How Do Small Players Deduce Beliefs about Uncertainty? A Look at Texas Shale Oil Investments

Thomas Fillebeen

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Chris Anderson, Chair

Robert Halvorsen, Chair

Eric Zivot

Program Authorized to Offer Degree:
Economics

University of Washington

Abstract

How Do Small Players Deduce Beliefs about Uncertainty?
A Look at Texas Shale Oil Investments

Thomas Fillebeen

Co-Chairs of the Supervisory Committee:
Associate Professor Chris Anderson
School of Aquatic and Fishery Sciences

Professor Robert Halvorsen
Department of Economics

Derivative markets enable firms to eliminate unwanted risk and, thereby, focus on their core competence. We ask whether small firms respond to changes in complex risk structures, using the emerging Texas shale oil industry. Unlike their conventional drilling counterparts, shale oil producers are small, with short production lead times and most revenue earned in two to three years. This exposes them to recent unprecedented volatility in product prices, production costs, and-in the Permian, pipeline transportation basis risk. We develop a real option model of their decision process, within which we measure firm responsiveness to volatility. To estimate revenue expectations, we use forward-looking time-varying beliefs deduced from futures and futures option prices. For costs, we use expectations about the evolution of production technology and forward-contracted transport capacity. Our results show that investment in new wells responds optimally not only to time-varying market price, but also to a second source of uncertainty: transportation basis. This demonstrates that even small firms are able to coordinate their production activity to integrate two sources of risk, impacting their investment returns. Lastly, we estimate the impact of the Nixon-era crude export ban, which became binding in August 2013, and caused US domestic crude price to fall below the world price, distorting US shale oil production revenue by $10.7B.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I'd like to express my sincere appreciation to the University of Washington, where I had the opportunity to learn from Professors Chris M. Anderson and Robert F. Halvorsen, whose guidance, supervision, and accessibility were invaluable. The numerous Anderson lab meetings throughout the years, and their participants, served as incubators for my research. Also, I'm grateful to Professor Levis A. Kochin, who helped shape and steer the topic through multiple long conversations, and to Professor Eric W. Zivot for his help selecting the methods and supporting the project early on and throughout. Lastly, I'd like to thank Professor Joe P. Mahoney for sharing his energy infrastructure knowledge.

I am appreciative of suggestions from participants of numerous seminars including, but not limited to, Ibrahim Demir, Anthony Sanford, Alexander Rodivilov, and Jorge Rojas. I thank Danyaal Khan and Phong Tran for their technical expertise in managing the data.

DEDICATION

to my lovely wife, parents and brother

Chapter 1

MANAGERIAL DECISION-MAKING

It is unclear how small firms—independent price-takers with short lead-times and production horizons—deduce beliefs about uncertainty when making irreversible investment decisions. As uncertainty increases, real options theory stipulates that the opportunity cost of postponing an investment decreases. This thesis argues that it is the integration of this time-dependent option to delay, which can incorporate multiple sources of uncertainty, that shows firm sophistication in estimating risk. We value a small[1] firm's investment project using a real option model and test whether its investment heuristics mimic those of an optimal decision-making process by measuring whether myopic constant or market time-varying uncertainty yields a better fit of firm investment behavior. Price and price volatility expectations can be central in estimating a real option forward-looking investment decision-making model. One centralized source of information about price expectations is futures and futures options (or derivative) markets. Standard practice among academics, policymakers, and businessmen is to interpret futures and futures options as the market expectation of the spot price and price volatility.[2] These derivative markets provide certain industries with unique access to expected price and uncertainty of each investment's return. The ability to recover the full distribution of expected investment returns signifies that these particular firms are uniquely positioned to deduce forward-looking uncertainty from markets. A small firm with a grasp of market or price risk—associated with unexpected changes in market price of commodities—can then fo-

[1]See Appendix A.1.2 for a discussion on the makeup of the Texas shale oil market.

[2]A subtlety is that, unlike speculators who require a risk-premium, producers can transact, or lock-in, the existing forward (or futures) price now to make their investment decision. We use the stronger risk-neutrality assumption, which states that without, e.g., risk-aversion the futures price is the conditional expectation of the spot price: $F_{\tau,t} = E[S_{\tau,t}]$. Since producers can hedge away risk using forward prices the assumption holds. However, we can also use the weaker assumption stating that it isn't necessary that the forward price equal expectation, but only that a producer be able to transact at that forward price.

cus on its core competence (Chew (2001)). However, the high upfront costs of acquiring technical expertise and software may serve as a barrier to estimating forward-looking beliefs. An alternative possibility is for small firms to incorporate information on the price of its major inputs acquired through daily business operation and not the market. The question then becomes, how sophisticated is a small firm's response to forward market traded uncertainty?

There is a lack of empirical work analyzing whether small leveraged firms are sophisticated enough to handle complex forms of time-varying risk when making irreversible investment decisions. Our thesis argues that appropriately measuring uncertainty allows a sophisticated small firm to accurately estimate the opportunity of waiting and learning more before actually investing. First, a firm estimates the real option value using expected time-varying price or market uncertainty.[3] Second, it measures time-varying basis risk expectations and its level of exposure to it, which is generally forecasted from spot dynamics.[4] Basis exposure results from the value of the commodity being tracked, and possibly hedged, not always covarying with the value of the derivative contract used to track price risk. Three common forms of basis risk are quality, transportation—different delivery points—and timing.

We contribute to a literature that, as Copeland et al. (2005) argue, needs more empirical research on the application of real options, as there is not enough evidence that they are priced correctly in the marketplace. Theoretical work, by Marschak (1949) and Dixit and Pindyck (1994) shows that it is suboptimal for investment decisions to be based solely on prevailing market prices and to disregard the volatility-induced optionality of waiting. A few empirical studies by Grenadier (1996), Aguerrevere (2003) and Dias (2005) measure the impact of uncertainty on irreversible investment decisions under various market structures. Also, in natural resources, Moel and Tufano (2002) and Dunne and Mu (2010) show the importance of realized volatility on mining and refining investments, respectively. However, all these studies apply backward-looking rather than

[3] A project with a real option value is equal to: Project Value = Net Present Value + Real Option Value.

[4] Forward contracts can be resold over-the-counter. The ease of resale signify that when forecasting basis spot dynamics are of primary importance.

forward-looking market measures of expected price volatility. Recent progress in the natural resource field by Kellogg (2014) captures market time-varying expectations for in-fill well investments in a stable price and cost environment from 1993-2003 in Texas. He shows the importance of forward-looking measures of price and price volatility expectations on firm decision-making by allowing the volatility process governing future oil prices to vary over time.[5]

Using the emerging Texas shale oil industry as a case study, our thesis estimates small firm responsiveness to complex sources of uncertainty. These small firms simultaneously drive and respond to advances in technological innovation, providing insight into the incentive structure of oil markets with energy policy implications. With technological innovation and access to capital and financial markets, these small leveraged firms are able to enter the oil industry, which was previously dominated by a few vertically integrated large firms. In particular, these are small debt- and equity-financed firms, which developed a new well production technology that exposes them to short-term risk. Unlike the large conventional incumbent firms that use production technology that exposes them to long-term risk, small firms with limited working capital and resources track their market or price risk using derivative markets. These cash flow-strapped small firms are more sensitive to production risk but benefit from access to developed derivative and capital markets. We empirically model them as independent shale oil producers with short lead-times and well production horizons, while the incumbents are large firms drilling conventional wells with long production horizons and lead-times. The entry of these new firms into global oil markets changed the way we conceive of energy policy. Previous literature by Hotelling (1931) views conventional oil production as an extraction rate choice, while recently Anderson et al. (2014) presents the production problem as a well investment choice, where extraction rate is determined by reservoir pressure. Both authors model an environment where uncertainty's impact on production decision is inelastic.

[5]Kellogg (2014) estimates market long-term option implied volatility by combining short-term implied volatility with a realized (parametric) estimation of volatility (or term structure of volatility) over the sample. However, we argue that imposing a constant term-structure is limiting and instead use a non-parametric estimation of the implied volatility surface to measure long-term implied volatility. This approach uses all available option data at every traded strike price and time-to-expiration applying methods that are widely used by practitioners (Chen (2011); Heston (1993)), see Appendix A.4.2.

However, as recently as 2009, with the entry of these shale oil producers, we find that the correspondence between production, well investment decision, and price risk is now more elastic in the short-run.

1.1 Shale Oil Environment

The thesis's 2009-2015 shale oil environment allows us to model firm well investment behavior using a single-agent framework. We then fit firm investment behavior onto Texas well-level drilling and completion data using responsiveness to time-varying volatility as a free parameter. This approach is possible because shale producers are atypical in the industry for a number of reasons. First, as small independent marginal producers with short lead-times and sharp well production decline rates, they track and respond to expected crude market price, volatility, forecasted transportation basis cost, and service cost per well. Second, they are not constrained by state or federal government in the way that other producers worldwide are. Third, these firms are price takers, as energy markets are global and competitive with no single shale oil producer capable of manipulating price. Fourth, they operate in a unique environment that doesn't constrain optimal decision-making, with lax credit, low expected corporate defaults, and mostly slack drilling and completion service supply. Fifth, because of the high price environment throughout most of the sample, these independents drill and complete homogeneously by adopting a 'factory drilling' approach, treating each prospect uniformly within the shale oil plays of a basin. Lastly, these shale oil agents don't diminish the real option value of waiting as strategic interactions are omitted, with land locked up early in the sample and few common pool resource dynamics that are typical of conventional wells.[6] As a result of these sample characteristics, we focus on a simple single-agent real option model, allowing the volatility process governing future oil prices to vary over time; we then measure whether a small firm's response to uncertainty is sophisticated.

The dominant crude oil futures and futures options contract is West Texas Intermediate (WTI), traded on Chicago Mercantile Exchange (CME) as a NYMEX product—these

[6]Shale oil rock has low porosity and permeability, which prevents it from flowing freely. Also, there is a cooperative incentive to increase lease values by sharing technological improvements to be sold on a secondary market.

are light and sweet crude oil contracts. The WTI contract is delivery-settled, calling for delivery of 1,000 barrels of crude oil in Cushing, Oklahoma. It is a primary source of price discovery in the international oil market. Many physical oil trades are quoted based on these futures and futures option instruments. Once a shale oil firm measures price risk from this derivative market, a common basis risk (or additional degree of sophistication) is transportation. This basis exposure results from the difference between oil production and delivery locations. For shale oil firms in certain basins, high, unexpected production growth early in the sample period increased congestion risk by overwhelming basin take-away capacity. We find that failure to respond to both price and transportation basis exposure by identifying and integrating them into firm-wide decision-making leads to a misestimation of real option value.[7]

As stated above, the first—and more standard—source of risk is price, which is obtained from derivative markets. In this paper, two out of the three independent variables used to capture small firm well investment responsiveness to expected price volatility (or uncertainty) are extracted from derivative market data. Expected prices are directly taken from futures markets, while expected price volatility is implied from futures option price data. Note that when long-term option prices are missing we estimate it using Heston (1993) and Chen's (2011) stochastic model.[8] Then using the Black (1976) model we estimate the implied volatilities for all available option prices. Finally, when short-term implied volatilities are missing we interpolate to estimate their values by applying a non-parametric cubic spline across maturities and strikes. Derivative markets are "efficient aggregators of information," incorporating more distributional information than backward-looking volatility measures. Work by Hilliard and Reis (1998), Szakmary et al. (2003), and Kellogg (2014) tests and supports the hypothesis that market-based volatility is a better predictor of expected uncertainty across multiple commodity and financial markets. These futures markets offer shale oil firms a measure of expected uncertainty

[7]Accurate real option valuation prevents over- or under-investment in good times, which allows a producer to make optimal capital expenditures without having, for example, to prematurely reenter capital markets in bad times.

[8]See Appendix A.4.1 and A.4.2 for a full discussion about parametric and non-parametric estimation of implied volatility from futures options data.

unavailable to firms in other industries. In only a few industries, such as energy production, airlines, and agriculture, are futures and futures options markets widely-traded and used by industry participants. With access to financial energy markets, a risk-averse producing firm has the option to hedge its price exposure. This optionality to hedge allows us to model shale oil producers as risk-neutral single agents.

Transportation accounts for the second, more complex source of risk, and is estimated using spot price dynamics[9] between production and delivery locations. The spot price differential between production and delivery points is known as (spot) transportation basis. Transportation basis is important when supply outstrips take-away capacity, resulting in congestion. Congestion leads to a temporary transport basis jump or, stated differently, a much lower price at the production location than at the delivery point. In Texas, we focus on the Permian and Eagle Ford basins; together, they represent more than 50% of shale oil production growth in the US. We look at the Permian to analyze well investment behavior that integrates both price and transportation risk—Eagle Ford isn't exposed to non-stationary transport costs. To estimate expected transportation basis, we forecast since there isn't a forward market in transportation costs or Permian Midland prices. Based on Schwartz (1997) and Hilliard and Reis (1998), we estimate a mean-reverting jump diffusion process on spot transportation basis data, for which we measure long-term transportation basis levels, as well as jump magnitudes and frequency, among other things. Once the model is fitted, we use Monte Carlo simulations to estimate[10] the expected long-term transportation basis between Permian Midland and WTI Cushing. By identifying and integrating both price and transportation basis uncertainty[11] into the real option model, this thesis concludes that small firms are sophisticated, as they integrate time-varying risk that can come from multiple sources, yielding a better fit of investment behavior.

[9] As stated earlier, forward-contracted and hedged transport capacity can be resold.

[10] When extrapolating 24-month transportation basis, we normalize for mean-reversion in volatility for further out maturities. Volatility term-structure mean-reversion is discussed in Appendix A.4

[11] We use a mixed-lognormal distribution and measure portfolio weights and a correlation structure over the sample to estimate complex volatility. Note that portfolio transportation basis weights increase in tandem with transportation congestion.

1.2 Energy Policy Implications

In August 2013, increased US shale oil production made the 1975 export ban binding and distorted Texas horizontal shale oil well investment. The ban prevents the export of crude but not crude product. This jump in production was driven by progress in drilling and completion technology, which allowed the US to extract large volumes of shale oil from tight rocks as of 2009, thereby altering the face of global energy geopolitics and, in particular, the economics of US energy policy. From 2009-2015, with heightened production, US shale oil producers became the marginal price setters.[12] Hence, to model current crude production dynamics and the full impact of the distortionary export ban policy,[13] we identify and estimate the supply response (elasticity) of shale oil producers.

To estimate the export ban's production impact, we measure the responsiveness of the shale oil portion of the supply curve. Hotelling's (1931) optimal exhaustible resource extraction framework, which controls for the production rate, isn't appropriate for modeling horizontal wells.[14] The production decline rate from these wells is sharp for the first two years and, once horizontal hydraulic fracturing is initiated, the flow rate cannot be intertemporally controlled. Anderson et al.'s (2014) modified Hotelling rule handles constrained flow rates by viewing production as an investment rather than a production problem. Well investments are thus key in modeling shale oil production. However, they apply this framework to conventional wells with low, and not high shale oil, production decline rates.

Our focus on irreversible well investments with sharp production decline rates marks a divergence from the literature, that up until now, found that the marginal crude supply was unresponsive in the short-run to price changes. Short-run shale oil production's responsiveness to price shocks is important to the macro-empirical literature epitomized by

---

[12] See Appendix C.1 for a depiction of the 2014 medium-run aggregate production cost curve. Shale oil production cost is located between that of onshore conventional and offshore deepwater.

[13] For of 2015, the ban was ineffective at depreciated prices of \$30-\$40; then on December 18th 2015 it was removed. Hence, for our Jan 2009-Jan 2015 sample, we estimate the full impact of the ban on well investments.

[14] Even then, structural econometric papers testing the "Hotelling rule" on the natural resource industry, including Halvorsen and Smith (1991), and Black and LaFrance (1998) find mixed results.

Hamilton (2009) and Kilian (2009) as well as, more recently, by Güntner (2014). These authors' work focuses on inelastic conventional oil supply, which responds gradually to crude price shocks. This new unconventional production response signifies that policy impacts, such as the binding export ban, lead to instantaneous well investments and production impacts.[15] For these wells with high decline rates, Kellogg (2014), with tertiary wells uses an irreversible investment approach to show that producers coordinate investment decisions based on time-varying market based expectations. The author finds that these producers are highly responsive to forward-looking price and volatility market signals.[16] Similarly, we deduce producers' beliefs about revenue and costs during the life of a new well from crude futures and futures options markets, transportation cost forecasts, technological trends, and well productivity gains, and use these beliefs as a basis for measuring well investment behavior.

The real option framework offers a new structural policy technique, which we use to estimate level, or price, and variance effect on forward-looking oil producer investment response. First, a real option allows a decision maker to not only assess current expected crude price, volatility, and total cost per well, but also to weigh the dynamic option to drill and complete now or to delay the investment decision.[17] Second, this approach, coupled with access to counterfactual state variables without the export ban, allows us to estimate the impact of the ban post-August 2013, as it became binding. Hence, once US domestic production outstripped demand and the export ban policy became binding, we identify its impact on production by estimating the number of well prospects that went un-drilled and -completed.

Empirically, there are three primary financial global benchmark crudes—West Texas Intermediary (WTI), Louisiana Light Sweet (LLS), and Brent—that are delivered in the

---

[15]The degree of production impact depends on the current production levels, the operational age of the producing wells, and proportion of price or market hedges.

[16]For example, when expected prices are much higher than current prices—known as contango—shale oil producers are incentivized to consider drilling and not completing their well—this storage option was primarily used as of mid-2015 when crude prices remained depressed.

[17]Shale oil producers may opt to delay their well investment decision to see if demand is high (or supply is low). If demand is high (or supply is low), they can decide to invest in a well, while with low demand (or high supply), a producer can avoid making a financial commitment by waiting further.

North Sea (Brent) and other water price points such as the US Gulf Coast where LLS is traded. WTI, however, is inland in Cushing, Oklahoma, connected by pipeline to the US Gulf coast. The WTI contract is a primary source of price discovery in the global oil market. Many physical oil trades are quoted based on these futures and futures option instruments. With access to price and volatility of these three financial market benchmarks, we can estimate the impact of the export ban on drilling and completion decision-making.[18] To do this, we use a counterfactual for the period when the ban was binding. We refer to this new counterfactual design that controls for both price and volatility as a real option-state switching policy estimation. The first state is the current domestic WTI environment with the export ban in effect, while the second is the (theoretical) international Brent state in which the ban doesn't impact prices and volatility. In fact, by controlling Brent international prices and volatility for Texas regional transportation dynamics, we estimate our counterfactual environment without an export ban.[19] Once we estimate and apply our real option well investment model on these two different sets of state variables—switching states as of 2013-08-27 when the ban became binding—we take their difference to measure the number of displaced wells.

As stated above, this counterfactual environment offers access to domestic price dynamics with and without the export ban. We use state price and volatility estimates from exchange traded WTI—the most traded crude benchmark—while a transportation cost-adjusted Brent serves as a counterfactual for global market fundamentals, illustrative of an environment without the binding export ban. Without the export ban, LLS is the US Gulf Coast domestic price, which at the margins is set by the imported Brent—North Sea international—price: the two prices are equal once we control for freight transportation basis.[20] WTI is the US Cushing, OK domestic inland price; the WTI and LLS price difference reflects transportation basis from the Gulf Coast to Cushing, OK. When the

[18]The counterfactual total cost per well is parametrically adjusted for higher price levels and well productivity gains; see Appendix C.2.

