bcf and the median is 6.000 bcf when considering the full sample. For comparison purposes note that the median forecasted storage change equals 40. The standard deviation of the forecast error is large relative

[39] The total of 102 firms is after consolidating observations where the same firm was listed under two or more slightly different names on different dates. In most cases, the same firm name was recorded slightly differently on different weeks, such as an abbreviated name some weeks and a full name in other weeks. In other cases, firm names changed slightly over our data period. There were 146 names prior to consolidation. Economist names were much more consistent.

to the mean. Together these statistics indicate there is variation of forecast errors across forecasters. The subperiod statistics show that in the first period (2003-2012) and the second period (2013-2021), the mean FEr is negative but is less negative during the second period. Likewise, both the mean and median Absolute Forecast Error (AFE) are smaller during the second period, as is the standard deviation of this variable.

Relative absolute forecast error (RAFE) is our study's primary measure of forecasting error. As mentioned above, the mean forecast error varies across the months of the year. The seasonal corresponds to a seasonal in total storage that is driven by a seasonal in natural gas consumption. Similar to AFE the mean RAFE is smaller during the second period as is the median. RAFE displays less skewness and kurtosis in the first period than in the second period.

We test whether the means and medians of the three variables are equal between the first and second subperiods. The results indicate that there are statistically significant differences between the two periods for both AFE and RAFE, the variable we focus on henceforth. Here and in subsequent tests we employ a conservative 1% significance level criterion. Both the mean and median RAFE are smaller during the second period, indicating by this measure that accuracy improved between the two subperiods. However, the standard deviation of RAFE relative to the mean indicates considerable variation across forecasters.[40] Mean RAFEs for the four least accurate forecasting firms are 1.9 times the four most accurate during the first subperiod, and 1.7 times greater in the second subperiod. The null that forecast accuracy does not differ across forecast issuing firms is rejected at the <1% level for both subperiods (Subperiod 1: F =14.46,

[40] Gay, Simkins, and Turac (2009) and Gu and Kurov (2018) explore natural gas forecaster accuracy for the 1997–2005 and 2008–2016 periods, respectively, and reject the null that forecast accuracy does not differ among forecasters. However, those authors do not explore why some forecasters are more accurate than others nor whether there have been changes over time.

p<.01; Subperiod 2: F =19.47, p<.01).[41] While these results do not control for other factors that could determine accuracy, they are consistent with an improvement in accuracy over time. We also investigate how persistent forecast accuracy was between the subperiods studied.[42] The correlation coefficient between the RAFEs in the two subperiods equals .64 and is significantly different from zero. We conclude that in a comparison of the two subperiods future forecast accuracy was positively correlated with past forecast accuracy, however, the relation is not conclusively perfect, suggesting the possibility that the second period was associated with changes in analyst behavior.

5. Herding and Anti-herding Behavior

The three measures we use to test for herding and anti-herding amongst forecasters are from Bernhardt, Campello, and Kutsoati (2006) (BCK), and Chen and Jiang (2006) (CJ_Dev, CJ_Prob). The development of each of these measures establishes a baseline from which herding or anti-herding can be inferred by first establishing how an optimizing forecaster would utilize public and private information when constructing a forecast. The Appendix B presents the development of each measure. Table 8 summarizes the link between the numerical value of each measure and the interpretation of whether herding, anti-herding or neither is indicated. Aside from the paper in which Bernhardt et al. (2006) develop and implement the BCK test, it has also been used by several researchers to test for herding, including Pierdzioch, Rulke, and Stadtmann (2010), Pierdzioch and Rulke (2013), and Broughton and Lobo (2018). Chen and Jiang (2006) compute their measures in a study of earnings forecasters. The BCK and CJ_Prob statistics are based upon similar assumptions and, as we will see, produce statistics that are similar in sign and

[41] The ANOVA test allows for unequal variances and includes all weekly observations for each firm.
[42] Brown (2001) and Michaely et al. (2023) find that the accuracy of earnings forecasters tends to be persistent.

magnitude. The authors show that the measures they propose are robust to alternative assumptions about their construction.

5.1 Herding and Anti-herding at the Firm Level

Table 9 presents the computed BCK and CJ measures in each subperiod for firms making at least 20 forecasts in each. Table 9 does not distinguish between individual forecasters. As described in the Appendix B and Table 8, each measure is based upon the week t forecast released by firm j, F_{jt} ,the consensus forecast at the time the forecast is released, $C_{t,}$, and the actual value of the variable being forecasted, A_t.[43] For forecasts released two days prior to the EIA announcement, we use the median forecast on the previous day (Day 3) as our measure off C_t. For forecasts released one day prior (Day 1), C_t is measured as the median of the forecasts on Days 2 and 3, and for forecasts released on the day of the EIA report (Day 0) (before the 10:30 AM ET, EIA announcement), C_t is the median of the forecasts on Days 1, 2, and 3. It should be noted that C_t cannot be measured for forecasts released on Day 3 or for forecasts released on Day 2 if there were no forecasts on Day 3. While the results presented in Table 9 are based upon firm level data, in later analyses in which we study the relation between herding/anti-herding and forecast accuracy, we present results based upon individual forecaster data.

As shown in Table 9, anti-herding predominates, but there is also evidence of optimal weighting of public and private information, that is, neither herding nor anti-herding. There is considerable variation in the measures across the sample firms for all three measures. The level of forecaster herding behavior is also persistent. The correlation between the two subperiods overall is greater than 60%. Roughly 78% of the firms did not deviate in behavior between the two subperiods based upon tests that the CJ_Dev measure was equal to 0 (the null hypothesis

[43] In unreported results we present statistics showing that over all forecasting firms each of the three herding measures confirms that anti-herding is the dominant behavior. The results are available from the authors upon request.

that a forecasting firm was neither a herder nor an anti-herder, Table 8). Of the firms that changed behavior, all were anti-herders during the first subperiod and switched to behavior consistent with optimal weighting, that is, they became neither herders nor remained anti-herders. We conclude that forecasting firms generally choose to anti-herd, and tend to stick to those strategies over time or choose the course of neither herding nor anti-herding. However, forecasting firms do choose different levels of anti-herding intensity. Likewise, while highly correlated between the two subperiods, the level of anti-herding behavior is slightly lower during the second period. We test whether the mean levels of each herding measure differ between the two subperiods. The test results are reported near the bottom of Table 9. The tests reject at the 1% level the null of equality of means between the two subperiods for two of the herding measures.[44] The results suggest that forecasters representing the sample firms shifted their behavior slightly between the two periods, but generally engaged in anti-herding.

6. The Relation Between Accuracy, Herding, Anti-herding and Forecast Timing

6.1 Between-Firm Estimation Results

We begin with between-firm regression results which emphasize forecast accuracy variation across forecasting firms. We standardized all variables and then computed the average values for each firm across the sample observations for that firm. The use of standardized variables enables us to assess the relative impact of each of the explanatory variables on forecast accuracy, as the raw variables are of different scales. We regress the average values for our standardized measure of RAFE (relative absolute forecast error), for each of the firms in the sample on: (1) the average of the standardized herding measure for that firm, (2) the average standardized number of days between the firm's forecast release and the EIA announcement and

[44] The p-values for tests of mean equality between subperiods: BCK (0.007), CJ_Dev (0.091), CJ_Prob (0.000).

(3) an interaction variable equal to the product of the herding and timing variables. We account for heteroskedasticity in coefficient tests by using Huber-White standard errors. A change of 1 standard deviation in an explanatory variable is associated with a change of θ standard deviations of RAFE where θ is the estimated linear regression coefficient.

Table 10 presents the results for models in which we test the null hypotheses that there is no relation between the herding measures and accuracy, and likewise the null that there is no relation between the timing of the forecast release and accuracy. We also test whether there were any changes in the relations between accuracy, the herding measures and timing between the two subperiods. We define a dummy variable D1 equal to 1 for observations during the period *1/1/2013-12/9/2021* and 0 otherwise.[45] The dummy variable is interacted with the other explanatory variables and is included as a constant effect as well. The estimated coefficients of the interaction variables allow us to test whether there was a statistically significant change in the impacts of timing and herding on accuracy between the subperiods. The estimation results reported in Table 10 are for models which also include interaction variables. We restrict the sample to those firms making 20 or more forecasts in either subperiod.

The estimated coefficients on each of the herding measures are positive and significantly different from zero at the 1% level. A more positive value for each of the herding measures implies anti-herding behavior, so we conclude that greater anti-herding intensity is associated with less accuracy. The estimated coefficients on the timing variable 'Days before EIA release' are also positive and generally significantly different from zero at the 1% level. Thus, early

[45] To address whether the results are sensitive to controlling for time using a dummy variable to distinguish observations after 2012 we estimated an alternative specification where the base model is that estimated in Table 10. The dummy variable was replaced with a linear time trend variable which increases linearly over the sample years. The estimation results for the time trend specification showed no significant effects of the time trend on forecasters' tendency to herd, timing choice, or the interaction between herding and timing. These results are for brevity not reported but are available from the authors upon request.

forecasters tend to be less accurate. Second, the estimated coefficients on the interaction variables equal to the period dummy times the herding measures are never significantly different from zero at the 1% level. These results suggest that there was no differential impact of herding (anti-herding) on accuracy between the two subperiods. Likewise, there was no differential impact of the relation between accuracy and the timing of the forecast release at the 1% level (variable 'Days before EIA release') between the two subperiods. These results suggest that on average the relation between accuracy, the herding measures and timing was stable across the two subperiods. A takeaway is that the general increase in information availability was not associated with a change in these relations. Results based upon the full sample, irrespective of how many observations a firm had reported in either subperiod, are found to be similar to those reported in Table 10. Further, the average adjusted r-squared across the models presented in Table 10 is equal to roughly 50%, indicating that timing and herding explain a large fraction of the variation in accuracy between firms.

The size of each estimated coefficient provides an indication of the relative importance of the variable. A one standard deviation change in the CJ_Dev herding measure (*Model 1* of Table 10) leads to roughly a change of .60 to .67 standard deviations in the mean RAFE measure during both subperiods. Correspondingly a one standard deviation change in the timing variable (Days before EIA release) for the same model leads to a change of .20 to .50 standard deviations in RAFE depending upon the measure of herding. We conclude that the forecaster's herding choice is more important to accuracy than the timing release choice, although both are individually important and statistically significant. The estimated coefficients on the interaction variable constructed as the product of a herding measure and the timing measure are not significantly different from zero for any of the models estimated, indicating there is no moderating or exacerbating effect of timing on herding. Likewise, this result does not shift between the subperiods.

Accuracy may also be influenced by forecaster ability. Ability should plausibly be positively related to experience. That is, if analysts who have made a large number of forecasts have used that experience to refine the models they use for forecasting, then we predict that more experienced analysts should on average be more accurate. *Model 2* includes a proxy for experience. The proxy variable we use is the number of forecasts produced by a firm which we label "Number of forecasts".[46] We add to the regression models the standardized number of forecasts issued by the firm over the sample period. The coefficient estimates for the variable measuring the number of forecasts are not significantly different from zero in either subperiod for two of the herding measure models but is positive and significant for the CJ_Dev model, however the overall results for herding and timing are consistent across all the models estimated under the *Model 2* structure.

The results presented in Table 10 as mentioned are between firm regressions utilizing individual firm averages. Those regressions do not capture any time-series variation within a single firm. We could potentially supplement those results by presenting panel regressions with fixed firm effects and weekly firm measures of timing and accuracy and an interaction variable for each subperiod. Table 10 suggests however that herding behavior is relatively invariant at the firm level. The expectation might therefore be that weekly measures of herding would likewise exhibit time invariance. The herding measures, however, cannot be estimated on a week-by-week basis. But even if they could, to capture within firm effects of anti-herding we would need to include both the herding measures and fixed firm effects. But because herding (anti-herding)

[46] We recognize that our measure of experience is a noisy proxy. A limitation of the measure is that some firms in the sample may have issued forecasts prior to the start date of our sample. This could potentially add noise and weaken any empirical relation between experience and accuracy in our data.

is relatively time invariant at the firm level, if we were to include weekly herding measures and firm fixed effects we would encounter a severe multicollinearity problem. Nevertheless, the results presented in Table 10 suggest that a large fraction of the accuracy variation is due to between-firm variation.

6.2 Individual Analyst Accuracy

Our data allow us to estimate a disaggregated model in which we focus on individual analysts in contrast to firms. While not perfect, this allows us to estimate a panel data model that allows for within-firm variation of accuracy by estimating a model using forecaster-level forecast errors. We regress individual *i*'s relative absolute forecast error in week t on firm-level herding measures (BCK(*i*), CJ_Dev(*i*), or CJ_Prob(*i*)), analyst *i*'s choice of timing in week *t* (Days before EIA release(i,t)), the interaction of the two variables, and the analyst's experience. All variables are standardized. The primary weakness of the model is our use of the firm level herding measures as proxies for the individual forecaster herding measures and the assumption that these measures are both constant during each subperiod and are reflective of the individual analyst's behavior. We justify the use of the firm level herding measures in the regressions based upon the observation that in general each firm is represented by generally only one or two analysts and that the measures are relatively invariant between the two subperiods for most firms. Our assumption is that the invariance property holds generally week-by-week during each subperiod. We cannot include firm fixed effects in the regressions due to the multicollinearity issue mentioned above. But the disaggregation implicitly captures some of the within-firm variation through the individual forecaster errors. However, as we are interested in time variation in accuracy we do include year fixed effects. Dummy variables for each year except year 2003 which serves as the base year, are included. The year-dummies allow us to control for the influence of year-level omitted variables on individual analyst forecast accuracy. These could for instance include fundamental shocks to the gas industry as well as any time trend in forecast accuracy. The variable “Number of forecasts” which proxies for forecasting experience, is now time-varying. We use the number of forecasts before week t of the analyst i instead of the total number of estimates for a firm as the measure of experience.

Table 11 presents the estimation results for the analyst-week specifications. In *Model 1*, we find results that are consistent with those presented in Table 11. Anti-herding and timing continue to have a positive and statistically significant influence on forecast accuracy. Earlier forecasts and those associated with greater anti-herders tend to be less accurate. The low R-squared values for the regressions however when compared to those for the between-firm regressions presented in Table 11, suggest that it is the cross-sectional relation reflected in the between-firm regressions that is the primary explanation for accuracy variation. The *Model 2* results which include year fixed effects support this conclusion. Although the coefficient on the variable “Number of forecasts” (our proxy for experience) is significant in *Model 1*, it is not significantly different from zero when year fixed effects are included in *Model 2*.

7. Additional Evidence on Herding and Anti-herding

The extent to which a forecaster anti-herds is a choice. The structure of the CJ_Dev measure offers an opportunity for us to explore anti-herding behavior at the level of the analyst. Our tests are designed to investigate several questions. First, was anti-herding the predominant form of behavior during the period we study? Second, did the level of anti-herding change over time, and was it different during the two subperiods we study? Third, were better forecasters greater anti-herders or conversely, were poorer forecasters greater anti-herders? Fourth, was anti-herding related to the timing of an analyst’s forecast release?

The CJ_Dev measure is given by the estimate of the coefficient β in the following regression as described in Table 8 and the Appendix B,

$$\mathrm{F_{it}} - \mathrm{A_t} = \alpha + \beta(\mathrm{F_{it}} - \mathrm{C_t}) \tag{1}$$

where the variables are as described earlier and the subscript i indexes the individual forecaster.

The CJ_Dev measure is estimated by, regressing the deviation of an analyst's forecast from the actual at time t $\mathrm{F_{it}} - \mathrm{A_t}$, on the deviation of the forecast from the consensus forecast, $\mathrm{F_{it}} - \mathrm{C_t}$. As outlined in Table 8, an estimated positive β coefficient is an indication of anti-herding, while a negative coefficient is an indication of herding. In this section we present results in which the focus is on the forecasts of the individual forecasters. We use the median of the prior forecasts as our estimate of the consensus. As this specification uses analyst-week level observations, we estimate two different specifications, one that includes firm-fixed effects, where the firm is the company that employs the analyst, and a separate model that includes analyst (individual) fixed effects.

Was anti-herding the predominant form of behavior? The estimation results are presented in Table 12. All models use the same dependent variable, the individual forecaster errors, $\mathrm{F_{it}} - \mathrm{A_t}$. These are unbalanced panel models because forecasters might issue forecasts in different weeks over our sample years 2003-2021. *Model 1* is our baseline specification including as a right-hand side variable the forecaster's deviation from the consensus $\mathrm{F_{it}} - \mathrm{C_t}$ and firm fixed effects. *Model 2* differs from *Model 1* in that the specification includes analyst fixed effects in place of firm fixed effects. In both *Models 1* and *2* the estimate of the coefficient on the analyst forecast deviation from the consensus $\mathrm{F_{it}} - \mathrm{C_t}$ is positive and statistically significantly different from zero consistent with the inference that the sample of forecasters on average displayed anti-herding behavior. These results support the conclusion drawn earlier that analysts in our sample tend towards anti-herding, that is, to present forecasts that overweight their private information.

Did the level of anti-herding change over time? Given the results for *Models 1* and *2*, *Model 3* introduces a dummy variable, D1, which equals 1 for observations after the end of 2012 and an interaction variable equal to D1 times $\mathrm{F_{it}} - \mathrm{C_t}$. The estimated coefficient on the interaction variable is an indicator of whether analysts tended to tilt more towards or away from anti-herding during the second subperiod. The coefficient on the dummy variable D1 has a negative sign but is not significantly different from zero. The coefficient on the interaction variable is both negative and statistically significantly different from zero. The net coefficient for the second period equals .698-.374 = .324 indicating that anti-herding by individual analysts during the second period was still present but less intense, similar to the firm-level results presented in Table 9 showing a decline in average anti-herding between the subperiods.

Were better forecasters greater anti-herders or conversely, were poorer forecasters greater anti-herders? *Model 4* turns to the question of whether the top (more accurate) forecasters tended to anti-herd, herd or neither, and likewise for the bottom (least accurate) forecasters. The model is streamlined and includes as explanatory variables $\mathrm{F_{it}} - \mathrm{C_t}$, firm fixed effects, and two interaction variables. If a forecaster was in the top 20 most accurate (forecasters who have the smallest average forecast error over all weeks), Top 20 equals 1 and otherwise 0. If a forecaster was in the bottom 20 least accurate then Bottom 20 (forecasters who have the largest average forecast errors over all weeks), equals 1 and otherwise 0. We interact both dummy variables with $\mathrm{F_{it}} - \mathrm{C_t}$. The results indicate that the coefficient on $\mathrm{F_{it}} - \mathrm{C_t}$ remains positive and significant and is equal to 0.547 indicating anti-herding is prevalent overall. However, the coefficient on the interaction variable [Top 20 × ($\mathrm{F_{it}} - \mathrm{C_t}$)] is negative and significant and equal to -.504. The sum of the two coefficients, 0.547+(-.504), is close to zero, although the sum is significantly different from 0, nevertheless indicating the top 20 forecasters tend more towards neither herding nor anti-herding but display choices consistent with acting optimally. Conversely the estimated coefficient on the interaction variable [Bottom 20 × $(\mathrm{F_{it}} - \mathrm{C_t})$] is positive and statistically significant, indicating that the worst forecasters display the greatest anti-herding behavior.

Was anti-herding related to the timing of an individual analyst's forecast release? Further did this relation change over time? *Model 5* extends the analysis by examining whether anti-herding differs across the days prior to the release of the EIA report. We define two dummy variables, the first, DA, takes the value 1 if the analyst's forecast is released 2 days prior to the EIA report and 0 otherwise. The second DB takes the value 1 if the forecast is released one day prior to the EIA report and 0 otherwise. We interact these dummies with $\mathrm{F_{it}} - \mathrm{C_t}$. The forecasts released 3 days prior to the EIA report is the base case. The estimated coefficients on the two interaction variables provide estimates of whether the intensity of anti-herding changed as the EIA report date approached. We also interact the dummies DA and DB with [D1 × $(\mathrm{F_{it}} - \mathrm{C_t})$] to test whether the relations changed during the second subperiod. The estimation results for *Model 5* indicate that anti-herding predominates as found for *Models 1-4*. The estimated coefficient on the variable $\mathrm{F_{it}} - \mathrm{C_t}$ is positive and statistically significantly different from 0 at the 1% level and equal to 1.047, indicating that forecasts issued 3 days prior to the EIA report are associated with intense anti-herding. Controlling for the timing of the forecast release and jointly for the sample subperiod, the coefficient on the base variable that accounts for a shift in anti-herding in general, the coefficient on [D1 × $(\mathrm{F_{it}} - \mathrm{C_t})$], while negative is not significantly different from zero at the 1% level. The result indicates that forecasts issued 3 days early during the second period show the same level of anti-herding as during the first period. However both coefficients on the variables that control for the day of the release are negative and statistically significant but are less in absolute value than the base coefficient on $F_{it} - C_t$ which equals 1.047. Anti-herding intensity decreases between the third day and the second day prior to the EIA report date but then increases. The net anti-herding coefficient two days out equals 1.047-0.901 = 0.146, showing a decrease in intensity but nevertheless that anti-herding is present, and one day out equals 1.047-0.319 = 0.728. The increase in intensity of anti-herding as the EIA release date approaches potentially is in response to an effort by forecasters to distinguish themselves as well as potentially satisfy the demand for unique forecasts as discussed in Section 6. These results do however support the overall conclusion that anti-herding is always the norm.

Controlling for the subperiod there is an incremental statistically significant decrease in anti-herding intensity one day out during the second period, the net intensity during the second subperiod equals 1.047-0.319-0.344 = 0.384, as compared with the net intensity during the first period of 1.047-.319 = .728. This was not the case for the forecasts released 2 and 3 days ahead of the EIA release. Thus, while the intensity of anti-herding did change for forecasts issued one day out, our conclusion that anti-herding continued to be the dominant form of behavior.

Model 6 presents the results of the full model. Here we see that the conclusions drawn for the disaggregated *Models 1-5* continue to be supported when all explanatory variables are included. Summarizing the results based upon *individual forecaster behavior,* 1) anti-herding is the dominant form of behavior in both subperiods studied, 2) anti-herding intensity first decreases and then increases as the EIA report date approaches relative to the earliest forecast dates, 3) the top 20 forecasters in terms of accuracy neither herd nor anti-herd while the bottom 20 forecasters tend to anti-herd. The result suggests that the most accurate forecasters behave as if they weight their private information and the public consensus view optimally, 4) during the second subperiod, anti-herding by early forecasters, those issuing forecasts 2 or 3 days prior to the EIA release, exhibits the same relative behavior as during the first subperiod, but as mentioned above anti-herding intensity declined on the day before the release of the EIA report.

8. Robustness Tests

To address whether the results are sensitive to controlling for time using a dummy variable to distinguish observations after 2012 we estimated an alternative specification where the base model is that estimated in Table 10 as highlighted in footnote 45. The dummy variable was replaced with a linear time trend variable. The estimation results for the time trend specification showed no significant effects of the time trend on forecasters' tendency to herd, timing choice, or the interaction between herding and timing. These results are for brevity not reported but are available from the authors upon request. In addition to the alternative controls for time, we estimated several additional model specifications. The base structure of the estimated models of accuracy is as shown in Table 11. First to test whether the basic results are sensitive to the choice of forecast when computing forecast errors, we repeat the estimation using only the one day ahead forecasts and find the basic results are unchanged. Second, to control for the potential that an information shock received prior to the forecast release might influence the results, we include the lagged standardized natural gas futures price change. Here the assumption we are making is that any information shock about the underlying asset could potentially lead to a shift in expectations. The basic results as suggested in Table 11 are unaffected. Finally, based upon information obtained from the source for the forecasts we study, i.e., Bloomberg, we know the order of the forecast releases for each date for the year 2021.

Based upon this subset of data, which represent within day releases, we find that our base results on herding and accuracy are further confirmed.[47]

9. Conclusions

There has been an explosion of information available to analysts over the last 20 years. We present evidence on one dimension of the potential impact of this information revolution, the accuracy of forecasts made by professional analysts that focus on the energy sector. The set of information about the energy sector and ease of access to that information has increased dramatically. Changes in natural gas in storage reported by the U.S. Energy Information Administration (EIA) are closely followed by energy market participants and are known to reflect changes in supply and demand conditions. A professional cadre of analysts engage in forecasting such changes which serve as the platform for our analysis. We present evidence that the accuracy of forecasts of the changes in natural gas storage levels is strongly associated with (1) the forecaster's tendency to either anti-herd or to neither herd nor anti-herd, and, (2) how long before the EIA announcement the forecast is released. Our evidence indicates that accuracy has improved over time, and while anti-herding is the norm, the intensity of anti-herding fell over the sample period we study. Anti-herders' forecasts tend to be much less accurate and early forecasts are less accurate than late forecasts. These two variables account for over 50% of the variation in forecast accuracy across forecasting firms. Our results indicate that the influence of anti-herding on the accuracy of the forecasts studied experienced no significant change between the two subperiods studied. A similar result is found for forecast timing. On the other hand, we find that the mean absolute forecast error (AFE) and relative absolute forecast error (RAFE) for our sample of forecasters declined during the post-2012 period as did the standard deviation of forecast errors. These results are broadly consistent with the hypothesis that the increase in

[47] Robustness test results are available from the authors upon request.

information available during the sample period led to an overall improvement in accuracy despite the fact that the influence of anti-herding and the release of early forecasts nevertheless continued to reduce accuracy.

