

Real Options, Financial Constraints, and Drilling Rig Rental Rates

ZEIGHAM KHOKHER, MOHAMMAD MOROVATI, SHERIDAN TITMAN, and
SHYAM VENKATESAN^a

^aFirst Version November 2013; Current Version November 2016

Abstract

We present a simple real options model that illustrates how changes in uncertainty can result in changes in equilibrium investment costs rather than in changes in investment levels. To empirically test the model, we examine a panel of oil rig rental rates. Our empirical analysis confirms that after we control for several relevant economic variables, price uncertainty negatively affects rig rates.

Keywords: Real Options, Commodity Market

Real options are generally defined as assets that provide their owners with economic opportunities, but not obligations. For example, a vacant lot provides its owner the opportunity to construct a building (Titman (1985)), and mineral rights provide the owner with the opportunity to extract natural resources (Brennan and Schwartz (1985)). The key insight from this analysis is that, *ceteris paribus*, increased uncertainty increases the value of the real option and decreases the tendency that it will get exercised (see also McDonald and Siegel (1986)).

In general, it is difficult to imagine situations where shocks to uncertainty do not affect other aspects of the real option exercise decision. For instance, consider the option to drill

[☆]Khokher and Venkatesan are at the A.B. Freeman School of Business, Tulane University, 7 McAlister Drive, New Orleans, LA 70118 USA. Morovati is at Stanford University. Titman is with the McCombs School of Business, The University of Texas at Austin, 2110 Speedway, Stop B6000 Austin, TX 78712 USA. We thank workshop participants at the 2015 Financial Research Association meeting for comments. Remaining errors are ours alone.

an oil well in the Gulf of Mexico, which is the focus of this paper. Because the supply of drilling rigs is at least temporarily fixed, a shock to uncertainty about oil prices, or even a shock to the level of oil prices, will not necessarily affect the amount of drilling. As our simple model illustrates, such a shock may, instead, cause the rental rates on the oil rigs to adjust. In the language of real options, shocks to option payoffs (the oil revenues) will affect the option exercise prices (the cost of drilling), but may not affect the exercise choices (whether or not to drill). If this situation is indeed the case, a more appropriate test of the real options model will examine the cost of drilling rather than the level of drilling.

In our simple model, the supply of drilling rigs is fixed and producers can extract their oil reserves immediately, or they can delay extraction until the next period. As uncertainty increases, this timing flexibility becomes more valuable and producers are less likely to extract reserves immediately. In this case, the producers' demand curve for oil rigs shifts downward and the equilibrium rig rental rate declines.

To test the model, we analyze the rental rates of drilling rigs in the Gulf of Mexico. Our dataset, which covers virtually all offshore drilling projects from 2000 to 2014 in the Gulf of Mexico, contains over 2,000 detailed rental contracts for drilling equipment. Our database contains 256 active oil companies, a large majority of which are not publicly listed. We have data on the rental rates of 250 separate drilling rigs: we control for rig-type specific fixed effects and identify the effect of macro variables on the rental rates.

As expected, rig rental rates are low when oil prices are low. In addition, as the water depth increases, the forward looking effect (reflected in the rental rates' responses to futures prices) becomes dominant, while the effect of financial constraints decreases. Furthermore, our empirical results show that rig rental rates are low during periods of increased uncertainty. We perform a host of econometric tests to establish that rig operators respond to changing market conditions by adjusting rig rental rates, and also to show that volatility has a causal impact on these investment costs. In our tests, we control for a variety of market variables

and the technical features of the rigs.

This paper adds to the literature on exhaustible resource investments. Litzenberger and Rabinowitz (LR 1995) use Tourinho’s (1979) characterization of reserves as call options to determine oil prices. Both Tourinho and LR take extraction costs as given. In contrast, our focus is on the determinants of these investment costs in equilibrium. This paper is also related to studies that examine granular capital expenditure data. For example, Paddock et al. (1988) apply option valuation techniques to data on offshore lease auctions in the Gulf of Mexico. Kellogg (2015) reports that in, onshore Texas oil fields, the response of drilling activity to price uncertainty is quantitatively consistent with his reduced-form model. Gilje and Taillard (2016) use natural gas drilling activity to compare the investment decisions of public and private firms. However, these studies do not focus on the drivers of investment costs per se; they rely on financial statement or “intensity-of-drilling” based measures to proxy for investment. Compared to these approaches, our rig rate data gives us a more precise measure of investment costs.

Section 1 presents a simple real options model of the market for oil rigs. Section 2 reviews the institutional setting of the offshore oil and gas industry in the Gulf of Mexico and describes the rental rate data. Section 3 presents the empirical results. Section 4 concludes.

1. The Model

In this section, we present a simple model to show that in some cases, a shock to the volatility of oil prices has no effect on the number of oil rigs in service or on the quantity of oil produced. Instead, a volatility shock causes the market to react by adjusting rig rental rates.

1.1. The Economy

We consider a two period economy ($t = 0, 1$) with a reserve of crude oil and N drilling rigs.¹ The oil is located underground at depths that lie within $(0, \Delta]$. At each depth $\delta \in (0, \Delta]$, an unlimited quantity of oil is available. All of the N rigs can be used in both periods. Even if a rig is used in period 0, it becomes available to be reused by period 1. The rigs are of different types; each type of rig is appropriate to extract reserves located at a particular depth. For example, rig type N can only be used to extract oil located at depth Δ .

The economy contains two types or categories of agents: oil producers (or drillers), and rig operators. Agents in each of these categories exist in a continuum. We identify a producer in the continuum by the depth (δ) at which her reserve is located. Each producer in the continuum owns equal quantities of reserves. Each rig operator in the continuum owns equal numbers of oil rigs.

Because the supply of rigs in the economy is limited (N), producers can only extract Q of the total reserve in any period. Producers rent rigs to extract reserves; to extract a quantity q of the reserve, producers need n rigs, where $n(q) = N * q/Q$.

We are modelling oil reserves in the Gulf of Mexico and these reserves constitute a small part of the total supply, so we take oil prices as exogenous in our model. All market participants observe the spot price of oil (P_0) at the initial period ($t = 0$) and the distribution of the oil price at the terminal period ($t = 1$). In particular, the time 1 price of oil (\tilde{P}_1) is lognormally distributed with $E(\ln \tilde{P}_1) = \mu$, and $Var(\ln \tilde{P}_1) = \sigma^2$. We denote the rate of interest between period 0 and 1 as r .

The producers have different extraction costs. The differences in extraction costs occur because each producer's reserve is located at a different depth and the costs of extracting reserves increase with the depths at which the reserves are located. The time t extraction

¹Figure 1 highlights the key model assumptions.

cost (x_t^δ) for producer δ lies within $(0, X_t]$. Producers extract low cost reserves first, and so $q(\delta) = Q * \delta / \Delta$, where q is aggregate production.

The producers' extraction costs consist of the costs of renting rigs and additional costs, for example labor. A producer δ pays $x_t^\delta = (R_t + C) \delta$ to extract her reserve at time t . The labor cost parameter C is a given scalar. We determine the specific values (R_t^e) of the rig rental rate variable R_t in equilibrium.

1.2. The Producers' Problem

At period 0, producers demand oil rigs to extract their oil. The producers' demand for rigs depends on the time 0 value of the rig rate variable (R_0) . The relation between the number of oil rigs that producers demand and R_0 constitutes the demand curve for oil rigs at period 0. The optimal number of rigs for a particular level of R_0 is given by,

$$\max_{\delta_m} \left[\int_0^{\delta_m} (P_0 - x_0) d\delta + e^{-r} \int_{\Delta - \delta_m}^{\Delta} E_0^Q(P_1 - x_1)^+ d\delta \right]. \quad (1)$$

Here, δ_m is the marginal producer for the specified level of R_0 , and $E_0^Q(\cdot)$ denotes expected value computed at time 0 using risk neutral probabilities.² The optimal number of rigs that producers demand for the specified level of R_0 is then immediate as, $N_D(R_0; P_0, \sigma) = \frac{\delta_m N}{\Delta}$.

At this initial period, each producer can extract her reserve for sale in the spot market or leave the reserve underground. Producers do not extract oil for above ground storage. If the producer leaves her reserve underground, she retains the option to extract her reserve in period 1. The producer's decision involves comparing the value of the extracted oil to the value of the underground oil reserve. If the payoff from immediate extraction is greater

²Agents at time 0 observe the distribution of oil prices and rig rates at time 1. They observe the distribution of rig rates at time 1, because we solve the model recursively starting at time 1. In this case, agents can make their time 0 decisions conditional upon the distribution of oil prices and rig rates at time 1.

than the value of the underground reserve, the producer extracts the oil. In this case, we will have a marginal producer, δ_m , such that all producers with lower extraction costs will extract their reserves at $t = 0$, and those with higher costs will not. This value of δ_m enters into the calculation of the optimal number of rigs demanded in period 0, $N_D(R_0; P_0, \sigma)$. This marginal producer, δ_m , will vary with the particular choice of R_0 . The demand for rigs in period 0, $N_D(R_0; P_0, \sigma)$, will, therefore, also vary with R_0 .