[19]Oil markets are global, hence firms are price takers and Texas shale oil as a whole represents a small percentage of worldwide oil activity.

[20]Remember from earlier, basis exposure results from the value of the commodity being tracked, and possibly hedged, not always covarying with the value of the derivative contract used to track price risk. Three common forms of basis risk are quality, transportation—different delivery points—and timing.

export ban became binding, WTI and LLS prices converged,[21] while LLS and Brent prices diverged. By correcting for transportation basis proxied by the LLS and WTI price difference, Brent is effectively converted into an accurate representation of a counterfactual WTI from August 2013-January 2015.

With current revenue and cost inputs and their counterfactual states, we estimate the diffusion process governing their future expected evolution. This helps measure the real option value, that is, the value of the right (but not the obligation) to invest in a well during the life of a shale oil lease. This value is modeled using a time-varying diffusion process (geometric Brownian motion). To numerically solve the real option, we use a dynamic programming Bellman Equation, and to estimate the diffusion process, we apply a Markov chain. The approach is first used by Kellogg (2014), with respect to tertiary wells for 1993-2003 in Texas. For shale oil, we use our non-parametric estimate of expected volatility and incorporate dual sources of risk—price (or market) and transportation basis—to estimate an empirically accurate well investment response in both the Permian and Eagle Ford. Note that Samis et al. (2005) find that a real option is the appropriate method to properly identify risk between different assets. We estimate the real option model and its diffusion process to measure drilling and completion activity in the Permian and Eagle Ford post-August 2013, capturing the well count difference with and without the export ban for August 2013 to January 2015. The difference between these option values is then converted into production output, severance and tax revenue, and forgone CO_2 emissions in Texas.

We divide this thesis into four chapters. Chapter 2 introduces the sample period of unexpected production and technological growth and motivates the real option model. It then characterizes the empirical data, which is extracted from multiple sources to estimate current and future expected revenue and cost. Finally it discusses the integration of these profit beliefs and volatility responsiveness into the real option model and transition dynamics. We also report and discuss the real option model empirical fitting results for Eagle Ford. Chapter 3 shows the estimation of transportation basis and its incorporation into investment decisions as well as reports and discusses the empirical results for the

[21] Increased transportation capacity and the export ban enhanced the convergence rate.

Permian. Finally, in chapter 4, we estimate the welfare costs of the export ban on investment behavior using a real option model and counterfactual expected price and price volatility.

Chapter 2
PRICE OR MARKET RISK

2.1 *Shale Oil Production*

Unlike the very long term investment decisions involved in conventional drilling, the technology of shale oil production allows a producer to observe and respond to relatively short term market conditions. We consider one of a large number of *ex ante* homogeneous prospects, indexed i, which will yield a series of monthly productions $\{x_{it}, \ldots, x_{it+T}\}$ if drilled at month t and produced for T periods; T is typically 36 months, with hyperbolic declining production over time and 80% of revenue earned in the first 24 months.

Production is sold in spot or futures markets at prices P_t, a vector of prices for delivery in periods $t, \ldots, t+T$, which vary with market conditions and volatility $\sigma_P(t)$, where the argument denotes the market price volatility expected at drilling time t—it is a producer's current expectation of volatility of the oil price. Prices expected in future time periods evolve according to a first order Markov process:

$$\ln P_{t+1} = \ln P_t + \mu(P_t, \sigma_P^2(t)) - \sigma_P^2(t)/2 + \sigma_P(t)\epsilon_{P_{t+1}} \tag{2.1}$$

where $\mu(P_t, \sigma_P^2(t))$ is a price trend (or drift) as a function of expected level and volatility[1] and $\epsilon_{P_{t+1}}$ price shock.

However, as exogenous demand conditions, global supply costs, and commodity speculation shifts, so too does the level of volatility. This moves according to a similar first order Markov process:

$$\ln \sigma_P(t+1) = \ln \sigma_P(t) - \gamma_t^2/2 + \gamma_t \epsilon_{\sigma_{t+1}} \tag{2.2}$$

[1] Heightened volatility affects the convenience yield of holding an additional unit of crude inventory—marginal value of storage. Expected storage value—the difference between injected spot and the expected withdrawal future price—increases with spot price volatility. Higher volatility increases the probability of spot and future spreads expanding. Furthermore, as price volatility increases there is higher inventory demand used to smooth crude supply. Therefore, price volatility leads to inventory build-up and raises prices.

where γ is a parameter capturing the volatility of inter-period volatility, and $\epsilon_{\sigma_{t+1}}$ is a shock to volatility.[2] Specifically, γ is the volatility of volatility process and $\epsilon_{\sigma_{t+1}}$ is an i.i.d standard normal random variable.[3]

Drilling and completion costs also experience volatility, as competition for drilling rigs and resources increases when higher price expectations stimulate new drilling activity. Therefore, we capture the evolution of costs with a volatility that is proportional to price volatility, $\sigma_{TC}(t) = \alpha\sigma_P(t)$, and a Markov process:

$$\ln TC_{t+1} = \ln TC_t + \hat{\mu}(TC_t, \hat{\sigma}^2_{TC}(t)) - \hat{\sigma}^2_{TC}(t)/2 + \hat{\sigma}_{TC}(t)\hat{\epsilon}_{t+1} \tag{2.3}$$

where $\hat{\mu}(TC_t, \hat{\sigma}^2_{TC}(t))$ is the trend (or drift) in costs as a function of expected level and volatility, and $\hat{\epsilon}_{t+1}$ is a cost shock drawn from a joint distribution with ϵ_{t+1} with correlation, ρ. TC_t is normalized by technological improvements, which steadily reduces cost.

2.1.1 Behavioral Model

At time t, the producer can construct an expected profit from investing in a well immediately based on current market conditions and volatility. Profit from drilling and completion is then conceptualized as:

$$\pi_i(P_t, TC_t, \sigma(t)) = \sum_{\tau=t}^{t+T} x_{i\tau}\mathrm{E}_{\sigma(t)}[P_\tau(\sigma(t))] - TC_t \tag{2.4}$$

where $x_{i\tau}$ is the production from well i in month τ, and T is the maximum life of a well, so the summation represents expected revenue given current beliefs about volatility; $P_\tau(\sigma(t))$ is the expectation of prices at future time τ, taken with respect to volatility at time t.[4] Further, drilling at t locks the production costs and associated drilling and completion technology available at t.

---

[2] For the Permian, to model current volatility we reestimate it to include price, or market (P), and transportation basis (B) risk: $\sigma^2(t)_{PB} = \sigma^2_B(t)w^2_B + \sigma^2_P(t)w^2_P + 2w_Bw_P\sigma_P(t)\sigma_B(t)\rho_{B,P}$

[3] We go over calibrating these two parameters in the empirical/data Section 2.2.4

[4] In practice, shale firms can lock in time-t pricing for most of their production in the future market, which is thickly traded up t+24 months, over which period most of their production occurs.

The choice a producer faces is whether to drill now and receive $\pi_i(P_t, TC_t, \sigma(t))$, or wait to exercise the option to drill and potentially benefit from more favorable realizations of the state variables later. The key state variables are $\{P_t, TC_t, \sigma(t)\}$. An optimizing producer will make the decision that maximizes the value function Equation 2.5:

$$V_i(P_t, TC_t, \sigma_P(t)) = \max\{\underbrace{\pi_i(P_t, TC_t, \sigma_P(t))}_{LHS}, \underbrace{\delta^t \mathrm{E}[V_i(P_{t+1}, TC_{t+1}, \sigma_P(t+1))]}_{RHS}\} \tag{2.5}$$

where the left-hand (LHS) term is Equation (2.4) and the right-hand (RHS) term is the expected value of waiting to observe the next period's state variables and then proceeding optimally. Waiting defers profits into the future, which are discounted at rate δ, but might confer several advantages. Our producer well drilling and completion investment decision can be differed indefinitely; to solve the dynamic programming problem we then use a stationary argument. This means that, conditional on the state, the value is independent of time. Hence, we'd theoretically remove the t index. Dropping the time-t index from the value function $V_i(P_t, TC_t, \sigma_P(t))$, it then becomes $V_i(P, TC, \sigma_P)$. Then next period's variables (RHS) are denoted with a prime, e.g. $V_i(P', TC', \sigma'_P)$.[5] However, this removes the drilling decision's temporal dimension, which is best explained, going forward, using a time-t index. First, drift in prices or costs may lead to higher prices, or lower costs. Second, volatility may mean P_{t+1} (or P') and TC_{t+1} (or T') are more advantageous. Third, time-varying volatility may yield a better volatility regime in which to drill in the future.

The real option tradeoff in Equation (2.5) is best framed as a dynamic programming optimal search and stop problem. Optimal stopping theory indicates that the next step is to solve for a well's time t critical productivity trigger, x_t^*, that leaves a producer indifferent between current profits (LHS) and the future expected evolution of profits (RHS) from drilling. At time t, only when $x_{it} \geq x_t^*$—expected productivity is bigger or equal to the productivity trigger—does a producer decide to drill and complete now.

In Figure 2.1, we show the effect of volatility $\sigma(t)$ for different expected price per barrel (bbl) scenarios ($40, $60, and $80/bbl) and fixed total cost per well ($6.8M) on solving for the critical well productivity x_{it} trigger. Figure 2.1 illustrates that as price

[5] With value function: $V_i(P, TC, \sigma_P) = \max\{\ \pi_i(P, TC, \sigma_P),\ \delta\mathrm{E}[V_i(P', TC', \sigma'_P]\}$

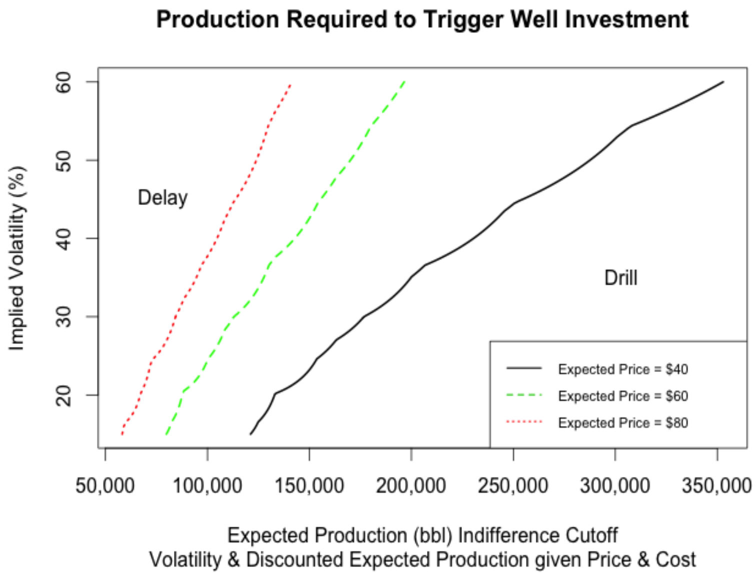

Figure 2.1: Drilling Trigger

increases, holding all else constant, the opportunity cost of delaying climbs.[6] Also, as volatility increases holding, costs and expected price constant, the well expected productivity necessary for a well to be drilled and completed goes up; volatility affects next period's expected price distribution, holding current price fixed.

Accurately assessing this tradeoff requires solving two complex dynamic inference problems. First, $\sigma_P(t)$ must be inferred based on available market information. Kellogg (2014) shows in a stable price, volatility, and cost environment, where on average only five wells are drilled per month, that small and large firms in his subsample draw on futures and futures options markets to make this inference. However, shale oil firms operate in a more dynamic price and volatility environment with high technological innovation and production productivity gains, where on average one-hundred and ninety wells are drilled a month, and it's unclear whether they are sophisticated enough to integrate these sources of risk.[7] Second, they must grapple with a dynamic inference problem: they must

[6] A producer's decision to drill and complete a well without integrating volatility leads to an uneconomical well investment, which is illustrated in Appendix A.7

[7] Industry participants state that shale firms use net present value (NPV) and value investment ratio (VIR) to analyze the value of a project. These approaches omit the real option value of the project:

infer γ to understand how volatility changes over time. There isn't a market to directly estimate volatility of volatility γ, so they must infer its value from futures options market implied volatility realizations. The empirical question is whether these small firms are able to heuristically supplement their well investment decision with similar analyses to make optimal decisions that reflect the dynamic complexity of their markets.

2.2 Data and Methods

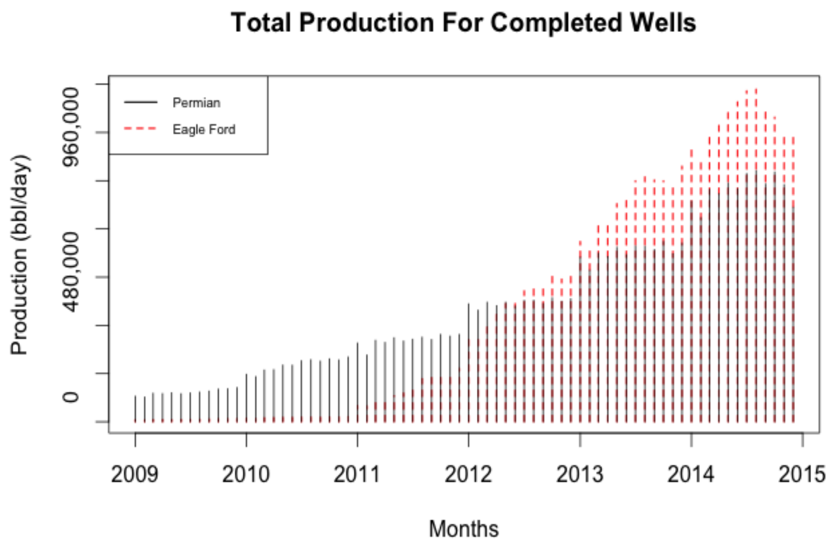

Figure 2.2: Eagle Ford Higher Production Growth than Permian

In order to test whether small firms are able to integrate complex risk information into their decision-making, we model the choices of firms to initiate drilling and completion of shale wells in west Texas during a period when the industry was most rapidly growing, as shown in Figure 2.2. Figure 2.3 shows the Eagle Ford and Permian basins where these firms are active. They present a strong test of this model, because they are mostly small, with the largest firm representing 5.7% of the market, and all but 2.6% are unaffiliated

Project Value = Net Present Value + Real Option Value. To compensate scenario analyses are used that approximate an output distribution. However, these approximations fail to measure the distribution of and relationship between inputs, and treat the investment horizon as finite, leading to approximation-errors, which widen when volatility increases.

with major producers. From 2009 to 2015, first month well productivity (bbl) increased by seven-fold, total cost (in $1,000,000) per well indexed to 2014 designs declined by 30%, expected prices averaged $84 per barrel, and implied volatility averaged 25%.[8] In contrast Kellogg's sample focused on infill wells[9] with sharp production decline curves and much lower volumes (bbl per initial production month) than shale. In his environment, price and implied volatility averaged $22 and 18%, respectively, and productivity per well prospect over time was stable with total cost per well (in $100,000) increased by 19%—well design (or input factor productivity) didn't change much for the 1993-2003 period.

As depicted in Figure 2.3, we focus on the Permian[10] and Eagle Ford basins, which differ in transportation basis risk, total cost per well, and well productivity gains. As shown in Figure 2.2, these differences lead to disproportionate shale oil production growth in both basins.[11] The Permian invests in both vertical (conventional) and horizontal shale (unconventional) wells; hence it takes longer for well productivity gains to make horizontal shale wells profitable.[12] These two basin specific differences allow us to test producer degree of sophistication. The earlier price risk specification is best tested on Eagle Ford, while with the Permian we strengthen the test by showing a more complex well project valuation, which requires different parameters and model structures.

Further, the Permian presents a one-step more complex volatility environment, because as, shown in Figure 2.3, it is more geographically isolated in the Texas Midlands and thus susceptible to transportation congestion. Usually, excess transportation capacity can be sold on a secondary market; however, as US shale oil production unexpectedly took off in 2009, and with unexpected pipeline construction delays[13], domestic production

[8]Where expected price and price volatility is proxied by the 24-month futures and futures options contract maturity.

[9]Wells drilled into the same reservoir as known producing wells to improve oil recovery by "filling in" areas of the reservoir that were not originally exploited.

[10]More specifically, Wolfcamp, Spraberry, and Bone Springs represent the horizontal shale oil fracking plays in the Permian Delaware region while Permian Midland accounts for the remainder.

[11]To estimate total production per day, we divide total production per month by twenty five days.

[12]Because of higher well investment activity in Eagle Ford, we can stipulate that producers benefited from more learning-by-doing and information sharing than in the Permian.

[13]Aside from state and government delays, a midstream company constructing a new pipeline must

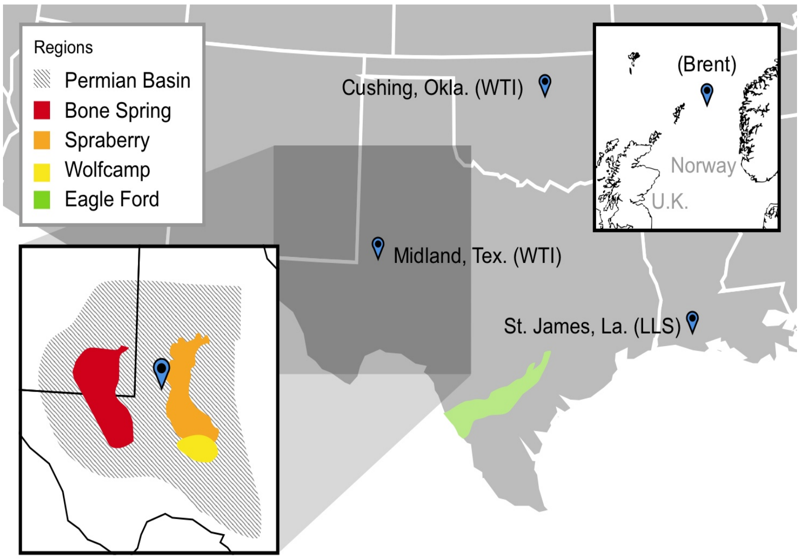

Figure 2.3: Texas Primary Shale Oil Basins

quickly overwhelmed takeaway capacity in the Permian.[14] This means that the Permian has a second, complexly-structured volatility component to its total cost term, which requires further analysis if small firms are going to time their development optimally.

To assess whether firms properly gauge and respond to complex volatility, we study how changes in the number of new wells put into production in each month varies with the underlying conditions in Equation (2.5). Building this model requires estimating, for each month, the expected revenue, costs, and volatility, as well as the evolution of volatility over time. The next section explains how we recover each of these factors from available data. We then test whether small producers appreciate the complexities of this environment by testing whether simple (static) or complex (time varying, or in the case of the Permian, multi-component) measures of volatility better explain production initiation decisions.

contract out up to 60% of its capacity before starting construction.

[14]Looking at producer revenue exposure, we find that at the beginning of the sample, price is the primary risk at 97%; however, towards the end of the sample, transportation basis contributes as much as 14%.