We find that forecaster accuracy varies greatly across forecasting firms and is highly persistent. The mean absolute forecast error of the least accurate forecasters is 1.9 times that of the most accurate during the pre-2013 period, but that ratio declines marginally during the post-2012 period, consistent with an overall reduction in forecast error. The less accurate forecasters rarely improve over time. This raises the question of how the least accurate forecasters continue to successfully sell their services; that is, what services do they provide to clients? Our results indicate two potential reasons for the continued viability of less accurate forecasters. First, they tend to release their forecasts early, possibly providing a first-mover advantage that clients demand. Second, they tend to base their forecasts on their private intelligence rather than herding, thereby disseminating information that clients find useful.

Our evidence indicates that anti-herding behavior is the norm, differs substantially across forecasters, and is strongly persistent. Specifically, anti-herders continue to anti-herd but there is little variation in this behavior between the pre-2013 and post-2012 periods, suggesting that this behavior was relatively stable. We do find evidence however that forecasters with the highest accuracy tend to neither herd nor anti-herd, while the worst forecasters tend to strongly anti-herd. Anti-herding exhibits a tendency to increase just prior to the EIA release, following a decrease in intensity, but this does not change between the subperiods. The increase in intensity of anti-herding as the EIA release date approaches is consistent with forecasters attempting to distinguish themselves as well as potentially strategically satisfy the demand for unique forecasts as discussed in Section 6.

Studies of forecaster accuracy differences in other markets have focused on cross-sectional determinants such as forecaster experience, specialization, and available resources. Our evidence indicates the two primary determinants for the sample studied here are the proclivity to anti-herd and the timing of when the forecast is released. Little is known about how the impacts of these determinants has changed over time. Our evidence indicates that the relations did not change over the course of the period we study, one during which there was a significant overall increase in information. While these relations persist, recall that overall forecast accuracy did improve during the post 2012 period. We do find a tendency for forecasters to issue forecasts earlier during the post 2012 period. One explanation for this change in behavior is that an increase in the available information set could have increased analysts' beliefs that their private information had improved in quality, leading them to issue more forecasts based upon the belief that they have acquired 'better' private information. Additional specifications involving alternative controls for time, the choice of forecast to evaluate and shocks to information prior to forecast release do not alter our basic inferences.

Although our study has focused solely on forecasters of changes in natural gas storage levels, our results potentially have implications for understanding forecaster accuracy for other economic variables and the influence of the shift in information, studies we leave for future research.

Chapter 3

Electronic Trading and Price Discovery in Commodity Futures

Coauthored with Kateryna Holland, Scott C. Linn, and Chitru S. Fernando

"Good riddance to open outcry," John Arnold, the billionaire former chief executive of natural-gas hedge fund Centaurus Energy, said in an email to *The Wall Street Journal*. "The transition from open-outcry to electronic trading happened so quickly because it provided far superior execution to the customer." (*The Wall Street Journal*, 12/29/2016, Alexander Osipovich, NYMEX "Trading Pits Shut Down, Marking End of an Era")

1. Introduction

The structure and design of financial markets have long been an area of interest to researchers, practitioners, and policy makers (Harris, 2003). As emphasized by O'Hara (2003), the two important functions of financial markets are liquidity and price discovery. The literature on the impact of electronic trading market structure change has largely focused on stock markets. A considerable literature has examined the liquidity and market quality effects of recent market structural changes, such as decimalization of bid-ask spreads, financialization, and moves to electronic trading (e.g. Bessembinder, 2003; Stoll, 2006; Lauter, Prokopczuk and Truck, 2023), drawing the conclusion that these events are associated with improvements in market liquidity and quality.[48] However, considerably less attention has been paid in the mainstream finance literature to the price discovery effects of significant structural changes in markets where both spot and futures prices are observed. This question is particularly important in commodity markets where physical spot and futures markets are less perfectly interconnected through arbitrage than in markets for financial assets (see, for example, Ederington et al., 2021). This

[48] Other examples of papers studying the implications of electronic trading in equity markets include Benveniste et al., (1992), Madhavan and Panchapagesan (2000), Gwilym and Alibo (2003), Venkataraman (2001); Handa, Schwartz, and Tiwari (2004), Jain (2005), Battalio et al. (2007), Shah and Brorsen (2011), Hendershott and Riordan (2013), Raman, Robe, and Yadav (2020). These papers examine electronic trading but do not examine the issue of where prices are discovered.

study helps fill this gap in the literature by examining how the switch from floor trading to electronic trading in the crude oil futures market impacted price discovery.

The switch to electronic trading can have a heterogeneous impact on price discovery. On the one hand, opponents of the switch to electronic trading, for example, George Gero, managing director at RBC Wealth Management, said "the floor made it possible to sense the collective mood of the market and relay it to clients—an impossibility in today's faceless electronic markets."[49] Moreover, while the observed increase in volume could be attributed to lower trading costs and attracting more informed traders into the futures market (potentially enhancing price discovery), it is also possible that lower costs attract an upsurge of noise traders, thereby hurting price discovery. Consequently, while electronification could have a beneficial effect from a liquidity standpoint, this benefit could be offset if there is a deleterious effect on price discovery, which in turn would have important implications for the myriad ways in which futures prices are used. Therefore, if a shift to electronic trading in a futures market leads to the expectation that futures prices will become less reliable, then the spot market may be viewed as providing an important complementary price signal and would then play a less important role in price discovery.

On the other hand, the transition to electronic trading was applauded by some longtime commodity traders, as in the opening quote by John Arnold who greeted the closing of open outcry trading on NYMEX with the sendoff "good riddance." Proponents of the change argued that electronification would slash costs and make markets more efficient (Fleming, Ostdiek, and Whaley, 1996; Raman, Robe, and Yadav, 2020). Therefore, if increased liquidity and market

[49] *The Wall Street Journal*, 12/29/2016, Alexander Osipovich, "NYMEX Trading Pits Shut Down, Marking End of an Era."

access following a shift to electronic trading is accompanied by an increase in informativeness of the futures market, then we would expect price discovery to be centered in that market.

Several recent studies, which reach different conclusions regarding the impact of electronic trading on market quality, are of relevance to this study on price discovery. Both market quality and price discovery are inherently linked to how securities are traded (e.g. O'Hara, 1995; Foucault, Pagano, and Roell, 2013) and as such are likely to be influenced by electronic trading through similar channels (e.g. the change in the mix of traders and relative importance of different contracts). Specifically, the takeaway from Brogaard, Ringgenberg and Roesch (2022) and Chung and Chuwonganant (2022) is that the elimination of floor traders was associated with an increase in pricing error and a decrease in market quality.[50] In contrast, Kye and Mizrach (2022), in a study of the same event, find no deterioration in market quality. Moreover, Lauter, Prokopczuk and Truck (2023) examine the influence of financialization and electronification on market quality of various commodities and find market quality improvements more specific to electronification, not financialization. The results presented in the aforementioned studies suggest that the question of where price discovery will occur after a shift to electronic trading remains an important empirical issue, which forms the central focus of this paper.

We believe the setting we investigate, the shift to electronic trading in oil futures markets, is tailormade for the analysis of price discovery effects that accompany a significant market structure change. We examine how the location of price discovery was impacted for a commodity traded in both the spot and futures markets when trading in the futures market shifted from the open-outcry (pit/floor) format to electronic trading. Specifically, on September 5, 2006,

[50] These two papers study the impact of a temporary suspension of floor trading on the NYSE starting on March 23, 2020, due to COVID-19, moving all trading to an electronic venue. See the release https://www.nyse.com/publicdocs/nyse/NYSE_Floor_Closure_FAQ_20200320.pdf

the New York Mercantile Exchange (NYMEX) commenced side-by-side trading (simultaneous electronic and open outcry trading) of the NYMEX physically delivered West Texas Intermediate (WTI, Cushing, OK, FOB) oil futures contract. This followed the move by the Intercontinental Exchange (ICE) to close its open-outcry trading floor and shift trading in the Brent oil futures contract to an electronic platform, which occurred on April 7, 2005.[51] There are several benefits from using this setting to examine the impact of a trading structure change on price discovery of crude oil. First, it is important to understand where prices are discovered and how changes in market architecture impact price discovery for users who rely on prices as signals of market fundamentals and oil price changes are known to have a significant impact on a wide-ranging collection of industries and the value of corporations.[52] Second, the underlying commodity is the same in both spot and futures markets, meaning that the fundamental value of the commodity underlying both markets is the same. Third, active spot and futures markets for oil exist in the U.S. and Europe and daily prices are available in long time series. Fourth, the shift to electronic trading in the futures market was complete and permanent. Fifth, following this shift there was an explosive increase in futures trading volume, as illustrated in Figure 3 (NYMEX) and Figure 4 (ICE). Such a dramatic increase in futures trading volumes raises the question of whether these structural changes also affected the relative influence of the spot

---

[51] See https://www.cmegroup.com/media-room/press-releases/2006/8/02/nymex_to_offer_sidebysidetradingofphysicallydeliveredenergyfutur.html and https://www.sec.gov/Archives/edgar/data/1174746/000095014406002072/g99070e10vk.htm for details.

[52] The U.S. Energy Information Administration (EIA) points out that "Oil plays a crucial role in the global economy—from the production of goods to the transportation of people and freight." The oil industry is becoming increasingly more important in the U.S. given the increase in domestic production from 2010 to 2022. Specifically, the value added by the oil and gas extraction industry in the U.S. amounted to $196 billion in 2021.
EIA: https://www.eia.gov/todayinenergy/detail.php?id=29972
Statista: https://www.statista.com/statistics/192910/value-added-by-the-us-oil-and-gas-extraction-industry-since-1998/#:~:text=The%20value%20added%20by%20the,billion%20U.S%20dollars%20in%202021.

versus futures market in oil price discovery, an empirical question that is the focus of this study.[53]

We examine changes in price discovery in the crude oil market caused by the introduction of electronic trading, which was a shock to structural trading arrangements, using daily spot and futures prices of WTI (West Texas Intermediate, Cushing, OK, FOB) and the corresponding data for Brent oil. To the best of our knowledge, this is the first paper to study the impact of the shift to electronic futures trading on the location of price discovery in commodity markets. We find that while the shift to electronic trading from floor trading in the futures market resulted in a significant ramp up in trading volume in both WTI and Brent markets, the two markets differed sharply in how electronic trading affected price discovery. In the WTI market, while both the spot price and the nearby futures contract price (first maturing contract) were significant contributors to price discovery prior to the change, the move to side-by-side electronic and floor trading in the futures market on September 5, 2006, shifted price discovery solely to the futures market. This finding suggests that in the WTI market, electronic trading was beneficial to both liquidity and price discovery. In contrast, while price discovery in the Brent market occurred in the futures market before the shift to electronic trading on April 7, 2005, this was not the case after the shift to electronic trading, with both the spot and futures markets contributing roughly equally to price discovery. This finding, which is the opposite of that we observe in the WTI market, suggests that while electronic futures trading was beneficial to liquidity as reflected in the consequent increase in volume, it resulted in a shift away from prices being solely discovered in the Brent futures market.

[53] As shown in Figures 3 and 4, the shift to electronic trading was associated with a large increase in trading volume and activity. Chordia, Roll and Subrahmanyam (2002) show that trading activity is positively associated with liquidity.

The theoretical basis for our empirical analysis stems from the equilibrium relation between spot and futures prices, driven by a no arbitrage condition and the potential existence of a convenience yield, and the dynamic equilibrium spot and futures price adjustment process developed in Figuerola-Ferretti and Gonzalo (2010) where the two prices exhibit a unit root and prices adjust when a deviation arises in the no arbitrage condition. We begin with a review of the former and then draw the parallel between the latter theoretical adjustment process and an empirical vector error correction model, VECM, in which there is a single long run relation between the spot price and the futures price equivalent to the theoretical no arbitrage relation. The VECM formulation allows the empirical estimation of the price adjustment parameters. Gonzalo and Granger (1995) show that in a setting like ours, each of the prices is linearly related to the fundamental intrinsic price of the commodity. This result allows us to compute the relative contributions to price discovery of each price as implied by the theoretical model. The primary measure of price discovery we compute, which follows from the economic theory, has been referred to in the literature as the Component Share, *CS* (see, for example, Harris et al., 1995).[54]

Our paper contributes to the literature on price discovery in general. Economic theory argues that futures markets can play an important role in price discovery and help make markets informationally complete by acting as a conduit for information about underlying fundamental values (Working, 1962; Black, 1976; Grossman, 1977). Our empirical results confirm the importance of the futures market in price discovery but, in addition, highlight that the spot market can also play an important role in price discovery, and that electronification of trading impacts where prices are discovered. We also contribute to the literature that compares price

[54] In the Appendix, we also compute two alternative measures, the Information Share measure of Hasbrouck (1995), and the Generalized Information Share of Lien and Shrestha (2014). Our focus on the Component Share measure stems from the fact that it is consistent with the economic theory we highlight and the theoretical statistical results in Gonzalo and Granger (1995).

discovery in traditional and electronic markets, such as Kurov and Lasser (2004), who examine floor traded and electronic E-mini equity futures prices but not spot prices, and Hendershott and Riordan (2013), who examine the algorithmic and human trades in the German equity market. Griinbichler, Longstaff, and Schwartz (1994) find that futures prices of an equity index lead spot prices more when futures are traded electronically. We extend this literature to commodity markets.[55]

Additionally, we contribute to the literature that examines price discovery in energy markets. The extant literature on price discovery in energy markets has (1) focused on a single time period and/or a single futures contract (Quan, 1992; Schwarz and Szakmary, 1994; Silvapulle and Moosa, 1999; Kim, 2011; Shrestha, 2014); (2) compared futures and spot prices and at times arrived at conflicting results while not considering an exogenous market structural change (Bopp and Sitzer, 1987; Bekiros and Diks, 2008; Shrestha, 2014; Chang and Lee, 2015; Holmes and Otero, 2019); (3) compared physically and financially settled contracts (Liu, Schultz, and Swieringa, 2015); and (4) concurrently compared different energy benchmarks (Lin and Tamvakis, 2001; Elder, Miao, and Ramchander, 2014; Liu, Schultz, and Swieringa, 2015).[56] These studies do not explore the implications of changes in market trading structure on price discovery. In contrast, our study directly explores the impact on price discovery of the important structural shift to electronic trading. Therefore, this study contributes both to the general issue of whether price discovery occurs in spot or futures markets, as well as the literature that examines the implications of structural breaks in crude oil prices (Fan and Xu, 2011; Chen, Lee, and Zeng, 2014). In addition, we examine how electronification of futures trading impacted both the WTI

[55] More generally, the determinants of price discovery have been examined in a variety of equity market settings (Hasbrouck, 1995; Shive, 2012; Boehmer and Wu, 2013; Brogaard, Hendershott, and Riordan, 2014 and 2019).
[56] Shrestha (2014) suggests that prices are discovered in the futures market for certain commodities, while in both futures and spot markets for others. Chang and Lee (2015) examine oil spot and futures prices and show that both contribute to price discovery in the long-run equilibrium.

and the Brent markets. Our focus on bivariate comparisons between spot and futures prices allows us to make statements about the relative contributions to price discovery of the spot and futures markets and, also, whether contract maturity is associated with price discovery.

The remainder of the paper is structured as follows. Section 2 presents an economic framework in which the relation between spot and futures prices, as well as its implications for the measurement of price discovery, are developed. Section 3 explains how we compute several measures of price discovery that are presented in our empirical study. Section 4 describes the data and results of unit root and cointegration tests. Section 5 describes the estimation results that form the basis for the computation and analysis of price discovery. Section 6 presents a discussion of the empirical results on price discovery. Section 7 presents our conclusions.

2. Futures Prices and Spot Prices

We begin by reviewing the relation between spot and futures prices for a storable commodity within two economic environments: the first, in which the convenience yield equals zero, and the second, in which the convenience yield is positive. A linear (in logs) relation between the spot and futures prices is the equilibrium outcome in both scenarios. We contrast the two resulting linear relations to highlight an important difference that emerges when the convenience yield is positive. These linear relations form the basis for a discussion of where price discovery occurs in response to a shock to the equilibrium relation.

2.1 Futures Prices with Unrestricted Arbitrage and a Zero Convenience Yield

Let S_t be the spot market price of a commodity in period t and s_t be its natural logarithm. Let F_t be the contemporaneous price of a futures contract maturing at $T, T - t > 0$, and let f_t be its natural logarithm. We make the following assumptions to establish the base unrestricted environment in which any frictions that might restrict arbitrage activity are absent:[57]

a) No taxes or transaction costs.

b) No limitations on borrowing at the risk-free rate.

c) No costs other than financing a (short or long) futures position and any costs associated with storing the commodity.

d) No limitations on the short sale of the commodity in the spot market.

e) Market participants can borrow and lend at the continuously compounded risk-free interest rate determined by the process $r_t = \bar{r} + \zeta$, where $\bar{r}$ is the mean and ζ is a stationary process with a mean of zero and finite positive variance.

f) The commodity can be stored at a continuously compounded storage cost per unit determined by the process $c_t = \bar{c} + \omega$, where $\bar{c}$ is the mean and is constant over the period *t* to *T* and ω is a stationary process with zero mean and finite variance.

g) The errors ζ and ω are independent.

h) s_t is integrated of order 1 and $\Delta s_t = s_t - s_{t-1}$is stationary, integrated of order 0.

For the moment, assume there is no convenience yield for the commodity. Under the above assumptions, futures and spot prices are related by a no arbitrage relation

$$F_t = S_t e^{(r_t + c_t)(T-t)} \,. \tag{4}$$

Assume for expositional purposes that $T - t = 1$. Taking the natural log of (1) gives

$$f_t = s_t + r_t + c_t \,, \tag{5}$$

and by substitution using (e) and (f) and rearrangement,

$$s_t = f_t + \bar{r} + \bar{c} + (\zeta + \omega) = \alpha + f_t + \varepsilon \,, \tag{6}$$

[57] Sections 2.1 and 2.2 draw on the developments in Figuerola-Ferretti and Gonzalo (2010).

with the implied coefficient on f_t equal to 1. The relation between the spot and futures prices will, however, be different if a nonzero convenience yield is present.

2.2 Futures Prices and the Convenience Yield

Restrictions on the ability to arbitrage can arise for a variety of reasons, such as storage capacity limitations, giving rise to a nonzero convenience yield. Commodities can confer real benefits to those who hold them in inventory. For example, owners of inventories may receive an added benefit or "convenience" when inventories in general are low because holding the commodity allows them to avoid stock-outs. The convenience yield as defined by Brennan and Schwartz (1985) is "…the flow of services that accrues to an owner of the physical commodity but not to an owner of a contract for future delivery of the commodity."[58] Under certain conditions (for example, when there is a shortage of the commodity in the spot market), the convenience yield can cause the futures price to be below the current spot price. A futures price that does not exceed the spot price by enough to cover "carrying cost" (interest plus storage cost) implies that inventory holders receive an implicit return from holding inventory. When supply falls, the convenience yield should rise and vice versa when supply becomes more abundant. The convenience yield may therefore be correlated with the spot or futures price as well as with other fundamentals, such as inventory availability or storage capacity or, for that matter, past values of the convenience yield.

The Kaldor-Working hypothesis postulates that the convenience yield depends inversely on the level of inventories. Bessembinder et al. (1995) and Casassus and Collin-Dufresne (2005) document a positive relation between convenience yields and prices, while Dincerler, Khokher, and Simin (2020) document an inverse relation between convenience yields and inventory. The

[58] See also Kaldor (1939) and Working (1949).

equilibrium model developed in Routledge, Seppi, and Spatt (2000) implies that the relation between convenience yields and spot prices varies with relative scarcity and, while the authors do not present a formal test of the model, their calibration exercise produces a relation consistent with the theory. Therefore, the convenience yield can potentially be related to prices as well as a vector of other variables, including, for instance, inventory level, which we represent with the notation X_t,

$$\delta_t = \delta(s_t, f_t, X_t) \,. \tag{7}$$

We return to the inclusion of other variables later in this section.

Introducing a convenience yield leads to the following equilibrium relation between the futures price and the spot price: $F_t = S_t e^{(r_t + c_t - \delta_t)(T-t)}$, where δ_t is the convenience yield. In terms of natural logs, the equilibrating relation now becomes

$$f_t + \delta_t = s_t + r_t + c_t \,. \tag{8}$$

As in Figuerola-Ferretti and Gonzalo (2010) and following the discussion above and equation (4), assume that the convenience yield is related to the spot price and the futures price and that it can be approximated by a linear relation[59]

$$\delta_t = \gamma_1 s_t - \gamma_2 f_t, \gamma_i \in (0,1), i = 1{,}2. \tag{9}$$

Substituting (6) into (5) and considering assumptions (e) to (g), the following relation between the spot and futures prices is obtained

$$s_t = \beta_2 f_t + \beta_3 + (\zeta + \omega) \,. \tag{10}$$

[59] This assumption is not without precedent. Routledge, Seppi, and Spatt (2000) model the function as a linear function of storage costs, interest rates, spot prices, and futures prices, while Gibson and Schwartz (1990) model the function as an autoregressive process with errors correlated with spot prices (see also Casassus and Collin-Dufresne, 2005).

By our earlier assumptions, the term in brackets is a stationary mean zero finite variance error. The result indicates that the spot and futures prices are linearly related to one another and in parlance we will use later are cointegrated with

$$\beta_2 \ = \ \frac{1 \, - \, \gamma_2}{1 \, - \, \gamma_1}; \beta_3 \ = \ \frac{-(\bar{r}+\bar{c})}{1 \, - \, \gamma_1} \, . \tag{11}$$

Equation (11) shows that the coefficients of the cointegrating vector can depend on the sizes of the parameters in (9) and the parameters $\bar{r} + \bar{c}$. The coefficient values shown in (11) will influence the slope of the futures curve, causing futures prices to be above or below the current spot price. The result is that the cointegrating relation is not the standard (1, -1) (unitary relation) often assumed in the literature on price discovery (Hasbrouck, 1995; Lehman, 2002). If $\beta_2 > 1$, then the market is in backwardation (inverted forward curve); if $\beta_2 < 1$, then the market is in contango (upward sloping forward curve). Also, notice that $\beta_3 < 1$ since $\bar{r} + \bar{c}$ is not negative.

The assumption that the convenience yield is linearly related to the spot and futures prices of the commodity can be extended to include other relevant variables such as inventory storage capacity. Suppose the convenience yield is also a function of another variable, x_t (Kim, 2011). For example, x_t could equal the available inventory storage capacity or the lagged value of the convenience yield. Assume the convenience yield reflects the value of x_t in the following manner

$$\delta_t = (\gamma_1 s_t - \gamma_2 f_t) + \gamma_3 x_t \ . \tag{12}$$

Substitution into (10) now shows that

$$\beta_2 = \frac{1 - \gamma_2}{1 - \gamma_1}; \beta_3 = \frac{\gamma_3 x_t - (\bar{r}+\bar{c})}{1 - \gamma_1}. \tag{13}$$

The result again is that $\beta_2$does not equal 1 unless $\gamma_1 = \gamma_2$. Further, the intercept,β_3 can have any sign depending on the relative magnitudes of $\gamma_3 x_t$ and $(\bar{r} + \bar{c})$.

3. Price Discovery Measures

Consider a setting in which a spot market and a futures market exist for a commodity. Assume that spot prices and futures prices for a unit of the commodity are observable, but the true intrinsic price (value) of the commodity is not. We hypothesize that each observable price reflects information about the intrinsic value, but potentially not to the same degree.[60] The degree to which the spot price or the futures price reflects the intrinsic value would therefore be an indicator of whether price discovery occurs in the spot or futures market. The market in which the observable price first reveals the intrinsic price would be considered the market in which prices are first "discovered." The discovery process evolves through the supply and demand decisions of market participants in the two markets. Figuerola-Ferretti and Gonzalo (2010) develop a dynamic equilibrium model of a spot and futures market for a storable commodity building on the framework presented by Garbade and Silber (1983). The authors show that in a dynamic equilibrium based upon supply and demand, spot and futures prices are related in the long run as in equation (10), and prices dynamically adjust to the random errors in the linear relation (10) based upon adjustment parameters that stem from fundamental supply and demand parameters. The analytical adjustment process the authors derive is empirically equivalent to a VECM. The empirical estimates of the adjustment parameters form the basis for the computation of a measure of price discovery and can be estimated from an empirical VECM involving the spot and futures prices.