1.3. The Rig Operators' Problem

At period 0, rig operators supply their rigs to producers. The producers use these rigs to extract their oil. Rig operators decide upon the number of rigs, $N_S(R_0; P_0, \sigma)$, they will supply. This number will, in general, depend upon the level of R_0 . The relation between $N_S(R_0; P_0, \sigma)$ and R_0 constitutes the supply curve for oil rigs at period 0.

At this initial period, each rig operator compares the revenue she can earn if she rents her rig immediately against her revenue if she chooses not to rent her rig. If she rents her rig, she earns δR_0 ; she earns nothing otherwise. The aggregate of each rig operator's optimal decision to rent or not will yield the optimal number of rigs supplied in period 0, $N_S(R_0; P_0, \sigma)$.

1.4. Equilibrium

At period 0, producers and rig operators bid on and supply rigs to determine the specific value R_0^e of the rig rental rate variable R_0 .³ The quantity of rigs demanded $[N_D(R_0; P_0, \sigma)]$ and supplied $[N_S(R_0; P_0, \sigma)]$ is a function of the rig rental rate variable, R_0 . In equilibrium, the quantity of rigs demanded will equal the quantity supplied, $(N_D(R; P_0, \sigma) = N_S(P; P_0, \sigma))$. This equilibrium is characterized by the value of the rental rate variable, $R_0^e(P_0, \sigma)$, that equates demand and supply. Proposition 1, which follows later, documents key characteristics of this equilibrium.

³We consider the current price of oil (P_0) is sufficiently high, so that it does not constrain this optimal choice.

At period 1, agents observe the realized price of oil (P_1). Producers and rig operators again bid on and supply rigs. This process determines the specific value R_1^e of the rig rental rate variable R_1 that equates demand and supply at this terminal period. The specific value R_1^e of the rig rental rate variable R_1 will depend on the realized price of oil, P_1 . At this time, producers will rent all rigs that are economically viable. Specifically, producers will rent rigs to extract reserves as long as extraction costs are below the realized price of oil, P_1 . Figure 2 explicitly relates producer δ 's payoff at time 1 to the price of oil at that time (P_1).

Proposition 1. *In this equilibrium, the economy fully utilizes the available supply of rigs (N) and aggregate production (q) is at capacity (Q). In particular, $N_D(R_0^e; P_0, \sigma) = N_S(R_0^e; P_0, \sigma) = N$, reserve Δ is extracted and $q(\Delta) = Q$. Furthermore,*

i. The equilibrium rig rental rate decreases in the riskiness of the spot price and increases in the spot price level. Specifically, given $\sigma' > \sigma$ and $P'_0 > P_0$, then $R_0^e(P_0, \sigma') < R_0^e(P_0, \sigma)$, and $R_0^e(P'_0, \sigma) > R_0^e(P_0, \sigma)$.

ii. The equilibrium rental rate adjusts to a shock to the riskiness of the spot price so that the economy fully utilizes the available supply of rigs, $N_D^e(R_0^e; P_0, \sigma') = N_S^e(R_0^e; P_0, \sigma') = N$, and continues to produce at capacity (Q).

Proof. See Appendix. □

The intuition for these results is as follows. First, the available supply of rigs will be fully utilized at the initial period because each rig operator will optimally choose to rent her rig. The full utilization of the rigs follows because, each rig operator can and will rent her rig in both periods. Because all rigs are used, the economy produces at capacity.

Furthermore, producers can extract their oil reserves immediately, or they can delay extraction until the following period. These reserves grant producers the option to delay their investment decisions. An increase in the riskiness of oil prices makes this timing flexibility more valuable and, thus, producers are less likely to extract reserves immediately. In

this case, the producers' demand curve for oil rigs shifts downward, and the equilibrium rig rental rate will decline. This result implies that, *ceteris paribus*, equilibrium rig rental rates should decrease with oil price volatility. Figure 3 illustrates this intuition. Analogously, an increase in the spot price of oil makes extracting producers' oil immediately attractive, thus raising the demand for rigs. Hence, equilibrium rental rates will rise. Whether the level or the volatility of the spot price changes, rig rental rates will adjust endogenously, but oil production will remain unchanged.

2. The Offshore Drilling Rig Data

2.1. Institutional Setting

Large reservoirs of hydrocarbons, such as crude oil and natural gas, are currently available beneath the surface of the earth. Extracting resources out of these reservoirs requires sophisticated techniques that depend on several factors, including the geology of the area and the type of resource. In conventional reservoirs, which are still by far the dominant source of energy in the world economy, extraction requires drilling one or several wells into the reservoirs.⁴ The wells often pass through several thousand feet of rocks to reach the reservoirs, and once the wells are completed, crude oil usually starts flowing out.⁵ Oil and gas professionals refer to a tract of sea or land that contains oil reservoirs as an oil field, these fields can extend several miles in length.⁶ Developing a field entails bringing it into production by drilling the appropriate wells.

⁴These stocks are often trapped within porous rocks under immense pressure. A majority of conventional oil and gas production projects involve drilling oil wells into these high-pressure reservoirs and the resources simply start flowing, at least during the early stages of production, due to their high pressure.

⁵Recent advances in drilling techniques has been a major enabling component of shale gas revolution. Horizontal drilling, once beyond imagination, requires directing the drilling bit thousands of feet under the surface to make a vertical turn and penetrate the hydrocarbon layers, often parallel to the earth's surface.

⁶Most of the hydrocarbon reservoirs contain both oil and gas. If a hydrocarbon reservoir mainly contains natural gas, the field is called a gas field. In the rest of this paper, we will simply refer to crude oil for much of our study, but our results are easily extended to natural gas, as well.

Offshore drilling targets the hydrocarbon resources that exist under the seabed. The Bureau of Ocean Energy Management (BOEM) is responsible for managing and administering petroleum production in the Federal regions of the Gulf of Mexico. BOEM has divided the Gulf of Mexico into block grids; each grid contains several square miles, usually in the form of rectangular tracts.⁷ BOEM offers a set of tracts for sale, and energy companies can compete in the bidding process to win the leases. Energy companies can perform seismic analyses in advance to evaluate the potential for hydrocarbon discovery in each tract. The winning bidder pays the sale price and obtains the right, but not the obligation, to start developing the petroleum field for a certain period of time, usually five or eight years. If a company starts developing a field, it keeps the right to production from that field for as long as it wishes to pay the associated fees and royalties. If a lease owner decides to abandon a tract and not to develop it, the tract will return to the stock of BOEM tracts after the lease expires and becomes available for potential future auctions. The tract lease is analogous to a call option. To own the lease, the winning bidder has to pay certain fees, similar to the price of a call option. Owning the lease gives the owner the right, but not the obligation, to start drilling to make the tract productive and create future streams of income. The owner will only start developing if she finds the development worthwhile or, in other words, finds the expected future cash flows more valuable than the development costs. The development costs are similar to the strike price of a call option, and the discounted expected future cash flows are analogous to the value of the underlying asset in a call option.

Offshore projects require large upfront investments, but yield uncertain production rates and cash flows to the investors. The major cost of developing an offshore oil field is the rental cost of drilling rigs; this may account for 60-70 percent of the total development cost. Offshore drilling rigs are huge mobile structures that firms use to drill wells in the seabed to

⁷For a precise map of the grid blocks, please download the following pdf file:
http://www.boem.gov/uploadedFiles/BOEM/Oil_and_Gas_Energy_Program/Mapping_and_Data/visual1.pdf

reach the petroleum reservoirs. Petroleum service companies (e.g. Transocean) usually own this equipment and rent it out to oil companies (e.g. BP). Often, oil companies contract rigs for short periods of time, ranging from several weeks to a few months. Included in the rig rental fees are costs of highly skilled labor provided by the service company, specialized equipment and material needed for the specific rig, and other overhead items like helicopters, which are necessary for employees to commute to and from the rigs. Once drilling is complete, the oil company installs the production facility and the petroleum service company transfers the rig to its next drilling location. Although companies can transfer drilling rigs virtually anywhere, the high opportunity cost of the time spent on long routes effectively creates a relatively competitive regional market. Superior drilling technology, environmental risks, and risks associated with severe weather are some of the other aspects that are relevant to offshore drilling.⁸

2.2. Data and Methods

Our investment cost data is from RigLogix, a leading energy industry data vendor. This unique dataset gives us precise estimates of investment costs in the Gulf of Mexico offshore drilling market. An important point to note is that we are mainly focusing on private sector petroleum investments, not on government investments. All of the investments we consider in this paper, in effect, contribute to capacity building within a competitive fringe that is a price taker in crude oil markets.

Analysts often use the capital expenditures of publicly listed oil and gas companies as a measure of investment costs. Government regulations require public firms to make their balance sheets available. However, using balance sheet data alone imposes serious limitations on the empirical analysis. First, the frequency of publicly available data on capital expen-

⁸The DeepWater Horizon disaster in April 2010 provides an example of the immense potential risks involved.

ditures is very low (annual, for most cases). In addition, the number of public exploration and production firms (E&P) is small. This small number limits the available data and, as a result, decreases the power of empirical tests. Second, the published data on capital expenditures is aggregated and does not relate to a specific investment opportunity. The data aggregates assorted types of investments. These investments may include building refineries, installing pipelines, expanding current production plants, or developing new fields. Because of this aggregation, identifying the factors that will impact new investment opportunities is difficult. Third, many active private companies do not publish their balance sheets.