2.2.1 Overview of Data Used in Study

Real option model inputs and information regarding their evolution come from multiple sources. First, the dependent variable, individual well prospects, comes from the Texas Rail Road Commission's (TRRC) "Drilling Permit Master and Trailer" dataset, with well drilling and completion dates from 1975 to 2015. This dataset provides field ID, operator name, lease number, and qualitative individual well information. Well productivity and total production per month are drawn from the TRRC's "Oil and Gas Annuals" dataset.[15] Second, the independent variables, expected futures and futures options, and spot Permian Midland price come from Bloomberg's CME/NYMEX product database for 2009-2015; for total (fixed) cost per well, we use insight from a recent Energy Information Administration (EIA)-commissioned IHS Global Inc. (IHS) study with yearly total cost estimates for 2009-2015,[16] in addition to two surveys with quarterly and monthly data. One of these surveys is for rig rental cost per well from the *Day Rate Report* for 2009-2015, and the other for fracking and miscellaneous[17] costs per well from the *Spears & Associates DCS Report* for 2012-2015. We interpolate to estimate monthly total cost data when there is only quarterly data available. Also, in the cost section, we discuss how total cost per well integrates, for example, the reduction in number of days a rig is rented—productivity of inputs—and our normalization of total cost by well productivity gains in the first month post completion over the period t indexed to 2009.[18]

2.2.2 Dependent: New Wells Completed

Changes in the fundamentals in the Bellman Equation (2.5) will result in changes in producers' decisions about whether to make new investments and complete additional wells. Our dependent variable is the number of wells completed in each month in each

[15] The production dataset is currently only supplied through the end of 2014, as there is a one-year lag. We divide monthly production by 30 days to get barrels per day.

[16] It measures both the average well (oil and gas) drilling and completion costs indexed, and non-indexed, to 2014 well design per basin. For our purposes, the estimates are broken down into Eagle Ford and Permian horizontal shale oil total cost per well.

[17] For example, hazardous waste disposal and insurance and consulting.

[18] We can either multiply expected unobserved productivity x_{it} by increases in realized production, or we can divide total cost per well by it, over time t.

basin out of all the available prospects. In shale oil, well completion is the process of making a drilled well ready for production by stimulating (or fracking) the rock formation for optimal flow. This is constructed from the Texas Railroad Commission's (TRRC) "Drilling Permit Master and Trailer" dataset, which provides drilling and completion dates for each well drilled from 1975 to 2015. For the Permian and Eagle Ford, we count the number of wells completed in each month.

Since the number of undeveloped prospects available for investment changes during the sample, we control for this by also considering the number of available prospects that were not developed. Using TRRC's number of permits issued to drill shale oil wells as of 2009, we know that the number of well prospects that could have been drilled for the sample was around 24,000 for 2009-2015. In Eagle Ford and the Permian, there were about 18,102 and 6,214 shale oil well prospects (or permits issued).[19] Out of these prospects, 10,353 shale oil wells were drilled and completed in Eagle Ford and 3,580 in the Permian for the sample.

2.2.3 Revenue: Constructing x_{it} and P_t

If a producer chooses to drill at t, it must construct an expectation of its revenue. In our model of west Texas, the prospects are *ex ante* homogenous, and thus the quantity produced at each time $t + \tau$ are exogenous, conditioned on drilling at t. However, Figure 2.4 shows that a firm will have a strong expectation about the rate at which production will arrive.[20] By fitting a hyperbolic function on single wells for Eagle Ford and the Permian, we find that the production decline rate is consistent over time across basins. Thus, we use this function to predict $x_{it+\tau}$ for values of τ. Initial production volume increases due to technological innovation.[21]

---

[19] Prospective wells are used to fit the real option model using a log-likelihood. The number doesn't impact the model fit but if it is unrealistically high, the expected productivity triggers (x_t^*) required to initiate well investments will be too low over time t.

[20] This figure and analysis are limited to the subsample of leases with only one well on them, because leases are flow metered rather than individual wells.

[21] We divide cost by production gains (bbl per initial production month) available at time t indexed to 2009 levels, which we then use to solve the real option model for the productivity trigger x_t^* well investment activity per month, which is equivalent to increasing expected unobserved productivity

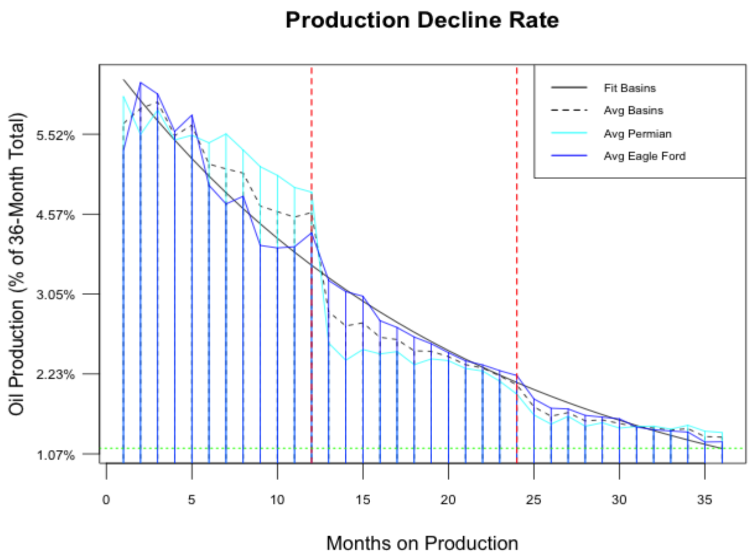

Figure 2.4: Eagle Ford & Permian Decline Rate per Well

As a consequence of the rapid, hyperbolic decline in production, we can establish an accurate revenue expectation based on price expectations within two years of drilling; Figure 2.4 shows that well production in the Permian and Eagle Ford declines by 60% in the first and 30% in the second year, so although the remaining years are not negligible, they don't drive the decision to invest in a well.

Producers can participate in the oil futures market to hedge revenue exposure—take short positions—using 1-24 month futures to lock in time t pricing for delivery at time-$(t+\tau)$[22] for half to two-thirds of their revenue.[23] In fact, shale oil producers can observe market expectations P_t in this market, which is heavily traded out to τ equal to 24-month. These expected futures price come from Bloomberg's CME/NYMEX product database for 2009-2015.

One feature of these markets, conjectured by Samuelson (1965) and empirically sup-

(x_{it}) by these same production gains over time.

[22]Analyzing individual company financial reports, we find that hedging activity is between 0-60%. Producers are incentivized to hedge by banks, or if their price forecast corresponds with that of the market.

[23]Note that the intuition holds, whether we use a real risk-free rate of 1% or weight-average cost of capital (WACC) of 15% to discount expected cash flows.

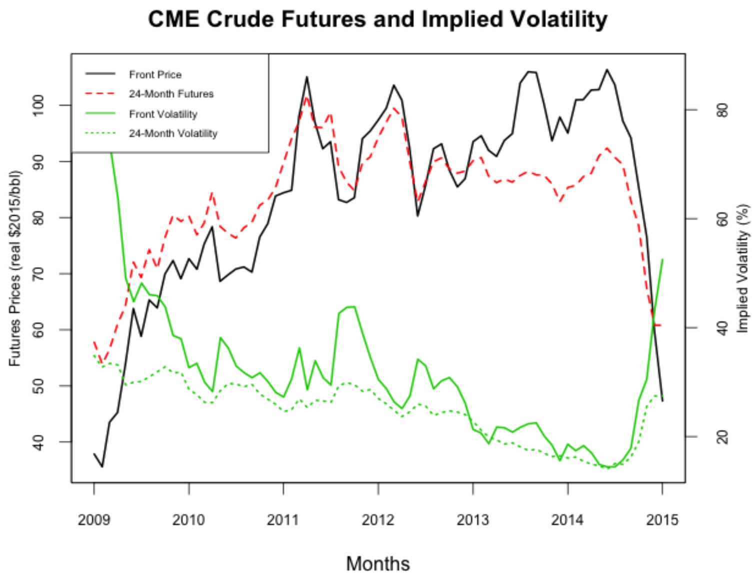

Figure 2.5: Futures Data & Implied Volatility Estimates

ported by Clewlow and Strickland (2000), is that futures price volatility increases as futures contracts approach expiration, when more information is available and uncertainty is removed. Proximity to the delivery date results in more trading, which leads to more volatility. For example, as physical traders price new weather information into the contracts closest to delivery, or reverse out of positions before delivery, volatility increases. Short-term, in particular 1-24 month, futures contracts are consistent with a mean-reverting process of expectations about the future oil price. From 2009 to 2011, we observe front-month mean-reverting tendency, as illustrated in Figure 2.5. When the 1-month oil price, the black line, is lower (at $40/bbl), the 24-month contract, the red line tends, to be higher (at $80/bbl). The reverse holds in 2013-2015, when the 1-month is higher than the 24-month. Specifically, towards the end of the sample, the front-month reverts to the 24-month contract at $60/bbl.[24] Hence, for current revenue, (LHS) of Equa-

[24] A producer's hedging strategy depends on the size of the derivative market position; it can either purchase the full 2-year strip—all the contracts from 1-24 months—the front-month contract and roll the position every month, or the 24-month contract only. Rolling its front month positions can be a costly proposition if adjacent months go up significantly before the rolls are put in place and the front-month expires. Although these hedging strategies differ because of the mean-reverting dynamics, all three options offer an empirically similar *ex ante* revenue expectation.

tion (2.5), the 24-month futures price (P_t) is an adequate proxy of the full 1-24 month strip. That is, we use the 24-month contract as an average for 1-24 month maturities and producer current price expectations because of the mean-reverting nature of volatility in crude oil markets.[25]

2.2.4 Revenue Volatility: Constructing $\sigma_P(t)$ and γ

We deduce producer current expected 24-month oil futures price volatility ($\sigma_P(t)$) from the 24-month volatility implied by the CME/NYMEX traded futures options prices. Arguably, if this oil futures options exchange is efficient,[26] then all the most current forward information is priced in at every period t.

Oil futures options are contracts, which give purchasers the option (but not the obligation) to buy or sell barrels of oil forward at the specified price, called the strike price, within a specific period of time τ (the time-to-maturity—TTM). To price an oil futures option, Black's (1976) model[27] is used for a given volatility, τ, strike price, underlying futures price, and risk-free rate—from T-Bills. With the oil futures option's price from derivative markets, we can infer the implied volatility from Black's formula. We then repeat the process, estimating the implied volatility for each τ and strike price; the result is called an implied volatility surface, shown in Appendix A.4.2.[28] As futures options are not as widely traded at TTM 24-month,[29] we estimate the volatility surface for 1-24 TTM to extrapolate missing 24-month data when necessary.

---

[25]Figure 2.5 shows that front-month volatility is higher than the 24-month contract. See Appendix A.4 for a full discussion on estimating the volatility term-structure and its mean-reversion dynamics.

[26]That is, if prices exist for everything, every producer faces the same prices and takes them as given.

[27]For example, it is a closed-form solution used to price a call option (C) on futures contracts using a cumulative probability distribution function for a standardized normal distribution—priced applying the arbitrage principle for a replicating portfolio. It is similar to the class of log-normal forward models epitomized by Black-Scholes except the spot price of the underlying is replaced by a discounted futures price. Furthermore, to recover the implied volatility, we take the observed option market price as given (C^*) and solve for the value of $\sigma_P(t)$ that equates Black's call option formula to the observed market price, such that: implied volatility is the value $\sigma^*_P(t)$ that makes $C^* = C(S, X, r, \sigma^*, T)$, with oil price (S), strike price (X), risk-free rate and time-to-maturity (T).

[28]We separately use a non-parametric interpolation in volatility space across strikes, and to interpolate in TTM, we use a cubic spline interpolation in total implied volatility space. Estimating volatility at every strike and TTM allows us to avoid the assumptions of constant volatility across all TTMs.

[29]See Appendix A.4 for a discussion about futures and future option volume.

To recover 24-month futures volatility belief $\sigma_P(t)$, we use all traded option price information for all strikes at t for 24-month TTM. For completeness, when 24-month option data isn't traded for every (strike) price, we estimate the implied volatilities for the 1-23 month maturities too. This allows us to extrapolate missing 24-month option data using adjacent τ-month contracts. For the sample, when the 24-month implied volatilities are estimated they are smoothed using a non-parametric spline across strikes. See Appendix A.4.2 for a detailed explanation on recovering the implied volatility surface to extrapolate the 24-month implied volatility when data is missing.[30] Briefly, the implied volatilities are estimated using Black's (1976) model on existing option price data. However, when option price data is missing, e.g. for the 24-month maturity, we extrapolate it using Chen's (2011) improved Heston (1993) methodology, and then use Black's (1976) model to estimate option prices and their corresponding implied volatilities.[31]

When extrapolating, instead of assuming that the option's underlying volatility is constant, Heston's stochastic volatility model integrates these long-observed characteristics of the implied volatility surface to vary over τ and strike price. Chen's approach improves the long-term volatility estimations by using a mixed-lognormal volatility extrapolation adjustment to compute the implied volatility at any strike price relatively close to 24-month futures price.[32] If the data isn't directly observable, this approach allows us to estimate the implied volatility of the 24-month futures price accurately using surrounding implied volatility information at different τ and strike price at every time t. The benefit of this approach is that it uses all available futures options price data—time-to-maturity, strike price, calls and puts—as inputs for extrapolating or interpolating[33] the 24-month,

[30]Previously, literature used a parametric realized volatility estimation to extrapolate long-term, 24-month, implied volatility. However, the approach only works if the term-structure of volatility is stable over time, as discussed in Appendix A.4.1 and A.4.2.

[31]Chen's (2011) method is a slight improvement on Heston's (1993) mixed-lognormal extrapolation of long-term implied volatilities, as it allows the extrapolated or interpolated long-term maturity not to be sensitive to short-term at-the-money dynamics.

[32]Mixed-lognormal method allows us to match implied volatilities beyond Heston's model last listed τ (or TTM) or strike price.

[33]When data is missing, we use the Heston-Chen stochastic model to extrapolate 24-month volatility from surrounding τs (or TTMs) and strikes, and interpolate between 24-month strikes. Also, interpolation and extrapolation are used when traded data is missing to construct the volatility surface for any strike price and TTM.

long-term, implied volatility. We emphasize that we update the estimated implied volatility surface, and thereby the 24-month contract dynamics, for every period t. Hence, we now have our 24-month current expected price and its corresponding volatility (P_t and $\sigma_P(t)$[34] on LHS) deduced from derivative markets.

Lastly, beyond 24 months, firms believe volatility $\sigma_P(t)$ will evolve into next month's implied volatility, in part because of a volatility of volatility process γ. There isn't a market to directly estimate volatility of volatility γ; hence we calibrate it on the sample's volatility realization.[35] It is the standard deviation of logged market volatility differences for the period, that is, the standard deviation of $\ln\sigma_P(t+1)$ - $\ln\sigma_P(t)$ over the sample.[36]

2.2.5 Cost: Total Cost per Well and Productivity Gains

Expected total cost per well is the last input we need to estimate our real option model. As stated above, for total cost per well, we use a recent Energy Information Administration (EIA)-commissioned IHS Global Inc. (IHS) study with yearly total cost estimates for 2009-2015, in addition to two surveys with quarterly and monthly data.[37] These surveys provide the main cost items when investing in a well: drilling rental and completion fracking costs.[38] They are the primary items driving variance in total cost over time t. Given yearly total cost per well data, we estimate total cost for each month t by interpolating that year's total cost, blending it with the corresponding monthly drilling and completion ratio indexed to that year. Early in the sample, when only drilling rental data is available, we similarly index drilling cost to that year's total cost per well

[34] We measure current expected price and volatility by averaging the daily traded price and implied volatility for every month t.

[35] We cannot accurately estimate the full distribution of volatility for strike prices distant from the traded 24-month futures price and τs beyond 24-months. Hence, to measure the volatility of volatility, we are unable to recover it from the distribution of volatility since we don't estimate an implied volatility surface for every strike price and τ (TTM).

[36] We estimate a γ of 0.0605 and 0.0740 for Eagle Ford and the Permian, respectively.

[37] Total cost per well integrates learning-by-doing, input productivity/efficiency, and increased well design size.

[38] Although the proportions of drilling rental and completion fracking costs as a function of total costs differs in the Permian and Eagle Ford.

and interpolate monthly data integrating drilling rental changes for that year.[39] When interpolating total cost per well between years, we control for the reduction in number of days a rig is rented and completion equipment used. As shown in Figure A.8 and A.9 in the Appendix A.5.1, rig rental days went from 32 to 15 in Eagle Ford, and 33 to 23 in the Permian, respectively. This approach allows us to measure not only the level of total cost per well, but also recover some of its variance for the sample. For further discussion, see Appendix A.5 and, for an illustration, Figure A.7. When drilling and completion data is quarterly, at the beginning of the sample, we interpolate using polynomial regression fitting.

Once we have total cost at a monthly frequency, we normalize by well productivity gains during the first month post-completion to estimate total cost per barrel. Specifically, we divide total cost per well by production gains (bbl per initial production month) available at time t indexed to 2009 levels.[40] Productivity (bbl) for the first month of a completed well increased seven-fold in Eagle Ford and five in the Permian, as illustrated in Figure A.10 in the Appendix A.5.2. This gives us a normalized total cost reduction per basin commensurate with the technological leaps modeled by Grübler (1990), Benkard (1999), and Gruber (1992) in the aerospace and semiconductor literature. This normalization of total cost per well—climbing down the cost curve over time—is illustrated in Figure A.11 and A.12 for Eagle Ford and the Permian, in the Appendix A.5.2. Given that a producer only observes current total cost per well, we use it as an expectation of current total cost per well (TC_t on LHS).

2.2.6 Expected Profit Function

With small shale oil producers' revenue and cost estimates, we formally identify the expected profit function. It is with the expected profit function and its future evolution that we measure a well project's value, and later fit the decision to drill empirically. The

---

[39] We also use a polynomial regression fit (or spline) to interpolate when completion cost data is missing for the year early in the sample; a discussion is in the Appendix A.5. However, the resulting estimates are not statistically different from each other.

[40] As stated initially in Section 2.2.3, this is equivalent to increasing expected unobserved productivity (x_{it}) by these same production gains over time.

project's real option value is the tradeoff between the decision to drill and complete now or wait to learn more before investing. For month t, total cost per well TC_t is decomposed into drilling and completion costs. First, drilling cost is proxied by both rig rental $d_i R_t$ and non-rental cost rc_i, where d_i is rig days with rental rate R_t.[41] Second, completion costs include fracking F_{it} and non-fracking cost fc_i. With expected current price P_t, total cost per well TC_t, and expected unobserved productivity x_{it}, we construct current profit π_{it} driving the decision to drill and complete a prospective well i at time t. More specifically, profit is a function of the expected 24-month oil price (P_t) times well-expected productivity (x_{it})[42] minus current total (drilling and completion) cost per well (TC_t):

$$\pi_{it} = \pi_i(P_t, TC_t) = x_{it}P_t - TC_t = x_{it}P_t - rc_i - fc_i - d_{it}R_t - F_{it} \tag{2.6}$$

Beyond 24 months, less information is traded and volume drops off—this results in a wide bid/ask spread.[43] We argue that the lack of traded information beyond then signifies that time-varying revenue expectations are modeled using a first-order Markov chain evolution, with random walk specification,[44] which we also use for expected current Permian price and price volatility evolution. Furthermore, total cost per well doesn't have a futures market and depends on oil's price evolution over time t. Hence, for total cost per well, we apply a first-order Markov chain evolution where cost volatility is scaled to crude volatility and cost shocks are correlated to crude oil standard normal shocks.[45]

As stated earlier and shown in Equations (2.1)-(2.3), as of $t+1$ the future expected evolution of profits is measured using memoryless Markov-chains[46] for price, volatility,

[41] See Appendix A.5 for a discussion on estimating total cost per well.