The price series we examine are empirically each integrated of order 1 and each pair of prices examined is related empirically by one cointegrating equation consistent with the existence of the equilibrium linear relation between spot and futures prices predicted by the

[60] This result is a product of what is commonly referred to as the Gonzalo-Granger decomposition, in which in our context an observable price is said to be composed of an underlying fundamental price and a random error (Gonzalo and Granger, 1995).

theory developed in Section 2.[61] The VECM representation used in our analysis therefore takes the following form:

$$\Delta P_t = c + \alpha(\varepsilon_{t-1}) + \sum_{k=1}^{K-1} \Gamma_k \Delta P_{t-k} + e_t, \tag{14}$$

where

$$E[\varepsilon_t \varepsilon_s^{'}] = \begin{cases} 0 & \text{if } t \neq s, \\ \Omega & \text{otherwise,} \end{cases}$$

ΔP_t is a 2 × 1 vector that contains the log spot price change and the log futures price change, $\begin{bmatrix} \Delta s_t \\ \Delta f_t \end{bmatrix}$ and to which in the tables displaying estimation results, we refer to as (Δs, $\Delta f1$). The variable ε_{t-1} is from the cointegration equation $\varepsilon_{t-1} = s_{t-1} - \beta_2 f_{t-1} - \beta_3$. Note that ε_{t-1} measures the extent to which the relation between the spot price and the futures price violates the equilibrium condition. The parameter vector $\alpha = \begin{bmatrix} \alpha_s \\ \alpha_f \end{bmatrix}$ contains the speed of adjustment coefficients that drive the two prices back towards the equilibrium relation whenever there is a shock that disrupts the equilibrium. The parameter K is the lag length of the unrestricted VAR of the bivariate model.

The developments in Gonzalo and Granger (1995) and Stock and Watson (1988) indicate that if the price vector in the VECM representation contains two variables, in our case, the spot price and a futures price, and there is one cointegrating equation, then the observed prices can each be expressed as a linear function of a single “common” but unobservable intrinsic price and stationary noise. The underlying intrinsic price in our setting is best thought of as the underlying fundamental value of the commodity. To identify the underlying common price, Gonzalo and Granger (1995) invoke two conditions. The first condition, that the intrinsic price is linearly related to the observable prices, is sensible if we turn things around and contend that each

[61] As discussed in Section 4.3.

observable price is itself related to the intrinsic value. The second condition is that observable prices can be decomposed into a permanent (long-run) component and a transitory component. This condition is also sensible if we think of non-zero deviations from the long-run relation between the spot and futures prices as transitory, as would be the case when market participants exploit arbitrage opportunities when they arise, causing these opportunities to soon vanish. This translates into the permanent long-run price being expressed as $P_{perm} = \alpha_{\perp}{}'P$, where $\alpha_{\perp}{}'\alpha = 0$ and the vector $\alpha_{\perp}$ is orthogonal to the estimated adjustment coefficient vector α in the VECM representation (equation (14)). Thus, we can compute the elements of the vector $\alpha_{\perp}$ using the estimated adjustment coefficients in α. The linear relation $P_{perm} = \alpha_{\perp}{}'P$ is unique up to a scalar multiplication. This allows one to identify which of the two observed prices reveals the most about the underlying fundamental price by examining the elements of the vector $\alpha_{\perp}$.

The vector $\alpha_{\perp}$ for the bivariate case can be computed directly from the estimated adjustment coefficients $\begin{bmatrix}\alpha_s \\ \alpha_f\end{bmatrix}$ in (14). Normalizing by the sum of the orthogonal values gives what the literature identifies as the Component Share, *CS* (Harris et al., 1995; Harris, McInish, and Wood, 2002) for each market in which we introduce the subscripts *s* and *f* to denote the spot and futures component shares respectively (CS_s, CS_f).[62] Specifically,

$$\left(\frac{\alpha_f}{\alpha_f - \alpha_s}, \quad \frac{-\alpha_s}{\alpha_f - \alpha_s}\right) = (CS_s, CS_f)\ . \tag{15}$$

To understand the *CS* measure, let's revisit the long-run relation between the spot and futures prices and the adjustment process in the VECM characterization (14), ignoring lagged changes in prices for illustrative purposes. Specifically, there are two adjustment equations, one for the spot price and one for the futures price:

[62] The measure *CS* was first proposed on an informal basis by Schwarz and Szakmary (1994); see also Theissen (2002), Lehman (2002), de Jong (2002), and Putniņš (2013).

$$\Delta s = \alpha_s(s_{t-1} - \beta_2 f_{t-1} - \beta_3)$$

$$\Delta f = \alpha_f(s_{t-1} - \beta_2 f_{t-1} - \beta_3) \qquad (16)$$

The α_i adjustment coefficients in equation (16) represent the sizes of the adjustments of the respective prices to a deviation in the long-run equilibrium relation. Suppose a shock occurs and the futures price immediately jumps to its new equilibrium level, but the spot price does not. The long-run relation is now out of equilibrium. To reestablish the equilibrium spot/futures relation, the spot price must adjust. Now, if the proportion of the *subsequent* adjustment in the long-run relation is all attributed to the spot price changing, then it would mean that 100% of the elimination of the error occurs in the spot market and α_s will be nonzero. Consequently, that implies that none of the subsequent adjustment to eliminate the error occurs in the futures market, or conversely, that 100% of the *initial shock* is immediately reflected in the futures price, or in other words, all the price discovery due to the shock occurred in the futures market and none in the spot market. Hence, α_f will equal 0. Referring to equation (15), we see that for this illustration $CS_s = 0$ and $CS_f = 1$. The inference would then be that 100% of price discovery occurs in the futures market.

To further illustrate the point, suppose the adjustment coefficients are $\alpha_s = -0.5$ and $\alpha_f = 0.01$. A disruption in the long- run relation $(s_{t-1} - \beta_2 f_{t-1} - \beta_3)$ results in a larger absolute adjustment in the spot price, implying that the futures price is where the initial response occurs. At the new futures price the disequilibrium error equals $(s_{Old} - \beta_2 f_{New} - \beta_3)$. Subsequently, this is then followed by an adjustment to the spot price which moves it to a new equilibrium level s_{New}. Consequently, we can infer the relative contribution of each market to price discovery by

$$\left(\frac{\alpha_f}{\alpha_f - \alpha_s}, \quad \frac{-\alpha_s}{\alpha_f - \alpha_s}\right) = \left(\frac{.01}{.01\text{-}(-.5)}, \quad \frac{\text{-}(-.5)}{.01\text{-}(-.5)}\right) = (CS_s, CS_f)$$

$$(CS_s, CS_f) \approx (0.02, \quad 0.98).$$

In the illustration, a larger value for CS_f indicates that price discovery occurs first in the futures, market followed by an adjustment in the spot market.

Gonzalo and Granger (1995) develop a likelihood ratio test for joint hypotheses on (CS_s, CS_f). We specifically test two alternative null hypotheses, which we label

H1: $\left(CS_s, CS_f\right) = (1, 0)$ and H2: $\left(CS_s, CS_f\right) = (0,1)$. H1 is a test of the hypothesis that the permanent long-run price is revealed solely by the spot price, and H2 is a test of the hypothesis that the price is revealed solely by the futures price. The test statistic is distributed χ^2 with 1 degree of freedom under the null hypothesis.[63]

4. Spot and Futures Prices for Crude Oil

4.1 The Data

The data we examine are daily spot and futures prices of West Texas Intermediate (WTI, Cushing, OK, FOB) and corresponding data for Brent oil. Daily spot prices are obtained from the website of the U.S. Energy Information Administration and cover 1/2/1986–1/31/2020 for WTI and 6/23/1988–1/31/2020 for Brent, sourced from Bloomberg and Thomson Reuters Datastream. The EIA defines the spot price as "the price for a one-time open market transaction for near-term delivery of a specific quantity of product at a specific location where the commodity is purchased at current market rates."[64]

We label the spot price S and its natural log, *s*. Daily settlement NYMEX futures prices

[63] We refer the reader to Gonzalo and Granger (1995) for a complete development of the test.

[64] U.S. Energy Information Administration, "Glossary": http://www.eia.gov/tools/glossary/index.cfm?id=S . Spot oil trades are typically done through direct one-to-one interaction of traders by phone, text, or through online information networks. These are bilateral agreements. To the best of our knowledge, this is the first paper to study the impact on the location of price discovery in commodity markets that experienced a shift from pit trading of futures contracts by 'floor traders' to electronic trading. Focusing on a setting in which a commodity is traded both in a spot market as well as futures market, allows us to directly address which market plays the more prominent role in price discovery.

are obtained from the archives of the EIA for WTI and ICE Brent futures prices are obtained from Bloomberg.[65] We examine prices of the four nearest term futures contracts (one-month, two-month, three-month, and four-month maturities) traded on CME/NYMEX. Longer dated contracts tend to have low volume and hence the reliability of settlement prices may be questionable. We define the futures prices for the first four contracts as F1, F2, F3, F4, respectively, and the natural logs as *f1*, *f2*, *f3*, and *f4*.

We separate the analysis of the WTI market into two time periods, 1/2/1986–9/4/2006 and 9/5/2006–1/31/2020, and the analysis of the Brent market into two time periods, 6/23/1988–4/6/2005 and 4/7/2005–1/31/2020. Our motivation for this demarcation, as discussed earlier, stems from the structural shift to electronic trading that occurred in these markets and our interest in whether this structural change influenced where prices are discovered. The period following 1/31/2020 has been both turbulent and unusual for the oil market. As the estimation methods we employ are asymptotically efficient and the number of observations for each subperiod just mentioned is large, but the number of observations after 1/31/2020 is small, we defer an examination of the post-1/31/2020 period for a future study after more observations become available.

4.2 Trading Activity Across Subperiods

Table 1 reports average daily open interest and volume for the WTI futures contracts for the two subperiods. The source of the volume and open interest data is DataStream.

Two results are clear from the data reported in Table 1. First, in both periods WTI open interest and volume fall as the time until maturity increases. Second, open interest and volume increase dramatically during the second period, as already illustrated in Figure 1. Average daily

[65] WTI NYMEX futures prices observed at 2:30 PM Central Time; Brent ICE futures prices observed at 19:30 PM London time.

volume of the first maturing contract during the second period is roughly 6.6 times as large as that value during Period 1. Test results reported in Panel C indicate that average daily volume as well as open interest between the two subperiods are statistically significantly different and that this is also true for median tests based upon medians. Hence, we conclude that there was a significant increase in trading activity following the adoption of electronic trading.

Companion data for the Brent futures contract are presented in Table 14 and mirror the results presented in Table 13. Both futures trading volume and open interest increased significantly in the Brent futures market following the shift to electronic futures trading. Average daily trading volume of the first maturing Brent contract also increased by a factor of 6.6, surprisingly similar to WTI.

4.3 Unit Root and Cointegration Tests: WTI and Brent

Unit root test statistics based upon the Augmented Dickey-Fuller (ADF) test and the Phillips-Perron (PP) test were computed for each log price series for both the WTI and Brent spot and futures price series for both subperiods of the respective oil contracts. The tests consistently do not reject the null hypothesis that a unit root is present in the log price variables at the 5% significance level.[66] The unit root hypothesis, however, is always rejected for the first differences (change in log price). We therefore infer that the level log price series are nonstationary and each exhibit a unit root. As mentioned in Section 3, the price discovery metric we focus on, *CS*, is computed from the estimated adjustment coefficients of the bivariate vector error correction model (equation (14)), The theory developed in Section 3 suggests that the spot/futures prices are linked via a long-run linear equilibrium relation (equation 10). Given the unit root test results, we tested the hypothesis that in the bivariate models the log price series

[66] The unit root test results are available from the authors upon request.

exhibit a long-run linear relation (are cointegrated; Greene, 2008)), as would be predicted by the economic theory discussed earlier.

We test the hypothesis that each spot/futures price pair exhibits a cointegrating relation using the methods developed in Johansen (1991, 1996). We use the Schwarz BIC criterion for lag determination. The cointegration tests we conduct impose the restriction that the cointegrating equation contains an intercept. The test statistics are computed for both subperiods.

The cointegration tests uniformly reject the hypothesis of no cointegrating equation between the spot price and each futures price at the 5% level for both the WTI and Brent series, but do not reject the hypothesis of at most one cointegrating relation. We conclude from these results that the pairs of variables we investigate are all cointegrated, consistent with the theoretical result discussed earlier regarding the equilibrium relation between spot and futures prices.[67]

5. Price Leadership

The inferences based upon the cointegration tests argue for specifying the relation between each spot/futures pair as a VECM for purposes of estimating the adjustment coefficients from which the component shares are computed. Lag lengths for the models are determined by the Schwarz BIC criteria. There is one cointegrating equation in each model and it contains an intercept. As we pointed out in Section 2.1, the general cointegrating relation between the spot price and a futures price in the presence of a convenience yield is $s_{t-1} = \beta_2 f_{t-1} + \beta_3 + \varepsilon_{t-1}$, where the coefficient multiplying the futures price is not equal to 1. Therefore, we do not impose that restriction. Likewise, as indicated by equations (11) and (13), the parameter β_3 can

[67] For brevity, we do not report the tables which present the unit root and cointegration test results. The results are available from the authors upon request.

be positive or negative, especially in the presence of a convenience yield, so we place no constraints on that coefficient as well.

5.1 Estimation Results

The WTI estimation results are presented in Table 15 and the Brent results are presented in Table 16.[68] The layout of these two tables is the same: Models 1-4 present the estimated parameters of equation (14) for each pairing of the spot and futures contracts for the first, second, third, and fourth maturing contracts. The Vector Error Correction Model representation, equation (14), of the relation between a spot price and a futures price expresses the left-hand side vector ΔP_t as the change in each of the prices considered (Δs,Δf). Model 1 involves the spot and the nearby futures price, (Δs,$\Delta f1$); Model 2 involves (Δs,$\Delta f2$); Model 3 involves (Δs,$\Delta f3$); and Model 4 involves (Δs,$\Delta f4$). The parameters α_s and α_f are the estimated adjustment coefficients, which reflect how each price adjusts in response to deviations from the long-run equilibrium (the 2 × 1 vector α in equation (14)). The results are reported for the sample period 01/02/1986–9/4/2006 (Panel A), and separately for the period 9/5/2006–1/31/2020 (Panel B) for WTI (Table 15) and for the periods 6/23/1988–4/6/2005 (Panel A), and 4/7/2005–1/31/2020 (Panel B) for Brent (Table 16). We first discuss the results for the normalized coefficients that multiply the futures prices, β_2, and the estimated intercepts, β_3, in the cointegrating equations and then turn to a discussion of the estimated adjustment coefficients in Section 5.2.

In terms of the model, the cointegrating equilibrium relation is specified as ($s_{t-1} = \beta_2 f_{t-1} + \beta_3 + \varepsilon_{t-1}$), so the estimated coefficient multiplying the futures price is β_2. While the estimates of β_2 are always significantly different from 0 in both subperiods and for both sets of

[68] From the estimation results for each VECM, we find that the intercept term in the VAR is never significantly different from zero, so all subsequent results are based on models suppressing that intercept. We do not suppress the intercept in the cointegrating relation, as it is always significantly different from zero.

oil prices (row β_2 of both tables), the estimated coefficients are significantly different from 1 during the second subperiod for both oil markets. Values different from 1 are consistent with the existence on average of a convenience yield. The estimates of the intercepts of the cointegrating equations (β_3) are not significantly different from zero during the first subperiod but are consistently negative and significantly different from zero during the second subperiod for both WTI and Brent. As pointed out earlier when discussing equation (13), the sign of the intercept depends on both the level of the interest rate and carrying cost and factors that influence the convenience yield. If the sign of the intercept is negative, it implies that the influence of the latter factors is slight, so that effectively equation (11) explains the sign of the intercept.

The estimates of the adjustment coefficients will be discussed more fully in the next section. At this stage, we note that the estimates of α_s (the spot price adjustment parameter) are consistently negative and significantly different from zero, which is the case for both WTI (Table 15) and Brent (Table 16). However, for WTI the estimates of α_f (the futures price adjustment parameter) are not significantly different from zero *except* for Model 1 involving the spot and a nearby futures price, (Δs,$\Delta f1$), during the period prior to the shift to electronic trading. Following the movement to electronic trading, the WTI estimated futures price adjustment coefficients are never significantly different from zero. In contrast, the estimates of α_f for Brent are not significantly different from zero prior to the shift to electronic trading except for Model 1 but are all positive and significantly different from zero following the shift. We explore these results in more detail in the next section.

5.2 Which Market Leads and Which Follows?

To motivate our discussion of price leadership, consider the following example. Suppose α_s is greater in absolute value than α_f in a bivariate version of equation (14) in which the variables are the log spot price and the log futures price for a commodity. This would imply that the spot market responds more to deviations (shocks) to the long-run equilibrium relation between the spot and futures prices. This situation would likewise imply that the futures price responds less to deviations in the long-run equilibrium relation. The implication is that more of the information that led to the deviation has already been impounded in the futures price, hence, price leadership would be said to occur in the futures market.

The first piece of evidence regarding which market responds first to a shock leading to a deviation from the long-run relation comes from examination of the adjustment coefficients that are presented in Table 15 and Table 16. The estimated cointegrating (long-run equilibrium) relation for each model has the general form $s_{t-1} = \beta_2 f_{t-1} + \beta_3 + \varepsilon_{t-1}$, where f is the log futures price being examined. Suppose there is an information shock, and the futures price moves to a new higher equilibrium level. This would generate an error with a negative sign, assuming the spot price has not yet adjusted. For the equilibrium between the spot and futures prices to be reestablished, the spot price would need to also rise. The per period adjustment to the spot price will depend on the adjustment coefficient α_s. If the estimated adjustment coefficient for the spot price model is significantly different from zero and negative, then the response of the spot price will be for it to change in the desired direction, a positive change in the example, to reestablish the equilibrium relation. Likewise, if the new information results in a drop in the futures price creating a positive error, the spot price would need to fall to reestablish the equilibrium. These examples characterize the case in which the futures price moves first. Conversely, if the spot price moves first, or if both prices contribute to price discovery, we would expect the adjustment coefficient for the futures price model to be positive and significantly different from zero, inducing adjustments to the futures price in response to a positive or negative deviation wrought from an increase or decrease in the spot price. Thus, by examining both the sign of the adjustment coefficients as well as whether they are significantly different from zero, we can infer which price is the leader and which price is the follower, or if both prices contribute to price discovery.

Table 15 shows the estimates of α_s for the spot WTI adjustment models, which are consistently negative and significantly different from zero at the 1% or 5% level for both subperiods. Note, however, that all estimates of α_s are less than 1, implying only partial adjustment in the spot price each period (day) going forward. Also, the adjustment coefficients for the futures prices are significantly smaller in absolute value than the adjustment coefficients for the spot price. However, the critical exception is the futures price adjustment coefficient in Model 1 involving the spot and a nearby futures price, (Δs,$\Delta f1$). During the first subperiod, the coefficient is both positive and significantly different from zero, indicating that both prices adjust to deviations in the long-run relation. However, this is not true for the remaining models, in which the adjustment coefficient for the futures price is never significantly different from 0. These results suggest that price leadership occurs largely in the futures market, i.e., an information shock results first in a change in the futures price followed by an adjustment to the spot price to reestablish the equilibrium. While this conclusion is universally true for WTI during the second subperiod following the adoption of side-by-side trading, it is not true for Model 1, (Δs,$\Delta f1$), during the first subperiod.

Recall that we are studying price data measured at a daily frequency. The results reported in Table 15 indicate that α_s, the adjustment coefficients for the spot price change model, are largest in absolute value in both subperiods for Model 1, (Δs,$\Delta f1$). The pre-shift period estimate equals -0.481 and the post-shift period estimate equals -0.453, both significantly different from zero. However, the estimates of α_s decline in absolute value as the futures maturity increases. Focusing on the first subperiod, the adjustment coefficient α_s equals -0.481, implying that 48.1% of any long-run error is adjusted in one day through an adjustment to the spot price. Coupling this with the adjustment coefficient for the futures price change model α_f , which equals 0.142 and is also significantly different from zero, the total one-day adjustment to eliminate the error was 62.3%. During the second subperiod, the spot price adjustment coefficient equals -0.453, which is statistically significant, but the futures price adjustment coefficient is not significantly different from zero, implying that the one-day adjustment was 45.3% during the second subperiod and was entirely concentrated in a change to the spot price. These results indicate that the speed of correction for the spot price model was large and significant, while the futures price adjustment coefficient fell between the two periods.

An interesting contrast involves comparison of the adjustment coefficients estimated for the first subperiod for Model 1, $(\Delta s,\Delta f_1)$, at (-0.481, 0.142) with those for Model 2, $(\Delta s,\Delta f_2)$, for the same subperiod, which equal (-0.081, -0.007).[69] Clearly, the size (and hence speed) of the spot price adjustment coefficient for the second model is much smaller in absolute value than for the first model. This is also true for the futures price adjustment coefficients (0.142 versus -0.007). In the latter case, we also see that the estimated coefficient (-0.007) is not significantly different from zero for Model 2. These results indicate that there is a much smaller spot price correction for deviations between spot prices and the second maturing contract. The difference in the spot price adjustment coefficients (-0.481 and -0.081) suggests that the spot price adjustment keys off the near-term futures price *f1,* and not *f2,* which is perhaps not a surprise. The surprise is really that the deviation of the long-run relation involving the second maturing contract elicits

[69] The nearby WTI oil contract is considered one of the most liquid contracts traded (Fett and Haynes, 2017), and its liquidity has been shown to be associated with price efficiency (Raman, Robe, and Yadav, 2020).

any spot price correction at all. One explanation is that the theoretical no-arbitrage relation *between* futures contracts does not, in fact, always hold.[70]

The results presented in Table 15 show that the adjustment coefficients for the futures price are generally (except Model 1) not significantly different from zero. Therefore, deviations in the long-run relation arise because the futures price moves and is followed by a change in the spot price to restore the equilibrium relation. This is consistent with futures playing an important price discovery role in the WTI market both in the period prior to and after the shift to electronic trading.

Turning to the Brent estimation results in Table 16 we see some parallels with the WTI results, but also some striking differences. First, the estimated spot price adjustment coefficients are consistently negative and statistically significant in both subperiods. These results indicate that, like the WTI results, the spot price adjusts to deviations in the equilibrium relation and in the expected direction. The result holds in both subperiods. Interesting however is the fact that the absolute sizes of the spot price adjustment coefficients are much smaller than for the WTI adjustment coefficients. Second, the adjustment coefficients for the futures price model, while consistently positive, are not significantly different from zero during the period prior to the shift to electronic trading except for Model 1 at 10% level. However, following the shift, the futures

---

[70] Specifically, assume two futures contracts are written on the same underlying asset. The first matures at time T and the current date is t, and the second contract matures at time $T+v$. Under no arbitrage, the futures price today of the first maturing contract is $F_{1,t} = S_t e^{(r_t + c_t - \delta_t)(T-t)}$, where $(r_t + c_t - \delta_t)$ equals the net of the continuously compounded interest rate, storage cost, and convenience yield. In the absence of arbitrage, the second contract's futures price at the same date t, assuming $(r_t + c_t - \delta_t)$remains fixed over the period T to $T+v$, should equal $F_{2,t} = S_t e^{(r_t + c_t - \delta_t)(T+v-t)} = F_{1,t} e^{(r_t + c_t - \delta_t)(v)}$. As such, if $F_{1,t}$ adjusts immediately to new news (that is prices are discovered in the futures market), then there should likewise be an immediate adjustment to $F_{2,t}$ as well. However, the spot price adjustment to the deviation between $F_{2,t}$ and the spot price could be influenced by a delayed residual adjustment. Ederington et al. (2021) present evidence that the no-arbitrage condition does sometimes fails to hold in the oil market and that adjustments to restore equilibrium can occur with a lag. In which case, if the longer-dated futures contract, that is, spot price errors are actually too small because the futures price has not fully adjusted, then we should see a smaller spot price adjustment. What we see in Table 3 are much smaller in magnitude spot price adjustment estimated coefficients for the models involving longer-dated futures, which would be consistent with this conjecture.

price adjustment coefficients are both positive and significantly different from zero. The latter result indicates that following the shift to electronic trading in the Brent futures market, deviations from the long-run equilibrium relation also resulted in subsequent adjustments to the futures price. This is in stark contrast to the results for WTI, suggesting that despite an increase in liquidity (as reflected in the increase in volume) in the Brent futures market, those prices do not initially fully adjust to deviations in the spot/futures equilibrium relation.