Our detailed dataset allows us to evaluate investments at the smallest unit of investment in the oil industry: an oil well. The rental cost of drilling equipment constitutes the majority of the cost of developing an oil field. The cost of drilling alone accounts for 60 to 70 percent of the cost involved in developing an offshore oil field.⁹ The dataset contains details of all of the offshore drilling contracts between service companies and oil companies from 2000 to 2014. The contracts include the names of the service companies that owned the rigs, the names of the oil companies that rented the equipment, the contracts' start and end dates, and the rig rental rates. The dataset also includes the fixture date (contract signing date) of each contract, cost and age of the equipment, and the rated water depth. The technical specifications and rig type allow us to follow each group of rigs with great precision.¹⁰

The dataset includes almost 2033 contracts for drilling in the Gulf of Mexico region. Our database contains 256 active oil companies, a majority of which are not publicly listed companies. The publicly listed companies operate mainly in deep water, primarily, because of the great technological sophistication and investment costs necessary to operate in deep waters. Both of these factors are prohibitive for small companies. Furthermore, large public

⁹<https://www.eia.gov/analysis/studies/drilling/pdf/upstream.pdf>

¹⁰A detailed description of the data fields present in our data is available at the data vendor's website: http://www.riglogix.com/RigLogix_Data.aspx

firms can access the capital markets more easily than small concerns, and they face fewer financial frictions. To evaluate the effect of financial constraints on the firms in our sample, we divide our sample into three different water depths. We run our tests on each bin separately to identify the effects of financial frictions.

The drilling data consists of more than 250 drilling rigs that have been repeatedly contracted. Hence, they are unavailable at times and become available again after their rental periods expire. The unit of observation is at the individual contract level. Accordingly, each rig appears several times in our data. We match each contract to the *OVX*, *VIX*, *spot*, *futures*, *treasury*, and *AAApremium* levels that existed on the first date of the contract signing month. We use the CBOE Crude Oil ETF Volatility Index (*OVX*) as a proxy for the volatility in the oil prices. *OVX* is a market estimate of expected 30-day volatility of crude oil futures prices and is calculated using the CBOE Volatility Index (*VIX*) methodology applied to options on the United States Oil Fund, LP (USO). Unfortunately, the *OVX* series goes back only until June 2007.

To increase the statistical power of our tests, we also use *VIX* as a proxy for price volatility. *VIX* is a key measure of the market's expectation of near-term volatility as reflected by the S&P 500 stock index option prices and, importantly, it is highly correlated with the *OVX*.¹¹ Significantly, unlike the *OVX* series, the *VIX* series spans the entire duration of our sample.

Spot and *futures* are the spot price of Brent Crude and the slope of the 12-month futures price of Brent Crude, respectively. Although we have the exact contract signing dates, businesses often make investment decisions using information from several months prior to and following considerable negotiations with the rig operating companies. To allow for this difference in timing, we take the past six-month moving average of volatility, spot

¹¹The correlation between *OVX* and *VIX* is greater than 0.83. More information about *OVX*, *VIX*, and their computational methodology is available at <http://www.cboe.com/>

price, and futures price. We use the yield on the treasury bill (*treasury*) and the excess yield of three month-to-maturity investment grade bonds over treasury bills of the same maturity (*AAApremium*) to capture different aspects of the financing cost. Bank of America/Merrill Lynch provide the US Corporate AAA Index (*AAAYield*). This index is the yield on dollar denominated investment grade corporate debt publicly issued in the US domestic market.

Table 1 provides a summary of the overall matched data. In Panel A, we provide the distributional details about the contract-related variables, and also of the other relevant financial variables. The median age of the rig is 27 years, and the median days for which it is rented is 39. Panel B presents summary information by the different rig types. Jackup and Semisub are the two most popular rig types that drillers use in the Gulf of Mexico region. Because the complexity of drilling increases with the water depth, the rigs are distinguished primarily by the maximum level of water depth at which they can operate. Commensurate with the challenges of deep water drilling, companies need superior and expensive technology to successfully drill at these depths. Therefore, rigs that can drill in the deep waters are expensive to construct and, hence, also have high daily rental rates. For example, the average rental rate for a Submersible rig, which operates in relatively shallow water, is \$61,646 per day. Compare this amount to \$406,343 for renting a Drillship, a rig that operates at depths of close to 10,000 feet. Also note the high correlation in operable water depths, construction cost of the rig, and the average daily rental rates.

We recognize that there exist great differences in the rental rates across the different types of rigs, and we deal with these differences in our multivariate regression specifications. In Panel C, we highlight the variation in the daily rental rates, one of the main points of our paper. The standard deviation in the rental rates, across each of the rig types, is at least 35 percent of the average rental rate. Standard real options models do not account for such variation in investment costs. In the following section, we identify the different factors that explain the variation in the daily rig rental rates.

3. Empirical Analysis

In this section, we empirically investigate several factors that affect the investment decisions of oil and gas companies, and we document the results of our tests.

3.1. Multiple Regression

To investigate the determinants of rig rental rates, we first estimate the following regression specification:

$$\begin{aligned} \log(\text{rental rate}_{i,j,t}) = & \beta_0 + \beta_1 \text{Volatility}_t + \beta_2 \text{Cost of capital}_t \\ & + \Pi_{i,j,t}\phi + Z_t\lambda + X_i\delta + \Gamma_j\theta + e_{i,j,t}. \end{aligned} \quad (2)$$

For every contract, i , and rig type, j , the rental rates are denominated in dollars per day. OVX and VIX are our primary measures of volatility in oil prices. Yield on the Treasury bill is our proxy for the cost of capital (*treasury*).¹² Furthermore, we include the spot price of Brent Crude oil (*spot*) and the slope of the futures contract price for Brent Crude to be delivered in 12 months (*futures*) as control variables.¹³

Equation (2) also uses some rig-specific characteristics as control variables. We include age of the rig (*Age*), and construction cost of the rig (*Construct*) as independent variables to explain the variation in daily rental rates. We expect the rental rates to decrease with the age of the rig and to increase with the construction cost of the rig. In addition, we also add contract-specific variable(s). A potentially important predictor of the daily rental rate should be the water depth at the location of the well (*LocDepth*). Based on our earlier

¹²We define the variable *treasury* as $\log(1 + \text{treasury bill yield})$.

¹³We compute the slope of the futures contracts as $\log(\text{futures}_{12m}/\text{futures}_{1m})$. Here, *futures*_{12m} is the futures price of the Brent Crude to be delivered in 12 months and, *futures*_{1m} is the futures prices of Brent Crude to be delivered in one month.

discussion, we expect the rental rates to increase with the water depth and hence, predict the coefficient of *LocDepth* to be positive.

To capture the differences across rig types, all regression specifications have a rig-type specific fixed effect. We cluster the standard errors by every calendar month, rig manager, and by rig owner to correct for any correlation in error terms. Further, we take the natural log of all of the explanatory variables. Taking the natural log helps us to interpret the marginal effects of the coefficients as a percent change in the dependent variable to a percent change in the explanatory variable.¹⁴ We also deal with any possible time trends in dollar denominated variables like *rentalrate* and *spot* by removing the effect of general inflation. We use the monthly consumer price index (CPI) to adjust prices for inflation.

Table 2 presents the results from a pooled OLS regression. Consistent with standard economic intuition, age of the rig is negatively related to rental rates and construction cost of the rig is positively related to rental rates. Further, we also find that the rental costs increase with the water depth at the drilling location. We observe this effect because the technical sophistication and risks of drilling increase as the water depth increases. Note, this effect captures the variation in rental rates after controlling for the rig-type.

The cost of capital and price volatility are the two key variables in Table 2. High cost of capital, as reflected by the effective yield on a treasury bill, is positively correlated with rig rental rates. As the cost of capital in the aggregate economy increases, the rig owners' costs of running the business will go up. In a perfectly competitive economy, we would expect the rig owners to pass these additional costs on to their eventual customers, the lessees. Our results are consistent with this expectation. Importantly, in column (I) of Table 2, which uses the whole sample, we also observe a negative relationship between future price volatility and rig rental rates.

¹⁴We use lowercase to denote the natural log of the variables.

The financial risks associated with drilling, and the nature of companies that engage in drilling, vary substantially with water depths and well locations. To further illustrate the different reasons for the variations in rig rental rates, we split the data into sub-samples, based on water depths. In columns (III), (IV), and (V) of Table 2, we report our results for rigs that operate in depths lower than 100 feet, between 100 and 1000 feet, and depths greater than 1000 feet, respectively. The results for low and medium water depths are not very different from our main results. However, for deep water, the results are in sharp contrast.