[42] $x_{it}P_t$: discounted (δ) expected well productivity times average state price. For simplicity and tractability, the cash-flow stream is aggregated into a single number $x_{it}P_t$, where expected price (P_t) is the 24-month futures price.

[43] Futures are more widely traded beyond 24 months than options; see Appendix A.4 for a discussion on futures and options traded volume. It is then option traded volume that limits expectations of current price (P_t) and volatility $(\sigma_P(t))$ to the 24-month maturity.

[44] See Appendix A.4 for an explanation of traded volume in the crude oil derivative market.

[45] The total cost crude oil scalar α ,used to link price to cost volatility, is equal to 1.696, and 1.10 when cost is normalized for productivity, while the correlation ρ between standard errors is equal to 0.20, and -0.19 when cost is normalized for productivity.

[46] With random walk drift, the mean-reverting drift specification didn't statistically change the results. See Appendix A.9 for reasons why a mean-reverting drift is theoretically interesting.

and total cost. Using the above specifications for current profit function (LHS) and expected future evolution (RHS) from Equation (2.5), we theoretically solve for the optimal productivity cutoffs x_t^*. All producers solve for the same cutoffs for all prospects at any given time t, as they face the same price, volatility, and total cost. In the following section, we introduce the volatility identification parameter β,[47] which helps test a producer's degree of sophistication when investing by determining the appropriate productivity cutoff x_t^* each period.

2.3 Model Empirical Fit

As discussed throughout, we model shale oil producers using a single-agent, risk-neutral and price-taker framework.[48] At every period t, a producer decides whether to drill and complete a prospect i or wait. In our empirical model, monthly aggregate well count is a function of the expected production threshold (x_t^*), which is conditional on previously not having drilled and completed a prospect. We take the net present value of a well's monthly expected production profile, using the firm's discount rate ($\delta = 0.89$) net state and federal withholding tax and lease royalty interest. We then sum discounted expected unobserved productivity and divide it by the mean total cost for the sample[49] and denote it as x_{it}.

2.3.1 β Identification ($x_{it} \geq x_t^*$)

In Section 2.2, we specified the state P_t, $\sigma_P(t)$ and TC_t transition functions, which we use to estimate the responsiveness to time-varying volatility (β), and expected unobserved productivity (x_{it}) of each prospect over time. First, we discuss the identification of volatility responsiveness using β sensitivity using Equation (2.7):

[47] In the theoretical specification, β is set to one; going forward we empirically estimate this β value.

[48] All shale oil firms are price-takers, as energy markets are competitive, with no single shale oil producer able to manipulate price and representing only a small percentage of worldwide oil production. Credit is lax during the sample period, with low expected corporate defaults and a slack drilling and completion service sector supply, with the only exception of a brief period in 2012. During oil production a producer's revenue is exposed to crude oil price and transportation basis risk. Furthermore, our agent is risk-neutral as he has the option to hedge risk-aversion by participating in derivative markets and signing take-or-pay agreements for price and transportation basis risk, respectively.

[49] Note that for the sample, Eagle Ford's mean total cost is \$8.2M and the Permian's is \$6.8M.

$$\ln \sigma_P(t) = \ln \bar{\sigma} + \beta(\ln \sigma_P^{Mkt}(t) - \ln \bar{\sigma}) \tag{2.7}$$

this β identification allows us to transform our volatility diffusion process from Equation (2.2) into an identification procedure. With Equation (2.8), we estimate a producer's degree of sophistication by measuring how producers believe volatility evolves over time t:

$$\ln \sigma_P^{Mkt}(t+1) = \ln \sigma_P^{Mkt}(t) - \gamma_t^2/2 + \gamma_t \epsilon_{\sigma_{t+1}} \tag{2.8}$$

the β behavioral parameter regulates the degree of firm responsiveness to changes in $\sigma_P^{Mkt}(t)$.[50]

The parameterization of the expected evolution of volatility, in Equation (2.8), allows us to measure the sophistication of a small firm by testing whether a shale oil producer's investment project heuristically mimics that of an optimal decision-making process. In particular, we estimate whether myopic constant or market time-varying uncertainty yields a better fit of firm investment behavior.[51] This decision to invest in a well is repeated at every time t using either a current expected myopic constant or market time-varying volatility for the sample. Furthermore, at t the chosen current expected volatility controls the evolution of volatility beliefs over τ as driven by either a constant or time-varying Markov process.

The estimated β value gauges a producer's degree of sophistication when making a well investment decision. Specifically, the value of β can identify both myopic constant ($\beta \neq 1$) and market time-varying ($\beta = 1$) beliefs. A β that empirically converges on 1 indicates that a firm optimizes using market time-varying volatility, while a β converging on 0 shows the usage of myopic restricted constant volatility. Volatility beliefs that deviate from time-varying market beliefs show small firm failure to sophisticatedly specify and coordinate investment decisions; see Appendix A.7 for real option value.

[50]Using this specification of the volatility evolution process on market volatility $\sigma_P^{Mkt}(t)$, we see that beliefs about the evolution of expected volatility σ_P are scaled by β such that we have $\beta\gamma$. Unit root cannot be rejected, random walk drift $\tilde{\mu}(\sigma_P(t)) = 0$).

[51]The idea is that if a real option value exists, a small firm incorporates it. However, it may either use a myopic (unsophisticated) or time-varying (sophisticated) measure of volatility.

To fit the real option model, we use firm responsiveness to market time-varying volatility as a free parameter—using β. By approximating the current volatility $\sigma_P(t)$ with β, we price the embedded value of the diffusion process in Equation (2.5) and determine the corresponding productivity cutoff x_t^* at each period. Second, with β we solve for the expected unobserved productivity x_{it} that puts a producer over the estimated productivity trigger x_t^* to invest in a well: we want the expected number of prospects that best fits $x_{it} \geq x_t^*$.

2.3.2 *Empirical Model Estimation (β, μ and ζ)*

Our empirical specification of the model is such that we want to fit the dependent variable number of prospects drilled and undrilled, given behavioral parameter β and the mean μ and variance ζ of the log of expected productivity $\log(x_{it})$ across both time t and prospects i—where μ is the average productivity and ζ controls for the variance of productivity across prospects and time.[52] Both of these parameters are estimated in the main model with β.[53,54] We construct a contemporaneous and temporal distribution around $\log(x_{it})$ to account for geological and engineering errors, and drilling and completion service supply delays. If we do not construct a distribution around x_{it}, then all prospects will be completed in an environment when prices, cost, and volatility are advantageous, while in the following period, which might be incrementally less advantageous, no prospects would be completed.

To predict the number of wells that will be drilled at any time t, we fit the real option model by estimating the β parameter using a maximum log-likelihood empirical fit on actual data with a Gaussian distribution with expected unobserved productivity mean μ

[52]We assume that the log of productivity is i.i.d. normal across prospects i and time t.

[53]The productivity distribution allows us to better empirically fit the model by allowing firm specific effects in the decision to invest in a well across prospects and time. If firms were identical, they'd invest at the same time, and if drilling decisions weren't distributed over time, they'd only drill in auspicious times and withhold otherwise. For example, if we do not account for geological and engineering errors, and drilling and completion service supply delays—stochastic lags between the decision to drill and final completion—then all prospects are completed in an environment when prices, cost, and volatility are advantageous, while in the following period, with incrementally less advantageous states, no prospects are completed.

[54]A more detailed explanation of firm volatility behavior identification is in Appendix A.2.

and variance ζ.[55] Given a choice set of prospective wells for 2009-2015 (N=24,316), we sum the probability that a well will be drilled, conditional on it not being drilled before and the probability that it will remain undrilled for the sample (N_T=10,383).[56,57]

The reported log-likelihood (LL) fit in Section 2.4's Table 2.1 and Section 3.3's Table 3.1 measures the probability that x_{it} exceeds the trigger productivity x_t^*, given time-varying realized state values (P, TC, σ), probabilities of transitioning between these realized price and cost states, and behavioral fitting parameters (β, μ, ζ).

$$\begin{aligned}\log \mathcal{L}\left(\beta, \mu, \zeta \mid P, TC, \sigma\right) = \sum_{t=1}^{T} \left[N_t \log \operatorname{Prob}\left(I_{it} = 1 \mid P, TC, \sigma; \beta, \mu, \zeta\right)\right] \\ + N_T \log \operatorname{Prob}\left(I_{it} = 0 \;\forall t \mid P, TC, \sigma; \beta, \mu, \zeta\right)\end{aligned} \tag{2.9}$$

I_{it} is equal to 1 if prospect i is drilled in month t and is 0 otherwise.
N_t is the number of wells actually drilled at t.
N_T is the number of prospects not drilled for the whole period T.
In Equation (2.9), we write the probability that a prospect is drilled at $t = 1$ as the conditional probability that it is drilled at $t = 1$ multiplied by the probability that it was not drilled at $t = 0$. The objective is to fit an estimate of β, μ and ζ that maximizes LL(β) by locating the optimal $\hat{\beta}$. Numerically, the maximum can be found at the point where there is no further increase for a predefined tolerance.[58] For a more complete visualizing of the algorithm's estimation procedure used to solve and fit the model, see appendix section A.10.

[55] We find the values of β, μ and ζ that maximize the log-likelihood.

[56] The idea is to measure the probability that x_{it} exceeds trigger productivity given our three state transition processes.

[57] When numerically solving the problem, as shown in Appendix A.10, μ and ζ are measured in the primary outer loop process, while β market volatility responsiveness is estimated in the inner loop. Appendix A.8 illustrates the empirical fitting of the log-likelihood of the productivity trigger over actual data on the number of prospects and wells drilled and completed.

[58] Each iteration estimates a new $\hat{\beta}$ and the log-likelihood value, which is higher than at the previous value until the desired tolerance between steps in achieved. Additionally, we apply the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method to approximate an arc-Hessian (Train (2009)) as the log-likelihood function can be non-quadratic. The method uses information at more than one point on the likelihood function to estimate the curvature.

2.4 Results and Discussion

2.4.1 Eagle Ford Firm Volatility Response

We fit the real option model for Eagle Ford, where mean total cost per well is $8.2M and the average expected productivity is estimated at 189,875 barrels (bbl) for the sample. The parameter of interest is a shale oil producer's response to time-varying volatility β.[59]

Table 2.1: Real Option Results for Eagle Ford

	Constant Volatility (1)	Time-Varying Volatility (2)
β (Volatility Responsiveness)	–	1.08
	–	(0.12)
μ (log expected productivity)	−5.97	− 6.79
	(2.01)	(2.40)
ζ (std of μ)	1.91	2.81
	(1.02)	(1.51)
Log Likelihood	−6,371.98	−6,362.99

Notes: μ is expressed as expected oil production (in bbl) divided by the total cost of drilling and completion. All estimates use 24 month futures prices and option volatilities. A random walk forecast is used for future volatility. Drilling and completion data are matched to drilling and completion likelihoods with a 2.4-month lag. Also, standard errors are clustered at the field level and corrected for auxiliary parameters.

In the restricted model (1) from Table 2.1, β is set equal to zero: firms do not respond to time-varying market volatility and instead apply constant volatility in their decision-making process. The resulting fit is displayed in Figure 2.6.[60] Both the distributional parameters μ and ζ are statistically significant. This model specification is basic and simpler than is expected to be used by industry participants. However, it provides a good baseline comparison to measure whether firms heuristically mimic an optimal decision process. Note that discounted expected productivity x_{it} is reported as production

[59] For Eagle Ford we motivate the state space and correlation structure by using a reduced form of the model as descriptive statistics. The results for this analysis are in Appendix A.6.

[60] For illustrative purposes, we use the first year average volatility in the sample as the constant volatility measure.

(barrels) per average total cost ($100,000).

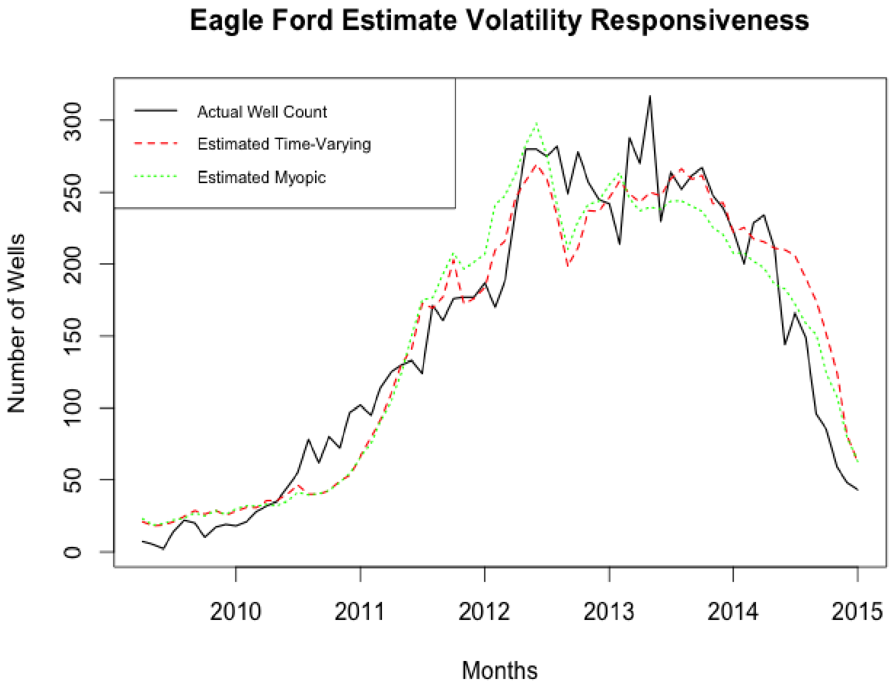

Figure 2.6: Eagle Ford Myopic $\beta = 0$ & Time-Varying $\beta = 1.08$

In the unrestricted model (2) from Table 2.1, the estimated behavioral fitting parameters β, μ, and ζ are all statistically significant with a β that is not statistically different from one. The reported log-likelihood (LL) fit measures the probability that discount well production x_{it} exceeds the solved trigger productivity given realized time-varying state values (P, D, σ). From Table 2.1, when we estimate β using market time-varying beliefs, $\beta = 1.08$, which in log-likelihood terms is a better fit than constant volatility.[61] As illustrated in Figure 2.6, a real option with time-varying volatility results in a better fit for most of the sample (−6,371.98 versus −6,362.99). This is particularly true starting in 2013, when expected volatility sharply declined and prices spiked (see Figure 2.5), and a constant volatility—average of high first year volatility—fit overestimates the wait on well investments. Thus, model (2) is a better fit and yields a β estimate statistically

[61]The likelihood ratio test rejects, with a p-value less than 1%, a null hypothesis that a firm does not respond at all to time-varying volatility $\beta = 0$.

different than zero, with standard errors that make it approximately equal to one. Hence, we find that a small shale oil producer employs a sophisticated decision-making process mimicking that of a real option approach.

In Figure 2.6, we notice that the time-varying volatility specification allows for a better overall fit of the data, particularly from 2012 to 2014, when myopic constant volatility is too high and is most strongly identified. Towards the end of the sample, the fit isn't as good. First, this can be explained by a short-term increase in wait on rig time: in late 2012 35% of operators reported, waiting longer than the previous three months. Second, around 2012-2013, leases signed in 2009, which are signed for an average of 3-5 years, began to expire. If the lease contract stipulates that a well must be drilled in 3-5 years, producers are incentivized to preemptively invest in a well to maintain or renew their leases. Lastly, given the sharp drop in prices and spike in volatility at the tail end of the sample, the decision to invest and complete may need more complex identification, which we simplified away for the sample by controlling for corporate structure (debt/equity), bank credit lines, and quality of the prospective well portfolio—high-grading.

The current results are important in that they indicate that small firms mimic a real option model with a time-varying implied volatility—from markets—used to initiate the future evolution of expected prices, volatility, and total cost. However, the failure to account for time-varying volatility leads to unsophisticated investments in uneconomical wells that don't properly account for the real option value of the project. Therefore, in a volatile and uncertain environment, these Eagle Ford high stakes producers are sensitive to volatility and do in fact use time-varying market-based expectations of volatility to coordinate their investment decisions.

Chapter 3

TRANSPORTATION BASIS RISK

3.1 Transportation Basis

By managing price and transportation basis, a producer can make well investment decisions that integrate risk from their adverse movements. Unlike price or market risk, however, there isn't a forward market in transportation basis. Therefore, to measure producer futures transportation basis, we simulate basis paths forward and take expectations. When modeling the spot transportation basis time series and its characteristics, we focus on price jumps and reversion to a long-term mean as a result of congestion and its relief. Then, when estimating forward basis we use the positive correlation structure between spot and futures prices, as well as the empirically well-documented finding that the volatility term structure of futures prices is a decreasing function of maturity.[1] Transportation basis risk comes from the difference in the delivery location of a barrel of crude oil and the market price contract.

Permian producers can hedge price risk by locking-in CME/NYME WTI (Cushing) futures. However, they are still exposed to the basis risk between Cushing—the pricing point for CME/NYMEX WTI oil—and the Permian spot market: with increased domestic supply transportation costs, from the Cushing to the Permian delivery point, were more susceptible to discontinuous increases. Throughout our sample, transportation cost for Eagle Ford was stable. Its proximity to the Gulf Coast offers multiple cheap transportation alternatives such as additional pipeline take-away capacity and freight—a competitive transport mode for short distance. As for the Permian, hedging by signing a take-or-pay agreement can reduce the transportation basis risk but does not eliminate it completely.

[1] The methodology for estimating forward transportation basis is well documented and discussed in the Appendix B.1.2.

Scant empirical work has been conducted to model and measure small firms' investment behavior under transportation cost uncertainty. Based on Schwartz's (1997) work on time series that mean revert and spike, we apply a mean-reverting jump diffusion process with term-structure dynamics, controlling for price levels—the higher the price, the higher likelihood of congestion—and transportation basis risk contribution as they affect investment decisions.[2] This approach is supported by Clewlow and Strickland (2000), who view firm management of enterprise risk as the systematic process of identifying, measuring, and incorporating primary risks across a company's operations to make informed managerial energy investment decisions.

Industry experts point out that typically, an upstream shale oil energy firm is exposed to midstream transportation risk through injection and withdrawal rates, as well as capacity and flow constraints. With access to three alternative modes of transportation, that risk can be reduced but not eliminated as their costs differ. Pipeline is the cheapest option, followed by rail and then freight transportation. As the primary transportation option is exhausted, a costlier alternative may then be used. As mentioned above, such a firm can hedge its price and transportation basis risk.[3] These contracts stipulate that, for a fixed tariff, a producer is responsible for supplying a predetermined amount of volume or pays a penalty; excess capacity can usually be sold on a secondary market.

3.2 Revenue Integrating Transportation Basis and Volatility

The movement of oil from the Permian Delaware and Midland regions to Cushing is constrained by pipeline take-away capacity. A producer can partially hedge its transportation basis by signing a pipeline agreement. The contract stipulates that, for a fixed tariff, a producer is responsible for supplying a predetermined number of barrels or pay a penalty. Producers usually hedge forward no more than 70% by signing three to nine year contract

[2] In Appendix B.1.2, we simulate Permian transportation basis level and volatility. First, we estimate the mean reversion rate, α, and long-term mean when it is above and below the sample mean, μ_t, using a linear regression. Second, we measure our jump-diffusion parameters using a recursive process until volatility estimates converge using a tolerance requirement, mainly the standard deviation of jump returns, frequency of jumps, and standard deviation of non-jump returns.