5.3 Are There Lagged Short-Term Relations Between Price Changes?

The general formulation of the VECM, repeated below in equation (17), shows that price changes, including those induced by adjustment to deviations in the long-run equilibrium relation, can also be related to lagged short-term price changes, through the coefficient matrices Γ_k , $k = 1 \ldots K\text{-}1$.

$$\Delta P_t = c + \alpha(\varepsilon_{t-1}) + \sum_{k=1}^{K-1} \Gamma_k \Delta P_{t-k} + e_t. \tag{17}$$

We test the null hypotheses that: 1) short-run changes in futures prices do not Granger-cause changes in spot prices, and 2) short-run changes in spot do not Granger-cause futures price changes.[71] We conduct such tests based upon estimated versions of equation (17), where the 2 × 1 vector ΔP_t contains the change in the log spot price and the change in the log futures price, respectively. The test results almost uniformly reject hypotheses 1) and 2) at the 1% level for all WTI spot/futures bivariate combinations for both subperiods. The one exception is Model 1 with (Δs,$\Delta f1$) during both the subperiods, where we cannot reject the null that changes in spot prices do not Granger-cause changes in futures prices. Therefore, while the WTI futures market is

[71] In a bivariate system with K lags, x_t is said to Granger-cause y_t if y_t can be better predicted using the histories of both x_t and y_t than it can by using the history of y_t alone. The classical test of the null hypothesis that x_t does not Granger-cause y_t is a test to confirm that the coefficients of the K lagged values of x_t are jointly equal to zero (a Wald test statistic) and is based on the assumption that the series being studied are stationary (Granger, 1969; Sims, 1972). The condition for the validity of the test in terms of equation (17) is that the *change* in the futures price and the *change* in the spot price are stationary, which we have confirmed (please refer to Section 4.3), and for specification purposes that the long-run error from the cointegrating relations is also stationary.

generally the market in which information is first reflected, inducing a futures price change and a consequent deviation in the long-run equilibrium relation, there are nevertheless also cross-market adjustments in response to short-term price changes.

We repeat the tests for causality for the Brent price models and find that once again we reject both hypotheses in both the periods before and after the shift to electronic futures trading in the Brent futures market for (Δs,$\Delta f1$). That is, our inference is that we reject the null that spot prices do not Granger-cause futures prices, but we also reject the null that futures prices do not Granger-cause spot prices. While we do not explore the relation between these short-run changes further, one possible explanation for the finding of such short-run effects is that these changes are the manifestation of trading frictions.[72]

6. Relative Price Discovery

Table 17 presents the *CS* measure and its statistics for the pairs of prices based on the results presented in Table 15 and 16.[73] The table is partitioned to show results for the two sample periods, before and after the introduction of electronic trading. Panel A presents results for WTI and Panel B for Brent. As mentioned earlier, a finding of $CS_f > CS_s$ would be evidence supporting the hypothesis that the futures market contributes more to price discovery than the spot market. Recall that, in the bivariate cases we examine, component shares are computed as

$$\left(\frac{\alpha_f}{\alpha_f - \alpha_s}, \quad \frac{-\alpha_s}{\alpha_f - \alpha_s}\right) = (CS_s, CS_f)\ . \tag{18}$$

[72] Note that there is an estimation issue that arises in a regression of price changes on lagged price changes when the two prices are cointegrated. The issue can be seen in equation (17) by recognizing that a regression involving first differences alone would suffer from omitted variable bias because the variable ε_{t-1} would be excluded from the model. This means that the estimated coefficients on the lagged changes will in most cases be biased, and it is possible that the standard errors will be smaller than the model that includes the omitted variable. For brevity we do not include the respective tables.

[73] Appendix Table C.1 presents alternative measures of price discovery, Information Share and Generalized Information Share.

Thus, the component shares are an alternative way in which to view price leadership. If the spot price adjusts to shocks to the long-run equilibrium relation that first arise from the futures price moving, then the contribution of the spot price to price discovery (CS_s) should be 0 since the futures price is the venue through which shocks are first manifested, followed by adjustments in the spot. Test results are reported for the null hypotheses H1: $(CS_s, CS_f) = (1,0)$ (null hypothesis: prices are discovered in the spot market), and H2: $(CS_s, CS_f) = (0,1)$ (null hypothesis: prices are discovered in the futures market) for each subperiod.[74]

Table 17 Panel A shows that for WTI, within both the first and second subperiods, the estimates of *CS* suggest that price discovery is concentrated in the futures market relative to the spot market for WTI. The test results uniformly reject H1 but do reject H2 for all spot price/futures price pairs involving the spot price and the second, third, and fourth maturing contract prices for both subperiods, and for the spot and first maturing contract during the period following the commencement of side-by-side trading. The result that both H1 and H2 are rejected at the .01 level during the first subperiod for the spot and nearby futures price pair support the alternative hypothesis that both prices contributed to price discovery before the shift to electronic trading. In contrast, H1 is rejected and H2 is not rejected during the second period for the spot and nearby futures price pair indicating that price discovery migrated to the futures market during the second subperiod, following the shift to electronic futures trading. While the result suggests that the spot and nearby futures contract both contributed to price discovery during the first subperiod, the contribution of the futures price became significantly more important during the second subperiod, indeed was the sole contributor to price discovery following the shift to electronic trading under the *CS* measure. One explanation is that trading in

[74] The *p*-values shown in Table 17 are based on the test statistic developed in Gonzalo and Granger (1995), as discussed in Section 3.

the futures market through increased access via electronic trading enhanced the information set, liquidity, and price discovery in the futures market.

Table 17 Panel B presents the results for Brent, where electronic trading commenced on 4/7/2005. In contrast to the results for WTI, price discovery in the Brent market did not shift exclusively to the futures market following the adoption of electronic trading, even though (as Figure 4 and Table 14 of the indicate) there was also a tremendous increase in Brent futures trading volume following the shift to electronic trading. For Brent, the futures market was where price discovery occurred primarily, by reference to the *CS* measure, prior to electronic trading. Likewise, while the test (H1) that the spot market was the sole source of price discovery is rejected, the test that the futures market was the sole source of price discovery (H2) is not rejected at the .01 level for the period prior to the shift to electronic trading. This changed following the shift to electronic futures trading on ICE. Unlike the WTI market, following the commencement of electronic futures trading, in the Brent market price discovery became divided between the spot and futures markets. Statistical tests of the hypotheses that the spot market is the sole source of price discovery and separately that the futures market is the sole source are both rejected following the shift to electronic trading for Brent.

During the post-shift period, both the spot and futures prices are significant contributors to price discovery in the Brent market for all spot/futures price pairs up through *spot*/*f3*, and that discovery occurs solely in the spot market for *spot*/*f4*. These results indicate that the spot market has played an increased role in price discovery in the Brent market following the shift to electronic trading by ICE, which is the opposite of what we find for the WTI market. These results are also the opposite of what we would expect to find if liquidity increases in the futures market are positively correlated with price discovery. Referring to Figure 4, these results indicate that a liquidity (volume) increase associated with a shift to electronic trading did not lead to a shift in price discovery to the market in which increased volume has occurred, that is, the futures market. While higher trade volume indicates greater overall market interest in a particular commodity, that volume may not imply improved price quality. For instance, if greater liquidity attracts more uninformed, or less well informed, traders, then price quality could deteriorate. Consequently, the location of price discovery could be impacted if, for instance in our case, the spot market is viewed as providing useful signals.[75]

We conclude that there are significant differences in how the switch to electronic trading impacted price discovery in the Brent and the WTI markets. For WTI, both the spot and futures markets contributed to price discovery prior to the shift to side-by-side electronic and open outcry trading, following the shift the futures market became dominant. For Brent, futures price was the dominant source of price discovery prior to the shift to electronic trading but following this shift both spot and futures become important to price discovery. In other words, while the futures market became the principal point for price discovery for WTI after the commencement of electronic trading, the importance of futures prices for price discovery declined in the Brent market after the switch to electronic trading. These conclusions reinforce our earlier discussion on price leadership and suggest that this structural change, while associated with a dramatic increase in trading volume, interestingly had a heterogeneous impact for the importance of futures price for price discovery.

[75] ICE (the primary exchange for Brent futures) claims that "Physical traders prefer the (Brent) oil complex by a wide margin, enabling prices based on the Brent benchmark to more accurately reflect global supply and demand, with less impact from speculators and market makers." https://www.theice.com/insights/market-pulse/brent-the-worlds-crude-benchmark.

7. Summary and Conclusions

O'Hara (2003) argues that price discovery is one of the two most important functions of a market and Black (1976) asserts that price discovery even dominates the risk management role played by futures markets. We study whether price discovery in the WTI oil market was influenced by the commencement of side-by-side electronic and open outcry trading of the NYMEX futures contract on 9/5/2006, and in the Brent oil market by the shift to electronic trading of the Brent futures contract by ICE on 4/7/2005. Focusing on a setting in which a commodity is traded both in a spot as well as a futures market allows us to directly address which market plays the more prominent role in price discovery and how that role was impacted by the shift to electronic trading. Daily crude oil spot and futures prices form the basis for the study. We examine price leadership through the direct analysis of coefficients of vector error correction models relating spot and futures prices, as well as measures of price discovery.

We conclude that price discovery was largely concentrated in the WTI futures market during the period we study, and that the implementation of electronic trading did not alter this result. In fact, our findings suggest that the commencement of electronic trading in WTI futures was associated with a total migration of price discovery to the futures market. This improvement in price discovery occurred despite the increased volume of trading that followed the change. These results provide "breathing room" in two respects. First, the increase in trading activity following the shift to electronic trading did not lead to a deterioration in the market's ability to discover prices in the WTI futures market. Second, those expressing regulatory or institutional concerns surrounding electronic trading of WTI oil futures in the United States can take comfort from the fact that increased trading activity following this institutional change did not lead to a deterioration in price discovery in the WTI oil futures market.

In contrast, however, while our evidence indicates that price discovery was concentrated in the Brent oil futures market prior to the shift to electronic trading by ICE on 4/7/2005, this was not the case in the aftermath (4/7/2005–5/26/2020). Following the change, we find that both the spot and futures markets contributed to price discovery, suggesting that the contribution of the Brent futures market to price discovery deteriorated after the shift to electronic trading. Our findings from the Brent market show that liquidity increases are not necessarily accompanied by price discovery improvements and suggest that the benefits of electronic trading need to be assessed from the standpoint of improvements in both liquidity and price discovery.

Explaining the differences between how WTI futures and Brent futures responded to electronification from the standpoint of price discovery is beyond the scope of this paper. While market participants look to both markets to assess the global oil price trajectory, there are important differences between these two markets that need to be accounted for in such a study, which we leave for future research.

References

Abel, A. B., Dixit, A. K., Eberly, J. C., & Pindyck, R. S. (1996). Options, the value of capital, and investment. *The Quarterly Journal of Economics*, *111*(3), 753-777.

Aguerrevere, F. L. (2009). Real options, product market competition, and asset returns. *The Journal of Finance*, *64*(2), 957-983.

Ahmadi, M., Manera, M., & Sadeghzadeh, M. (2019). The investment-uncertainty relationship in the oil and gas industry. *Resources Policy*, *63*, 101439.

Aizenman, J., & Marion, N. P. (1993a). Policy uncertainty, persistence and growth. *Review of International Economics*, *1*(2), 145-163.

Aizenman, J., & Marion, N. P. (1993b). Macroeconomic uncertainty and private investment. *Economics Letters*, *41*(2), 207-210.

Al-Harthy, M. H. (2007). Stochastic oil price models: Comparison and impact. *The Engineering Economist*, 52(3), 269-284.

Anderson, S. T., Kellogg, R., & Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3), 984-1026.

Ap Gwilym, O., & Alibo, E. (2003). Decreased price clustering in FTSE100 futures contracts following a transfer from floor to electronic trading. *The Journal of Futures Markets*, 23(7), 647-659.

Arnold, T., & Crack, T. F. (2000). Option pricing in the real world: a generalized binomial model with applications to real options. *Available at SSRN 240554*.

Arouri, M. E. H., & Nguyen, D. K. (2010). Oil prices, stock markets and portfolio investment: Evidence from sector analysis in Europe over the last decade. *Energy Policy*, *38*(8), 4528-4539.

Asmussen, S., Avram, F., & Pistorius, M. R. (2004). Russian and American put options under exponential phase-type Lévy models. *Stochastic Processes and Their Applications*, *109*(1), 79-111.

Bachmann, R., & Bayer, C. (2014). Investment dispersion and the business cycle. *American Economic Review*, *104*(4), 1392-1416.

Backus, D. K., Foresi, S., & Wu, L. (2004). Accounting for biases in Black-Scholes. *Available at SSRN 585623*.

Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, *16*(1), 101-143.

Baldwin, C. Y. (1982). Optimal sequential investment when capital is not readily reversible. *The Journal of Finance*, *37*(3), 763-782.

Bali, T. G., & Murray, S. (2013). Does risk-neutral skewness predict the cross section of equity option portfolio returns?. *Journal of Financial and Quantitative Analysis*, 1145-1171.

Bates, D. S. (1997). The skewness premium: Option pricing under asymmetric processes. *Advances in Futures and Options Research*, *9*, 51-82.

Battalio, R., Ellul, A., & Jennings, R. (2007). Reputation effects in trading on the New York Stock Exchange. *The Journal of Finance*, 62(3), 1243-1271.

Bekiros, S. D., & Diks, C. G. (2008). The relationship between crude oil spot and futures prices: Cointegration, linear and nonlinear causality. *Energy Economics*, 30(5), 2673-2685.

Bennouna, K., Meredith, G. G., & Marchant, T. (2010). Improved capital budgeting decision making: evidence from Canada. *Management Decision*.

Benveniste, L. M., Marcus, A. J., & Wilhelm, W. J. (1992). What's special about the specialist?. *Journal of Financial Economics*, 32(1), 61-86.

Berger, P. G., Ofek, E., & Swary, I. (1996). Investor valuation of the abandonment option. *Journal of Financial Economics*, *42*(2), 259-287.

Bernales, A., Valenzuela, M., & Zer, I. (2023). Effects of information overload on financial markets: How much is too much?. *Board of Governors of the Federal Reserve System International Finance Discussion Paper*, (1372).

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, *98*(1), 85-106.

Bernhardt, D., Campello, M., & Kutsoati, E. (2006). Who herds?. *Journal of Financial Economics*, 80(3), 657-675.

Bessembinder, H. (2003). Trade execution costs and market quality after decimalization. *Journal of Financial and Quantitative Analysis*, 38(4), 747-777.

Bessembinder, H., Coughenour, J. F., Seguin, P. J., & Smoller, M. M. (1995). Mean reversion in equilibrium asset prices: Evidence from the futures term structure. *The Journal of Finance*, 50(1), 361-375.

Bewley, R., & Fiebig, D. G. (2002). On the herding instinct of interest rate forecasters. *Empirical Economics*, 27, 403-425.

Bikhchandani, S., Hirshleifer, D., Tamuz, O., & Welch, I. (2021). Information cascades and social learning, Working paper, *National Bureau of Economic Research*.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3(1-2), 167-179.

Block, S. (2007). Are “real options” actually used in the real world?. *The Engineering Economist*, *52*(3), 255-267.

Bloom, N., Bond, S., & Van Reenen, J. (2007). Uncertainty and investment dynamics. *The Review of Economic Studies*, *74*(2), 391-415.

Boehmer, E., & Wu, J. (2013). Short selling and the price discovery process. The *Review of Financial Studies*, 26(2), 287-322.

Boomhower, J. (2019). Drilling like there's no tomorrow: bankruptcy, insurance, and environmental risk. *American Economic Review*, 109(2), 391-426.

Bopp, A. E., & Sitzer, S. (1987). Are petroleum futures prices good predictors of cash value?. *The Journal of Futures Markets*, 19(4), 705-719.

Borland, L., & Bouchaud, J. P. (2004). A non-Gaussian option pricing model with skew. *Quantitative Finance*, *4*(5), 499-514.

Boyarchenko, S. (2004). Irreversible decisions and record-setting news principles. *American Economic Review*, 94(3), 557-568.

Boyarchenko, S. I., & Levendorskiĭ, S. Z. (2000). Option pricing for truncated Lévy processes. *International Journal of Theoretical and Applied Finance*, 3(03), 549-552.

Boyarchenko, S. I., & Levendorskiĭ, S. Z. (2002). Perpetual American options under Lévy processes. *SIAM Journal on Control and Optimization*, 40(6), 1663-1696.

Boyer, B., Mitton, T., & Vorkink, K. (2010). Expected idiosyncratic skewness. *The Review of Financial Studies*, *23*(1), 169-202.

Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory*, 26(2), 318-339.

Bray, M. M., & Savin, N. E. (1986). Rational expectations equilibria, learning, and model specification. *Econometrica*, 1129-1160.

Brennan, M. J., & Schwartz, E. S. (1985). Evaluating natural resource investments. *Journal of Business*, 58, 135-157.

Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *The Review of Financial Studies*, 27(8), 2267-2306.

Brogaard, J., Hendershott, T., & Riordan, R. (2019). Price discovery without trading: Evidence from limit orders. *The Journal of Finance*, 74(4), 1621-1658.

Brogaard, J., Ringgenberg, M. C., & Roesch, D. (2022). Does floor trading matter?. Available at SSRN 3609007.

Broughton, J. B., & Lobo, B. J. (2018). Herding and anchoring in macroeconomic forecasts: The case of the PMI. *Empirical Economics*, 55, 1337-1355.

Brown, C. A., & Robinson, D. M. (2002). Skewness and kurtosis implied by option prices: A correction. *Journal of Financial Research*, *25*(2), 279-282.

Brown, L. D. (2001). A temporal analysis of earnings surprises: Profits versus losses. *Journal of Accounting Research*, 39(2), 221-241.

Brown, M. B., & Forsythe, A. B. (1974). Robust tests for the equality of variances. *Journal of the American Statistical Association*, 69(346), 364-367.

Brunnermeier, M. K., & Parker, J. A. (2005). Optimal expectations. *American Economic Review*, *95*(4), 1092-1118.

Caballero, R. J. (1999). Aggregate investment. *Handbook of Macroeconomics*, *1*, 813-862.

Cameron, C. & A. Trivedi (2005). Microeconometrics: Methods and Applications. Cambridge University Press, Cambridge, UK.

Cao, H., Guo, L., & Zhang, L. (2020). Does oil price uncertainty affect renewable energy firms' investment? Evidence from listed firms in China. *Finance Research Letters*, *33*, 101205.

Carruth, A., Dickerson, A., & Henley, A. (2000). What do we know about investment under uncertainty?. *Journal of Economic Surveys*, *14*(2), 119-154.

Casassus, J., & Collin-Dufresne, P. (2005). Stochastic convenience yield implied from commodity futures and interest rates. *The Journal of Finance*, 60(5), 2283-2331.

Chang, B. Y., Christoffersen, P., & Jacobs, K. (2013). Market skewness risk and the cross section of stock returns. *Journal of Financial Economics*, *107*(1), 46-68.

Chang, B. Y., Christoffersen, P., Jacobs, K., & Vainberg, G. (2012). Option-implied measures of equity risk. *Review of Finance*, 16(2), 385-428.

Chang, C. P., & Lee, C. C. (2015). Do oil spot and futures prices move together?. *Energy Economics*, 50, 379-390.

Chen, F., & Linn, S. C. (2017). Investment and operating choice: Oil and natural gas futures prices and drilling activity. *Energy Economics*, 66, 54-68.

Chen, J., Hong, H., & Stein, J. C. (2001). Forecasting crashes: Trading volume, past returns, and conditional skewness in stock prices. *Journal of Financial Economics*, *61*(3), 345-381.

Chen, P. F., Lee, C. C., & Zeng, J. H. (2014). The relationship between spot and futures oil prices: Do structural breaks matter?. *Energy Economics*, 43, 206-217.

Chen, Q., & Jiang, W. (2006). Analysts' weighting of private and public information. *The Review of Financial Studies*, 19(1), 319-355.

Chiou-Wei, S. Z., Linn, S. C., & Zhu, Z. (2014). The response of US natural gas futures and spot prices to storage change surprises: Fundamental information and the effect of escalating physical gas production. *Journal of International Money and Finance*, 42, 156-173.

Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *The Journal of Finance*, 56(2), 501-530.

Chung, K. H., & Chuwonganant, C. (2022). Covid-19 and the stock market: Liquidity, pricing efficiency, and trading, Forthcoming *Journal of Financial Markets*.

Claerbout, J., & Muir, F. (2008). Hubbert math. *Sepstanford.edu*.

Clement, M. B. (1999). Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter?. *Journal of Accounting and Economics*, 27(3), 285-303.

Clement, M. B., & Tse, S. Y. (2005). Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance*, 60(1), 307-341.

Clement, M. B., Hales, J., & Xue, Y. (2011). Understanding analysts' use of stock returns and other analysts' revisions when forecasting earnings. *Journal of Accounting and Economics*, 51(3), 279-299.

Conrad, J., Dittmar, R. F., & Ghysels, E. (2013). Ex ante skewness and expected stock returns. *The Journal of Finance*, *68*(1), 85-124.

Cooper, I. (2006). Asset pricing implications of nonconvex adjustment costs and irreversibility of investment. *The Journal of Finance*, *61*(1), 139-170.

Cooper, R. A., Day, T. E., & Lewis, C. M. (2001). Following the leader: A study of individual analysts' earnings forecasts. *Journal of Financial Economics*, 61(3), 383-416.

Corrado, C. J., & Su, T. (1996a). Skewness and kurtosis in S&P 500 index returns implied by option prices. *Journal of Financial Research*, *19*(2), 175-192.

Corrado, C. J., & Su, T. (1996b). S&P 500 index option tests of Jarrow and Rudd's approximate option valuation formula. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, *16*(6), 611-629.

Covert, T. (2015). Experiential and social learning in firms: the case of hydraulic fracturing in the Bakken Shale. *Available at SSRN 2481321*.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2), 187-202.

Cox, J. C., & Ross, S. A. (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3(1-2), 145-166.

Dai, Z., Zhou, H., Kang, J., & Wen, F. (2021). The skewness of oil price returns and equity premium predictability. *Energy Economics*, *94*, 105069.

Datta, D. D., Londono, J. M., & Ross, L. J. (2017). Generating options-implied probability densities to understand oil market events. *Energy Economics*, 64, 440-457.

Davis, G. A., & Cairns, R. D. (2017). The odd notion of "reversible investment". *Journal of Banking & Finance*, *81*, 172-180.

De Jong, F. (2002). Measures of contributions to price discovery: A comparison. *Journal of Financial Markets*, 5(3), 323-327.

Décaire, P. H., Gilje, E. P., & Taillard, J. P. (2020). Real option exercise: Empirical evidence. *The Review of Financial Studies*, *33*(7), 3250-3306.

Dechow, P. M., Hutton, A. P., & Sloan, R. G. (2000). The relation between analysts' forecasts of long-term earnings growth and stock price performance following equity offerings. *Contemporary Accounting Research*, 17(1), 1-32.

Dias, J. C., Larguinhob, M., & Braumannc, C. A. (2015). Entry and Exit Decisions under Uncertainty for a Generalized Class of One-Dimensional Diffusions. In *Real Options Conference. Athens and Monemvasia*.

Diavatopoulos, D., Doran, J. S., Fodor, A., & Peterson, D. R. (2012). The information content of implied skewness and kurtosis changes prior to earnings announcements for stock and option returns. *Journal of Banking & Finance*, *36*(3), 786-802.

Dincerler, C., Khokher, Z., & Simin, T. (2020). An empirical analysis of commodity convenience yields. *Quarterly Journal of Finance*, 10(02), 2050009.

Dittmar, R. F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *The Journal of Finance*, *57*(1), 369-403.

Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of Political Economy*, *97*(3), 620-638.

Dixit, A. (1991). Irreversible investment with price ceilings. *Journal of Political Economy*, *99*(3), 541-557.

Dixit, R. K., & Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton university press.

Doshi, H., Kumar, P., & Yerramilli, V. (2018). Uncertainty, capital investment, and risk management. *Management Science*, *64*(12), 5769-5786.

Drakos, K., & Konstantinou, P. T. (2013). Investment decisions in manufacturing: assessing the effects of real oil prices and their uncertainty. *Journal of Applied Econometrics*, *28*(1), 151-165.

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(4), 1350-1380.

Ederington, L. H., Lin, F., Linn, S. C., & Yang, L. Z. (2019). EIA storage announcements, analyst storage forecasts, and energy prices. *The Energy Journal*, 40(5), 121-142.

Ehrbeck, T., & Waldmann, R. (1996). Why are professional forecasters biased? Agency versus behavioral explanations. *The Quarterly Journal of Economics*, 111(1), 21-40.

Elder, J., & Serletis, A. (2010). Oil price uncertainty. *Journal of Money, Credit and Banking*, *42*(6), 1137-1159.

Elder, J., Miao, H., & Ramchander, S. (2014). Price discovery in crude oil futures. *Energy Economics*, 46, 518-527.

Eppler, M. J., & Mengis, J. (2008). The concept of information overload: A review of literature from organization science, accounting, marketing, mis, and related disciplines, *The Information Society*, 20(5), 325-344.

Fan, Y., & Xu, J. H. (2011). What has driven oil prices since 2000? A structural change perspective. *Energy Economics*, 33(6), 1082-1094.

Ferderer, J. P. (1996). Oil price volatility and the macroeconomy. *Journal of Macroeconomics*, *18*(1), 1-26.

Fernandez-Perez, A., Garel, A., & Indriawan, I. (2020). Natural gas storage forecasts: Is the crowd wiser?. *The Energy Journal*, 41(5), 213-238.

Fett, N., & Haynes, R. (2017). Liquidity in select futures markets. White paper. Office of the Chief Economist, United States Commodity Futures Trading Commission.

Figlewski, S. (2016). What goes into risk-neutral volatility? Empirical estimates of risk and subjective risk preferences. *The Journal of Portfolio Management*, 43(1), 29-42.

Figuerola-Ferretti, I., & Gonzalo, J. (2010). Modelling and measuring price discovery in commodity markets. *Journal of Econometrics*, 158(1), 95-107.