Firms that operate in deep water are large public firms that face little or no financial frictions. As a result, such firms are mostly indifferent to marginal changes in price volatilities. Further, rigs that operate in deep water, Semisubs and Drillships, are significantly more expensive equipment to rent than shallow water rigs. Therefore, the lessor firms that lease deep water rigs are also, on average, large public firms.¹⁵ These large lessor firms do not exhibit much sensitivity to changes in cost of capital because the daily rig rental rates do not change in response to changes in treasury bill rates. However, given the large amount of initial investment necessary to operate in deep water areas, future expectations about oil prices matter a great deal for rental rates. Among the three depth categories, the slope of futures prices is positive and most significant for the rigs that operate in deep water. *post1* is an indicator variable for contracts signed between January 1, 2004 and April 20, 2010, and *post2* is an indicator variable for contracts signed after April 20, 2010. Having these two variables in the regression controls for time trends in rig rental rates. The negative coefficients for both of these variables indicate that rig rental rates declined in these periods compared to the base group, which consists of contracts signed prior to 2004.

To further substantiate our results, we recreate the results presented in Table 2 by using

¹⁵Well over 63% of the rig owners in the deep water category were publicly listed firms. Compare this percent with only 30% for leased rigs in the below 100 feet category.

ovx, an alternate measure of the market’s expected volatility. By using *ovx*, we lose nearly half of our sample and therefore, also lose considerable statistical power. Nevertheless, despite using *ovx*, we find that the results remain qualitatively similar to those we presented earlier.¹⁶

We also perform a falsification test to provide robustness to our results. To each rental contract, we randomly assign a volatility number. We pick this volatility number from the empirical distribution of *ovx* and *vix*. After randomizing the associated volatility, we should not expect to find any relationship between such volatility and rig rental rates. We run the same specification as in equation (2). Consistent with our expectations, we do not find any relationship between rig rental rates and these volatilities.

3.2. *Endogenous rental rates*

One of the central goals of this article is to establish unequivocally that rig rental rates are an endogenous variable. To further illustrate this point, we use the exogenous event of the British Petroleum (BP Deepwater Horizon) oil spill that occurred on April 20, 2010 to examine how rig rental rates changed around this event.

On April 20, 2010, Deepwater Horizon, a Semisub rig, exploded while drilling a well in the Gulf of Mexico region. Although the company finally capped the well, by the time it completed the process, the explosion had resulted in the largest offshore oil spill in U.S. history. This catastrophic accident led to substantial loss of human life and threatened the ecological balance in all of the southern states that surround the Gulf of Mexico. In response to the incident, the U.S. Department of the Interior issued a moratorium on new deepwater oil drilling permits in 500 feet of water or more. Soon after, Hornbeck Offshore Services and several other companies that were engaged in offshore drilling challenged this order. A

¹⁶Because no OVX data was available prior to 2007, our regression specification cannot have both *post1* and *post2*. Therefore, we included only *post2* in the specification. The coefficient on *post2* now shows the marginal change in the rental rates after April 20, 2010, compared to earlier dates.

Federal judge, Judge Martin Feldman, overturned the moratorium, and an appellate court subsequently upheld his ruling. The Interior Secretary then issued a second ban in June, 2010 that was scheduled to expire in November, 2010.¹⁷

The Deepwater Horizon explosion and the subsequent litigation activity created a great deal of uncertainty about future drilling activities and their associated cash flows. We predict that as a response to these events, companies would lower rig rental rates for deep water rigs to induce manufacturers to drill. Because the moratorium did not apply to depths lower than 500 feet, we should not expect to see any effect on the rental rates for such rigs. However, in an effort to use the available capacity, we might find that some firms move to drilling in shallower water (less than 500 feet) and therefore, demand shallow water rigs. In these cases, this substitution could raise the rental rates of the shallow water rigs.

We use the following regression specification to test our prediction:

$$\log(\text{rental rate}_{i,j,t}) = \beta_0 + \text{post} + \beta_2 \text{deepwater} + \beta_3 \text{post} * \text{deepwater} + X\Gamma + e_{i,j,t}, \quad (3)$$

where *post* is an indicator variable that takes the value one if the contract signing date is in the post-event period, and zero otherwise. *deepwater* is also an indicator variable that takes the value one if the water depth at the well location is greater than 500 feet, and zero otherwise. *Post * deepwater* is the interaction of the two indicator variables mentioned above. The above specification includes all of the covariates we used earlier in Table 2. The difference-in-difference estimator of the “treatment” effect, (β_3) , is more efficient in the presence of these exogenous controls.¹⁸ The coefficient of interest here is β_3 because it

¹⁷Additional details about Hornbeck Offshore Services LLC v. Salazar are available at https://en.wikipedia.org/wiki/Hornbeck_Offshore_Services_LLC_v._Salazar

¹⁸See Roberts and Whited (2013) for a discussion of difference-in-difference estimators and their efficiency with additional exogenous controls.

captures the marginal effect of the event on the rig rental rates after we control for any inherent differences between the regions (shallow or deep), and for any systematic changes in the economic conditions over time. A negative β_3 would be consistent with our previous results.

Table 3 presents the values for the difference-in-difference estimator from the estimating equation (3). For column (I), April 20, 2010 to April 19, 2011 represents the post-event period. The coefficient on *deepwater* is positive, confirming that rigs that work in deeper water have higher rental rates than shallower water rigs. Most importantly, the coefficient on the interaction term, as predicted, is negative and highly statistically significant. This finding suggests that the rig rental rates endogenously adjusted to accommodate for the economic environment that resulted from the Deepwater oil spill.

Amidst all of the legal controversy, on October 13, 2010, the Interior Secretary finally lifted the moratorium on deep water oil drilling. Strict new rules, including one-on-one worker training, accompanied this relaxation. On account of the new compliance requirements and because of increased scrutiny in the permit review process, companies encountered substantial delays in obtaining permits. The delays were such that even two months after the Interior Secretary lifted the ban, no new permits were awarded.¹⁹ Overall, in the aftermath of the oil spill, a great deal of uncertainty surrounded deep water drilling and experts expected this uncertainty to continue well into early 2012. In column (II) of Table 3, we redefine our post-period to be two years, ranging from April 20, 2010 to April 19, 2012. We define the pre-period, in both columns (I) and (II), as the two years from April 20, 2008 to April 19, 2010. The results of column (II) are consistent with our earlier results from the one-year horizon.

In column III of Table 3, we provide the results from a falsification/placebo test. Instead

¹⁹See the following report: <http://www.wsj.com/articles/SB10001424052970204204004576050451696859780>

of the actual date of the BP oil spill, we assign April 20, 2009 as the exogenous event date. The result in column III does not show the same negative coefficient for *interact* as in columns I and II. Further, we also provide bootstrap results we obtain from applying the difference-in-difference estimator on randomly picked dates. We perform 500 iterations and in each iteration, we pick a random date between January 1, 2004 and April 19, 2009. We avoid the period before January 2004 because the number of contracts was pretty sparse. Then, for each selected date, we define the post-event period as the one year period after such a date. We run the same specification as in Table 3 and collect the coefficient on *interact*. Because we used these randomly picked dates, our prediction is that we should not find any significance on the *interact* variable. The mean of the point estimates over the 500 iterations is 0.220. The 5th and the 95th percentiles of the distribution are -0.009 and 0.509, respectively, and the inverse quantile function of zero is 0.072. Collectively, these figures imply that close to 93 percent of the estimates were positive, which validates that pre-existing differences in the rig rental rates did not drive our results. Overall, these results reinforce our findings regarding the endogeneity of rig rental rates.

3.3. Predicting rig utilization rates

The real options literature argues that in the presence of irreversibility and uncertainty, firms should consider the value of continuing to hold options and, perhaps, delay investments. In this context, we should naturally test the response of drilling activity to changes in price volatility. However, an important assumption underlies the expected relation between these two variables. In particular, this relation assumes investment costs are constant and inelastic. Above, we have demonstrated that considerable variation exists in rig rental rates and that they dynamically adjust to the macro-economic environment. Therefore, our thesis is that in such an environment, drilling activity and price volatility might not be negatively related. We now design an empirical scheme to test our prediction.

In addition to having contract level details, our dataset also provides month-level information on the utilization rates of the different rig types. In other words, it gives us information about the percent of available rigs for GOM exploration that producers are actually using. The utilization rates tell us a great deal about the drilling activity in each month. Figure 5 plots the time series of the utilization rates for the two most-used rig types in the GOM.²⁰ The average utilization rates are extremely high, in excess of 80 percent, and show limited variation.²¹ Table 4 provides results from an OLS regression, where the (log) rig utilization rate is the dependent variable. Utilization rates are very sticky; therefore, we also include a lagged utilization rate in the specification. Additionally, we include the average rental rate of the specific rig type in the previous month as one of the explanatory variables. We lag other regressors by three months to account for the delay between the decision to drill and the actual commencement of drilling. Columns (I) and (II) of Table 4 use different measures of price volatility. However, in neither of these cases can we reject the null hypothesis that the coefficient on the volatility term is equal to zero.

These results are in sharp contrast to the findings of Kellogg (2014). The dataset in Kellogg (2014) is for onshore drilling activity and, importantly, it lacks information on the amount of investment needed for the projects he studies. In other words, it does not have data on the cost of drilling. In fact, the inherent differences between the onshore and offshore drilling industries could also be an important reason for the differences in the results. First,

²⁰We have made no adjustments for seasonality or any kind of smoothening to the raw utilization rates.