[3] Industry experts contend that pipeline capacity is modeled as a stair-step function where rail transport is used to smooth out transitions from step-to-step. In fact, if a pipeline gets congested there are two, more expensive, alternative modes of transportation: rail and trucks.

terms.[4] They can also forgo hedging altogether by relying exclusively on excess pipeline capacity sold on the spot market—resale is conducted on a secondary market. However, as demand for pipeline capacity outstrips supply, more expensive modes of transportation, such as rail and freight, are used. In our sample, usage of alternative, more costly modes of transportation leads to a spike (or jump[5]) in (spot) transportation basis—the spot price differential between the production (Permian) and delivery point (Cushing)—that results in Permian crude selling at a severe discount to the WTI, Cushing price.

As discussed above, we proxy producer expected revenue by price and price volatility over two years using a 24-month futures and futures options contract. In the Permian, however, transportation basis is an additional input into our real option well investment model. Over time t, we adjust the current expected 24-month market price (WTI, Cushing) by transportation basis to infer the expected current 24-month price observed by producers in the Permian. As the primary transportation option is exhausted, the costlier alternatives may then be used. To estimate current expected transportation basis dynamics, we use a mean-reverting jump diffusion model that computes the reversion characteristics of transportation basis to a long-term mean after an initial congestion generated jump. See Appendix B.1.2 for a full discussion of the econometric model.

To estimate current Permian revenue, there is no expected current 24-month transportation basis, or forward Permian price derivative market, and we don't have access to take-or-pay agreements. Therefore, we estimate what producers know by measuring the expected Permian Midland price ($\mathrm{E}(P_t^{PTB})$) in Equation (3.1)—WTI Cushing futures price (P_t) adjusted by expected transportation basis ($\mathrm{E}(TB_t)$).

We also integrate an additional source of volatility derived from our estimate of current expected transportation basis. That is, we integrate transportation basis risk ($\sigma^2_{TB}(t)$) into expected time-varying volatility to give us current expected price and volatility estimates ($\sigma^2_{PTB}(t)$ in Equation (3.2)). For this we use a calibrated correlation measure

[4] The length of the contract is a function of the amount of capacity being hedged. Habitually, shorter term contracts are offered when less capacity is expected to be used. Beyond 70% producers are exposed to excess capacity commitments, in which they may end up paying for unused capacity.

[5] Modeled with a Poisson process. The jump process $\mathrm{d}q$ is assumed to be independent of the continuous stochastic increment dz, see Appendix B.1.2 for notation and estimation.

($\rho_{TB,P}$), estimated on realized price and transportation basis with and without jumps.[6]

$$\mathrm{E}(P_t^{PTB}) = P_t - E(TB_t) \tag{3.1}$$

where P_t is the 24-month market price, while the PTB superscript is short for price and transportation basis. E(TB_t) is the expected current 24-month transportation basis from simulating forward a mean-reverting jump diffusion model; for the specifics of the econometric model see Appendix B.1.2. The process is first calibrated on spot transportation basis dynamics, and second simulated forward—capturing the correlation structure between spot and 24-month contracts—to estimate the expected current 24-month transportation basis. With a measure of 24-month transportation basis, we estimate the adjusted expected current 24-month price observed by Permian producers—estimated by discounting the 24-month WTI Cushing price.

To measure Permian producer expected current 24-month volatility, we first compute price (P_t) and transportation basis (E(TB_t)) contribution to revenue exposure as w_P and w_{TB}, and their correlation ($\rho_{TB,P}$)[7] with and without jumps. Second, we use the realized standard volatility.[8] A producer's belief about current 24-month complex volatility ($\sigma_{PTB}(t)$) is then measured by combining futures option implied volatility and estimated expected transportation basis risk[9]:

$$\sigma_{PTB}^2(t) = \sigma_{TB}^2(t)w_{TB}^2 + \sigma_P^2(t)w_P^2 + 2w_{TB}w_P\sigma_P(t)\sigma_{TB}(t)\rho_{TB,P} \tag{3.2}$$

where $\rho_{TB,P}$ is the correlation coefficient between transportation basis (TB_t) and market

[6]In the Appendix we first motivate the integration of transportation basis (Appendix B.1.1). Second, we show the general mean-reverting jump diffusion model to be used to measure spot transportation basis dynamics (Appendix B.1.2). Third, we illustrate the estimation of mean-reversion and jumps (Appendix B.1.3). Finally, we walk through the extrapolation of spot transportation basis time series to compute the corresponding $E(TB_t)$ time series (Appendix B.1.4).

[7]Correlation of the price and basis returns for the period. As basis increases, Permian price decreases; therefore towards the end of the sample the correlation structure between expected Permian and transportation basis returns was more negative with a $\rho_{TB,P} = -0.35$.

[8]Basis volatility is estimated as the standard deviation of $\ln TB_{t+1}$ - $\ln TB_t$ with and without jumps for the sample.

[9]The bivariate relationship between price and transportation basis returns is approximately normal, with a slightly negative correlation structure for the sample. If their relationship had high kurtosis and/or skewness we'd apply a bivariate Gaussian copula to control for the non-normal marginal distributions.

price (P_t). w_{TB} and w_P are basis and price component weights, i.e. producer proportion of transportation basis and price and their contribution to revenue exposure.

In summary, we first calibrate a mean-reverting jump diffusion model on spot transportation basis data.[10] Second, we extrapolate—using simulations—current expected 24-month transportation basis controlling for the correlation structure between spot and 24-month contracts, as well as new expected pipeline capacity. Third, we estimate the expected current 24-month Permian price by discounting the 24-month WTI (Cushing) futures contract by our measure of expected current 24-month transportation basis (Equation 3.1). This result is illustrated in Figure 3.1. Lastly, we apply a mixed-lognormal distribution (Equation 3.2), and calibrated time-varying portfolio weights—percent of contribution between price and transportation basis—and the correlation structure to compute expected 24-month complex volatility—dual source price and transportation basis volatility.[11]

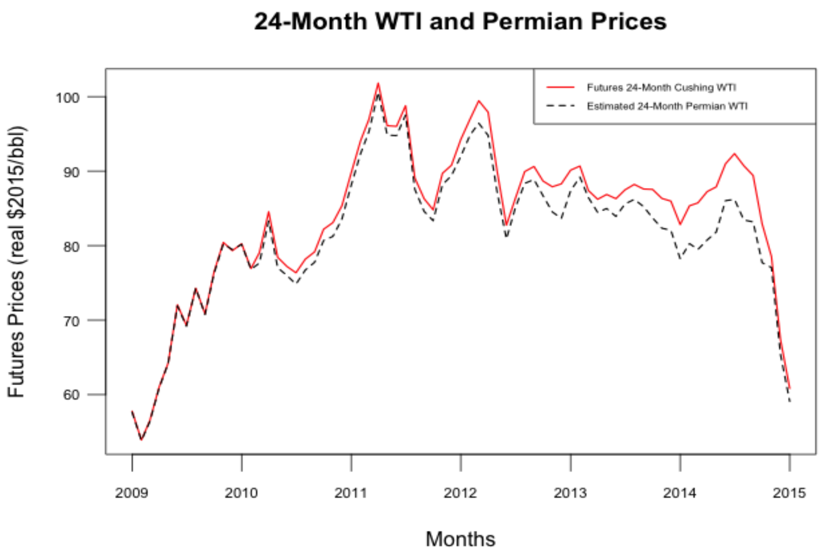

Figure 3.1: Basis Adjusted 24-Month Permian Price

[10]A discussion of the mean-reverting jump diffusion model and its calibration are in Appendix B.1.2 and B.1.3.

[11]Note that, unsurprisingly, portfolio transportation basis weights increase in tandem with transportation congestion.

Estimated expected 24-month Permian Midland price and transportation basis are illustrated in Figure 3.1 and, in the Appendix B.1.4, Figure B.5, respectively. The majority of transportation basis jumps occur towards the end of the sample, when high expected oil futures price lead to increased production, overwhelming the existing pipelines' takeaway infrastructure. We notice that in 2009, the likelihood of congestion was low. However as prices and production increased over the sample period, the likelihood of congestion went up and in 2010, 2012, 2013, and 2014 jumps in expected transportation basis ranged from \$8/bbl to \$1/bbl, with an expected reversion to a long-term mean of \$4.93/bbl when $P_t > \bar{P}$ and \$2.52/bbl when $P_t < \bar{P}$.[12] For example, in 2014, as shown in Figure 3.1, the 24-month WTI and 24-month Permian WTI price traded at \$87.26 and \$80.76, respectively. We now have an approximate measure of a Permian producer's expected current Permian price and volatility (P_t^{PTB} and $\sigma_{PTB}(t)$ on LHS).

As before, we argue that once the expected Permian price and volatility are estimated, the lack of traded information beyond 24-months allows us to model multi-risk (or complex) time-varying revenue expectations using a first-order Markov chain evolution. In particular, the belief about current 24-month complex volatility is estimated by combining futures options implied volatility and estimated current expected transportation basis risk.[13] Once current 24-month expected complex volatility is measured, we use a time-varying first-order Markov chain to estimate the expected evolution of the Permian price to value the real option. By integrating 24-month Permian price level, volatility, and their expected evolution, we more accurately mimic Permian producer response to sophisticated levels of risk.

3.3 Permian Firm Complex Volatility Response

We fit the real option model for the Permian, where mean total cost per well is \$6.8M, and the average expected productivity is estimated at 82,831 barrels (bbl) for the sample. The parameter of interest is a shale oil producer's response to complex time-

[12] The estimation of long-term means are in the Appendix B.1.2. Note that $\bar{P}$ is the sample mean.

[13] The bivariate relationship between price and transportation basis returns is approximately normal. If their relationship had high kurtosis and/or skewness we'd apply a bivariate Gaussian copula to control for the non-normal marginal distributions.

varying—integrating price and transportation basis—volatility β. We find that a better fit to the myopic model is the complex time-varying specification—even if there isn't a clear statistical difference between time-varying and complex volatility. Hence, investments in the Permian reflect an understanding of complex risk, multi-source time-varying volatility.

Table 3.1: Real Option Results for Permian

	Constant (1)	Time-Varying (2)	Complex (3)
β (Volatility Responsiveness)	–	0.93	1.04
	–	(0.169)	(0.14)
μ (log expected productivity)	−5.14	− 6.26	− 6.30
	(2.14)	(3.21)	(3.33)
ζ (std of μ)	1.79	2.01	2.04
	(0.98)	(2.25)	(2.63)
Log Likelihood	−2,173.04	−2,169.56	−2,168.55

Notes: μ is expressed as expected oil production (in bbl) divided by the total cost of drilling and completion. All estimates use 24 month futures prices and option volatilities. A random walk forecast is used for future volatility. Drilling and completion data are matched to drilling and completion likelihoods with a 2-month lag. Also, standard errors are clustered at the field level and corrected for auxiliary parameters.

In the restricted model (1) from Table 3.1, β is set to zero: producers do not respond to time-varying market volatility and instead apply constant volatility in their decision-making process. The resulting fit is displayed in Figure B.7 in Appendix B.2.[14] Both the distributional parameters μ and ζ are statistically significant. Note that discounted expected productivity x_{it} is reported as production (barrels) per average total cost ($100,000).

Model (2) uses time-varying price volatility, while the complex model (3) integrates both time-varying price and transportation volatility. In the unrestricted models (2) and (3) from Table 3.1, the estimated behavioral fitting parameters β, μ, and ζ are all statistically significant with βs equal to one. In Table 3.1, we estimate β using complex

[14]For illustrative purposes, we use the first year average volatility in the sample as the constant volatility measure.

time-varying beliefs where $\beta = 1.04$, which has a better fit log-likelihood measure than constant volatility. On the other hand, the market time-varying volatility model where $\beta = 0.93$ also yields a better fit than the restricted myopic model. Although the complex (3) specification has higher LL than the time-varying one ($-2{,}168.55$ versus $-2{,}173.04$), we cannot conclude with more than 10% significance that this is statistically better.[15] Thus, model (3) is a better fit and yields a β estimate statistically different than zero, with standard errors that make it approximately equal to one. Hence, we find that a small shale oil producer mimics a sophisticated decision-making process.

The current results show that in, the Permian, small firm response to complex risk is consistent with an optimal decision process. Failure to respond to price and transportation basis risk leads to uneconomical well investments that don't properly value the real option value.[16] That is, we find that failure to respond to both price and transportation basis exposure by identifying and integrating them into firm-wide decision-making leads to a misestimation of real option value. In a volatile and uncertain environment with transportation exposure, small leveraged producers do in fact have a sophisticated response to complex uncertainty, necessary to coordinate their well investment decisions.

3.4 Conclusion

In parametrizing the expected evolution of volatility, we estimate the sophistication of a small firm by testing whether a shale oil producer's investment project heuristically mimics that of an optimal decision-making process. In fact, we test whether myopic constant or market time-varying uncertainty yields a better fit of firm investment behavior. This approach allows us to measure whether a small firm's response is sophisticated and consistent with an optimal decision-making process that integrates commonly held beliefs about prices an transportation basis inputs.We find that small firms account for forward time-varying volatility when making irreversible investment decisions. By de-

[15] After identifying $\ln \sigma_P(t) = \ln \sigma_P^{Mkt}(t) + \beta(\ln \sigma^{PTB}(t) - \ln \sigma_P^{Mkt}(t))$, with $\sigma_P^{Mkt}(t)$ identification, a log-likelihood ratio test concludes at the 10% level that complex risk is better than time-varying market risk. Once complex risk is parameterized and the model is nested we draw a statistical comparison between model (2) and (3).

[16] See Appendix B.2 for a detailed discussion about model specifications, and Appendix A.7 for the consequences of investing using an inaccurate measure of volatility.

ducing beliefs from futures and futures options markets we show that failure to measure the real option value leads to uneconomical well investment decision-making. Furthermore, the integration of both price and transportation basis risk deduced from derivative markets and forward-contracts—such as take-or-pay agreement—with a more complex distribution yields a better fit of investment behavior. Thus, our Texas case study shows that small firm drilling and completion activity from 2009-2015 in the highly productive shale oil fields—Permian and Eagle Ford—has a sophisticated response to uncertainty by integrating time-varying price and transportation basis risk.

3.4.1 Discussion

Small firms within our sample hedge between 0% to 60% of their price exposure. We are interested in what is driving this heterogeneous hedging behavior. Are differing input price forecasts reducing hedging activity and therefore a form of speculation? In this thesis, we show that the setting of common beliefs using all available information from futures, options, and over-the-counter markets enables a small firm to respond sophisticatedly to uncertainty using time-varying expectation of volatility. In the current literature, it's unclear how a small firm incorporates information on the price of its major inputs: does it use information acquired through daily business operations or coordinate supply chain activity using common market beliefs? Given the high upfront cost of acquiring technical expertise and software, firms may put off deducing beliefs from forward-looking financial markets and contracts. However, we find that whether or not these firms' risk forecasts differ from those of the market, they will set common beliefs and we conjecture that they will then selectively hedge and speculate based on firm-specific variance reduction preferences. Hence, the role of futures and options markets, as a coordinator of investment strategies and not necessarily an active-management tool, should be analyzed further.

In the current sample, the decision to drill and complete is a single decision and driven by three factors: sharp production decline curves, technological innovation, and volatility swings. Furthermore, the analysis was conducted in an environment without common-pool dynamics, where homogeneous firms are price takers—with no single shale oil producer able to manipulate price—credit is lax, there are low expected corporate defaults, and drilling and completion supply is slack. Hence, the study encompasses a volatile environment absent corporate structure effects and bank credit constraints. Currently, however, new crude curve dynamics decoupled the decision to drill and complete, with producers reserving the right to complete for a later date. Expected increases in cost translate into preemptive drilling, while high short-term expected prices result in completion delays. Additionally, the drop in oil prices from \$100 to \$40 increased corporate defaults with company downgrades by Fitch, Standard and Poor, and Moody's. Hence, a natural extension of this work is to model the small firm responsiveness given cheap oil, capital market tightening, equity and debt corporate structures, and new high grading strategies. Finally, towards the end of the sample, costly-lower-tail outcomes became likelier and resulted in financial distress and bankruptcy. As of 2015-2016 small firms are unable to carry out their well investment strategies and make interest rate payments. To model this high-order risk we need to extract the full distribution of future expected prices; we will then be able to estimate extreme events that go unaccounted for in a risk-neutral valuation. This work will serve to explain small firm response to uncertainty in a distressed market and estimate the role of corporate structure and bank lending.

Chapter 4

ENERGY POLICY: COST OF THE CRUDE EXPORT BAN

4.1 Unconventional Oil Response in Texas

In August 2013, increased US shale oil production made the 1975 export ban binding and distorted Texas horizontal well investment. The ban prevents the export of crude but not crude product. This jump in production was driven by progress in drilling and completion technology, which allowed the US to extract large volumes of shale oil from tight rocks as of 2009, thereby altering the face of global energy geopolitics and, in particular, the economics of US energy policy. From 2009-2015, with heightened production, US shale oil producers became the marginal price setters. Hence, to model current crude production dynamics and the full impact of the distortionary export ban policy,[1] we identify and estimate the supply response (elasticity) of shale oil producers.

4.2 Counterfactual and Policy Valuation Data

We use drilling and completion, and production data for 2009-2015 from the Texas Railroad Commission (TRRC). In particular, production comes from the "Oil and Gas Annuals" dataset and is used to estimate both well productivity and displaced number of barrels from the export ban as of August 2013. To estimate total displaced production during the ban, we multiply the average first month well production (bbl) profile for the corresponding month by the consistent well production decline rate.

[1] As of 2015, the ban was ineffective at depreciated prices of \$30-\$40; on December 18$^{\text{th}}$ 2015 it was removed. Hence, for our Jan 2009-Jan 2015 sample, we estimate the full impact of the ban on well investments.

4.2.1 Revenue: Counterfactual Price and Price Volatility

Expected prices are directly taken from futures markets, while expected price volatility is implied from futures option price data using the Black (1976) model and methods that are widely used by practitioners (Heston (1993); Chen (2011)). These derivative markets serve as an efficient aggregator of beliefs used to hedge[2] and coordinate downstream, midstream and upstream energy investment activity. Our WTI and Brent futures and options price data come from Bloomberg's Chicago Mercantile Exchange (CME) and Intercontinental Exchange (ICE) product database for 2009-2015, which provides us with both domestic and international expected future oil prices and allows us to estimate the corresponding implied volatilities.[3]

4.2.2 Cost: Counterfactual Total Cost per Well

The final input for our real option model is current expected total cost per well, normalized for productivity.[4] As discussed and estimated earlier, cost of a well falls into two primary line-items categories: drilling and completion, with drilling averaging between 42% and 52% of total cost. Within the drilling category, rig rental is the costliest line item—between 15% in Eagle Ford and 25%, in the Permian[5]—of total cost per well, while within completion, the costliest line item, fracking stimulation, is 20% for the Permian and 39% for Eagle Ford—as a percentage of total cost per well. Out of all the components making up total cost per well, there is a direct relationship between rig rental rates and

[2] Industry experts indicate that producers hedge 30-60% and 75% of their price and, when applicable, transportation cost exposure, respectively. Hence, before initiating drilling and completion, the 24-month market price expectations are partially locked in to mitigate half to two-thirds of revenue exposure. The Permian is exposed to both price and transportation cost exposure, with transportation cost risk hedged using pipeline take-or-pay agreements. As discussed earlier, we simulate this transportation risk using a mean-reverting jump (MRJD) diffusion process, which controls for congestion expectations, and show that it is more consistent than using only derivative market price expectations.