Fisher, I. (1930). Theory of interest: as determined by impatience to spend income and opportunity to invest it. Augustusm Kelly Publishers, Clifton.

Fleming, J., Ostdiek, B., & Whaley, R. E. (1996). Trading costs and the relative rates of price discovery in stock, futures, and option markets. *The Journal of Futures Markets*, 16(4), 353-387.

Foucault, T., Pagano, M., & Röell, A. (2013). Market liquidity: Theory, evidence, and policy. Oxford University Press, USA.

Gallo, G. M., Granger, C. W., & Jeon, Y. (2002). Copycats and common swings: The impact of the use of forecasts in information sets. *IMF staff Papers*, 49(1), 1-18.

Garbade, K. D., & Silber, W. L. (1983). Price movements and price discovery in futures and cash markets. *The Review of Economics and Statistics*, 289-297.

Gay, G. D., Simkins, B. J., & Turac, M. (2009). Analyst forecasts and price discovery in futures markets: The case of natural gas storage. *The Journal of Futures Markets*, 29(5), 451-477.

Gibson, R., & Schwartz, E. S. (1990). Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance,* 45(3), 959-976.

Gilje, E. P. (2019). Does local access to finance matter? Evidence from US oil and natural gas shale booms. *Management Science*, *65*(1), 1-18.

Girardi, D. (2021). The neoclassical theory of aggregate investment and its criticisms.

Gonzalo, J., & Granger, C. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1), 27-36.

Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 424-438.

Greene, W. H. (2003). *Econometric analysis*. Pearson Education India.

Greene, W. H. (2008). Econometric analysis, 6th ed, Pearson Prentice Hall, Upper Saddle River, NJ.

Grossman, S. J. (1977). The existence of futures markets, noisy rational expectations and informational externalities. *The Review of Economic Studies*, 44(3), 431-449.

Gu, C., & Kurov, A. (2018). What drives informed trading before public releases? Evidence from natural gas inventory announcements. *The Journal of Futures Markets*, 38(9), 1079-1096.

Gukhal, C. R. (2001). Analytical valuation of American options on jump-diffusion processes. *Mathematical Finance*, *11*(1), 97-115.

Gulen, H., & Ion, M. (2016). Policy uncertainty and corporate investment. *The Review of Financial Studies*, *29*(3), 523-564.

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

Handa, P., Schwartz, R., & Tiwari, A. (2006). The economic value of a trading floor: Evidence from the American Stock Exchange, *Journal of Business*, 77, 331-355.

Harris, F. H. D., McInish, T. H., & Wood, R. A. (2002). Common factor components versus information shares: A reply. *Journal of Financial Markets*, 5(3), 341-348.

Harris, F. H. D., McInish, T. H., Shoesmith, G. L., & Wood, R. A. (1995). Cointegration, error correction, and price discovery on informationally linked security markets. *Journal of Financial and Quantitative Analysis*, 30(4), 563-579.

Harris, L. (2003). Trading and exchanges: Market microstructure for practitioners. Oxford University Press, New York, NY.

Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, *55*(3), 1263-1295.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *The Journal of Finance*, 50(4), 1175-1199.

Hausman, J. (1999). The effect of sunk costs in telecommunications regulation. In *The new investment theory of real options and its implication for telecommunications economics* (pp. 191-204). Springer, Boston, MA.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica: Journal of the Econometric Society,* 213-224.

Hendershott, T., & Riordan, R. (2013). Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*, 48(4), 1001-1024.

Hilbert, M. (2012). How to measure "How Much Information"? Theoretical, methodological, and statistical challenges for the social sciences. *International Journal of Communication,* 6, 1042-1056.

Hirshleifer, J. (1958). On the theory of optimal investment decision. *Journal of Political Economy, 66*(4), 329-352.

Hirshleifer, J. (1970). *Investment, interest, and capital.*

Hirst, D. E., Hopkins, P. E., & Wahlen, J. M. (2004). Fair values, income measurement, and bank analysts' risk and valuation judgments. *The Accounting Review*, 79(2), 453-472.

Holland, A. S., Ott, S. H., & Riddiough, T. J. (2000). The role of uncertainty in investment: An examination of competing investment models using commercial real estate data. *Real Estate Economics*, *28*(1), 33-64.

Holmes, M. J., & Otero, J. (2019). Re-examining the movements of crude oil spot and futures prices over time. *Energy Economics*, 82, 224-236.

Hong, H., Kubik, J. D., & Solomon, A. (2000). Security analysts' career concerns and herding of earnings forecasts. *The Rand Journal of Economics*, 31, 121-144.

Huang, R., Krishnan, M., Shon, J., & Zhou, P. (2017). Who herds? Who doesn't? Estimates of analysts' herding propensity in forecasting earnings. *Contemporary Accounting Research*, 34(1), 374-399.

Ivković, Z., & Jegadeesh, N. (2004). The timing and value of forecast and recommendation revisions. *Journal of Financial Economics*, 73(3), 433-463.

Jain, P. K. (2005). Financial market design and the equity premium: Electronic versus floor trading. *The Journal of Finance*, 60(6), 2955-2985.

Jarrow, R., & Rudd, A. (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics*, *10*(3), 347-369.

Jegadeesh, N., & Kim, W. (2010). Do analysts herd? An analysis of recommendations and market reactions. *The Review of Financial Studies*, 23(2), 901-937.

Jens, C. E. (2017). Political uncertainty and investment: Causal evidence from US gubernatorial elections. *Journal of Financial Economics*, *124*(3), 563-579.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 1551-1580.

Johansen, S. (1995). Likelihood-based inference in cointegrated vector autoregressive models, 2nd ed, Oxford University Press, Oxford.

Johnston, J. C., & Pashler, H. (1998). Additional limitations in dual-task performance. *Attention.* Taylor & Francis, Hove, England, 155-189.

Jurczenko, E., Maillet, B., & Negréa, B. (2004). A note on skewness and kurtosis adjusted option pricing models under the Martingale restriction. *Quantitative Finance*, *4*(5), 479-488.

Kahneman, D. (1973). Attention and effort (Vol. 1063, pp. 218-226). Englewood Cliffs, NJ: Prentice-Hall.

Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies*, 7(1), 1-27.

Kaminski, V. (2012), *Energy Markets*. Risk Books, Incisive Media, London, UK.

Kang, W., Lee, K., & Ratti, R. A. (2014). Economic policy uncertainty and firm-level investment. *Journal of Macroeconomics*, *39*, 42-53.

Ke, B., & Yu, Y. (2006). The effect of issuing biased earnings forecasts on analysts' access to management and survival. *Journal of Accounting Research*, 44(5), 965-999.

Kellogg, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the Texas oilpatch. *The Quarterly Journal of Economics*, 126(4), 1961-2004.

Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from Texas oil drilling. *American Economic Review*, *104*(6), 1698-1734.

Kim, C. H. (2011). Price discovery in crude oil prices. Available at SSRN 1859550.

Kogan, L. (2001). An equilibrium model of irreversible investment. *Journal of Financial Economics*, *62*(2), 201-245.

Kothari, S. P., So, E., & Verdi, R. (2016). Analysts' forecasts and asset pricing: A survey. *Annual Review of Financial Economics*, 8, 197-219.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, *48*(8), 1086-1101.

Kou, S. G., & Wang, H. (2004). Option pricing under a double exponential jump diffusion model. *Management Science*, *50*(9), 1178-1192.

Kraus, A., & Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *The Journal of Finance, 31*(4), 1085-1100.

Kurov, A., & Lasser, D. J. (2004). Price dynamics in the regular and E-mini futures markets. *Journal of Financial and Quantitative Analysis*, 39(2), 365-384.

Kye, H., & Mizrach, B. (2022). Nyse floor shutdown and market quality into the close, Working paper, Rutgers University.

Lambrecht, B. M. (2017). Real options in finance. *Journal of Banking & Finance*, *81*, 166-171.

Lamont, O. A. (2002). Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior & Organization*, 48(3), 265-280.

Lang, A. (2000). The limited capacity model of mediated message processing. *Journal of Communication*, 50(1), 46-70.

Laster, D., Bennett, P., & Geoum, I. S. (1999). Rational bias in macroeconomic forecasts. *The Quarterly Journal of Economics*, 114(1), 293-318.

Leahy, J. V., & Whited, T. (1995). The effect of uncertainty on investment: Some stylized facts.

Lehmann, B. N. (2002). Some desiderata for the measurement of price discovery across markets. *Journal of Financial markets*, 5(3), 259-276.

Lensink, R., & Morrissey, O. (2000). Aid instability as a measure of uncertainty and the positive impact of aid on growth. *The Journal of Development Studies*, *36*(3), 31-49.

Levendorskiĭ, S. (2005). Perpetual American options and real options under mean-reverting processes. *Available at SSRN 714321*.

Lewis, A. L. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. *Available at SSRN 282110*.

Lien, D., & Shrestha, K. (2009). A new information share measure. *The Journal of Futures Markets*, 29(4), 377-395.

Lien, D., & Shrestha, K. (2014). Price discovery in interrelated markets. *The Journal of Futures Markets*, 34(3), 203-219.

Lin, H. W., & McNichols, M. F. (1998). Underwriting relationships, analysts' earnings forecasts and investment recommendations. *Journal of Accounting and Economics*, 25(1), 101-127.

Lin, S. X., & Tamvakis, M. N. (2001). Spillover effects in energy futures markets. *Energy Economics*, 23(1), 43-56.

Linn, S. C., & Zhu, Z. (2004). Natural gas prices and the gas storage report: Public news and volatility in energy futures markets. *The Journal of Futures Markets*, 24(3), 283-313.

List, J. A., & Haigh, M. S. (2010). Investment under uncertainty: Testing the options model with professional traders. *The Review of Economics and Statistics*, 92(4), 974-984.

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

Lucas, R. E. (1986). Adaptive behavior and economic theory. *Journal of Business*, 59, S401-S426.

Lukawski, M. Z., Anderson, B. J., Augustine, C., Capuano Jr, L. E., Beckers, K. F., Livesay, B., & Tester, J. W. (2014). Cost analysis of oil, gas, and geothermal well drilling. *Journal of Petroleum Science and Engineering*, *118*, 1-14.

Madhavan, A., & Panchapagesan, V. (2000). Price discovery in auction markets: A look inside the black box. *The Review of Financial Studies*, 13(3), 627-658.

Maghyereh, A., & Abdoh, H. (2020). Asymmetric effects of oil price uncertainty on corporate investment. *Energy Economics*, *86*, 104622.

Majd, S., & Pindyck, R. S. (1987). Time to build, option value, and investment decisions. *Journal of Financial Economics*, *18*(1), 7-27.

Marcet, A., & Sargent, T. J. (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48(2), 337-368.

Mayers, D. (1998). Why firms issue convertible bonds: the matching of financial and real investment options. *Journal of Financial Economics*, *47*(1), 83-102.

McDonald, R., & Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, *101*(4), 707-727.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, *3*(1-2), 125-144.

Miao, J., & Wang, N. (2007). Investment, consumption, and hedging under incomplete markets. *Journal of Financial Economics*, 86(3), 608-642.

Michaely, R., Rubin, A., Segal, D., & Vedrashko, A. (2023). Do differences in analyst quality matter for investors relying on consensus information?. *Management Science.*

Mincer, J., & Zarnowitz, V. (1969). The evaluation of economic forecasts, in J. Mincer, ed., *Economic Forecasts and Expectations* (New York: National Bureau of Economic Research).

Moel, A., & Tufano, P. (2002). When are real options exercised? An empirical study of mine closings. *The Review of Financial Studies*, 15(1), 35-64.

Mohn, K., & Misund, B. (2009). Investment and uncertainty in the international oil and gas industry. *Energy Economics*, *31*(2), 240-248.

Molls, B. M. (2001). *An empirical investigation of investment under uncertainty with sunk costs: Oil production in Oklahoma.* The Pennsylvania State University.

Muehlenbachs, L. (2015). A dynamic model of cleanup: Estimating sunk costs in oil and gas production. *International Economic Review*, 56(1), 155-185.

Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, *5*(2), 147-175.

O'hara, M. (1998). Market microstructure theory. Basil Blackwell, Cambridge, MA.

O'Hara, M. (2003). Presidential address: Liquidity and price discovery. *The Journal of Finance*, 58(4), 1335-1354.

Pástor, Ľ., Stambaugh, R. F., & Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, 146(2), 403-424.

Pew Research Center (2021) Internet/Broadband Fact Sheet (https://www.pewresearch.org/internet/fact-sheet/internet-broadband/)

Phan, D. H. B., Tran, V. T., & Nguyen, D. T. (2019). Crude oil price uncertainty and corporate investment: New global evidence. *Energy Economics*, *77*, 54-65.

Pierdzioch, C., & Rülke, J. C. (2013b). A note on the anti-herding instinct of interest rate forecasters. *Empirical Economics*, 45, 665-673.

Pierdzioch, C., Rülke, J. C., & Stadtmann, G. (2010). New evidence of anti-herding of oil-price forecasters. *Energy Economics*, 32(6), 1456-1459.

Pierdzioch, C., Rülke, J. C., & Stadtmann, G. (2013a). Forecasting metal prices: Do forecasters herd?. *Journal of Banking & Finance*, 37(1), 150-158.

Pindyck, R. S. (1986). Irreversible investment, capacity choice, and the value of the firm. *National Bureau of Economic Research.*

Pindyck, R. S. (1990). Irreversibility, uncertainty, and investment. *National Bureau of Economic Research.*

Pindyck, R. S. (1993). Investments of uncertain cost. *Journal of Financial Economics*, *34*(1), 53-76.

Potters, M., Cont, R., & Bouchaud, J. P. (1998). Financial markets as adaptive systems. *Europhysics Letters*, *41*(3), 239.

Prokopczuk, M., Simen, C. W., & Wichmann, R. (2021). The natural gas announcement day puzzle. *The Energy Journal*, 42(2), 91-112.

Putniņš, T. J. (2013). What do price discovery metrics really measure?. *Journal of Empirical Finance*, 23, 68-83.

Quan, J. (1992). Two-step testing procedure for price discovery role of futures prices. *The Journal of Futures Markets*, 43(1), 139-149.

Raimi, D., Krupnick, A. J., Shah, J. S., & Thompson, A. (2021). Decommissioning orphaned and abandoned oil and gas wells: New estimates and cost drivers. *Environmental Science & Technology*, 55(15), 10224-10230.

Raman, V., Robe, M. A., & Yadav, P. K. (2019). The third dimension of financialization: Electronification, intraday institutional trading, and commodity market quality, Working paper, University of Oklahoma.

Ramnath, S., Rock, S., & Shane, P. (2008). The financial analyst forecasting literature: A taxonomy with suggestions for further research. *International Journal of Forecasting*, 24(1), 34-75.

Richardson, S., Teoh, S. H., & Wysocki, P. D. (2004). The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives. *Contemporary Accounting Research*, 21(4), 885-924.

Rosen, S. (1981). The economics of superstars. *The American Economic Review*, 71(5), 845-858.

Routledge, B. R., Seppi, D. J., & Spatt, C. S. (2000). Equilibrium forward curves for commodities. *The Journal of Finance*, 55(3), 1297-1338.

Ruan, X., & Zhang, J. E. (2018). Risk-neutral moments in the crude oil market. *Energy Economics*, *72*, 583-600.

Samis, M., Davis, G. A., Laughton, D., & Poulin, R. (2005). Valuing uncertain asset cash flows when there are no options: A real options approach. *Resources Policy*, *30*(4), 285-298.

Scharfstein, D. S., & Stein, J. C. (1990). Herd behavior and investment. The *American Economic Review*, 80, 465-479.

Schneider, C., & Spalt, O. (2016). Conglomerate investment, skewness, and the CEO long-shot bias. *The Journal of Finance*, *71*(2), 635-672.

Schwarz, T. V., & Szakmary, A. C. (1994). Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. *The Journal of Futures Markets*, 14(2), 147-167.

Sepp, A. (2008). VIX option pricing in a jump-diffusion model. *Risk Magazine*, 84-89.

Shah, S., & Brorsen, B. W. (2011). Electronic vs. open outcry: side-by-side trading of KCBT wheat futures. *Journal of Agricultural and Resource Economics*, 48-62.

Shive, S. (2012). Local investors, price discovery, and market efficiency. *Journal of Financial Economics*, 104(1), 145-161.

Shrestha, K. (2014). Price discovery in energy markets. *Energy Economics*, 45, 229-233.

Silvapulle, P., & Moosa, I. A. (1999). The relationship between spot and futures prices: Evidence from the crude oil market. *Journal of Futures Markets*, 19(2), 157-193.

Sims, C. A. (1972). Money, income, and causality. *The American Economic Review*, 62(4), 540-552.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665-690.

Sinha, P., Brown, L. D., & Das, S. (1997). A re-examination of financial analysts' differential earnings forecast accuracy. *Contemporary Accounting Research*, 14(1), 1-42.

Slade, M. E. (2001). Valuing managerial flexibility: An application of real-option theory to mining investments. *Journal of Environmental Economics and Management*, 41(2), 193-233.

Spyrou, S. (2013). Herding in financial markets: A review of the literature. *Review of Behavioral Finance*, 5(2), 175-194.

Stickel, S. E. (1992). Reputation and performance among security analysts. *The Journal of Finance*, 47(5), 1811-1836.

Stock, J. H., & Watson, M. W. (1988). Testing for common trends. *Journal of the American Statistical Association*, 83(404), 1097-1107.

Stoll, H. R. (2006). Electronic trading in stock markets. *Journal of Economic Perspectives*, 20(1), 153-174.

Theissen, E. (2002). Price discovery for european cross-listed stocks. *Journal of Empirical Finance*, 9(4), 455-474.

Tobin, J., & Brainard, W. C. (1976). Asset markets and the cost of capital. *Cowles Foundation Discussion Papers*. 659.

Trueman, B. (1994). Analyst forecasts and herding behavior. *The Review of Financial studies*, 7(1), 97-124.

Tsekrekos, A. E. (2010). The effect of mean reversion on entry and exit decisions under uncertainty. *Journal of Economic Dynamics and Control*, *34*(4), 725-742.

Venkataraman, K. (2001). Automated versus floor trading: An analysis of execution costs on the Paris and New York exchanges. *The Journal of Finance*, 56(4), 1445-1485.

Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics*, 58(3), 369-396.

Woodford, M. (1990). Learning to believe in sunspots. *Econometrica*, 58, 277-307.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, *20*(1), 39-54.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*, 2nd ed. MIT press.

Working, H. (1949). The theory of price of storage. *The American Economic Review*, 39(6), 1254-1262.

Working, H. (1961). New concepts concerning futures markets and prices. *The American Economic Review*, 51(2), 431-459.

World Economic Forum (2020). How has technology changed-and changed us-in the past 20 years. In World Economic Forum, (https://www.weforum.org/agenda/2020/11/heres-how-technology-has-changed-and-changed-us-over-the-past-20-years/).

Yan, B., & Zivot, E. (2010). A structural analysis of price discovery measures. *Journal of Financial Markets*, 13(1), 1-19.

Yang, M., Blyth, W., Bradley, R., Bunn, D., Clarke, C., & Wilson, T. (2008). Evaluating the power investment options with uncertainty in climate policy. *Energy Economics*, *30*(4), 1933-1950.

Yoon, K. H., & Ratti, R. A. (2011). Energy price uncertainty, energy intensity and firm investment. *Energy Economics*, *33*(1), 67-78.

Zhang, X. (1997). Valuation of American options in a jump-diffusion model. *Numerical Methods in Finance*, 93-114.

Zitzewitz, E. (2003). Measuring herding and exaggeration by equity analysts and other opinion sellers, Working Paper, Stanford University.

Appendix A. Computation of BKM Option-Implied Risk-Neutral Central Moments

I obtain crude oil (symbol: CL) options on futures prices and futures prices from the CME Group Datasets (End of Day Complete database). Only option prices with maturities between 10 and 180 days are used, which are the most liquid options contracts. I manually dropped some outliers when the prices is 100 times greater than average, which is identified as data input error. Risk-free rates are obtained from the OptionMetrics files. Risk-free rates are extrapolated and interpolated to ensure sufficient risk-free rates for maturities of 10-180 days. I follow Chang, Christoffersen, and Jacobs (2013) and Ruan and Zhang (2018), when computing the risk-neutral central moments. I filter data by dropping option prices lower than 3/8 (minimum tick) and deep-in-the-money options (put options with an exercise price higher than 103% of futures price and call options with an exercise price lower than 97% of futures price) as well as drop days with fewer than 2 puts or calls prices, and the options prices violating the spot-futures arbitrage condition. I expand moneyness for each date and maturity to the range between 0.0001 and 3 and use the implied volatility (from CME's modified Black-Scholes option pricing model) to interpolate and extrapolate implied volatilities in that range. Next, I use the implied volatility to infer option prices using the Black-Scholes options on futures model to obtain a smooth option price in the moneyness 0.0001 and 3 for each date and maturity. With futures price and option prices in the moneynesses range [0.0001,3], I calculate option implied volatility, skewness, and kurtosis following Bakshi, Kapadia, and Madan (2003) using the trapezoidal integration following Chang, Christoffersen, and Jacobs (2013) and Ruan and Zhang (2018). I then compute the 18-month BKM risk-neutral central moments. I use futures prices to estimate the term structure of oil return realized central moments, then use the 10-180 days central moments to interpolate the one-month maturity central moments. I then use the one-month central moments and term structure to extrapolate 18-month maturity central moments. The last step is to take three-month lags of the central moments for drilling and shutdown model moments. The methods are similar to those employed by Kellogg (2014) when computing forward looking implied volatilities.

Appendix B. The Herding Measures

B.1 The Bernhardt, Campello, and Kutsoati (2006) (BCK) Measure of Herding

BCK assume that a rational forecaster seeking to maximize the accuracy of his forecast rationally incorporates into his forecast F_{jt} all information available at time (t), including previous forecasts that are available, which we summarize as the consensus forecast C_t. That is, the rational forecaster seeking to maximize accuracy weights all relevant available information, including the forecasts of others, in an optimal fashion.[76] BCK define a herder as a forecaster who overweighs the information in previous forecasts issued by others in forming his forecast, and an anti-herder as a forecaster who underweights such information. In other words, the BCK herder classification includes not just those who simply mimic the prevailing consensus, but also those who base their forecast on both their private information and the information in existing forecasts but overweight the latter relative to the most accuracy-maximizing weighting. Similarly, in addition to those who ignore the information in existing forecasts in forming their own, the BCK anti-herder classification includes those who consider both the prevailing consensus and their private information in forming their forecast but underweight the prevailing consensus. We now lay out the details of the measure and its assessment.

The BCK measure is based on the following reasoning. Let A_t be the actual storage level change based upon the EIA report, which, of course, is unknown at the time the analyst is preparing his forecast. Let E_{jt} be forecaster j's estimate of A_t based solely on his private information, that is, before seeing the forecasts of others. If j's private information forecast is unbiased, the probability of $E_{jt}>A_t$ and the probability of $E_{jt}<A_t$ should both be about .5.

[76] In BCK, rationality within the context of the development of their herding measure means optimal Bayesian updating.

Let C_t be the consensus forecast of other forecasters observable by j before publishing his forecast and let F_{jt} be the forecast released by j after observing the previous consensus forecast. If other forecasters are unbiased, $\Pr(C_t<A_t) \approx .5$ and $\Pr(C_t>A_t) \approx .5$. However, for any particular date t, the consensus C_t may deviate from the actual A_t and from the private information forecast E_{jt}, even if both E_{jt} and C_t are unbiased over many experimental trials. Suppose that the consensus, C_t, is below the private information forecast E_{jt} of forecaster j. If forecaster j issues unbiased forecasts, the probability that the ultimate forecast issued, F_{jt}, exceeds A_t should be independent of all conditioning information outside j's private information available at the time of his forecast, including C_t, and the probability should equal .5. The same holds for the probability that F_{jt} is less than A_t. This will be the case if the forecaster neither herds nor anti-herds.

On the other hand, if $E_{jt}>C_t$ and the analyst herds, the actual forecast issued, F_{jt}, will be shifted toward C_t and away from E_{jt} by more than would be called for by the optimal weighting of private and public information. That is, if herding occurs, $F_{jt}>C_t$. As a result, if herding does occur when $E_{jt}>C_t$, the probability that $F_{jt}>A_t$ should be influenced by the fact that the forecaster has shifted away from E_{jt} towards C_t. Thus, the probability that $F_{jt}>A_t$ should be less than .5.

Similarly, suppose $E_{jt}<C_t$ - that is, that the prevailing consensus forecast exceeds j's forecast based on his private information. If $E_{jt}<C_t$ and j herds, and analyst j herds, then j will adjust his forecast toward C_t, but if all the weight is not placed on C_t, then $F_{jt}<C_t$. If F_{jt} was unbiased, the probability that $F_{jt}<A_t$ should equal about .5. However, if F_{jt} is being shifted away from E_{jt} upwards towards the observed C_t, this will no longer be the case, and the probability that $F_{jt}<A_t$ will be less than .5.