²¹We chose to keep our model simple and focus on full utilization at date 0. Because most of the literature stresses the waiting to drill effect, we wanted to provide a simple model where all of the volatility effect goes through the rig rate channel and drilling activity is unaffected. Of course, in reality, increases in volatility are likely to result in a lower utilization rate as well as lower rig rates in some markets. To match the utilization data, we will need rigs that have different qualities – the low quality rigs require high labor costs and are thus taken out of service when oil prices are low. Including this feature in the model would add further complications that would make it less tractable. In reality, rigs may not be utilized for a short period because they are being transported from one location to another. Another plausible consideration is that 80% utilization may be close to full utilization, in the same sense as 5% unemployment is considered full employment.

because the offshore rig rental costs are extremely high, both in absolute dollar terms and also in terms of their proportions of total costs, negotiating these rates is a binding constraint on companies' investment decisions. Second, in the short run, the supply of offshore rigs is fixed. Commissioning a new rig happens rarely; besides, building new rigs is a long process. Additionally, offshore rigs are massive, making moving them from one place to another very expensive. Therefore, their geographical mobility is restricted to a very small area. Given these constraints on the supply of rigs, companies sign contracts to rent them several months before the actual start day of drilling. Typically, these contracts also have severe penalties for cancellation. As a result, the rate of drilling cannot easily adjust to changes in volatility. Third, the barriers to entering the onshore rig rental market are substantially lower than those of the offshore market, resulting in an excess supply of onshore rigs. Because of this situation, the rental market for onshore rigs is extremely competitive, leaving very low margins for any further price drops. Finally, the opportunity costs of keeping expensive offshore rigs idle are prohibitive and encourage the rig owners to enter the bargaining game. Overall, the above results are consistent with our view that supply side dynamics matter a great deal to firms' drilling decisions.

3.4. Causal Effect of Volatility

To provide additional support to our primary hypothesis, we ask whether volatility has a causal effect on the rental rates of rigs. Ideally, we would like to compare the contracts that companies make in times of high volatility to those they make during other times. However, a potential selection bias confounds our efforts to identify the causal effect by simply comparing sample means. Furthermore, the fundamental characteristics of the contracts that companies make during times of “high” volatility and “normal” volatility might be different, making the assignment to “treatment” and “control” groups nonrandom.

We overcome this problem by comparing the rental rates of contracts that companies

make during a “high” volatility regime (the treatment group) with those of matched samples of contracts made in a “normal” volatility regime (the control group). For the purposes of these tests, we label the contracts as “treated” if companies entered into them while volatility was in the highest quintile (above 80%) of the distribution. All other contracts are the control group.²² We use the contract characteristics and other exogenous macro variables to match the treated and control groups. We use the nearest-neighbor, as well as the optimal match algorithm.²³

We match the contracts on age of the rig (*age*), construction cost of the rig (*construct*), depth at the location of the well (*locdepth*), six month moving average of Brent Crude price (*brentavg*), spot price of Brent (*spot*), yield of the treasury bill (*treasury*), and rig type. We include an indicator variable for contracts signed between January 1, 2004 and April 20, 2010 (*post1*), and also an indicator variable for contracts signed after April 20, 2010 (*post2*). We can reasonably assume that conditional on these observable characteristics, assignments to treatment and control groups are random (unconfoundedness). Further, in the matching process, we enforce an exact match regarding the rig type and for the period in which the company signs the contract (*post1* and *post2*). Figure 4 is a love plot that provides the extent of balance between the two groups. Readers can see that after we complete the matching process, the two groups are very similar in the observed dimensions. Following the matching process, we proceed to compute the overall *average treatment effect* (*ATE*). This estimate points to the causal effect of price volatility on rig rental rates. We compute *ATE* as $\frac{1}{N} \sum_i (Y_i(1) - \hat{Y}_i(0)) + \frac{1}{M} \sum_j (\hat{Y}_j(1) - Y_j(0))$, where N and M are the number of treated and

²²A simple t-test for the differences in mean of the two sample shows that the mean for the treated group is 0.164 lower than for the control group. Because we are using log rental rates, this finding implies that the rental rates in a high volatility regime are, on an average, 16.4% lower than in regular times.

²³The nearest-neighbor matching method is a “greedy” matching algorithm, this algorithm chooses the closest control match for each treated unit individually, without trying to minimize a global distance measure. In contrast, “optimal” matching finds the matched samples with the smallest average absolute distance across all of the matched pairs.

controls in the matched sample, and $\hat{Y}_j(0)$ and $\hat{Y}_j(1)$ are the imputed potential outcomes for each observation j under the counterfactual condition.

Table 5 reports the *ATE* for the different matching methods that we used. The first row of Table 5 refers to the results we obtained by using the nearest-neighbor matching method. We use the inverse of the variance matrix as the weighting scheme for the covariates.²⁴ Companies made 374 contracts during the times of high volatility. In our sample, this number refers to contracts companies made in the months where the VIX was 23.6 percent or higher. Before undergoing any matching procedure, the control sample had 1431 contracts. In the second specification, we impose stricter matching requirements than those imposed previously. We stipulate that all matches are equal to or within 0.25 standard deviations of each covariate. This restriction obviously reduces the matched number of contracts. The first two rows report the *ATE* after adjusting for bias (see Abadie and Imbens (2012)) and report the correct Abadie-Imbens standard error. Both of the specifications lead to a negative coefficient that is also highly statistically significant. The point estimates for *ATE* are -0.047 and -0.042, respectively. This negative sign on the coefficient is consistent with our prediction regarding volatility and confirms that it has a negative causal effect on the rig rental rates.

4. Final remarks

The exhaustibility of natural resources, like crude oil, brings the inter-temporal decision to extract the resource to the core of a firm's investment decision. McDonald and Siegel (1986) show that because the developer of an unexploited oil field has discretion over the timing of her opportunity to invest, this waiting option creates value that is central to the cost-benefit decision. In addition, the capital expenditure for developing an offshore oil field

²⁴In unreported results, we also use Mahalanobis distance as the covariate weighting scheme. The results are not qualitatively different than those reported here.

is almost entirely irreversible because the scrap value of these projects is close to zero. This irreversibility feature and the timing option inherent in exhaustible resource investments has made the real options approach standard in the literature (for example Dixit and Pindyck, 1994) and ubiquitous in practice.

Our evidence suggests an important change in the way that oil and gas firms should approach their investment decisions. The rental costs of offshore drilling equipment constitute a major share of the costs of developing new wells. Consequently, rental rates may affect the decisions of oil and gas companies to either act on particular investment opportunities or to wait for better market conditions. We find significant variation in rig rental rates. An important conclusion to draw from this evidence is that standard real options techniques, which do not consider variations in investment costs, are likely to produce sub-optimal investment decisions. To address this issue, researchers should build and test real option models that can accommodate variations in investment costs.

Table 1: Summary of the Data

This table provides the summary of our data. Panel A provides details about the contract parameters and about the other economic variables we used in this article. Geographical region represents the different parts of the world where the drilling activity occurred. Contract length is the amount of time for which the rig is leased. Location water depth is the actual depth of water where the offshore drilling activity is going to take place. *Age* is the time, in years, since the rig was constructed. *AAAyield* is the yield on AAA rated US corporate bonds provided by Bank of America Merrill Lynch. Panel B reports the different types of rigs that are leased out, the number of contracts for each of these rig types, and a mean for different variables relating to the rigs. Daily rental rate is the rent that is charged by the rig owner for every day that the rig is rented out. Max WD is the technically specified maximum water depth in which each of the rigs can operate. The dollar denominated construction cost of the rig is provided in the final column. Panel C provides the summary statistics of the daily rental rate distribution for each of the different rig types.

Panel A: Summary Statistics

	Mean	Median	Standard Deviation
Number of contracts	2033		
Contract length (days)	50.75	39.00	45.76
Location water depth (feet)	892.70	114.00	1770.77
<i>Age</i> (years)	24.35	27.08	10.58
<i>OVX</i> (%)	37.78	34.82	13.34
<i>VIX</i> (%)	18.87	16.27	8.44
<i>AAAyield</i> (%)	4.24	4.57	1.24

Panel B: Types of Rigs

Rig Type	# of contracts	Daily rental rate (\$)	Max WD (feet)	Construction cost (\$ million)
Submersible	66	61,646.21	77.42	29.61
Jackup	1491	80,184.18	260.66	36.38
Semisub	368	284,738.17	6330.30	172.38
Drillship	108	406,342.60	10,675.92	486.71

Panel C: Variations in Rental Rates

Rig Type	Mean (\$)	Std Dev (\$)	Min Rate (\$)	Max Rate (\$)
Submersible	61,646	21,377	26,500	115,000
Jackup	80,184	35,943	6,500	222,000
Semisub	284,738	157,958	40,000	607,000
Drillship	406,342	155,967	103,000	681,000

Table 2: Explaining the Variation in the Daily Rental Rates.