[3] In all industries, expectations of prices and price volatility expectations influence decision-making. In only a few sectors, such as oil, gas, and wheat, do industry participants hedge and coordinate their activity extensively using financial futures and options markets. Those futures and option markets offer these industry participants a unique measure of expectations and volatility, unavailable elsewhere.

[4] Productivity (bbl) for the first month of a completed well increased by a multiple of 9 in Eagle Ford and 7 in the Permian.

[5] For the Permian most shale oil activity is in three sub-formation: Wolfcamp, Spraberry, Bone Spring.

oil prices. Rig rental is contracted from specialized oil service companies that measure their rental rates based on expected oil price levels. Hence, in our counterfactual policy estimation, we adjust total cost per well upward to control for the higher price levels from an environment without the export ban.[6] We obtain quarterly data on rental prices for drilling rigs from the *The Day Rate Report* survey for Q1 2009-Q2 2015 and on fracking during the completion phase from *Spears & Associates DCS Report*.

4.3 Real Option Model and Previous Results

A dynamic programming approach allows us to empirically estimate the structural parameters β, μ, and ζ of the underlying real option model. This is a stationary infinite-horizon model, where the value functions are computed using contraction mappings. The structural methods' strong link to data makes it clear which variation in state parameters is driving the results. Firm beliefs about the evolution of expected volatility σ is regulated using β volatility responsiveness such that $\ln \sigma_P(t) = \ln \bar{\sigma} + \beta(\ln \sigma_P^{Mkt}(t)$ - $\ln \bar{\sigma})$, where $\bar{\sigma}$ is the first year average volatility. The parameter β measures a firm's responsiveness to the future expected evolution of time-varying market price volatility used to measure the real option value of waiting to drill and complete a well.[7] In particular, we test whether, when coordinating well investment decisions and integrating transportation basis risk, firms are sophisticated enough to use market time-varying volatility. The restricted model with $\beta = 0$ indicates a firm is myopic, as it measures future uncertainty using constant volatility, while $\beta = 1$ recognizes the informational advantage of using a time-varying expectation deduced from futures options prices to estimate the initial implied volatility $(\sigma_P^{Mkt}(t))$, as well as simulating and integrating transportation basis risk expectations when necessary. Both models, however, estimate the 'wait and see' real option value resulting from integrating uncertainty. Also, we assume that the log of expected production, $\log x_{it}$, is i.i.d across prospect i and time t. To fit simulated drilling and completion activity and rationalize the data, we estimate the distributional

[6] See Appendix C.2 for a discussion on a simple time-series adjustment of total cost per well.

[7] β is the free parameter we estimate. We model the option that a firm has to wait as a random walk belief about future price, cost and volatility.

parameters mean μ and standard deviation ζ for $\log x_{it}$.[8] Once β, μ and ζ are estimated on Texas shale oil firm-level data, we measure the counterfactual policy impact of the export ban on drilling and completion investment activity by identifying the number of well prospects that went un-drilled and -completed.

For shale oil producers, the 2009-2015 period is an unconstrained environment, with homogeneous factory drilling where each well is treated uniformly within the basin and drilling and completion service sector supply is slack.[9] Furthermore, these homogeneous firms are price takers—with no single shale oil producer able to manipulate price—credit is lax, and there are low expected corporate defaults. A producing firm can either hedge fully or, if it has beliefs that differ from that of the market, it can partially hedge its price risk, leaving speculative biases unhedged.[10] Hence, the agent is risk-neutral as it can hedge away any price risk-averseness by participating in energy derivative markets. These simplifying assumptions allow us to apply a single-agent real option model for shale oil well investment decision-making at the individual producer level. As discussed earlier, we use our real option counterfactual framework on a large sample of small producers operating in the Permian and Eagle Ford during a particularly volatile price and cost environment.

The real option trade-off is best restated as a dynamic programming optimal search and stop problem—also known as dynamic programming Markov chain approach. The Bellman Equation (4.1) value function V_i captures the state-contingent, stationary, and infinite-horizon nature of the problem. We are then able to converge on a stable value function solution. Price and cost changes beyond 24-months are not drawn from markets, but measured using a Markov chain.[11] We can then estimate the state space transition

[8] There are distributional lags between the decision to drill and complete and production. Furthermore, there are idiosyncratic differences between firms on how geological data is interpreted over time.

[9] During a few months in 2012-2013, 35% of surveyed producers reported they had to wait a month or two before getting a rig. However, the time window is low and there isn't any data to control for waitlists systematically, aside from adding a lag between the decision to and commencement of drilling.

[10] As shown by Kellogg (2014), the market is an "efficient aggregator of information" and better than any spot or forecasting technique that derives its measure from price history. Policy and price exposure is derived from financial markets, which are efficient aggregators of information, incorporating more distributional information than backward-looking volatility measures.

[11] We showed earlier that the 24-month maturity rate of new information entering the market drops

dynamics using multivariate normal distribution.

Oil markets are global, with Texas shale oil producers output representing only a small percentage of worldwide oil activity. Hence, these producers are price takers. As outlined earlier, the value function encapsulates the trade-off between expected current and future profitability for a given unobserved expected well productivity x_{it},[12] discounted by δ^t, as a function of the state parameters price (P), price volatility (σ), and total cost per well (TC).

$$V_i(P_t, TC_t, \beta\sigma_P(t)) = \max\{\pi_i(P_t, TC_t), \delta^t \mathrm{E}[V_i(P_{t+1}, TC_{t+1}, \beta\sigma_P(t+1))]\} \tag{4.1}$$

Empirically, for month t, total cost per well TC_t is decomposed first into drilling costs, then by rig rental $d_i R_t$ and non-rental cost rc_i, where d_i is rig days for a rig rental rate R_t. Days to drill are not constant throughout the sample decreasing at a somewhat[13] linear rate from 33 days at the beginning of the sample to 15 days at the end. Completion costs include fracking F_{it} and non-fracking cost fc_i, where fracking cost is interpolated when data is missing. With expected current price P_t, total cost per well TC_t, and productivity x_{it}, we estimate current profit π_{it} driving the decision to drill and complete a prospective well i at time t. The relationship between price P_t, total cost per well TC_t, and expected productivity x_{it} is summarized in the profit π_{it} Equation (B.1):

$$\pi_{it} = \pi_i(P_t, TC_t) = x_{it}P_t - TC_t = x_{it}P_t - rc_i - fc_i - d_{it}R_t - F_{it} \tag{4.2}$$

$x_{it}P_t$: discounted (δ) expected well productivity times average state price

rc_i: non-rig costs

rc_i: non-fracking costs

and volatility beyond it is constant. We therefore designate contracts less than or equal to 24 months as short-term—modeled using a mean-reverting dynamic—and contracts longer than 24 months as long-term—using a random walk process without memory and with constant volatility. Kellogg (2014) makes a similar argument but focuses on a parametric estimation of an 18-month contract instead of a non-parametric 24-month contract because of data 'limitations.'

[12] We solve for the expected well productivity $\overset{*}{x_t}$ that leaves a producer indifferent between investing in a well now or later.

[13] From 2009 to 2015, the number of days to drill decays quickly and then slows down (e.g. 33 (2009), 22 (2010), 21 (2011), 20 (2012), 18 (2013), 16 (2014) and 15 (2015) days).

$d_i R_t$: number of days in operation times rig costs

F_{it}: total cost per well normalized for productivity

Table A.1 reports the convergent real option model fit using optimal structural parameters β, μ, and ζ. There are fewer producers operating on productive and homogenous shale oil fields in the Permian, which helps explain the tighter distribution (ζ) around expected productivity (μ) used to smooth the real option model's predicted well investment activity and rationalize the data to estimate a best fit. With both βs statistically equal to one, we conclude that Eagle Ford and Permian producers use time-varying market beliefs and a heuristic version of the real option framework to decide whether to drill and complete a well or wait and see.[14] Furthermore, Eagle Ford is exposed only to crude price risk, while the Permian is exposed to both price and transportation basis risk. We apply a mean-reverting jump process to simulate the 24-month price discount between the Permian Midland and West Texas Intermediary (WTI) Cushing market prices, which represents transportation basis (or cost).[15] This transportation basis tends to mean-revert to its long-term mean, except when transport takeaway capacity is congested and it temporarily jumps. For producers exposed to transportation basis risk, derivative market futures price and implied volatility aren't precise enough and need to integrate both price and transportation basis volatility to explain producer response to uncertainty and the resulting well investment decisions.

Similarly, for the counterfactual international Brent price, when necessary, we incorporate basin-specific transportation dynamics into price and volatility to account for the role of transportation basis congestion risk in the current inputs of the real option model. Beyond 24 months we still apply a memoryless diffusion process to estimate the real option value. Hence, for our policy estimation, we use the statistically significant parameter estimates from Table A.1 to predict well count in an environment with and without the export ban in both the Permian and Eagle Ford. To estimate the welfare impacts, we

---

[14]Eagle Ford is located close to the coast and doesn't suffer from any transportation congestion for the period.

[15]The geographic delivery points of these crude benchmarks is shown in Figure 4.3.

Table 4.1: Permian and Eagle Ford Model Results

	Permian (1)	Eagle Ford (2)
β (volatility responsiveness)	1.04	1.08
	(0.14)	(0.12)
μ (log expected productivity)	-6.30	$-$ 6.79
	(3.33)	(2.40)
ζ (std of μ)	2.04	2.81
	(2.63)	(1.51)

Notes: μ is expressed as expected oil production (in bbl) divided by the total cost of drilling and completion. All estimates use 24 month futures prices and option volatilities. A random walk forecast is used for future volatility. Drilling and completion data are matched to drilling and completion likelihoods with a 2-month lag.

then adjust the counterfactual price for transportation basis: freight and pipeline.

4.4 Investing Under the Export Ban

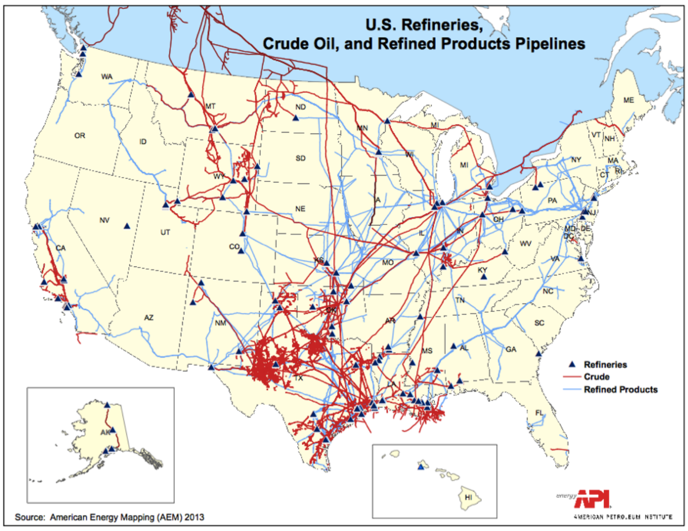

Figure 4.1: US Crude Oil and Product Pipeline Map

source: API

The United States' drilling and completion technological progress made the extraction of large volumes of shale oil feasible as of 2009, thereby altering the face of global energy geopolitics and, in particular, the economics of US energy policy. The primary shale oil fields, representing more than 50% of US shale oil production, are the Permian and Eagle Ford in Texas; with North Dakota's Bakken and Colorado's Niobrasa fields, they account for more than 95% of shale oil production in the US. As production of light oil from shale increased in Texas, where the majority of storage, transportation, and refining capacity is located, it displaced light oil imports and altered the US crude oil profit structure. Figure 4.1 shows a map of US crude and refined product pipelines, where blue illustrates refinery pipelines, red is for crude pipelines, and the triangles are refinery capacity. It is immediately clear that US refining capacity and crude pipelines are primarily located in Texas. Figure 4.1 also illustrates the concentration of takeaway capacity in the Gulf, as well as how isolated the eastern and western states are from it. In the context of increased domestic production, this takeaway infrastructure is a limiting factor, which the export ban aggravated by allowing an oil supply glut to form in Texas, depreciating domestic prices.

During the 2009-2015 period, the US shifted its source of crude oil from imported light oil to US-sourced light oil from shale. However, the presence of the originally slack 1975 export ban on crude oil created the risk of a light oil glut in the US. The ban prevents the export of crude oil but not crude oil product. Hence, as shale oil production began to take off in 2009, displacing imports, the export ban remained slack and steps were taken to avert transportation congestion. In particular, transportation import flows that were originally aimed inwards were reversed and pipeline capacity increased, storage filled up, and refiners invested in distillation capacity primarily aimed at circumventing the ban by increasing exportable pseudo-refined crude oil product volumes. Initially, US light oil only displaced imports, but as domestic production outstripped demand, new international outlets were necessary. Applying a 5-year seasonal analysis, Figure B.6 shows the decrease in Gulf and east coast imports of light crude product.[16] The shifting energy structure effectively calls into question the previously non-binding ban on oil exports dating back to the 1970s, from which crude products such as jet fuel and

[16]Gulf coast refineries produce gasoline from low-priced heavier crude, while on the East Coast they rely on high priced light sweet crude. However, given a lack of transportation infrastructure, the glut of light oil from shale oil was not easily moveable from Gulf to East coast.

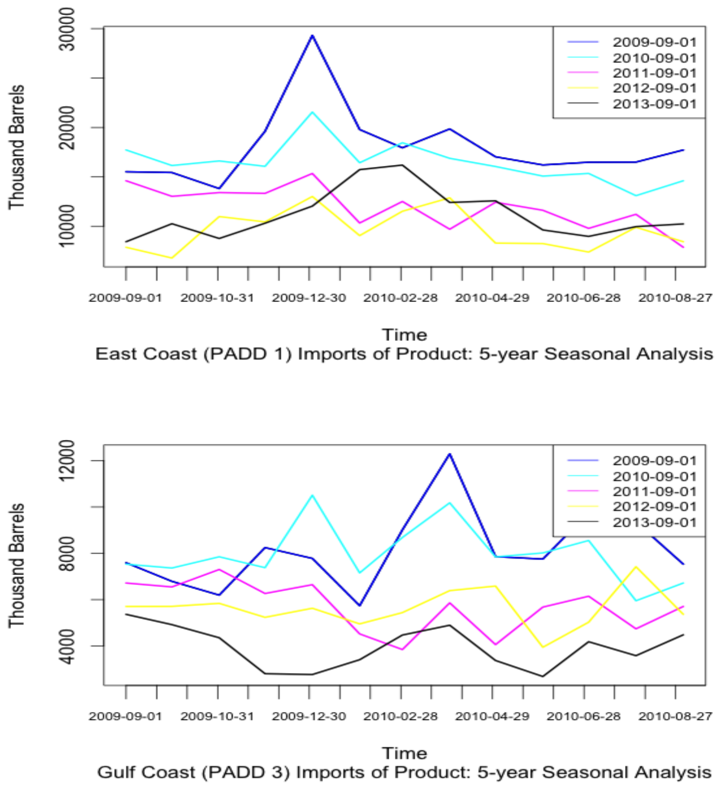

Figure 4.2: PADD Petroleum Products Imports

source: EIA

gasoline are exempt.

As we show using a time series structural break approach, by August 2013, the export ban effectively became binding,[17] resulting in a glut of light oil in Texas and domestic crude oil trading at a discount to the international price. Refiners benefited from the ban by buying cheaper domestic crude oil and selling refined product at the higher world price. We estimate that producers, on the other hand, received a lower price signal, and therefore invested in fewer wells than they would have if the ban had not become binding. Lower production resulted in lower state severance tax revenue, as well as lower CO_2 emissions from the displaced barrels of oil.

[17]There are two primary sanctioned export outlets to circumvent the ban: 1) US crude oil exports to Canada and 2) expansion of distillation towers to convert crude oil into exportable product.

4.4.1 Domestic and International Financial Markets

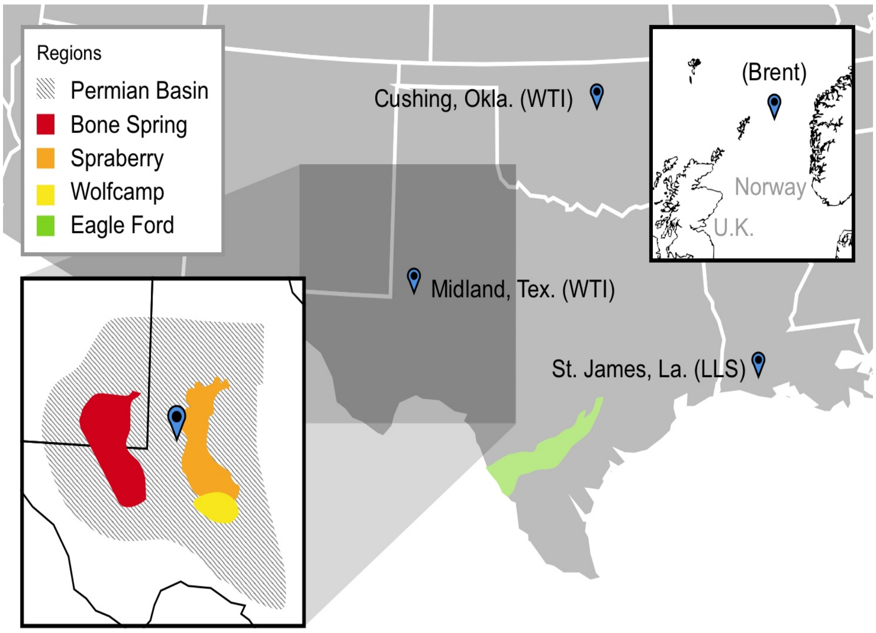

Figure 4.3: Benchmark Crude Oils

As illustrated in Figure 4.3, the three primary global benchmark crudes—West Texas Intermediary (WTI), Louisiana Light Sweet (LLS), and Brent—are delivered in the North Sea (Brent) and the US Gulf Coast, where LLS is traded. WTI, however, is inland in Cushing, Oklahoma, connected by pipeline to the US Gulf coast. With access to a unique energy derivative market, we can estimate the impact of the export ban on drilling and completion decision-making by using counterfactual expected prices, volatility and price adjusted total cost per well.[18] To capture the export ban's policy impact, we use two highly traded crude benchmarks, namely West Texas Intermediary (WTI) and North Sea Brent, coupled with the Louisiana Light Sweet (LLS) and WTI spread.[19] For purposes of estimation, we employ contracts that are traded two years out. Hence, we do not apply the standalone LLS contract but instead use the the LLS-WTI spread, which is heavily

[18] We modify total cost per well by calibrating rig rental cost for the change in price levels; see Appendix C.2.

[19] WTI and the LLS-WTI spread are traded on the Chicago Mercantile Exchange (CME), while Brent has a higher traded volume on the Intercontinental Exchange (ICE).

traded, to infer the LLS price.