Hence, the assessment using the BCK measure proceeds as follows. First, we separate the forecasts into two groups where (1) $F_{jt}>C_t$ and (2) $F_{jt}<C_t$, and we compute the proportion of times $F_{jt}>A_t$ is less than .5 in case 1 and the proportion of times $F_{jt}<A_t$ is less than .5 in case 2. If these proportions are both statistically significantly less than .5, it constitutes evidence of herding under the BCK measure. Similarly, a finding that the proportions are greater than .5 in these two cases implies that forecasters anti-herd.

B.2 The Chen and Jiang (2006) Measures of Herding

We continue with the notation used to describe the BCK measure for expository purposes but suppress the individual forecaster identifier j. Chen and Jiang (2006) posit a setting in which an analyst is privy to private as well as public information about the unknown value of a variable to be forecasted, which henceforth we define following our earlier convention as the change in storage A. A sufficient statistic for the private information is assumed to exist and is defined as $\mathrm{Y} = \mathrm{A} + \varepsilon_{\mathrm{y}}$ where $\varepsilon_{\mathrm{y}} \sim \mathrm{N}(0, \sigma_{\mathrm{y}}^2)$. Similarly, the public information is characterized as the consensus forecast $\mathrm{C} = \mathrm{A} + \varepsilon_{\mathrm{c}}$, $\varepsilon_{\mathrm{c}} \sim \mathrm{N}(0, \sigma_{\mathrm{c}}^2)$. Each information signal is characterized by a level of precision $\left(\frac{1}{\sigma_{\mathrm{i}}^2}\right)$ where i= (Y, C). The analyst's optimal Bayesian forecast of the unknown quantity A utilizes both private information as well as public information. Specifically, the conditional expectation of A is given by

$$\mathrm{E}(\mathrm{A}|\mathrm{Y},\mathrm{C}) = (h \bullet \mathrm{Y}) + (1-h) \bullet \mathrm{C} \tag{B1}$$

where *h* is the forecaster's optimal weight on the private information if the forecaster acts as a rational Bayesian.

Suppose analysts construct their forecasts of A conditional on Y and C using a weighting factor *k* on private information that is possibly different from *h*. Specifically, the analyst's actual forecast is constructed by

$$\mathrm{F} = (k \bullet \mathrm{Y}) + (1 - k) \bullet \mathrm{C} \tag{B2}$$

Combining the two expressions and incorporating the assumptions regarding Y and C, the conditional expectation of the *difference* between F and A (the forecast error) is given by

$$\mathrm{E}(\mathrm{F} - \mathrm{A} \mid \mathrm{Y}, \mathrm{C}) = \frac{k - h}{k}(\mathrm{F} - \mathrm{C}) \tag{B3}$$

Herding occurs if the analyst underweights the private information relative to what would be optimal, or conversely, overweights the public information. If the forecaster employs the optimal weight, h, then k=h and $\frac{k-h}{k}$ will equal zero. If the forecaster overweights her private information when making forecasts, then $k > h$ and the multiplier $\frac{k-h}{k}$ will be positive, an indication of anti-herding. In contrast, if the forecaster underweights her private information, $k < h$ and the multiplier $\frac{k-h}{k}$ will be negative, an indication of herding. A regression of the actual forecast error, F–A, on the forecast deviation from the consensus, F–C, will provide an estimate of $\frac{k-h}{k}$ and permit a test of herding behavior, specifically if $\frac{k-h}{k} < 0$, the analyst herds, if $\frac{k-h}{k} = 0$ the analysts neither herds nor anti-herds, and if $\frac{k-h}{k} > 0$ the analysts anti-herds. We refer to the test statistic as CJ_Dev.

The dependent variable in the estimation, the forecast error F–A, is the forecaster's estimate of the natural gas storage change in week t minus the actual storage change for week t; the regressor is the forecaster's estimate of the storage change minus the consensus forecast (F–

C). For forecasts issued on Day 0 (the day of the EIA release), the consensus is the median of the forecasts issued on Day 1, Day 2, and Day 3; for forecasts issued on Day 1, the consensus is the median of the forecasts issued on Day 2 and Day 3; and for forecasts issued on Day 2, the consensus is the median of forecasts issued on Day 3.

The intuition behind the second measure proposed by CJ runs as follows. If the analyst utilizes the efficient forecasting parameter h, her forecast is equally likely to deviate positively or negatively from the actual value of A. Because the consensus forecast is unbiased, it is also equally likely that an analyst's forecast, if constructed efficiently using h, will deviate positively or negatively from the consensus. Because of distributional symmetry, there is therefore a 50% chance that the sign of the forecast error will be the same as that of the deviation from the consensus, assuming that the forecasts are generated using the parameter h. This logic serves as the threshold for the null hypothesis: If the analyst uses the optimal factor h, the proportion of cases in which the sign of the forecast error equals the sign of the forecast–consensus deviation will equal .5. Conversely, if the actual forecast parameter k does not equal h (i.e., if $k > h$ or $k < h$), this will no longer be true. If $k > h$, the analyst overweights the private information, and the proportion of cases in which the deviation of the forecast from C matches the forecast error (the deviation of the forecast from the actual) will exceed .5, and if $k < h$, the proportion of cases with matching signs will be less than .5 (CJ), as the analyst underweights private information. We refer to this test as CJ_Prob.[77] CJ_Prob exceeds .5 if the analyst anti-herds and is less than .5 if the analyst herds.

Table 8 of the main text presents the formulas used in the calculation of each herding measure along with the interpretations of the resulting calculated measures.

[77] The BCK and CJ_Prob statistics are based upon similar assumptions and, as we will see, produce statistics that are similar in sign and magnitude. CJ show that the measures they propose are robust to alternative assumptions.

Appendix C. Other Measures of Price Discovery

Our primary measure of price discovery is the Component Share, *CS*, because it is consistent with the economic theory developed by Figuerola-Ferretti and Gonzalo (2010) and it allows for a statistical differentiation between the two hypothesis through the likelihood ratio test designed by Gonzalo and Granger (1995). However, alternative measures of price discovery have also been proposed in the literature and we estimate them here. Hasbrouck (1995) proposes a measure he labels the Information Share (*IS*). The *IS* measures the contribution of the innovation of series j to the total variance of the long-run impact of an innovation on series i, where for illustration the bivariate system contains series i and j. The *IS* measure assumes that $\beta_2 = 1$ in equation (13). However, as the theory described above suggests, this condition may not be satisfied for a storable commodity.[78] Also, the *IS* measure depends upon the ordering of the series in its computation because of the way the covariance matrix of the innovations is factored in the measure's computation (Cholesky factorization). This results in an upper and lower bound for the *IS* measure that depends upon the ordering. A common approach in the literature is to average the *IS* measures but Lien and Shrestha (2014) propose an alternative to the *IS* measure that they label the Generalized Information Share (*GIS*).[79]

The *GIS* measure does not impose the unitary cointegrating relation assumption and allows for nonzero correlations in the innovations in the bivariate version of equation (17). Lien and Shrestha (2014) demonstrate that the *GIS* is independent of the ordering of the variables. We present results for our principal measure of interest, the *CS* measure in Table 17, as it is most

[78] The developments in Yan and Zivot (2010) are also built on the assumption of a unitary cointegrating relation.
[79] A predecessor of *GIS* developed by Lien and Shrestha (2009), which the authors label *MIS*, also avoids the upper and lower bound issue of the *IS*, but the *MIS* measure presumes the unitary cointegrating relation.

consistent with the economic theory and allows for a test of significance. We also present the calculated *IS* and *GIS* measures in Table C.1.

Table C.1 results with the *GIS* measure of Lien and Shrestha (2009, 2014) corroborate those presented in Table 17 for the *CS* measure. Moreover, the numerical estimates of *GIS* during each subperiod provide further support for the conclusion we draw from the results based upon *CS* and our tests of H1 and H2.[80]

[80] For WTI, the result that both H1 and H2 are rejected at the .01 level during the first subperiod for the spot and nearby futures price pair support the alternative hypothesis that both prices contributed to price discovery before the shift to electronic trading. In contrast, H1 is rejected and H2 is not rejected during the second period for the WTI spot and nearby futures price pair indicating that price discovery migrated to the futures market during the second subperiod, following the shift to electronic futures trading. For Brent, H1 and H2

Table C.1
Other Measures for Price Discovery: Information Share and Generalized Information Share

This table provides additional measures of price discovery to those reported in Table 17. Panel A presents results for WTI and Panel B from Brent crude oil contracts. Futures Price *i* = Futures Price for Contract *i*, *i*= 1, 2, 3, 4 (maturity month). Information Shares (Hasbrouck, 1995), where *IS* is based on the ordering (spot, futures) and *IS*' is based on the ordering (futures, spot). Average *IS* values are shown. Generalized Information Shares, *GIS*, (Lein and Shrestha, 2009, 2014), are also reported. Data are from the archives of the U.S. Energy Information Administration, Bloomberg, and Datastream.

Panel A: WTI

	01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020	
	Spot Price	Futures Price 1	Spot Price	Futures Price 1	Spot Price	Futures Price 2	Spot Price	Futures Price 2	Spot Price	Futures Price 3	Spot Price	Futures Price 3	Spot Price	Futures Price 4	Spot Price	Futures Price 4
IS	0.896	0.104	0.919	0.081	0.699	0.301	0.755	0.245	0.665	0.335	0.779	0.221	0.640	0.360	0.802	0.198
IS'	0.009	0.991	0.000	1.000	0.002	0.998	0.019	0.981	0.001	0.999	0.007	0.993	0.001	0.999	0.002	0.998
Avg. *IS*	0.453	0.547	0.459	0.541	0.350	0.650	0.387	0.613	0.333	0.667	0.393	0.607	0.321	0.679	0.402	0.598
GIS	0.386	0.614	0.356	0.644	0.202	0.798	0.182	0.818	0.186	0.814	0.218	0.782	0.176	0.824	0.250	0.750

Panel B: Brent

	06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020	
	Spot Price	Futures Price 1	Spot Price	Futures Price 1	Spot Price	Futures Price 2	Spot Price	Futures Price 2	Spot Price	Futures Price 3	Spot Price	Futures Price 3	Spot Price	Futures Price 4	Spot Price	Futures Price 4
IS	0.767	0.233	0.808	0.192	0.711	0.289	0.850	0.150	0.653	0.347	0.896	0.104	0.618	0.382	0.928	0.072
IS'	0.033	0.967	0.056	0.944	0.014	0.986	0.093	0.907	0.006	0.994	0.138	0.862	0.006	0.994	0.178	0.822
Avg. *IS*	0.400	0.600	0.432	0.568	0.363	0.637	0.472	0.528	0.330	0.670	0.517	0.483	0.312	0.688	0.553	0.447
GIS	0.338	0.662	0.381	0.619	0.289	0.711	0.448	0.552	0.247	0.753	0.524	0.476	0.234	0.766	0.584	0.416

Figure C.1
Timeline of the Introduction of Electronic Trading in WTI and Brent Crude Oil Markets

Figure C.1 presents the timeline of important dates related to the introduction of electronic trading in WTI (Panel A) and Brent (Panel B) crude oil markets.

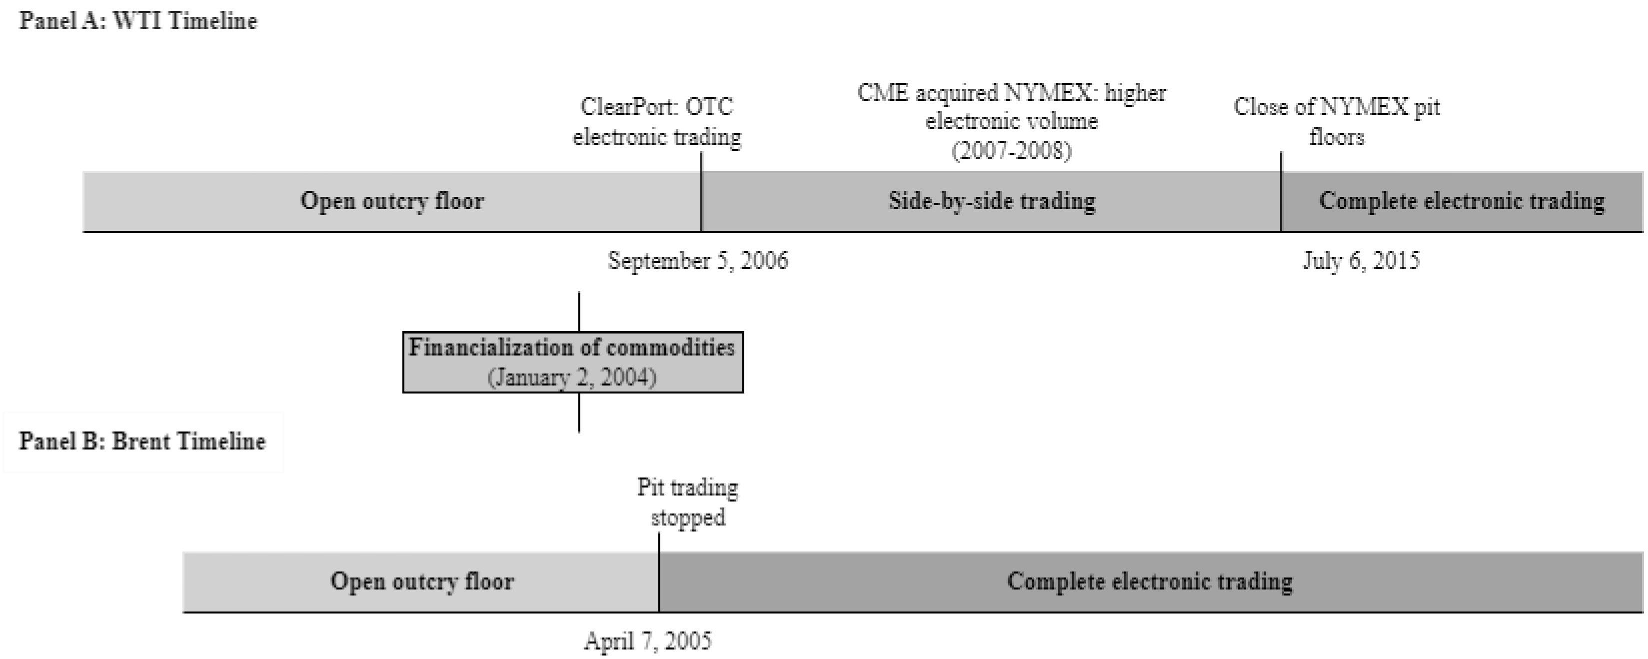

FIGURE 1

Optimal Exercise Prices of A Call Option Upon Changes in Skewness and Kurtosis in Underlying Security Price Return Distribution

Figure 1 shows the optimal exercise prices of a call option for different values of skewness and kurtosis in the underlying security price return distribution. The call option value is calculated using the modified Black-Scholes option pricing model. Immediate exercise value is calculated using the conventional Black-Scholes option pricing model. On the x-axis, skewness changes from -0.5 to +0.5 in 0.01 tick, and on the y-axis, kurtosis increases from 2.5 to 3.5 for the changes in excess kurtosis from -0.5 to +0.5 in 0.01 tick. Each fixed skewnss-kurtosis values is simulated for 10,000 times for underlying price changes from 92.0001 to 112.00. The exercise threshold is when the difference between call option value and immediate exercise value is lower than 0.01. Z-axis shows the optimal exercise price of the call option. The optimal exercise price of the call option is simulated using Stata and graphed in Matlab. The figure shows that the optimal exercise of the call option declines as skewness in underlying security prices increases. The optimal exercise price does not have an apparent change as kurtosis in underlying price increases. I assume a constant annual risk-free rate 2%, annual standard deviation in underlying security prices 20%, time interval one month, dividend yield 6%, and strike price 92.00.

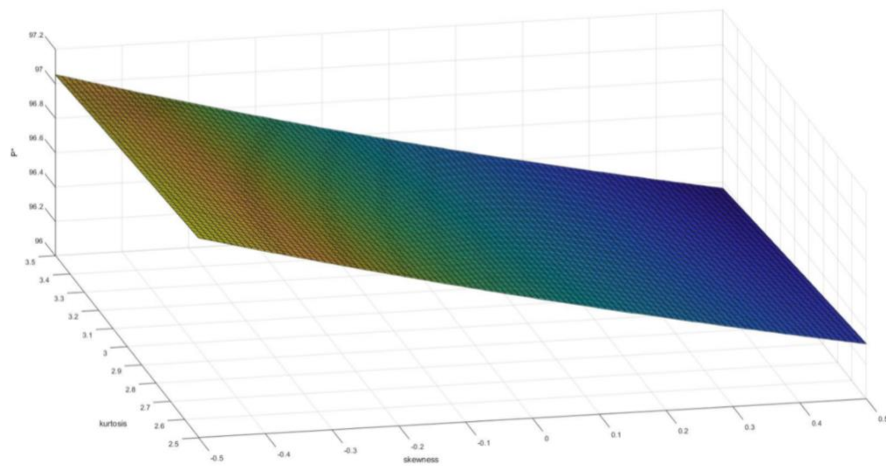

FIGURE 2

Oil Price and Central Moments

Figure 2 (1) and (2) show the histogram of three-month lagged 18-month oil futures price and graph of the price to date, respectively; (3) and (4) shows the histogram of 18-month three-month lagged option-implied volatility and graph of the volatility to date, respectively; (5), (6), (7), and (8) shows the histograms of 18-month three-month lagged option-implied skewness and kurtosis and the graphs of the central moments to dates, respectively. Oil price is scaled by 100 and deflated by CPI.

(1)

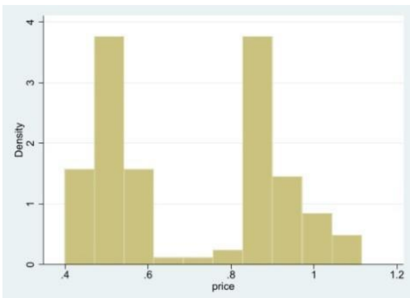

(2)

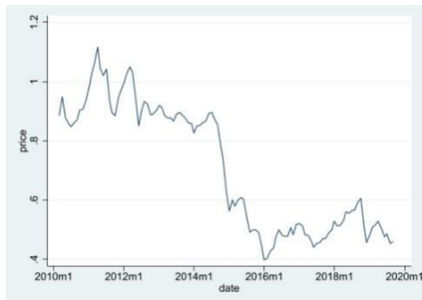

(3)

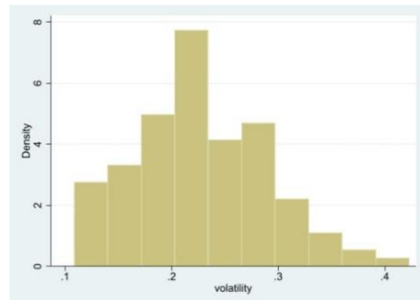

(4)

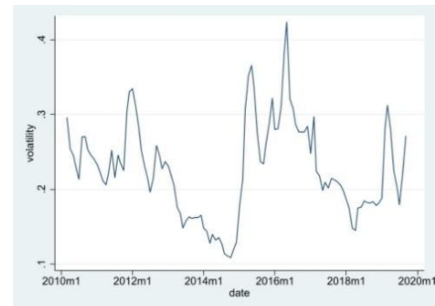

(5)

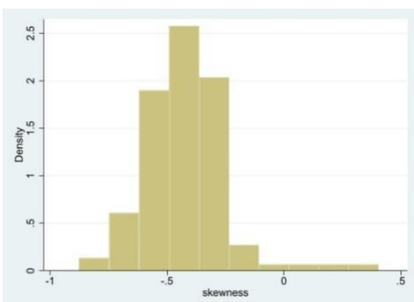

(6)

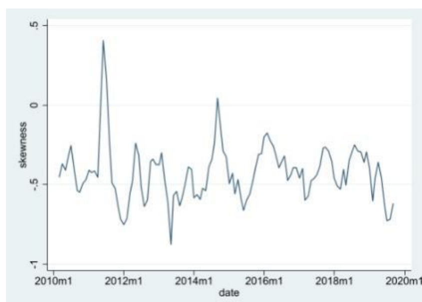

(7)

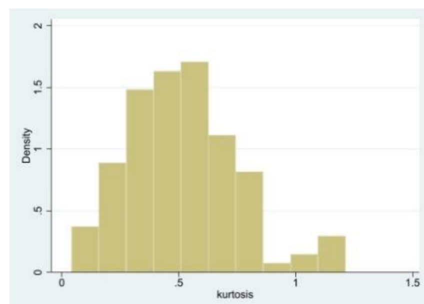

(8)

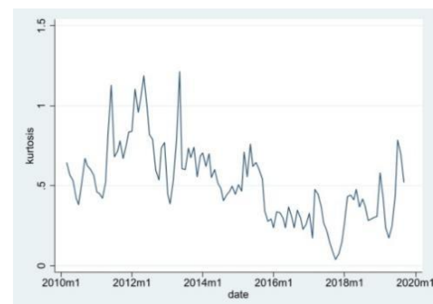

FIGURE 3

Total Daily Trading Volume WTI Futures Contracts

Figure 3 reports the total daily trading volume for the first four maturing NYMEX physically delivered traded WTI futures contracts for the period 4/2/1991 through 1/31/2020. Demarcation shows date of shift to side-by-side floor and electronic trading which occurred on 9/5/2006.The source of the data is DataStream.

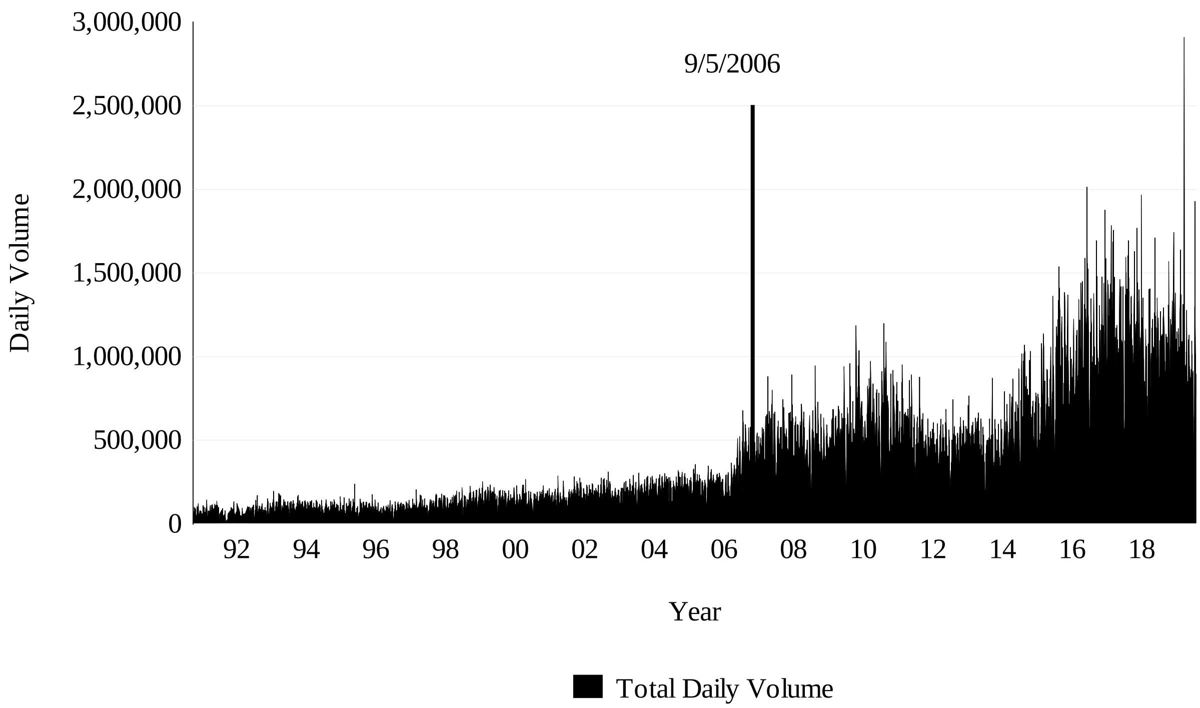

FIGURE 4

Total Daily Trading Volume Brent Futures Contracts

Figure 4 reports the total daily trading volume for the first four maturing ICE physically delivered traded Brent futures contracts for the period 4/2/1991 through 1/31/2020. Demarcation shows date of shift to side-by-side floor and electronic trading which occurred on 9/5/2006.The source of the data is DataStream.