This table shows the results from the pooled OLS regression. The dependent variable, in all of the columns, is the natural log of the daily rental rate. The variable *volatility* is the log of the six month moving average of CBOE Volatility Index computed from the S&P 500 stock index option prices (VIX); *spot* is the log of the six-month moving average of the spot price of Brent Crude; *futures* is the log of the six-month moving average of the slope of the futures contract on Brent Crude expiring in 12 months; *treasury* is $\log(1+\text{yield})$, where the yield is that of the treasury bill; *age* is the log of the time, in years, since the rig was constructed; *construct* is the log of the dollar denominated construction cost of the rig; *aaapremium* is the log of the difference between yield on AAA rated US corporate bonds and yield on a treasury bill; *locdepth* is the log of water depth at the location of the well; *post1* is an indicator variable for contracts signed between January 1, 2004 and April 20, 2010; *post2* is an indicator variable for contracts signed after April 20, 2010. All of the specifications include rig-type fixed effect. The standard errors are clustered by rig managers, by rig owners, and by each calendar month. The significance levels are denoted by *, **, and *** and indicate whether the results are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively.

	(I)	(II)	(III)	(IV)
	All Rigs	Depth < 100 ft	100 ft < Depth < 1000 ft	Depth > 1000 ft
<i>volatility</i>	-0.475*** (0.084)	-0.613*** (0.134)	-0.46*** (0.109)	-0.153* (0.085)
<i>spot</i>	1.13*** (0.114)	0.894*** (0.129)	1.134*** (0.082)	1.55*** (0.115)
<i>futures</i>	0.945* (0.545)	0.501 (0.832)	0.99** (0.501)	1.548*** (0.361)
<i>treasury</i>	0.141** (0.064)	0.136* (0.079)	0.164*** (0.061)	-0.054 (0.092)
<i>age</i>	-0.078** (0.037)	-0.187*** (0.036)	-0.055 (0.042)	0.024 (0.022)
<i>construct</i>	0.125** (0.059)	0.058 (0.052)	0.158*** (0.052)	0.173*** (0.050)
<i>aaa premium</i>	-0.014 (0.054)	0.007 (0.084)	-0.011 (0.052)	-0.044 (0.108)
<i>locdepth</i>	0.096*** (0.018)	0.031 (0.025)	0.296*** (0.077)	0.117 (0.080)
<i>post1</i>	-0.269* (0.141)	-0.069 (0.193)		-0.692*** (0.081)
<i>post2</i>	-0.365** (0.173)	-0.14 (0.267)	-0.058 (0.049)	-1.115*** (0.199)
<i>Observations</i>	1819	844	632	343
<i>Adj R²</i>	0.78	0.49	0.70	0.70

Table 3: Endogenous Rental Rates.

This table provides the results from a difference-in-difference analysis. The dependent variable is the natural log of the daily rental rate. The variable *volatility* is the log of the six-month moving average of the CBOE Volatility Index computed from the S&P 500 stock index option prices (VIX); *spot* is the log of the six-month moving average of the spot price of Brent Crude; *futures* is the log of the six-month moving average of the slope of the futures contract on Brent Crude expiring in 12 months; *treasury* is $\log(1+\text{yield})$, where the yield is that of the treasury bill; *age* is the log of the time, in years, since the rig was constructed; *construct* is the log of the dollar denominated construction cost of the rig; *aaapremium* is the log of the difference between yield on AAA rated US corporate bonds and yield on a Treasury Bill. *post* is a dummy variable that takes the value one when the contract signing date is in the post-event period, and zero if the contract signing date is in the pre-event period. For column (I), the post-event period is one year from April 20, 2010 and for column (II) post-event period is two years from April 20, 2010. The pre-event period is the two years prior to April 20, 2010. For the falsification test in column (III), the post-event period is one year after April 20, 2009. *deepwater* is a dummy variable that takes the value one when the water depth at the well is greater than 500 feet, and takes the value zero otherwise. *post * deepwater* is the product of two dummy variables, *post* and *deepwater*. All of the specifications include rig-type specific fixed effects. The standard errors are clustered in two dimensions, by each calendar month, and by rig owner. The significance levels are denoted by *, **, and ***, and indicate whether the results are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively.

	(I) 1 yr	(II) 2 yr	(III) Falsification
<i>post * deepwater</i>	-0.407*** (0.103)	-0.239*** (0.078)	0.389*** (0.144)
<i>post</i>	0.050 (0.080)	-0.007 (0.077)	-0.384*** (0.130)
<i>deepwater</i>	0.536*** (0.133)	0.476*** (0.098)	0.144* (0.075)
<i>volatility</i>	0.607*** (0.177)	0.726*** (0.096)	-0.024 (0.167)
<i>spot</i>	1.063*** (0.251)	1.140*** (0.224)	0.727*** (0.263)
<i>futures</i>	0.043 (1.345)	-0.342 (1.026)	1.969*** (0.607)
<i>treasury</i>	0.328** (0.142)	0.247*** (0.091)	0.230* (0.127)
<i>age</i>	-0.070** (0.029)	-0.074*** (0.021)	-0.134*** (0.026)
<i>construct</i>	0.135*** (0.025)	0.131*** (0.021)	0.111*** (0.028)
<i>aaa premium</i>	0.207 (0.162)	0.106 (0.150)	0.099 (0.097)
Observations	504	672	609
<i>Adj R</i> ²	0.78	0.82	0.83

Table 4: Rate of Drilling

This table shows the results from the pooled OLS regression where the log monthly utilization rates is the dependent variable. In column I, *volatility* is *ovx*, which is the log of the six-month moving average of the CBOE Crude Oil ETF Volatility Index. In column II, *volatility* is *vix*, which is the log of the six-month moving average of CBOE Volatility Index computed from the S&P 500 stock index option prices. The variable *laggedUtil* is the log utilization rate of the previous month; *spot* is the log of the six-month moving average of the spot price of Brent Crude; *futures* is the log of the six-month moving average of the slope of the futures contract on Brent Crude expiring in 12 months; *treasury* is $\log(1+\text{yield})$, where the yield is that of the treasury bill; *aaapremium* is the log of the difference between yield on AAA rated US corporate bonds and yield on a treasury bill; *avgrate* is the log of the average rig-specific rental rate in the previous month; *post1* is an indicator variable for the period between January 1, 2004 and April 30, 2010; *post2* is an indicator variable for months after May 1, 2010. *ovx*, *vix*, *spot*, *futures*, *treasury*, and *aaapremium* are all lagged by three months. All of the specifications include rig-type fixed effects. The standard errors are clustered by rig type and by each calendar quarter. The significance levels are denoted by *, **, and ***, and indicate whether the results are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively.

	(I)	(II)
<i>volatility</i>	0.060 (0.048)	0.001 (0.029)
<i>laggedUtil</i>	0.559*** (0.165)	0.639*** (0.085)
<i>spot</i>	-0.047 (0.080)	0.049* (0.024)
<i>futures</i>	-1.027*** (0.183)	-0.321* (0.176)
<i>treasury</i>	-0.034 (0.070)	-0.039 (0.047)
<i>aaa premium</i>	0.021 (0.021)	-0.028 (0.023)
<i>avgrate</i>	-0.046 (0.032)	-0.015 (0.013)
<i>post1</i>		-0.018 (0.024)
<i>post2</i>	-0.049 (0.072)	-0.078 (0.064)
<i>Observations</i>	298	190
<i>Adj R²</i>	0.63	0.56

Table 5: Causal Effect of Volatility on Rental Rates

This table provides estimates of the mean difference between the rental rates on contracts made during "high" volatility times and rental rates on contracts made at all other times. Contracts made during "high" volatility (treated group) are the contracts that were made in the months when the volatility was in the top most quintile (above 80%) of the volatility distribution. We match the two groups, treatment and control, in many dimensions. The observable covariates used are *age*, which is the age of the rig in years; *construct*, which is the dollar value of constructing the rig; *locDepth*, which is the water depth at the site of the drill. *aaapremium* is the log of the difference between yield on AAA rated US corporate bonds and yield on a treasury bill; *spot* is the log of the spot price of Brent Crude; *futures* is the log slope of the futures contract on Brent Crude expiring in 12 months; *treasury* is $\log(1+\text{yield})$, where the yield is that of the treasury bill; *post1* is an indicator variable for contracts signed between January 1, 2004 and April 20, 2010; *post2* is an indicator variable for contracts signed after April 20, 2010; and *rigtype* is the type of rig. We use two different matching algorithms, Nearest Neighbor and Optimal match. Results are provided for both matching schemes. We report the estimate of the overall average treatment effect (*ATE*) along with the estimates of Abadie-Imbens standard error in paranthesis below. Row 2 reports the *ATE* after the treatment and control group covariates have been constrained to be within 0.25 standard deviations of one another. The significance levels are denoted by *, **, and ***, and indicate whether the results are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively.

Estimator	Average Treatment Effect (ATE)	Number of Treated	Number of Controls
Nearest Neighbor (Inverse Variance)	-0.047*** (0.01)	374	1431
Nearest Neighbor (Stricter)	-0.042*** (0.01)	128	1431
Optimal Match	-0.107** (0.052)	359	1364

Figure 1: Model Assumptions and Timeline

This figure summarizes model assumptions and the agents choices.