The different futures prices between these three oil types reflect distinct grades, calendar (or temporal) dynamics,[20] and transportation bases. As shown in Figure B.1, WTI and Brent are cointegrated;[21] historically Brent and WTI tracked each other closely, however, with the export ban and inland transportation congestion, the spread expanded. Before the export ban became binding, a widening spread, through arbitrage activities, quickly converged back to its long-run equilibrium level a few times a day.[22] The one-month and 24-month Brent differentials in Figure B.1, are in some periods mainly the result of transportation costs and in others the consequence of the export ban and lack of capacity to transport the glut of new oil from Cushing Oklahoma or Houston Texas to international markets. In fact, prior to 2013, because of LLS's relative proximity to global markets, it equaled Brent crude, while the WTI-Brent spread reflected transportation costs between the Gulf and Cushing, OK. Post-August 2013, the export ban prevented Louisiana from turning import into export capacity, widening the LLS-Brent spread.

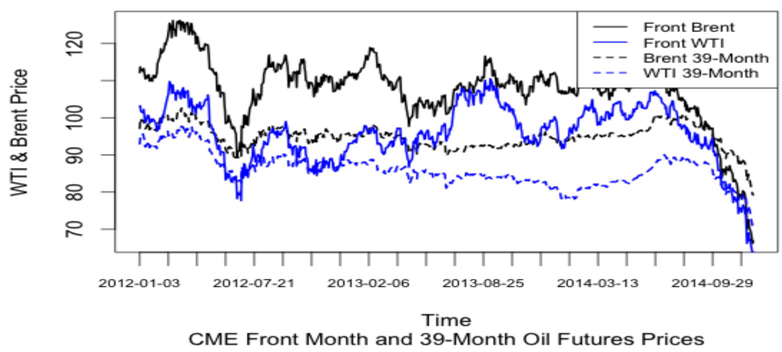

Figure 4.4: Crude Front and 39-Month Futures Prices

[20]Different contracts maturities are priced relative to each other as one is always more widely traded, offering more information, than the other. Also, future contracts are priced based on the cost of carry—cost of storage—and convenience yield—the benefit of having crude on hand—which are time-specific constraints applying to all contract maturities.

[21]They are integrated of the first order, i.e. the first difference of the time series is stationary.

[22]We used an error correction model (ECM) to confirm the reduction in the convergence rate of the spread from 2005-2015.

In our analysis, the expansion in the LLS-Brent differential proxies for the effect of the export ban while WTI-Brent is the combination of the export ban and transportation costs. We use Brent as the international price and show that US production only marginally impacts—on the order of less than 10% —global production, while WTI represents domestic prices observed by Texas shale oil producers. We control for transportation costs using the WTI Midland-WTI and LLS-WTI differentials for the Permian and Eagle Ford respectively. Hence, with access to domestically constrained WTI prices and the unconstrained Brent international prices, we can estimate the distortionary impact of the export ban on the decision to drill and complete wells.

4.4.2 Export Ban Structural Break

As shown in Figure 4.5, the LLS-WTI spread was initially wide, but the addition of transport capacity, reversal of pipeline flows, and use of rail and freight transport post-August 2013 between Texas and Louisiana tightened the spread.[23] With an LLS-WTI spread (or transportation basis) adjustment, the WTI-Brent differential represents the impact of the export ban policy on investment decision-making. Although the export ban was in place since 1975, energy derivative markets did not expect production to outstrip domestic demand by 2013. Mănescu and Nuño (2015), and Baumeister and Kilian (2016) estimate that financial markets did not anticipate the shale oil supply glut and the oil price drop in 2014-2015. As illustrated in Figures 4.5 and 4.6, in August 2013 the front-, 10-, and 24-month LLS-Brent contract expanded and LLS-WTI prices converged contemporaneously. Therefore, producer investment decisions, which use the 24-month contract as a input, were impacted by the export ban only when it became binding in August 2013.

The convergence of LLS-WTI spread, in Figure 4.5, was the result of pipeline capacity expansion and flow reversals, which reduced the risk of bottlenecks, while the LLS-Brent differential expansion, in Figure 4.5 and 4.6, came from domestic production outstripping demand without the option to export light oil directly. We conduct a linear level test to

[23]For illustrative purposes, we start with 2012, as production took off then. Also, the LLS-WTI is not traded consistently over 24 months; therefore, we focus on a 10-month contract.

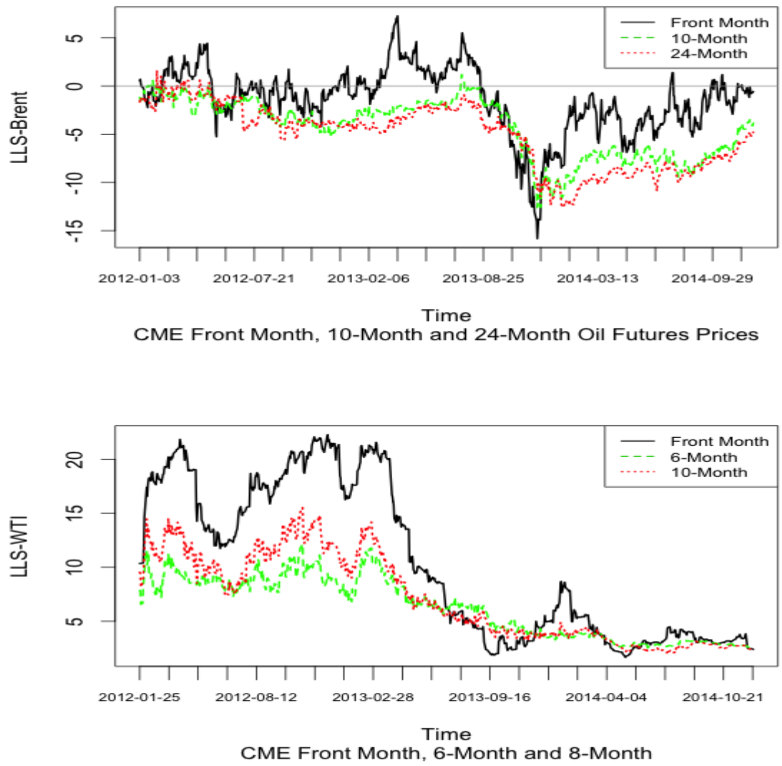

Figure 4.5: Futures Price Differential

support the theoretical claim that, as of August 2013, LLS-WTI differential converged, while LLS-Brent spread expanded.

The cointegrated nature of the WTI, LLS, and Brent time series allows us to estimate a structural break by looking either at WTI-Brent or LLS-Brent spreads. We use the LLS-Brent spread because of its stable distribution pre-2013[24] and apply a simple empirical fluctuation process to pinpoint the structural change. We have a ballpark range of when the ban became binding, as light sweet crude imports significantly decreased and east coast storage volume increased from 2012-2014. Additionally, given that all contract maturities covary in unison, as the ban became binding, we use the front-month to estimate the date.

[24]The LLS-Brent is exposed to a stable freight basis, while LLS-WTI integrates both transportation (e.g. pipeline) and freight basis risk.

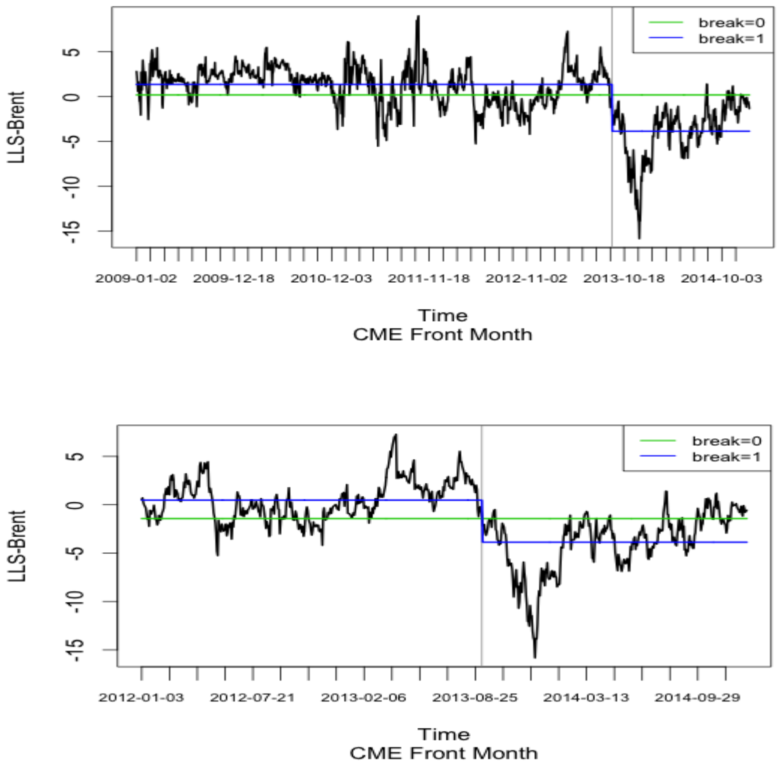

Figure 4.6: Identifying the Binding Export Ban

We use an OLS-based CUSUM structural change test (Ploberger and Krämer (1992)) to estimate the binding export ban's structural break date. The standardized cumulative sum is set up in Equation (4.4):

$$y_i = x_i^T \beta_i + u_i \tag{4.3}$$

$$\mathrm{W}_n^0(t) = \frac{1}{\hat{\sigma}\sqrt{n}} \sum_{i=1}^{t} \hat{u}_i \tag{4.4}$$

The limiting process for $\mathrm{W}_n^0(t)$ is the standard Brownian bridge $W^0(t) = W(t) - tW(1)$, where $W(.)$ denotes Brownian motion.

In the case of a single-shift alternative, the process peaks around the breakpoint. We are concerned with testing the hypothesis that the regression coefficient remains constant against the alternative that at least one coefficient varies over time. The result from

Equation (4.4) is graphically depicted in Figures 4.5 and 4.6. We estimate using a mean-reverting process that the LLS-Brent spread pre-2013 reverts to a mean of \$0. Post-2013, we measure a new long-term mean of minus \$5. We find, therefore, that for the sample, which has a long-term mean of \$0, the mean-reversion rate of the LLS-Brent spread is lower post-2013. Furthermore, the level change test finds that a significant structural change occurred on the date of 2013-08-27. Hence, this simple regression error test supports the claim that the export ban became binding in August 2013, which pinpoints the breakpoint date for our real option switching policy estimation.

The August 2013 binding export ban constrains transportation of US-sourced shale oil production in the Gulf region. As of 2013, the combined effect of the export policy coupled with the long lead time necessary to implement transportation infrastructure projects resulted in a widening of the crude price differential between domestic and international crude prices, represented by Louisiana Light Sweet (LLS) and Brent financial futures, respectively. This price differential was near zero prior to the ban becoming binding. Hence, to quantify the policy impact, we adjust Brent and WTI expected price and volatility by transportation basis risk—modifying rig rental costs for the higher oil price environment. Brent and WTI are then used as control and treatment, respectively, to measure the effect of the ban on shale oil well investment decisions.

US shale producers are atypical in the industry; as small independent players with sharp production decline curves, they are responsive to expected commodity prices, future uncertainty, and expected total costs per well, and are not directly impacted by state or federal government in the way that other producers worldwide are. Thus, we adopt a real option framework to estimate the number of displaced wells in the Permian and Eagle Ford as a result of the export ban. These displaced wells are then converted into lost barrels, producer surplus, state tax revenue, and CO_2 emissions.

4.4.3 Real Option Switching Methodology and Results

We adopt a real option state-switching framework, whereby we specify two different sets of state variables: pre-and post-August 2013. For the Permian and Eagle Ford, we apply an estimated real option model and its price responsive β, and μ and ζ expected productivity distributional parameters from Table A.2. This allows us to capture the dynamic effect of price, volatility, and cost on the decision the drill and complete a well. We use transportation cost-adjusted Brent as our counterfactual state, illustrative of an environment without the binding export ban. As stated above, by correcting for transportation exposure, Brent is effectively converted into an accurate representation of a counterfactual WTI, which is located inland. In particular, for the Permian, we estimate expected transportation basis risk by forecasting the WTI Midland-WTI (Cushing) differential. To estimate the inland and gulf coast transportation basis, we use the expected LLS-WTI futures price and implied volatility from derivative markets. For Eagle Ford, we only control the expected LLS-WTI spread.

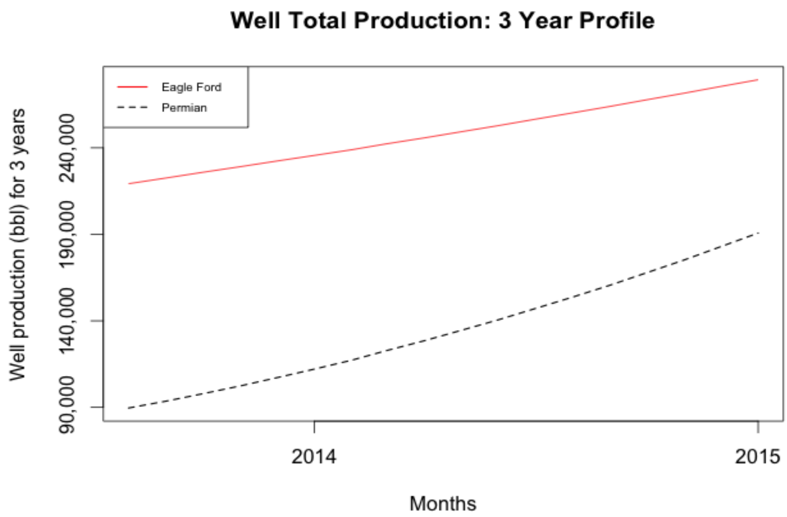

Figure 4.7: Well Three Year Total Production

Once the number of displaced wells is estimated from the ban, we use the fixed well decline rate with the corresponding initial month production to measure the number of forgone barrels. For both the Permian and Eagle Ford, we only use a well's three-year discounted accumulated production. Note that although Permian productivity growth was slow to start, towards the end of the sample it begins to converge on Eagle Ford's level, as shown in Figure 4.7. This accumulated three-year well production profile gives us a lower bound estimate. Forgone production also allows us to estimate the tons of CO_2 emissions displaced by shale oil production using the Environment Protection Agency's (EPA) 0.43 metric tons CO_2/barrel conversion. Furthermore, applying the expected prices from futures derivative markets for the corresponding month and well productivity, we measure producer-forgone revenue and surplus. Finally, knowing Texas' severance, levied at 4.6%, and oil regulation tax, levied at 3/16th of one cent per barrel of oil produced, we estimate forgone state tax revenue. The real option-switching measure between two states—one in which the export ban is binding and one in which it is not—is one of the first policy impact measures to integrate both price and variance into a dynamic programming behavioral response to empirically measure the distortion in investment decisions.

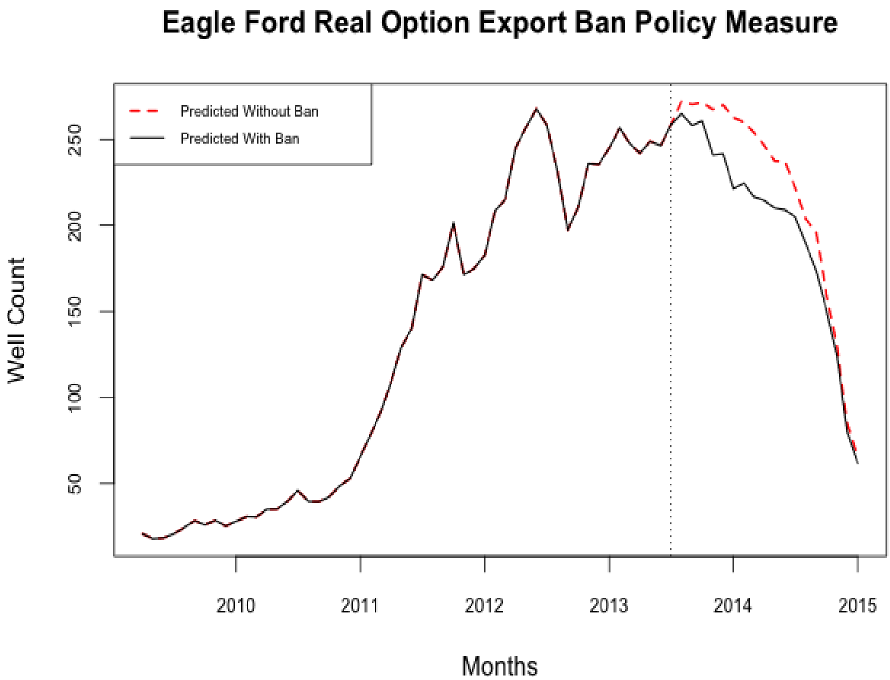

Figure 4.8: Eagle Ford Real Option Results

Export Ban Policy Results

The first state is the current environment with the export ban in effect, while the second is a theoretical state in which the ban doesn't exist. We estimate and apply our real option model on these two different sets of state variables, switching states as of 2013-08-27 when the ban became binding. Taking the difference in well investment decisions between these two states, we measure the number of displaced wells in the Permian and Eagle Ford as a consequence of the ban. As illustrated in Figures 4.8 and 4.9, as the ban becomes binding, the well count between environments with and without a ban diverges. Figure 4.9 shows slight divergence in the Permian prior to August 2013; the slight discrepancy results from changes in the correlation structure between WTI, WTI Midland, and Brent price and variance states.[25]

From August 2013 to January 2015, we estimate that 581 wells were displaced 364 from Eagle Ford and 217 from the Permian. These undrilled prospective wells account for 3% of total well investments, and reduced oil production by 116,213,859 barrels, representing about 7% of Texas shale oil production with the ban for the sample period. To estimate discount well production, we take each initial month's productivity profile and apply the estimated three-month well decline profile.[26] Note that the same Texas well production decline rate is consistent between basins and periods.[27] We then apply the forward curve going out three-years to estimate the corresponding producer expected forgone revenue of \$10,742,342,415.[28] The expected average for the sample and three years out is around \$96.69 per barrel, which is in line with the realized price for the sample period. To measure forgone producer surplus, we apply the corresponding costs per month and find \$6,853,080,941. Displaced production comes as reduced Texas severance

[25] The transportation cost and market price correlation structure is estimated over the sample. Hence, the concatenation of price and volatility time series indubitably leads to small estimation discrepancies in correlation between a state with and without the ban. As outlined earlier, this is a consequence of bundling mean-reverting jump-diffusion (MRJD) simulations using portfolio mixed log-normal techniques to capture producer sensitivity to Permian basin transportation basis risk.

[26] Note that accumulated production per well is increasing over time as producers are able to get more out of their wells.

[27] The three-year profile is used, as it represents two-thirds of a well's production capability.

[28] Producers coordinate and hedge their activity using expected price and volatility beliefs deduced from financial markets.