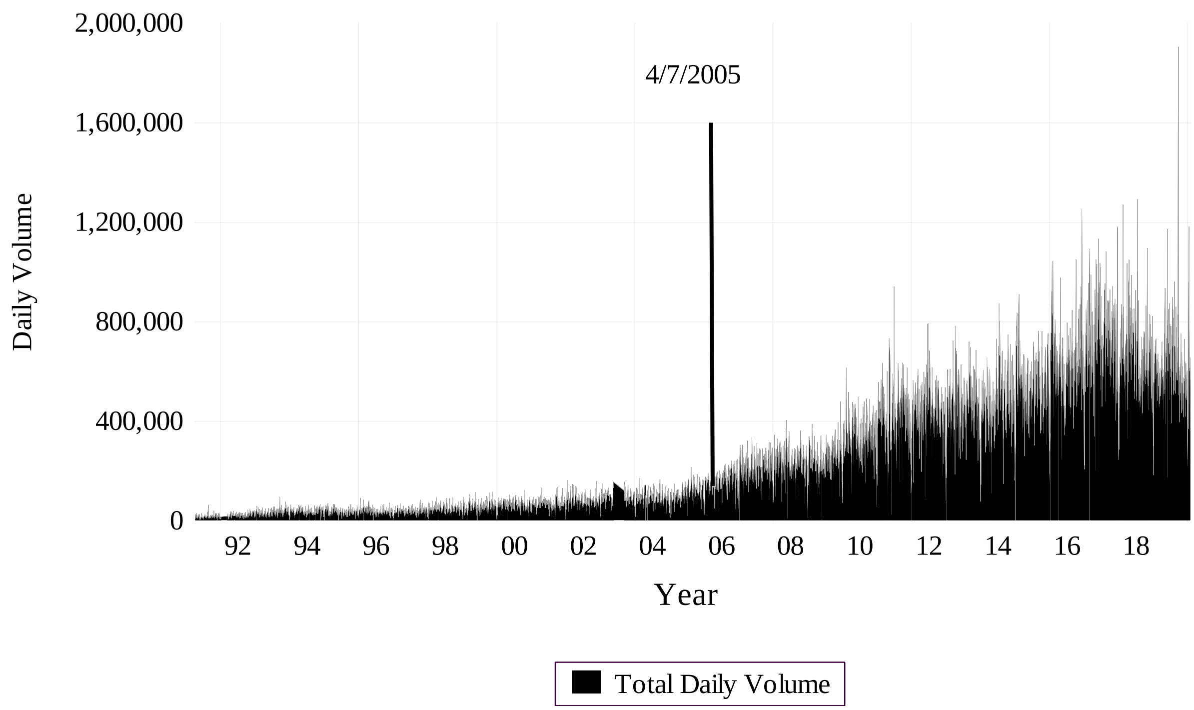

TABLE 1

Coefficient Sign Predictions of Choice Determinants for Choice Models

Table 1 shows the predicted relationships between the empirical model variables as discussed in sections 4.3.1, 4.4.1, and 4.5.1. and the probability of oil producers choosing "producing" ("drilling" or "shutdown") status at t according to the developed economic and empirical models in 4.3, 4.4, and 4.5. These relationships are indicated by the modified Black-Scholes option pricing model and Real Options Valuation as discussed in section 4.1. The first column shows the names of the variables in these empirical models (as in equations (1), (2), and (3)). Columns 3-5 show the predicted relationships between the variables and the probability of choosing "producing" as in the production model, and the probability to choose "drilling" in the drilling model, as well as the probability to choose to close the oil well in the shutdown model. In particular, the relationship between the probability and price moments for the production model depends on lagged status, so column 2 gives the lagged status of oil wells. In the second column, lagged producing status is given as signs of predicted coefficients depending on lagged status.

Model	Production		Drilling	Shutdown
Variable	Lagged Production State at t-1	Probability of Selecting "Producing" State at t	Probability of Selecting "Drilling" at t	Probability of Selecting "Shutdown" at t
Lagged Status(= "Producing")		Positive		
Price		Positive	Positive	Negative
Volatility	*Producing*	Positive	Negative	Negative
	Mothballing	Negative		
Skewness	*Producing*	Positive	Positive	Positive
	Mothballing	Negative		
Kurtosis	*Producing* *Mothballing*	Positive or Negative	Positive or Negative	Positive or Negative
Horizontal Drilling		Positive or Negative	Positive or Negative	Positive or Negative
Directional Drilling		Positive or Negative	Positive or Negative	Positive or Negative
Undetermined Drilling		Positive or Negative	Positive or Negative	Positive or Negative
Basins		Positive or Negative	Positive or Negative	Positive or Negative
Age		Negative		
Depth		Negative	Negative	Positive
Productivity		Positive	Positive	Negative
Total Costs		Positive	Negative	Negative

TABLE 2

Life Cycle of An Oil Well

Table 2 describes the six stages of the life cycle of an oil well – 1. Seismic survey; 2. Drilling; 3. Producing; 4. Mothballing, 5, Producing (Resuming), and 6. Shutdown. The second column describes the definition of each stage and the last column explains the situations in which each stage would be chosen.

Oil well life cycle stage	Define	When to choose the stage
1. Seismic Survey	Tests of oil well expected total production and operating cost	Oil producers to test oil reserve and reservoir depth
2. Drilling	Use rigs to open a new well and stabilize wellbore with cement and steel	Reservoir is detected; maintaining leasing a field
3. Producing	Pump up crude oil products and sell at market (future) prices	Expected revenue from selling is higher than operating cost; well is not exhausted or dry
4. Mothballing	Temporarily stop production of an oil well but maintain the option to re-open it later	Expected revenue is lower than operating cost but may becomes greater than cost in the future
5. Producing (Resuming)	Resume production after mothballing when prices are better	Resume production of oil well if expected revenue becomes greater than costs
6. Shutdown	Permanently close a well by plugging in wellhead with cement	Well is exhausted; too costly to maintain an active oil well

TABLE 3
Statistical Summary of Oil Price Moments and Controls

Data consists of oil wells in California, Pennsylvania, North Dakota, Oklahoma, and Texas. Dates range monthly from Jan 2010 to Jun 2019. “Prod”=1 for producing well-month observations and “Prod”=0 for mothballing. Age is the number of months in production from the completion date. Cum. oil production indicates the cumulative oil production measured in bbl. Price is the one-month maturity WTI futures prices (scaled by 100 and deflated by CPI). Volatility, skewness, and kurtosis are BKM’s (Bakshi et al. (2003)) option-implied risk-neutral central moments computed from WTI crude oil futures prices and option prices. *Reserve* is the total tested oil potential in bbl. *Depth* is in tens of thousands of feet. Drilling type dummies include directional, horizontal, vertical, and undetermined drilling types. *Last12* is the oil well’s last 12-months’ production rate and it is only for the shutdown dataset. *Exer(drilling) (exer(shutdown))* is the drilling varible in the drilling(shutdown) dataset and it is equal to one for spud (shutdown) month. The variables *age*, *cum. oil production, reserve*, *depth*, and *last12* are taken natural logarithm. The last panel summarizes the dummies for the wellheads’ located basins.

Variable	Obs	Mean	Std. dev.	Min	Max
prod	53,135,284	0.6239	0.4844	0.0000	1.0000
age	53,135,284	4.8450	1.0263	0.0000	6.6644
cum. oil production	42,251,592	8.8826	2.8138	-73.4215	16.4272
price	53,135,284	0.7147	0.2097	0.3977	1.1064
kurtosis	53,135,284	0.1636	0.3108	-0.4402	1.0916
skewness	53,135,284	-0.2318	0.1882	-0.6808	0.6169
volatility	53,135,284	0.2145	0.0603	0.1040	0.4042
reserve	6,932,426	4.2573	1.6864	-4.6052	9.7365
depth	21,521,972	8.2495	0.8085	2.9957	10.1266
last12	25,722,321	6.0047	2.1597	0.0000	13.9951
exer(drilling)	5,084,240	0.3281	0.4695	0.0000	1.0000
exer(shutdown)	27,391,704	0.0089	0.0941	0.0000	1.0000

Drilling Types	Obs	Percent	Cumulative
directional	2,127,199	4.00	4.00
horizontal	7,166,473	13.49	17.49
undetermined	170,977	0.32	17.81
vertical	43,670,635	82.19	100.00

State	Obs	Percent	Cumulative
California	5,790,727	10.90	10.90
North Dakota	1,245,491	2.34	13.24
Oklahoma	5,833,824	10.98	24.22
Pennsylvania	9,273,059	17.45	41.67
Texas	30,992,183	58.33	100.0

Basins	Obs	Percent	Cumulative
Anadarko	3,507,419	7.04	7.04
Appalachian	9,273,059	18.60	25.64
Ardmore	135,499	0.27	25.91
Ark-la-tx	2,772,800	5.56	31.47
Arkoma	1,002,489	2.01	33.48
Cherokee platform	832,007	1.67	35.15
Delaware	672,646	1.35	36.50
Fort worth	5,427,499	10.89	47.39
Gom offshore	70,442	0.14	47.53
Los angeles	444,786	0.89	48.42
Marietta	66,390	0.13	48.56
Mid-continent other	383,672	0.77	49.32
Midland	4,383,562	8.79	58.12
Palo duro	36,534	0.07	58.19
Permian other	7,546,061	15.14	73.33
Sacramento	116,535	0.23	73.56
Santa maria	115,984	0.23	73.80
Ventura	235,107	0.47	74.27
Western gulf	6,812,855	13.67	87.93
Western us other	4,769,578	9.57	97.50
Williston	1,245,491	2.50	100.00

TABLE 4

The Choice to Produce or Mothball

Table 4 Estimation results of producing status on price, central moments, and controls using the random-effects panel probit model. The regressions examine the impacts of the determinants on exercise of real options of oil well production. The model specification in Equation (1). The dependent variable, *prod*, =1 for "producing" and =0 for "mothballing". L.prod =1: lagged producing status is "producing" and =0: lagged producing status is "mothballing". The option-implied risk-neutral central moments (volatility, skewness, and kurtosis) are computed using BKM (2003). Well depth is measured in tens of thousands of feet, and well age is measured in months. "horizontal drilling" and "vertical drilling" are drilling types and "directional drilling" is the benchmark. "Ardmore", "Delaware", … are basin dummies. The variables, *age*, *depth*, *reserve*, and *cumulative oil production* are taken natural logarithms. The reported ρ is the proportion of panel-level variance in the total variance of dependent variable. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.

	(1)			(2)		
*prod=1 for producing and prod=0 for mothballing						
dependent variable: prod	coefficient	robust SE	p	coefficient	robust SE	p
l.prod	0.412	0.028	0.000 ***	0.376	0.029	0.000 ***
price	0.266	0.031	0.000 ***	0.229	0.033	0.000 ***
{l.prod=0}×vol	0.030	0.091	0.739	-0.115	0.083	0.166
{l.prod=1}×vol	0.266	0.056	0.000 ***	0.428	0.061	0.000 ***
{l.prod=0}×skewness				0.093	0.015	0.000 ***
{l.prod=1}×skewness				-0.007	0.013	0.580
{l.prod=0}×kurtosis				0.173	0.019	0.000 ***
{l.prod=1}×kurtosis				-0.192	0.014	0.000 ***
age	-0.307	0.021	0.000 ***	-0.318	0.021	0.000 ***
depth	-0.427	0.042	0.000 ***	-0.424	0.042	0.000 ***
reserve	-0.008	0.017	0.625	-0.009	0.017	0.602
cumulative oil production	0.198	0.029	0.000 ***	0.197	0.029	0.000 ***
horizontal drilling	0.828	0.067	0.000 ***	0.811	0.067	0.000 ***
vertical drilling	-0.280	0.063	0.000 ***	-0.738	0.197	0.000 ***
undetermined	-0.750	0.194	0.000 ***	-0.270	0.063	0.000 ***
Ardmore	-0.207	0.052	0.000 ***	-0.204	0.051	0.000 ***
Ark-la-TX	1.938	0.097	0.000 ***	1.929	0.097	0.000 ***
Arkoma	-0.725	0.084	0.000 ***	-0.725	0.085	0.000 ***

(continued on next page)

TABLE 4

The Choice to Produce or Mothball (continued)

	(1)			(2)		
*prod=1 for producing and prod=0 for mothballing						
dependent variable: prod	coefficient	robust SE	p	coefficient	robust SE	p
Cherokee platform	-0.479	0.039	0.000 ***	-0.478	0.039	0.000 ***
Delaware	1.679	0.052	0.000 ***	1.663	0.051	0.000 ***
Fort Worth	2.003	0.059	0.000 ***	1.989	0.059	0.000 ***
Gom offshore	2.116	0.384	0.000 ***	2.118	0.386	0.000 ***
Marietta	0.016	0.096	0.864	0.022	0.096	0.821
mid-continent other	-0.316	0.036	0.000 ***	-0.316	0.036	0.000 ***
Midland	3.025	0.043	0.000 ***	3.001	0.043	0.000 ***
Palo Duro	1.533	0.191	0.000 ***	1.520	0.191	0.000 ***
Permian other	3.159	0.059	0.000 ***	3.144	0.060	0.000 ***
western gulf	1.789	0.036	0.000 ***	1.775	0.036	0.000 ***
Williston	0.140	0.029	0.000 ***	0.161	0.029	0.000 ***
constant	3.391	0.333	0.000 ***	3.444	0.334	0.000 ***
ρ	0.650			0.648		
pseudo R-squared	0.058			0.059		
likelihood test ratio	0.000			0.000		
obs	6,452,102			6,452,102		
# of wells	99,041			99,041		
log likelihood	-766,151.54			-765,221.33		
marginal effects	*dy/dx*	*SE*	*p*			
price	0.015	0.002	0.000 ***			
{l.prod=0}×vol	-0.008	0.006	0.165			
{l.prod=1}×vol	0.037	0.005	0.000 ***			
{l.prod=0}×skewness	0.006	0.001	0.000 ***			
{l.prod=1}×skewness	-0.007	0.001	0.000 ***			
{l.prod=0}×kurtosis	0.012	0.001	0.000 ***			
{l.prod=1}×kurtosis	-0.025	0.002	0.000 ***			

TABLE 5

The Choice to Drill or Defer

Table 5 panels (3)-(4) show results of the regression of the real option exercise on price, central moments, and controls using the Cox hazard model. The regressions examine the impacts of the determinants on exercise of real options of oil well entry. The regression equation is Equation (2). The dependent variable "exer" dummy=1 for "drilling" months and =0 for the months before exercising of the options. The option-implied risk-neutral central moments (volatility, skewness, and kurtosis) are computed using BKM (2003). Well depth is measured in tens of thousands of feet. "horizontal drilling" and "directional drilling" are drilling types. "Arkoma", "east Texas", … are basin dummies. Drilling costs, *cost_def*, are sums of drilling, completion, tie-in, and transportation costs and are deflated. The variables, *depth*, *cost_def*, and *reserve* are taken natural logarithms. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.

	(3)				(4)			
*drilling=1 for spud month and =0 for prior months								
dependent variable: drilling	coefficient	robust SE	p		coefficient	robust SE	p	
price	1.267	0.076	0.000	***	1.000	0.101	0.000	***
volatility	-5.453	0.211	0.000	***	-5.587	0.216	0.000	***
skewness					0.039	0.083	0.636	
kurtosis					0.367	0.090	0.000	***
depth	-0.286	0.059	0.000	***	-0.291	0.069	0.000	***
cost_def	0.032	0.048	0.000	***	0.036	0.048	0.453	
reserve	0.141	0.011	0.002	***	0.142	0.011	0.000	***
horizontal drilling	-0.038	0.090	0.085	*	-0.039	0.090	0.665	
vertical drilling	0.153	0.083	0.000	***	0.151	0.083	0.070	*
Arkoma	-1.037	0.192	0.000	***	-1.034	0.192	0.000	***
Cherokee platform	-0.456	0.078	0.000	***	-0.454	0.078	0.000	***
Fort Worth	0.194	0.064	0.002	***	0.195	0.064	0.002	***
Hollis-Hardeman	0.230	0.133	0.085	*	0.234	0.134	0.081	*
Central basin platform	0.835	0.058	0.000	***	0.839	0.058	0.000	***
Delaware	0.752	0.052	0.000	***	0.750	0.052	0.000	***
East Texas	-0.263	0.116	0.023	**	-0.263	0.116	0.023	**

(continued on next page)

TABLE 5

The Choice to Drill or Defer (continued)

	(3)				(4)			
*drilling=1 for spud month and =0 for prior months								
dependent variable:		robust				robust		
drilling	coefficient	SE	p		coefficient	SE	p	
Eastern shelf	0.161	0.092	0.079	*	0.162	0.092	0.077	*
Gulf coast central	-0.020	0.077	0.794		-0.020	0.077	0.799	
Gulf coast west	0.558	0040	0.000	***	0.554	0.040	0.000	***
Kerr	1.055	0.250	0.000	***	1.064	0.253	0.000	***
Llano uplift	0.149	1.067	0.889		0.159	1.072	0.882	
Midland	0.738	0.040	0.000	***	0.740	0.040	0.000	***
Northwest shelf	0.149	0.218	0.494		0.152	0.218	0.487	
Palo Duro	0.194	0.211	0.358		0.194	0.211	0.358	
Val Verde	0.859	0.378	0.023	**	0.864	0.379	0.023	**
pseudo R-squared	0.016				0.016			
likelihood test ratio	0.000				0.000			
obs	1,203,762				1,203,762			
# of wells	32,817				32,817			
log likelihood	-120,328.90				-120,307.99			

TABLE 6

The Choice to Shutdown or Continue Producing

Table 6 panels (5)-(6) show results of the regression of the real option exercise on price, central moments, and controls using the Cox hazard model. The regressions examine the impacts of the determinants on exercise of real options of oil well exits. The regression equation is Equation (3). The dependent variable "exer" dummy=1 for "shutdown" months and =0 for the months before exercising of the options. The option-implied risk-neutral central moments (volatility, skewness, and kurtosis) are computed using BKM (2003). Well depth is measured in tens of thousands of feet and taken natural logarithm. "horizontal drilling" and "directional drilling" are drilling types. "Arkoma", "east Texas", … are basin dummies. Drilling costs, *cost_def*, are sums of drilling, completion, tie-in, and transportation costs and are deflated. The variables *cost_def* and *last12* are taken natural logarithms. *, **, and *** indicate significance at 10%, 5%, and 1%, respectively.

	(5)				(6)			
*shutdown=1 for shutdown month and =0 for prior months								
dependent variable:		robust				robust		
shutdown	coefficient	SE	p		coefficient	SE	p	
price	-4.782	0.098	0.000	***	-4.498	0.147	0.000	***
volatility	-8.010	0.300	0.000	***	-7.389	0.316	0.000	***
skewness					0.696	0.139	0.000	***
kurtosis					-0.083	0.109	0.444	
depth	0.235	0.034	0.000	***	0.234	0.034	0.000	***
cost_def	-0.091	0.020	0.000		-0.090	0.020	0.000	***
last12	-0.090	0.009	0.000	***	-0.090	0.009	0.000	***
horizontal drilling	-1.317	0.110	0.000	***	-1.316	0.110	0.000	***
vertical drilling	-0.229	0.080	0.004	***	-0.228	0.080	0.004	***
undetermined	0.013	0.114	0.909		0.013	0.114	0.909	
central basin platform	-0.151	0.107	0.157		-0.151	0.107	0.157	
Delaware	0.217	0.115	0.060	*	0.217	0.115	0.059	*
east Texas	0.280	0.082	0.001	***	0.280	0.082	0.001	***
east Texas costal	1.148	0.244	0.000	***	1.148	0.244	0.000	***
eastern shelf	-0.092	0.131	0.484		-0.092	0.131	0.481	
Fort Worth	0.925	0.259	0.000	***	0.923	0.158	0.000	***
gulf coast central	0.817	0.154	0.000	***	0.816	0.154	0.000	***
gulf coast west	0.671	0.127	0.000	***	0.671	0.127	0.000	***
Hollis-Hardeman	0.437	0.169	0.010	**	0.437	0.169	0.010	**
Kerr	-0.267	0.240	0.313		-0.267	0.264	0.312	
Llano uplift	-0.378	0.457	0.408		-0.380	0.457	0.406	
midland	0.005	0.100	0.957		0.005	0.100	0.961	

(continued on next page)

TABLE 6

The Choice to Shutdown or Continue Producing (continued)

	(5)				(6)			
*shutdown=1 for shutdown month and =0 for prior months								
dependent variable:		robust				robust		
shutdown	coefficient	SE	p		coefficient	SE	p	
northwest shelf	-0.419	0.119	0.000	***	-0.419	0.119	0.000	***
Palo Duro	0.093	0.248	0.709		0.092	0.248	0.710	
Val Verde	1.764	0.57	0.002	***	1.768	0.571	0.002	***
pseudo R-squared		0.064				0.064		
likelihood test ratio		0.000				0.000		
obs		10,513,982				10,513,982		
# of wells		135,034				135,034		
log likelihood		-227,430.25				-227,318.36		

Table 7
Descriptive Statistics for Forecasts and Forecast Errors for All Individual Forecasters

The table reports statistics for 19,677 individual forecasts by 102 different forecasters of the weekly EIA announcement of the change in total natural gas in storage in the lower 48 U.S. states in billions cubic feet (Bcf). The forecasts are compiled by Bloomberg and span 3/4/2003 through 12/9/2021. Forecast Error (Fer) is measured as the difference between the individual forecasted change and the actual change subsequently announced by the EIA. Absolute Forecast error is the absolute value of FEr. Relative absolute forecast error (RAFE) is the absolute forecast error standardized by monthly means. In Panel A, we report full sample statistics, in Panel B and Panel C, respectively, we summarize forecasts and errors from 2003 to 2012 and from 2013 to 2021. In the Equality tests panel, we test if the mean errors are equal between the two periods and we report p-values of the statistics. In the brackets we report p-values for equality of median tests. *** indicates rejection of the null that the means are equal between subperiods at the 1% level.

	Mean (1)	Median (2)	Standard Deviation (3)	Skewness (4)	Kurtosis (5)	Mininum (6)	Maximu m (7)
Panel A: Full sample (2003-2021): 19,677 obs							
Forecast storage change (Bcf)	1.620	40.000	95.876	-0.996	2.971	-350.000	135.000
Forecast error (FEr)	-0.368	0.000	12.149	-0.076	9.658	-145.000	97.000
Absolute forecast error (AFE)	8.542	6.000	8.646	2.904	19.421	0.000	145.000
Relative absolute forecast error (RAFE)	1.000	0.757	0.930	2.253	11.798	0.000	10.744
Panel B: First period (2003-2012): 11,200 obs							
Forecast storage change (Bcf)	3.224	38.000	91.222	-0.956	2.789	-279.000	131.000
Forecast error (FEr)	-0.483	0.000	13.119	-0.220	9.587	-145.000	97.000
Absolute forecast error (AFE)	9.235	7.000	9.330	2.857	19.483	0.000	145.000
Relative absolute forecast error (RAFE)	1.078	0.823	0.992	2.043	9.978	0.000	10.744

(continued on next page)

Table 7
Descriptive Statistics for Forecasts and Forecast Errors for All Individual Forecasters

Panel C: Second period (2013-2021): 8,477 obs							
Forecast storage change (Bcf)	-0.500	42.000	101.665	-1.014	3.035	-350.000	135.000
Forecast error (FEr)	-0.216	0.000	10.739	0.293	8.501	-82.000	81.000
Absolute forecast error (AFE)	7.631	6.000	7.558	2.806	16.068	0.000	82.000
Relative absolute forecast error (RAFE)	0.898	0.688	0.829	2.587	15.607	0.000	8.626

Equality tests (First period=Second period)	FEr	AFE	RAFE
p-value: Mean equality test (Median equality test)	0.127 (0.095*)	0.000***(0.000***)	0.000*** (0.000***)

Table 8
Herding Measures and Interpretation

BCK is due to Bernhardt, Campello, and Kutsoati (2006), CJ_Dev and CJ_Prob are due to Chen and Jiang (2006). Refer to Section 5 for background details. The variable F_{jt} represents the forecast issued by forecasting firm j of the change in the natural gas storage level at time t reported in Table 9 and used for the results presented in Table 10. F_{jt} represents the forecast issued by the individual forecaster j of the change in the natural gas storage level at time t used for the results presented in Tables 11 and 12. A_t denotes the actual storage level change as announced by the EIA, and C_t denotes the median of all forecasts prior to j's forecast. # denotes 'the number of'. Herding in the case of the BCK measure can be identified by either of the first parts of the two multiplications, hence we employ a weighted average of the two in our analyses similar to what is employed by Bernhardt, Campello, and Kutsoati (2006).