PANEL A: KEY ASSUMPTIONS

1. We consider a two period economy ($t = 0$ and 1) with an underground oil reserve and N drilling rigs.
 - The cost of extracting the reserve is x_t^δ and $x_t^\delta = (R_t + C) \delta$;
 - $\delta \in (0, \Delta)$ is the underground depth of the reserve. R_t is a rig rate variable; it is endogenous. C is a labor cost parameter; it is a given scalar.
2. We take oil prices as exogenous since the oil in the Gulf of Mexico is a small part of the total oil supply.
 - The spot price of oil is denoted P_0 .
 - The price of oil in period 1 (P_1) is lognormally distributed with $E(\ln P_1) = \mu$ and $Var(\ln P_1) = \sigma^2$.
3. Producers rent rigs to extract reserves. Because the supply of rigs is limited, only Q of the reserve can be extracted per period.
 - To increase aggregate production (q), deeper reserves must be extracted: $q(\delta) = Q * \delta / \Delta$.
 - The number of rigs required increases with production: $n(q) = N * q / Q$.

PANEL B: MODEL TIMELINE

- **Period 0:** Producers and rig operators bid on and supply rigs.
 - The quantity of rigs demanded and supplied is a function of the rig rate variable R_0 .
 - In equilibrium, the quantity of rigs demanded will equal the quantity supplied $N_D(R_0; P_0, \sigma) = N_S(R_0; P_0, \sigma)$.
 - The specific value, R_0^e , of R_0 that equates the demand and supply of rigs characterizes this equilibrium.
- **Period 1:** The realized price of oil (P_1) is observed.
 - Producers and rig operators again bid on and supply rigs. This process determines the value R_1^e of the rig rate variable R_1 .
 - Rigs with current rental rates below P_1 will be rented out and others will not.

Figure 2: The Oil Producers' Payoffs with Endogenous Rig Rates

This figure illustrates the producers' payoffs at time 1. The payoffs at time 1 for producer δ is defined as $Payoff_1^\delta = P_1 - (R_1 + C) \delta$. Here, R_1 is the time 1 value of the rig rental rate variable. We use the equilibrium condition equating the quantity of rigs demanded and supplied to determine R_1 ; we can then express R_1 as a function of the oil price at time 1, P_1 .




Figure 3: Equilibrium in the Market for Oil Rigs and Oil Price Volatility

This figure illustrates equilibrium at $t=0$ in the market for oil rigs. The figure also shows the effect of changing the exogenous oil price volatility variable σ on the endogenous rig rate variable R_0 .




Figure 4: Covariate Balance

The graph below plots the covariate balance between the control group and the treated group. The treated group contains the contracts that companies made in the months when the volatility was in the top most quintile of the volatility distribution (above 80%). The black points are the pre-matching differences between the covariates. The red points are the differences in the covariates in the matched sample.

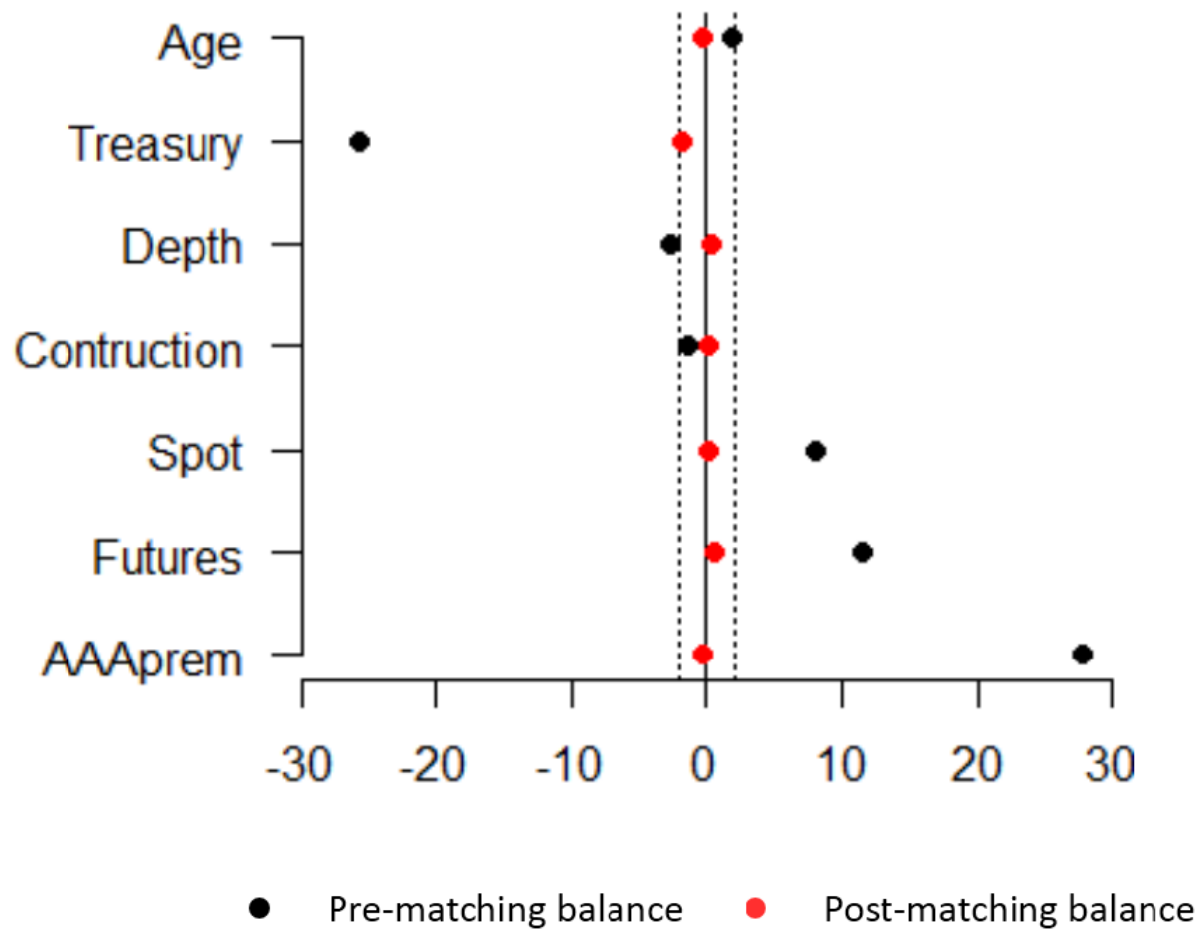




Figure 5: Utilization Rates

The graph below plots the utilization rates of two main rig types in our data. The utilization rate represents the number of rigs in use during a given month as a percentage of total available rigs.



	Mean (%)	Std (%)
Jackup	82	10.6
Semisub	84.04	6.61

Appendix A. Proofs

At $t = 1$, agents observe the realized price of oil (P_1). At this time, each one of the producers $\delta \in (0, \Delta]$ can spend $x_1^\delta = (R_1 + C)\delta$ and extract her reserve. Or, she can leave her reserve underground. A particular producer $\delta \in (0, \Delta]$ will rent a rig to extract her reserve if her reserve is economically viable ($P_1 - x_1^\delta > 0$). She will not rent a rig to extract her reserve if the current price of oil is below her extraction cost, ($P_1 - x_1^\delta < 0$). There is then a marginal producer with extraction cost $x_1^{\delta_m}$ who is indifferent between extracting the oil and leaving it underground. For this marginal producer, $P_1 - x_1^{\delta_m} = 0$. Hence, this marginal producer can be uniquely identified as, $\delta_m = \frac{P_1}{R_1 + C}$.

The marginal cost producer at $t = 1$ depends on the value of the rig rental rate variable R_1 . If the rig rental cost variable (R_1) is sufficiently high, $\delta_m < \Delta$ and not all producers will extract their reserve. With a decrease in R_1 it will be optimal for more producers to extract their reserve, this raises the value of δ_m .

Given the choice of the marginal producer, δ_m , we derive the optimal quantity of rigs demanded at $t = 1$, $N_D^1(R_1; P_1, \sigma)$. In particular, $N_D^1(R_1; P_1, \sigma) = \frac{\delta_m N}{\Delta} = \frac{N}{\Delta} \frac{P_1}{R_1 + C}$. If $\delta_m < \Delta$, the number of rigs used is $N\delta_m/\Delta$, and this number of rigs is less than the number of rigs available (N). In this case, the economy-wide production ($\delta_m Q/\Delta$) is below capacity (Q). A decrease in R_1 raises the value of δ_m and hence this also raises $N_D^1(R_1; P_1, \sigma) = \frac{\delta_m N}{\Delta}$. Therefore, other things equal, the quantity of rigs demanded, $N_D^1(R_1; P_1, \sigma)$, is decreasing in the rental cost parameter R_1 .

At $t = 1$, all rig operators will choose to rent their rigs; this follows because if they do not rent their rig they earn nothing. Therefore, the optimal quantity of rigs supplied at $t = 1$, $N_S(R_1; P_1, \sigma)$ follows immediately; $N_S(R_1; P_1, \sigma) = N$, for all $R_1 > 0$.

The equilibrium condition requires the quantity of rigs demanded and supplied to equate, so $N_S^1 = N_D^1$. Because $N_S^1 = N$, this equilibrium condition requires that all rigs are rented, $N_S^1 = N_D^1 = N$. In particular, $N_D^1 = \frac{N}{\Delta} \cdot \frac{P_1}{R_1 + C} = N$. In this case, the highest cost producer, Δ ,

will extract her reserve and the economy produces at capacity. Furthermore, denoting the equilibrium rig rental rate variable at $t = 1$ as R_1^e , we obtain that $R_1^e = \frac{P_1}{\Delta} - C$. However, because rig rates are positive, we note that $R_1^e = (\frac{P_1}{\Delta} - C)^+$.

As we show next, producer δ can use the time 1 value of the rig rate variable, R_1^e , to compute the value of her underground reserves at period 0.

At $t = 0$, agents observe the distribution of the terminal period oil price (P_1). At this initial period, each producer $\delta \in (0, \Delta]$ computes her time 1 extraction costs to be, $x_1^\delta = (R_1^e + C)\delta = \delta (\frac{P_1}{\Delta} - C)^+ + C\delta$ or:

$$x_1^\delta = \begin{cases} \delta \frac{P_1}{\Delta} & P_1 > C\Delta \\ C\delta & P_1 < C\Delta \end{cases}$$

Thus, at $t = 0$, each producer δ understands that her extraction costs at $t = 1$ will increase with the price of oil realized at that time (P_1). Moreover, extraction costs (x_1^δ) increase with the depth of the reserve(δ); $x_1^{\delta'} > x_1^\delta$ for $\delta' > \delta$.

At $t = 0$, producer δ also understands that because she has the option but not the obligation to extract her reserve at period 1, her time 1 payoff is $(P_1 - x_1^\delta)^+$ or:

$$(P_1 - x_1^\delta)^+ = \begin{cases} (1 - \frac{\delta}{\Delta})P_1 & P_1 > C\Delta \\ P_1 - C\delta & C\delta < P_1 < C\Delta \\ 0 & P_1 < C\delta \end{cases}$$

Producer δ 's reserve at $t = 0$ is equivalent to a long position in a call with an exercise price of $C\delta$ and a short position in $\frac{\delta}{\Delta}$ calls with an exercise price $C\Delta$. Producer δ can compute the value of her underground reserve (V_δ) at $t = 0$ by comparison with a schedule of option prices; these option prices follow from the model parameters. For example, call options for the pur-

chase of oil at $t = 1$ with exercise price E are traded at a price C_E . Therefore, $V_\delta = C_{C\delta} - \frac{\delta}{\Delta} C_{C\Delta}$. In addition, with $P_1 \sim N(\mu, \sigma^2)$ the reserve value V_δ can be computed directly as follows:

$$V_\delta = \left(\frac{(P_0 - e^{-r}C\delta)}{2} + \frac{P_0(\Delta - \delta)}{2\Delta} \right) erf \left(\frac{C\Delta - P_0 e^r}{\sqrt{2}\sigma} \right) - \frac{(P_0 - e^{-r}C\delta)}{2} erf \left(\frac{C\delta - P_0 e^r}{\sqrt{2}\sigma} \right) + k.$$

Here, $erf(x) = \int_0^x \frac{e^{-t^2}}{\sqrt{\pi}} dt$ and $k = \sqrt{2\pi} \left(e^{-\left(r + \frac{P_0^2 e^{2r} - 2C\delta e^r P_0 + C\delta^2}{2\sigma^2}\right)} - e^{-\left(r + \frac{P_0^2 e^{2r} - 2C\Delta e^r P_0 + C\Delta^2}{2\sigma^2}\right)} \right) + \frac{\sigma(\Delta - \delta)}{\Delta\sqrt{2\pi}} e^{-\left(r + \frac{P_0 e^{2r} - 2C\Delta P_0 e^r + C\Delta^2}{2\sigma^2}\right)} - \frac{P_0(\Delta - \delta)}{2\Delta}.$

At $t = 0$, each of the producers $\delta \in (0, \Delta]$ can spend $x_0^\delta = (R_0 + C)\delta$ and extract their reserve or leave it underground. A particular producer $\delta \in (0, \Delta]$ makes this decision at this time as follows:

if $P_0 - x_0^\delta > V^\delta$ extract her reserve

if $P_0 - x_0^\delta \leq V^\delta$ leave her reserve underground

There is a marginal producer with extraction cost $x_0^{\delta_m}$ who is indifferent between extracting the oil and leaving it underground. For this marginal producer, $P_0 - x_0^{\delta_m} = V_{\delta_m}$ and hence this marginal producer can be uniquely identified as, $\delta_m = \frac{P_0 - V_\delta}{R_0 + C}$.

The marginal cost producer at $t = 0$ depends on the value of the rig rental cost variable R_0 . If the rig rental cost variable (R_0) is sufficiently high, $\delta_m < \Delta$ and not all producers will extract their reserve. With a decrease in R_0 it will be optimal for more producers to extract their reserve, this raises the value of δ_m .

Given the choice of the marginal producer, δ_m , we derive the optimal quantity of rigs demanded at $t = 0$, $N_D(R_0; P_0, \sigma)$. In particular, $N_D(R_0; P_0, \sigma) = \frac{\delta_m N}{\Delta} = \frac{N}{\Delta} \frac{P_0 - V_\delta}{R_0 + C}$. A portion of the available supply of rigs (N) can remain unused at time 1. If $\delta_m < \Delta$, the

number of rigs used is $N\delta_m/\Delta$, and this number of rigs is less than the number of rigs available (N). In this case, the economy-wide production ($\delta_m Q/\Delta$) is below capacity (Q). A decrease in R_0 raises the value of δ_m and hence this also raises $N_D(R_0; P_0, \sigma) = \frac{\delta_m N}{\Delta}$. Therefore, other things equal, the quantity of rigs demanded, $N_D(R_0; P_0, \sigma)$, is decreasing in the rental cost parameter R_0 .

At $t = 0$, all rig operators will choose to rent their rigs; this follows because if they do not rent their rig they earn nothing. Therefore, the optimal quantity of rigs supplied at $t = 0$, $N_S(R_0; P_0, \sigma)$ follows immediately; $N_S(R_0; P_0, \sigma) = N$, for all $R_0 > 0$.

The equilibrium condition requires the quantity of rigs demanded and supplied to equate, so $N_S = N_D$. Because $N_S = N$, this equilibrium condition requires that all rigs are rented, $N_S = N_D = N$. In particular, $\frac{N}{\Delta} \cdot \frac{P_0 - V_\delta}{R_0 + C} = N$. In this case, the highest cost producer, Δ , will extract her reserve and the economy produces at capacity (Q). Denoting the equilibrium rig rental rate parameter at $t = 0$ as R_0^e , we obtain that $R_0^e = \frac{P_0 - V_\Delta}{\Delta} - C$. We abstract from corner solutions where R_0^e can be negative, by assuming that the current price of oil (P_0) is sufficiently high ($P_0 > C\Delta + V_\Delta$).

In equilibrium the market for rigs at $t = 0$ will clear, that is the quantity of rigs demanded will equal the quantity supplied ($N_D^0(R_0; P_0, \sigma) = N_S^0(R_0; 0, \sigma)$). This equilibrium is characterized by the time 0 rig rental rate of the marginal producer, $R_0^e(P_0, \sigma) = \frac{P_0 - V_\Delta}{\Delta} - C$. In this equilibrium, the highest cost rig operator will clear the market for rigs, $\delta_m = \Delta$. The economy in this case fully utilizes the available supply of rigs (N) and produces at full capacity (Q).

An increase in the oil price volatility to σ' raises V_δ . Each reserve increases in value because the likelihood that it will be economically feasible at the terminal date increases with volatility. Recall that the marginal producer is determined as, $\delta_m = \frac{P_0 - V_\delta}{R_0 + C}$. When a rise in volatility increases the option value of the reserves, V_δ , then the value of δ_m will decline. In this case, it will no longer be optimal for high cost producers to extract their

reserves. As a result, the optimal quantity of rigs demanded at $t = 0$, $N_D(R_0; P_0, \sigma) = \frac{\delta_m N}{\Delta}$, will also decline. This will be reflected in a downward shift in the demand for oil rigs, $N_D(R_0; P_0, \sigma)$, and a lower equilibrium rig rental rate, $R_0^e(P_0, \sigma')$. The decline in rental rate will be sufficient so that economy again fully utilizes the available supply of rigs (N) and produces at full capacity (Q). Analogously, an increase in the spot price of oil (P_0) makes it attractive for producers to extract their oil, this raises the demand for rigs and their rental rates.

References

- Abadie, A., and G. W. Imbens. 2012. Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*.
- Brennan, M. J., and E. S. Schwartz. 1985. Evaluating natural resource investments. *Journal of Business* pp. 135–157.
- Kellogg, R. 2014. The effect of uncertainty on investment: evidence from Texas Oil Drilling. *The American Economic Review* 104:1698–1734.
- McDonald, R., and D. Siegel. 1986. The Value of Waiting to Invest. *The Quarterly Journal of Economics* 101:707–27.
- Roberts, M. R., and T. M. Whited. 2013. Chapter 7 - Endogeneity in Empirical Corporate Finance. pp. 493 – 572. Elsevier.
- Titman, S. 1985. Urban land prices under uncertainty. *The American Economic Review* 75:505–514.