Herding Measure	Calculation Formulae	Herding Inference Indicator	Anti-Herding Inference Indicator	Neither Herding nor Anti-Herding
BCK	$= \left(\frac{N(F_{jt} > A_t \text{ and } F_{jt} > C_t)}{N(F_{jt} > C_t)}\right)\left(\frac{N(F_{jt} > C_t)}{\{N(F_{jt} > C_t) + N(F_{jt} < C_t)\}}\right) + \left(\frac{N(F_{jt} < A_t \text{ and } F_{jt} < C_t)}{N(F_{jt} < C_t)}\right)\left(\frac{N(F_{jt} < C_t)}{\{N(F_{jt} > C_t) + N(F_{jt} < C_t)\}}\right)$	< 0.5	> 0.5	$= 0.5$
CJ_Dev	$\left(\frac{k-h}{k}\right)_j$ in: $F_{jt} - A_t = \left(\frac{k-h}{k}\right)_j (F_{jt} - C_t)$ The empirical estimate of CJ_Dev is the estimate β in the following regression model $F_{jt} - A_t = \alpha + \beta(F_{jt} - C_t)$	$\left(\frac{k-h}{k}\right)_j < 0$	$\left(\frac{k-h}{k}\right)_j > 0$	$\left(\frac{k-h}{k}\right)_j = 0$
CJ_Prob	$= \frac{\#[(F_{jt} - A_t) \cdot (F_{jt} - C_t) > 0]}{\#[(F_{jt} - A_t) \cdot (F_{jt} - C_t) > 0] + \#[(F_{jt} - A_t) \cdot (F_{jt} - C_t) < 0]}$	< 0.5	> 0.5	$= 0.5$

Table 9
Forecasting Firm Herding Behavior by Subperiod

BCK and CJ herding measures for various forecasting firms are reported for two subperiods: 3/4/2003-12/31/2012 and 1/1/2013-12/9/2021. BCK Herding Measure: BCK < .5 → Herding; > .5 → Anti-herding. CJ_Dev Herding Measure: CJ_Dev < 0 → Herding; > 0 → Anti-herding; = 0 → Neither. CJ_Prob Herding Measure: CJ_Prob < .5 → Herding; > .5 → Anti-herding. Values are computed across weeks for each firm displayed. We restrict the sample to forecasting firms with at least 20 forecasts in each subperiod. The last row of the table reports the correlation coefficient between the subperiods. Refer to Section 5 of the text for details on the computation of the herding measures. ***, ** denotes significantly different from 0.5 for BCK, 0 for CJ_Dev, and 0.5 for CJ_Prob at the 1%, 5%, level, respectively.

	BCK				CJ_Dev				CJ_Prob			
Forecasting Firm	3/4/2003-12/31/2012		1/1/2013-12/9/2021		3/4/2003-12/31/2012		1/1/2013-12/9/2021		3/4/2003-12/31/2012		1/1/2013-12/9/2021	
C.H. Guernsey	0.5296		0.5359		0.7726	***	0.7093	***	0.7763	***	0.7778	***
Capital One Securities Inc	0.6863	**	0.5354		1.0980	***	0.7005	”	0.7500	***	0.6667	
Citi Futures Perspective	0.6366	***	0.5690	**	0.8234	***	0.8045	***	0.8387	***	0.7419	***
Citigroup Global Markets	0.5946		0.5173		0.2124		0.1487		0.6000		0.5289	
ConocoPhillips	0.5780	***	0.5540	**	0.4456	***	0.1162	**	0.6189	***	0.5561	**
Cormark Securities Inc	0.7419	***	0.6342	***	0.4562	***	0.7261	***	0.7576	***	0.6429	***
Credit Suisse	0.5514		0.5517		0.2701	**	0.2116		0.5955		0.5741	
Energy Overview	0.7753	***	0.7105	***	0.8262	***	0.7360	***	0.8000	***	0.7778	***
Fellon-McCord	0.7581	***	0.7067	***	0.8263	***	1.4690	***	0.7541	***	0.7027	***
FirstEnergy Capital	0.5509	*	0.5466		0.3144	***	0.2259	**	0.6220	**	0.6505	***
Gelber & Associates	0.6274	***	0.6721	***	0.6331	***	0.6190	***	0.6827	***	0.6895	***
Hencorp	0.7428	***	0.8019	***	1.0327	***	1.2347	***	0.7506	***	0.8100	***
ICAP Energy	0.6286	***	0.5385		0.5282	***	0.1228		0.6815	***	0.5517	
ION Energy	0.5000		0.4708		0.1152		-0.0238		0.5444		0.4645	
J.P. Morgan	0.5517	**	0.4781		0.1152		-0.0428		0.6091	***	0.4755	
KLR Group LLC	0.5781		0.5194		1.1378	***	0.2965	**	0.6667		0.5161	
Macquarie	0.4945		0.5065		0.2824	**	-0.0278		0.5636		0.5024	
Noble Americas Corp.	0.5211		0.5556		0.1897		0.2063		0.5082		0.6000	
Prestige Economics	0.6687	***	0.6136		0.7014	***	0.7021	***	0.7231	***	0.6250	
Ritterbusch & Assoc	0.7339	***	0.6471	***	0.8442	***	0.6895	***	0.7608	***	0.6526	***
Societe Generale	0.5538	**	0.4722		0.1968	***	0.2125		0.5804	**	0.4571	

(continued on next page)

Table 9
Forecasting Firm Herding Behavior by Subperiod

Forecasting Firm	BCK 3/4/2003-12/31/2012		BCK 1/1/2013-12/9/2021		CJ_Dev 3/4/2003-12/31/2012		CJ_Dev 1/1/2013-12/9/2021		CJ_Prob 3/4/2003-12/31/2012		CJ_Prob 1/1/2013-12/9/2021	
Summit Energy	0.6022	***	0.6211	***	0.4296	***	0.6160	***	0.6985	***	0.6639	***
SunTrust Robinson Humphrey	0.5979	***	0.7100	***	0.2399	**	0.5852	***	0.6118	***	0.7113	***
Tuohy Brothers	0.6769	***	0.6109	***	0.7732	***	0.4274	**	0.7500	***	0.6364	***
UBS	0.5874	***	0.5800		0.4107	***	0.1490		0.7123	***	0.5952	
Wells Fargo Securities	0.7133	***	0.5085		0.8357	***	0.0791		0.7409	***	0.5068	
Means	0.6223		0.5833		0.5581		0.4497		0.6806		0.6183	
Medians	0.6000		0.5548		0.4922		0.3619		0.6906		0.6307	
Test for Equality of Means (p-value)	(0.007)				(0.096)				(0.000)			
Correlation Coefficient	0.681***				0.617***				0.671***			

Table 10

Forecast Accuracy, Herding Strategy, and Forecast Timing

Full sample (3/4/2003-12/9/2021), period interaction specification. Dependent variable firm-level average of Standardized Relative Absolute Forecast Error. Raw variables are standardized by subtracting the sample mean and dividing by the sample standard deviation. Results are 'between' regressions in which average values are regressed on average values, and averages are computed across non-zero observations. In Model 1, we regress the firm average standardized value of our measure of forecast accuracy RAFE on the firm average standardized values of the three herding measures, the standardized number of days the forecast is released prior to the EIA report release, and the interaction of the two explanatory variables, D1=1 for second subperiod (after 12/31/2012), the interaction of the dummy and herding measures, the interaction of the dummy and days prior to EIA release, and the interaction of the dummy, herding measures, and days before EIA release. The sample is the 81 firms with 20 or more forecasts in either subperiod. Model 2 includes the firm average standardized value of the number of forecasts made by the forecasters, and an interaction variable with D1. The p-value for the test of the null that a coefficient equals zero is shown in parentheses and is calculated using the Huber-White standard error. *** indicates rejection of the null that the estimated coefficient equals zero at the 1% level.

	Model 1			*Model 2*		
Variable	BCK	CJ_Dev	CJ_Prob	BCK	CJ_Dev	CJ_Prob
Constant	0.047	0.003	-0.002	0.047	0.003	-0.002
	(0.771)	(0.977)	(0.983)	(0.773)	(0.971)	(0.983)
Herding measure	0.600***	0.674***	0.633***	0.600***	0.696***	0.634***
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
Days before EIA release	0.589***	0.252***	0.200	0.591***	0.239***	0.199
	(0.001)	(0.002)	(0.020)	(0.001)	(0.002)	(0.027)
Herding measure x Days before EIA release	0.080	0.078	0.039	0.079	0.098	0.040
	(0.601)	(0.355)	(0.594)	(0.611)	(0.200)	(0.571)
D1	-0.079	-0.005	0.004	-0.076	-0.006	0.002
	(0.656)	(0.967)	(0.973)	(0.671)	(0.962)	(0.989)
(Herding measure) x D1	-0.022	-0.018	0.011	-0.012	-0.044	0.021
	(0.913)	(0.895)	(0.924)	(0.951)	(0.746)	(0.861)
(Days before EIA release) x D1	0.149	0.239	0.304	0.160	0.253	0.304
	(0.423)	(0.059)	(0.019)	(0.414)	(0.044)	(0.024)

(continued on next page)

Table 10
Forecast Accuracy, Herding Strategy, and Forecast Timing

(Herding measure x Days before EIA release) x D1	-0.196	-0.011	-0.058	-0.183	-0.029	-0.036
	(0.272)	(0.925)	(0.536)	(0.307)	(0.798)	(0.702)
Number of forecasts				-0.024	0.124	0.018
				(0.820)	(0.070)	(0.832)
(Number of forecasts) x D1				-0.096	-0.161	-0.157
				(0.467)	(0.124)	(0.169)
Adjusted r-square	0.435	0.595	0.566	0.431	0.595	0.567
Observations	107	107	107	107	107	107

Table 11
Analyst-Week Forecast Accuracy, Herding Strategy, and Forecast Timing

Full Sample ((3/4/2003-12/9/2021), analyst-week panel specification. Dependent variable is analyst-week Standardized Relative Absolute Forecast Error. Herding measures are standardized by subtracting the sample mean and dividing by the sample standard deviation. In Model 1, we regress the standardized value of forecast accuracy RAFE on the standardized values of the three herding measures, the number of days prior to the EIA release, the interaction of the two variables, and the experience of the forecast firms (the number of forecasts prior to week t). The sample consists of the 56 firms with 100 or more forecasts. Model 2 includes the year fixed-effect dummies. The p-value for the test of the null that a coefficient equals zero is shown in parentheses and are calculated using Huber-White standard errors. ***indicates rejection of the null that the estimated coefficient equals zero at the 1% level.

	Model 1			*Model 2*		
Variable	BCK	CJ_Dev	CJ_Prob	BCK	CJ_Dev	CJ_Prob
Intercept	0.001	0.012	0.005	0.487***	0.370***	0.370***
	(0.985)	(0.615)	(0.818)	(0.000)	(0.000)	(0.000)
Herding measure	0.104***	0.240***	0.249***	0.127***	0.193***	0.205***
	(0.004)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Days before EIA release	0.068***	0.081***	0.072***	0.058***	0.066***	0.056***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Herding measure ×	-0.018	0.016	0.002	-0.002	0.018	0.007
Days before EIA Release	(0.435)	(0.508)	(0.949)	(0.920)	(0.301)	(0.699)
Number of forecasts	-0.089***	-0.085***	-0.088***	0.044	-0.005	-0.009
	(0.000)	(0.000)	(0.000)	(0.109)	(0.830)	(0.615)
Year FE				YES	YES	YES
Adjusted r-square	0.036	0.069	0.071	0.074	0.092	0.097
Observations	17,465	17,465	17,465	17,465	17,465	17,465

Table 12
Analyst Herding Strategy

Full Sample ((3/4/2003-12/9/2021), analyst-week panel herding specification. Dependent variable analyst-level weekly forecast error, (F(i,t)-A(t)). Individual analysts forecast errors are regressed on the deviations of their forecasts from the prior consensus forecast. Analyst forecast error is calculated as the difference between the forecast of analyst i at date t, F(i,t) and the actual change in natural gas storage in week t, A(t). The forecast deviation from the consensus is calculated as (F(i,t)-C(t)) where the consensus C(t) equals the median of existing forecasts for actual in week t. Model 1 includes firm fixed effect dummies Model 2, replaces the firm fixed effects with economist fixed effects. The dummy variable D1 equals 1 if the forecast was issued during the second subperiod and 0 otherwise. Model 3 includes D1, and the interaction of D1 and (F(i,t)-C(t)). 'Top 20' equals 1 if the forecaster is in the top 20 in terms of accuracy and 0 otherwise, while 'Bottom 20' equals 1 for the bottom 20 forecasters in terms of accuracy. Model 4 includes interaction variables that are based on interacting Top 20 and Bottom 20 with (F(i,t)-C(t)). The dummy variable DA equals 1 if the forecast was released 2 days prior to the EIA report release and 0 otherwise, and DB equals 1 if the forecast was release 1 day prior and 0 otherwise. Model 5 includes the interaction of DA and (F(i,t)-C(t)) and the interaction of DB and (F(i,t)-C(t)) and the three-way interactions of D1 and these two interaction variables. Model 6 includes all variables. Models 4 and 6 5 exclude firm fixed effects due to the multicollinearity of fixed effects with Top 20, Bottom 20. In the regressions including the firm fixed effect dummies, only forecasts from firms making more than 20 forecasts within 2 days of the EIA actual storage announcement days are used. In the regressions including the economist fixed effect dummies, only forecasts from economists who made more than 20 forecasts within 2 days of the EIA announcements are used. The p-value for the test of the null that a coefficient equals zero is shown in parentheses and are calculated using Huber-White standard errors. *** indicates rejection of the null that the estimated coefficient equals zero, at the1% level.

(continued on next page)

Table 12
Analyst Herding Strategy

Variable	*Model 1*	*Model 2*	*Model 3*	*Model 4*	*Model 5*	*Model 6*
Intercept	-0.315***	-0.301***	-0.275***	-0.370***	-0.347***	-0.392***
	(0.000)	(0.000)	(0.003)	(0.000)	(0.000)	(0.000)
Analyst Forecast minus Consensus	0.524***	0.529***	0.698***	0.547***	1.047***	0.984***
F(i,t)-C(t)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
D1			-0.054		0.120	0.108
			(0.763)		(0.424)	(0.524)
D1 × (F(i,t)-C(t))			-0.374***		-0.131	-0.025
			(0.000)		(0.213)	(0.780)
D1 × Top 20						0.176
						(0.484)
D1 × Bottom 20						-0.964**
						(0.020)
Top 20 × (F(i,t)-C(t))				-0.504***		-0.380***
				(0.000)		(0.000)
Bottom 20 × (F(i,t)-C(t))				0.453***		0.354***
				(0.000)		(0.000)
DA × (F(i,t)-C(t))					-0.901***	-0.745***
					(0.000)	(0.000)
DB × (F(i,t)-C(t))					-0.319***	-0.337***
					(0.000)	(0.000)

(continued on next page)

Table 12

Analyst Herding Strategy

Variable	*Model 1*	*Model 2*	*Model 3*	*Model 4*	*Model 5*	*Model 6*
D1 × DA × (F(i,t)-C(t))					0.285	0.136
					(0.025)	(0.213)
D1 × DB × (F(i,t)-C(t))					-0.344^{***}	-0.256^{***}
					(0.001)	(0.002)
D1 × Top 20 × (F(i,t)-C(t))						-0.012
						(0.859)
D1 × Bottom 20 × (F(i,t)-C(t))						0.081
						(0.357)
Firm FE	YES	NO	YES	NO	YES	YES
Analyst FE	NO	YES	NO	NO	NO	NO
Adj. R-squared	0.181	0.183	0.203	0.241	0.241	0.273
Observations	13,630	12,885	13,630	13,630	13,630	13,630

TABLE 13

Open Interest and Volume: WTI Futures

Table 13 reports descriptive statistics for daily open interest (Open Int.) and volume (Vol.) on four NYMEX WTI futures contracts with maturity of one month (*f1*), two months (*f2*), three months (*f3*), and four months (*f4*) over the period from 01/02/1986 to 1/31//2020. The data sources are Bloomberg and Datastream.

Panel A. 01/02/1986–9/4/2006

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Mean	90,570	55,202	80,180	40,455	42,061	12,982	28,040	5,782
Median	79,419	49,487	68,207	32,792	38,819	10,953	25,317	4,832
Std. Dev.	54,402	32,358	43,104	28,032	17,700	8,879	13,843	4,201
Observations	5,161	4,534	5,170	4,703	5,148	4,641	5,126	4,506

Panel B. 9/5/2006–1/31/2020

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Mean	267,102	363,220	262,568	174,324	143,450	64,680	107701	37,204
Median	281,029	303,357	245,358	141,037	132,647	56,497	91,423	30,315
Std. Dev.	148,366	212,095	109,134	127,371	63,015	36,904	61,769	25,886
Observations	3,364	3,378	3,394	3,379	3,394	3,379	3,394	3,379

Panel C. Differences between subperiods

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Diff. in Means Test *p*-value	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001
Diff. in Medians Test *p*-value	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001

TABLE 14

Open Interest and Volume: Brent Futures

Table 14 reports descriptive statistics for daily open interest (Open Int.) and volume (Vol.) on four Brent futures contracts with maturity of one month (*f1*), two months (*f2*), three months (*f3*), and four months (*f4*) over the period from 06/23/1988 to 1/31/2020. The data source is Bloomberg and Datastream.

Panel A. 6/23/1988–4/6/2005

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Mean	61,114	26,294	57,386	19,040	25,852	5,522	15,816	2,132
Median	63,042	24,304	56,998	14,906	22,855	3,846	14,135	1,579
Std. Dev.	22,015	11,863	21,372	14,777	12,691	5,223	7,821	2,003
Observations	3,124	3,473	3,126	3,508	3,126	3,485	3,126	3,430

Panel B. 4/7/2005–1/31/2020

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Mean	199,577	173,569	252,467	134,676	143,009	63,733	93,696	35,223
Median	164,323	166,599	221,345	118,002	123,011	55,579	74,896	27,894
Std. Dev.	139,140	96,911	132,202	84,822	88,243	41,317	73,110	29,885
Observations	3,788	3,784	3,794	3,785	3,794	3,785	3,793	3,785

Panel C. Differences between subperiods

	f1 Open Int.	*f1* Vol.	*f2* Open Int.	*f2* Vol.	*f3* Open Int.	*f3* Vol.	*f4* Open Int.	*f4* Vol.
Diff. in Means Test *p*-value	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001
Diff. in Medians Test *p*-value	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001

TABLE 15

Vector Error Correction Model: WTI Estimates

Table 15 reports estimation results for the bivariate Vector Error Correction model, equation (14), in which the endogenous variables are the WTI spot price and futures price for futures contract maturities 1 through 4.The 2 × 1 vector ΔP_t in equation (14) contains the first differences of the spot and futures prices (Δs,Δf_t). The cointegrating equation is given by $s_{t-1} = \beta_2 f_{t-1} + \beta_3 + \varepsilon_{t-1}$. The parameters α_s and α_f are the estimated speed of adjustment coefficients as in equation (16) for the spot and future price models respectively. The parameters β_2 and β_3 measure the equilibrium relation between the spot price and the futures price as shown in equations (11) and (13). Lag length for the VECM is determined using the Schwarz BIC. Standard errors are presented in parentheses. ***, ** indicates significantly different from zero at the .01, .05 level. The symbol § indicates significantly different from 1 at the .01 level. The data are for the period from 01/02/1986 to 1/31/2020. The data source is the archives of the U.S. Energy Information Administration.

Panel A: 1/2/1986–9/4/2006

	(Δs,$\Delta f1$) Model 1	(Δs,$\Delta f2$) Model 2	(Δs,$\Delta f3$) Model 3	(Δs,$\Delta f4$) Model 4
α_s	-0.481***	-0.081***	-0.046***	-0.032***
	(0.040)	(0.014)	(0.009)	(0.007)
α_f	0.142***	-0.007	-0.004	-0.002
	(0.039)	(0.012)	(0.007)	(0.005)
β_2	0.999***	0.995***	0.990***	0.985***
	(0.001)	(0.006)	(0.011)	(0.017)
β_3	0.003	0.021	0.041	0.061
	(0.002)	(0.019)	(0.035)	(0.053)

Panel B: 9/5/2006–1/31/2020

	Model 1	Model 2	Model 3	Model 4
α_s	-0.453***	-0.097***	-0.046***	-0.029**
	(0.051)	(0.023)	(0.015)	(0.012)
α_f	-0.004	-0.031	-0.010	-0.004
	(0.051)	(0.022)	(0.013)	(0.010)
β_2	1.005***§	1.029***§	1.047***§	1.060***§
	(0.001)	(0.007)	(0.014)	(0.021)
β_3	-0.021***	-0.134***	-0.215***	-0.276***
	(0.003)	(0.029)	(0.060)	(0.090)

TABLE 16

Vector Error Correction Model: Brent Estimates

Table 16 reports estimation results for the bivariate Vector Error Correction model, equation (14), in which the endogenous variables are the Brent spot price and futures price for futures contract maturities 1 through 4.The 2 × 1 vector ΔP_t in equation (14) contains the first differences of the spot and futures prices (Δs,Δf_t). The cointegrating equation is given by $s_{t-1} = \beta_2 f_{t-1} + \beta_3 + \varepsilon_{t-1}$. The parameters α_s and α_f are the estimated speed of adjustment coefficients as in equation (16) for the spot and future price models respectively. The parameters β_2 and β_3 measure the equilibrium relation between the spot price and the futures price as shown in equations (11) and (13). Lag length for the VECM is determined using the Schwarz BIC. Standard errors are presented in parentheses. ***, **, * indicates significantly different from zero at the .01, .05, 0.10 level. The symbol § indicates significantly different from 1 at the .01 level. The data are for the period from 01/02/1986 to 1/31/2020. The data source is the archives of the U.S. Energy Information Administration, Bloomberg, and Datastream.

Panel A. 6/23/1988–4/6/2005

	(Δs,$\Delta f1$) Model 1	(Δs,$\Delta f2$) Model 2	(Δs,$\Delta f3$) Model 3	(Δs,$\Delta f4$) Model 4
α_s	-0.074***	-0.041***	-0.029***	-0.024***
	(0.014)	(0.009)	(0.007)	(0.006)
α_f	0.027*	0.009	0.004	0.003
	(0.015)	(0.009)	(0.006)	(0.005)
β_2	1.025***§	1.030***	1.029***	1.035***
	(0.008)	(0.016)	(0.025)	(0.033)
β_3	-0.080***	-0.087*	-0.082	-0.094
	(0.023)	(0.048)	(0.076)	(0.100)

Panel B. 4/7/2005–1/31/2020

	Model 1	Model 2	Model 3	Model 4
α_s	-0.122***	-0.050***	-0.026**	-0.015*
	(0.022)	(0.015)	(0.011)	(0.009)
α_f	0.064***	0.040***	0.030***	0.024***
	(0.024)	(0.015)	(0.011)	(0.009)
β_2	1.028***§	1.046***§	1.061***§	1.074***§
	(0.004)	(0.008)	(0.013)	(0.019)
β_3	-0.130***	-0.210***	-0.280***	-0.340***
	(0.015)	(0.034)	(0.056)	(0.081)

TABLE 17

Component Shares and Hypothesis Test Results

Table 17 reports calculations of Component Share, *CS*, values of relative price discovery for WTI based on the estimated parameters of the bivariate vector error correction model form equation (11). *CS* is computed as $\left(\frac{\alpha_f}{\alpha_f-\alpha_s}, \frac{-\alpha_s}{\alpha_f-\alpha_s}\right) = (CS_s, CS_f)$ for the spot and futures respectively, where the α_i coefficients are from the estimated VECM and the subscripts identify the spot and futures models. Futures Price *i* = Futures Price for Contract *i*, *i*= 1, 2, 3, 4 (maturity month). Two hypotheses are tested. H1: (1,0) denotes the test to confirm that the spot market contributes 100% to price discovery and the futures market contributes 0%; H2: (0,1) denotes the test to confirm that the spot market contributes 0% to price discovery and the futures market contributes 100%. A *p*-value less than .01 indicates rejection of the hypothesis at the 1% level. Test statistics for the two hypotheses are computed following the development in Gonzalo and Granger (1995). In those cases where the estimated values for $\hat{\alpha}_f$ are not significantly different from zero at the 1% level, we set the coefficient equal to zero for the computation of CS. However, the statistical tests of H1 and H2 do not rely on this restriction. The WTI data are for the period from 01/02/1986 to 1/31/2020 and are sourced from the archives of the U.S. Energy Information Administration. The Brent data are for the period from 06/23/1988 to 1/31/2020 and are sourced from the archives of the U.S. Energy Information Administration, Bloomberg, and Datastream. Panel A presents results for WTI and Panel B for Brent crude oil contracts.

Panel A: WTI

	01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020		01/02/1986–9/4/2006		9/5/2006–1/31/2020	
	Spot Price	Futures Price 1	Spot Price	Futures Price 1	Spot Price	Futures Price 2	Spot Price	Futures Price 2	Spot Price	Futures Price 3	Spot Price	Futures Price 3	Spot Price	Futures Price 4	Spot Price	Futures Price 4
CS	0.228	0.772	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000
H1:(1,0) *p*-value	0.000		0.000		0.000		0.000		0.000		0.000		0.000		0.003	
H2:(0,1) *p*-value	0.000		0.944		0.553		0.154		0.575		0.469		0.637		0.719	

Panel B: Brent

	06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020		06/23/1988–04/06/2005		04/07/2005–1/31/2020	
	Spot Price	Futures Price 1	Spot Price	Futures Price 1	Spot Price	Futures Price 2	Spot Price	Futures Price 2	Spot Price	Futures Price 3	Spot Price	Futures Price 3	Spot Price	Futures Price 4	Spot Price	Futures Price 4
*CS**	0.269	0.731	0.344	0.656	0.189	0.811	0.439	0.561	0.127	0.873	0.536	0.464	0.124	0.876	0.611	0.389
H1:(1,0) *p*-value	0.000		0.000		0.000		0.000		0.000		0.014		0.000		0.074	
H2:(0,1) *p*-value	0.051		0.004		0.264		0.005		0.485		0.005		0.485		0.005	