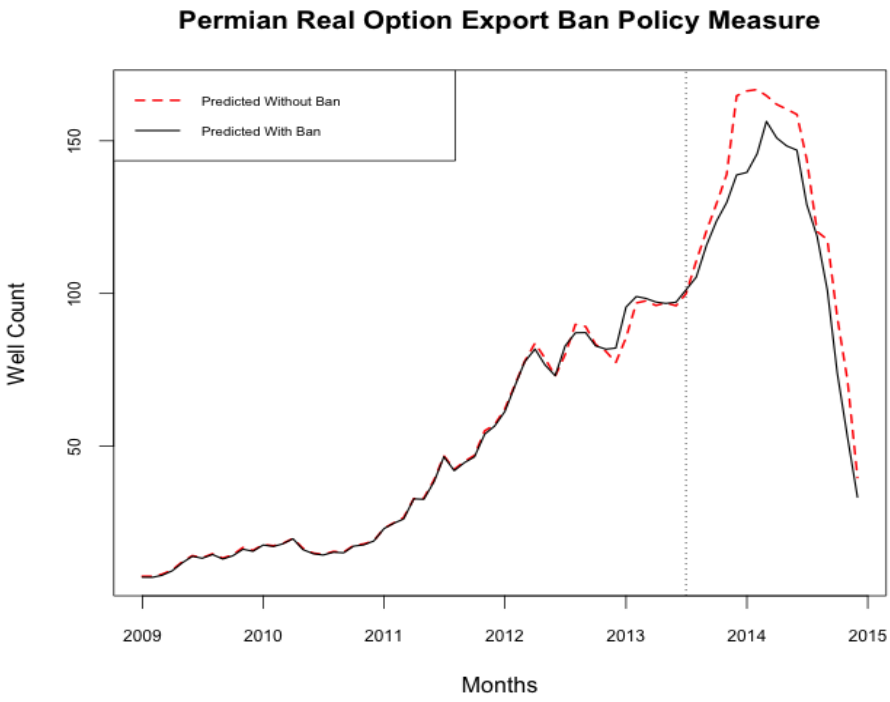

Figure 4.9: Permian Real Option Results

and oil regulation tax revenue of $500,861,715. However, the export ban policy prevented 49,971,959 metric tons of CO_2 from being burned. These estimates are based on a three-year well production profile and are therefore lower bound estimates of lost income and CO_2 emissions avoidance.

Lost producer revenue from the export ban is a direct transfer to Gulf Coast refiners. Without the ban, the 116,213,859 barrels would have gone into international markets, circumventing Gulf Coast refineries. However, with the ban, domestic is lower than the international price. These refineries buy crude domestically and produce and sell gasoline and heating oil, which is priced internationally. We simplify the estimation by assuming that refineries produce a product pegged to the international Brent price and purchase domestic WTI. Given the depreciable effect of the ban on domestic price, we measure that for all the barrels produced under the ban, refineries earned $10,136,189,800 in revenue.[29]

[29]This estimation does not account for operational expenses, barge transportation costs, and cracking-efficiency—utilization and capacity—of the refinery.

This revenue approximation shows that the ban effectively served to transfer revenue from producers to refiners and lowered the Texas severance tax base.

4.5 Conclusion

4.5.1 Policy Implications

We estimate well investment decisions in the shale oil sector by adopting an irreversible investment real option model. The decision to invest in a well is highly responsive to a firm's beliefs regarding expected price, volatility, and cost, which are deduced from crude futures and option markets, transportation cost forecasts, and technological trends. Additionally, the sharp production decline rate of these wells allows production flows to respond to price and volatility changes within months. By modeling a new small, innovative, and independent shale oil entrant, we diverge from previous literature that analyzes the inelastic responsiveness of super majors[30] or state run producing firms to oil shocks. These new supply dynamics have macro-economic ramifications for estimating the impact of crude oil shocks on economic decisions. In fact, crude price shocks now affect expectations about the future path of the crude price and drilling and completion costs over a shorter 24-month horizon. Hence, to accurately measure the impact of a policy, we need to analyze the impact of market variables on the decision to invest in a well or not.

The real option framework offers a new structural policy technique, which we use to estimate level and variance effects of a ban on forward-looking domestic producers. Incorporating uncertainty allows us to model the critical option value of investing now or waiting for a future project's uncertainty to be resolved. This approach, coupled with access to global market fundamentals without the export ban, allows us to estimate the impact of the ban post-August 2013, as the ban became binding. We find that once US domestic production outstripped demand and the export ban policy became binding, the Federal government effectively transferred revenue from producers to Gulf coast refiners. These refiners benefitted from the depreciated domestic input crude price and a higher

[30] Firms with global reach energy portfolios.

international output crude product price.[31] The model indicates that US production revenue would have been \$10.7B higher had investment decisions been based on global market fundamentals. In addition, the state of Texas would have earned severance and oil regulation tax revenues of \$500,861,715.

With an estimated real option model for 2009-2015, we find that shale oil firms are sensitive to distortionary policies such as the export ban. If it weren't for the constant progress in cost reduction and the decline in crude oil prices towards the end of the sample, the impact of the ban would be costlier. The decision to remove the ban when prices were below \$40 was insignificant as the ban was ineffective at that point, illustrated by the contraction in the LLS-Brent spread in 2015.

4.5.2 Current Market Trends

A consequence of this short-run well investment response is that oil price shocks —unanticipated changes in supply and demand—now impact production and macroeconomic outcomes within the same period. Also, as of 2015, the US economy became more reliant on a heavily indebted energy sector. In this newly responsive and indebted producer environment, the benefits of low 2015-2016 crude prices aren't clear. On the one hand, consumers have more disposable income, which may be saved; on the other hand, some producers are unable to service their debt, with all the financial ramifications that implies.

Currently, new crude market dynamics, with a severely depreciated spot price, have decoupled the decision to drill and complete, with producers reserving the right to complete at a later date. There have also been increased corporate defaults, with company downgrades by Fitch and Moody's and lower bank reevaluation of crude reserve value. In this stressed environment, the decision to drill and complete is impacted by a firm's debt to equity corporate structure. This is a concern, as credit limits are set based on the value of shale oil reserves in the ground; hence, banks are starting to have trouble making loans under this energy price collapse. In fact, producer credit lines were cut by

[31] Although producers circumvented the ban by increasing exportable pseudo-refined crude oil product and exporting to Canada, we estimate that without the ban, the 116M barrels would have gone into crude oil international markets, circumventing Gulf Coast refineries. However, with the ban, domestic prices are depreciated with respect to international prices.

around 39% in Fall 2015.

These are complex times, as banks are pressured to cut debt financing to producers living beyond their expected cash flows. These pressures risk reducing production, which will also hamper firms' ability to meet their financial obligations. This could potentially lead to sector-wide bankruptcy. Such an energy downturn could spill over and impact regional economic development by initially compromising the industrial and transportation sector. However, the welfare effect is unclear, as a lower crude price has positive effects too: consumers are better off from the fall in refined product costs. In the past, well-capitalized super majors or state firms with refining businesses underwent crude oil investments; now, with a negative oil price movement and small independent oil firms, we risk seeing an increase in bankruptcy and consolidation. Hence, a natural extension of this thesis is to model the fundamental financial market shift that occurred in 2015, assessing its impact on heterogeneous firms through capital market tightening, new operational innovation, and high grading strategies in a low price environment.

BIBLIOGRAPHY

Aguerrevere, F. L. (2003). Equilibrium investment strategies and output price behavior: A real-options approach. *Review of Financial Studies*, 16(4):1239–1272.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2014). Hotelling under pressure. Technical report, National Bureau of Economic Research.

Baumeister, C. and Kilian, L. (2016). Understanding the decline in the price of oil since June 2014. *Journal of the Association of Environmental and Resource Economists*, 3(1):131–158.

Benkard, C. L. (1999). Learning and forgetting: The dynamics of aircraft production. Technical report, National Bureau of Economic Research.

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

Black, G. and LaFrance, J. T. (1998). Is hotelling's rule relevant to domestic oil production? *Journal of Environmental Economics and Management*, 36(2):149–169.

Bloomberg, L.P. (1/1/09-6/1/15). Daily CME Futures and Options Curve. Bloomberg database. University of Washington Foster Business Library, Seattle, WA. (June 9, 2015).

Chen, T. (2011). Improve OVDV long-term volatilities. *Bloomberg Research*, pages 1–7.

Chew, D. (2001). University of Texas roundtable on energy derivatives and the transformation of the US corporate energy sector. *Journal of Applied Corporate Finance*, 13(4; SEAS WIN):50–75.

Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management.* Lacima Publ.

Copeland, T., Weston, J., and Shastri, K. (2005). Financial theory and capital policy. *Boston, MA: Addison-Wesely.*

Dias, M. (2005). Hybrid real options with petroleum applications. *Departamento de Engenharia Industrial. Rio de Janeiro: Pontifícia Universidade Católica do Rio de Janeiro.*

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

Dunne, T. and Mu, X. (2010). Investment spikes and uncertainty in the petroleum refining industry. *The Journal of Industrial Economics*, 58(1):190–213.

Energy Information Administration (EIA) (March 30, 2016). EIA Report Shows Decline in Cost of U.S. Oil & Gas Wells Since 2012. http://www.eia.gov/todayinenergy/detail.cfm?id=25592.

Grenadier, S. R. (1996). The strategic exercise of options: Development cascades and overbuilding in real estate markets. *The Journal of Finance*, 51(5):1653–1679.

Gruber, H. (1992). The learning curve in the production of semiconductor memory chips. *Applied Economics*, 24(8):885–894.

Grübler, A. (1990). *The Rise and Fall of Infrastructures: Dynamics of Evolution and Technological Change in Transport.*

Güntner, J. H. (2014). How do oil producers respond to oil demand shocks? *Energy Economics*, 44:1–13.

Halvorsen, R. and Smith, T. R. (1991). A test of the theory of exhaustible resources. *The Quarterly Journal of Economics*, pages 123–140.

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

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343.

Hilliard, J. E. and Reis, J. (1998). Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot. *Journal of Financial and Quantitative Analysis*, 33(01):61–86.

Hotelling, H. (1931). The economics of exhaustible resources. *The Journal of Political Economy*, pages 137–175.

Kellogg, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch. *The Quarterly Journal of Economics*, page qjr039.

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

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.

Mănescu, C. B. and Nuño, G. (2015). Quantitative effects of the shale oil revolution. *Energy Policy*, 86:855–866.

Marschak, J. (1949). Role of liquidity under complete and incomplete information. *The American Economic Review*, pages 182–195.

Moel, A. and Tufano, P. (2002). When are real options exercised? an empirical study of mine closings. *Review of Financial Studies*, 15(1):35–64.

Ploberger, W. and Krämer, W. (1992). The cusum test with ols residuals. *Econometrica: Journal of the Econometric Society*, pages 271–285.

RigData (2Q2010 1Q2012 and JAN2015). Day Rate Report.
http://rigdata.com/news-and-analysis/day-rate-report.aspx. (Feb 12, 2015).

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

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

Samuelson, P. A. (1965). *Proof that properly anticipated prices fluctuate randomly.*

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973.

Spears Associates, Inc. (Q2 2015). Spears and Associates DCS Report. http://spearsresearch.com/index.cfm. (Sept 24, 2015).

Szakmary, A., Ors, E., Kim, J. K., and Davidson, W. N. (2003). The predictive power of implied volatility: Evidence from 35 futures markets. *Journal of Banking & Finance*, 27(11):2151–2175.

Texas Railroad Commission (1977-2015). Drilling Permit Master and Trailer. http://rrc.state.tx.us/data/datasets/DrillingPermitData.php (March 6, 2015).

Texas Railroad Commission (2009-2013). Final Oil and Gas Annuals. http://rrc.state.tx.us/data/datasets/WellData.php (March 6, 2015).

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge university press.

Tufano, P. (1996). Who manages risk? an empirical examination of risk management practices in the gold mining industry. *The Journal of Finance*, 51(4):1097–1137.

Appendix A
TECHNICAL DETAILS & DATA, CHAPTER 2

A.1 Well Investment in the Permian and Eagle Ford

A.1.1 Conventional and Shale Oil

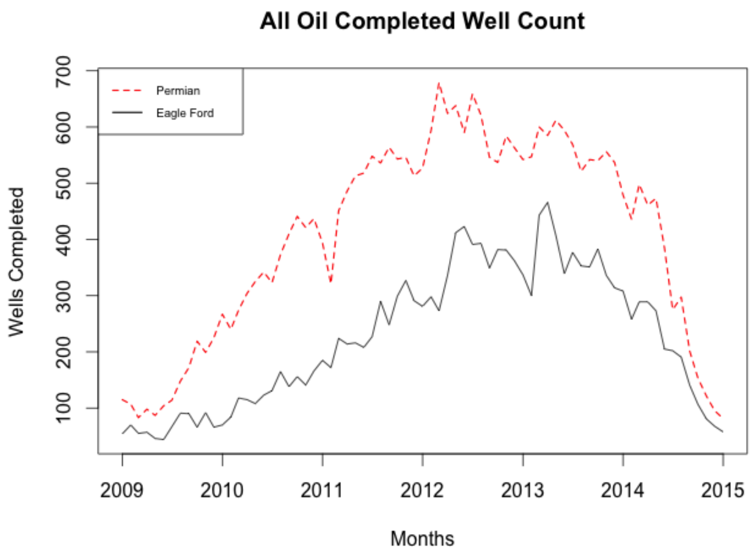

Figure A.1: Number of Onshore Oil Wells: Conventional and Shale Oil

Well activity data come from the TRRC's "Drilling Permit Master" dataset, which provides the drilling and completion date, district, county, and lease name for every well drilled and completed in Texas. As illustrated in Figures A.1 and A.2, once we filter for horizontal drilling and completed shale oil wells within particular counties/districts, we get 10,353, and 3,578 wells completed in the Permian and Eagle Ford, respectively. With the shale oil boom came an increased well completion rate per month. At first, even though costs were lower, Permian horizontal shale oil activity stagnated as a result of lower production efficiency, high transportation cost exposure, and preexisting vertical fracking activity.[1] For the period, most leases are privately owned mineral rights, paying severance and regulation taxes once completed—only 1.5-1.9% of leases are held publicly.

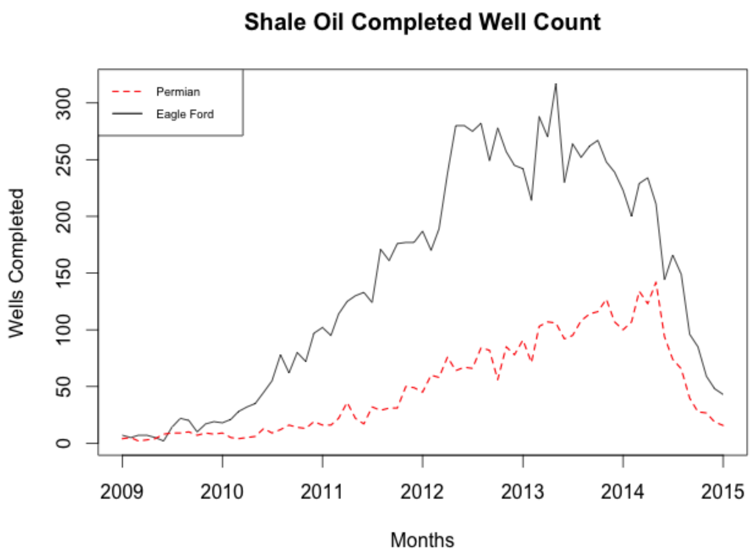

Figure A.2: Number of Onshore Oil Wells: Shale Oil

A.1.2 Shale Oil Number of Wells

In order to test whether small firms are able to integrate complex risk information into their decision-making, we model the choices of firms to initiate drilling and completion of shale wells in west Texas during a period of high shale oil production growth. Table A.1 shows Eagle Ford and Permian[2] basins well investment concentration per largest producer. From the TRRC data we find that per basin the largest firm represents 5.7% of the market, and all but 2.6%[3] unaffiliated with major producers.

[1] As shown in Figure 2.3, well investment in the three Permian low-permeability formations—Wolfcamp, Spraberry, Bone Spring—picked up in early 2012.

[2] More specifically, Wolfcamp, Spraberry and Bone Springs represent the horizontal shale oil fracking plays in the Permian Delaware region while Permian Midland represents the remainder.

[3] On average, per year there were 340 and 172 shale oil producers operating in Eagle Ford and the Permian respectively. There are 6 super-majors: Statoil, Exxon, ConocoPhillips, Chevron, Shell, and BP. Industry participants state that departments, such as shale oil, within these large companies compete for capital, acting as independent units.

Table A.1: Market Concentration: Wells Drilled and Completed

	Permian		Eagle Ford	
Date	% of Wells	Producer Name	% of Wells	Producer Name
2009-2015	5.7%	EOG RESOURCES	5.7%	APACHE CORPORATION
2009	3.7%	ENERVEST OPERATING	–	–
2010	7.5%	EOG RESOURCES	5.6%	DEVON ENERGY PRODUCTION
2011	7.5%	CHESAPEAKE OPERATING	5.1%	CHESAPEAKE OPERATING
2012	6.9%	CHESAPEAKE OPERATING	7.4%	EOG RESOURCES
2013	5.5%	EOG RESOURCES	6.6%	APACHE CORPORATION
2014	5.8%	EOG RESOURCES	7.3%	APACHE CORPORATION
2015	–	APACHE CORPORATION	–	–

Notes: "–" used when the number of firms operating is lower than 45 or there is insufficient data

For example, producing firms in the sample, such as EOG Resources, Devon Energy, Chesapeake Operating, or Apache Corporation, have a market capitalization in the order of tens versus hundreds of billion of dollars—typical of the super-majors. These producers, however, are not fully integrated across the oil supply chain or involved in multi-billion dollar offshore or conventional oil production projects like super-majors such as Exxon, Shell, BP or Chevron. EOG, Chesapeake, and Apache are recurring producers, but they represent a small proportion of total well investments. Furthermore, they have often been tied for first, with second place lagging only one or two well investments behind.

A.2 Firm Volatility Behavior Identification

Identification using β to distinguish between time varying market $\sigma_P^{Mkt}(t)$, constant $\bar{\sigma}$. We distinguish between an unsophisticated response to volatility (or differing producer beliefs) and those of futures and futures options markets in order to measure the response to market $\sigma_P^{Mkt}(t)$ by parameterizing producer beliefs β. We estimate that $\bar{\sigma}$ is equal to the first year average (30%) of the sample. Hence, we estimate market volatility as a function of the average volatility plus a firm's volatility belief:

$$\ln \sigma_P(t) = \ln \bar{\sigma} + \beta(\ln \sigma_P^{Mkt}(t) - \ln \bar{\sigma}) \tag{A.1}$$

A.3 *State Transition Process: Future Expected Evolution*

Picking up on long run random walk time series properties, the uninformed Wiener process reflects the appropriate modeling process. In particular, for long-run dynamics, beyond 24-month maturities, we focus on a Geometric Brownian Motion (GBM) description of future oil price evolution. The GBM follows stationary increments making it time-homogeneous, thereby likening it to a Markov process. Once the GBM stochastic differential equation—where proportional changes in crude price are assumed to have constant instantaneous drift and volatility—is discretized, we adapt the resulting Markov chain to our particular crude oil environment.[4]

Given our interest in modeling far out maturities, we adopt the GBM framework and use Itô's lemma, in discretized time, to derive the process for the change in the price, total cost, and price volatility state variables. For example, the crude price P in continuous time follows:

$$dP = (\mu - \sigma^2/2)dt + \sigma dz \tag{A.2}$$

σ: price volatility

μ: instantaneous drift

dP represents the increment in the crude price process during a time interval dt, and dz is the underlying uncertainty driving the model and represents an increment in a Wiener process during dt.

A.4 *Implied Volatility Term Structure*

Mean-Reverting Volatility

Volatility also follows a strong mean-reverting process: as less information is present in the early months, volatility declines sharply over time, as shown in Figure B.4.[5] This signifies that the 24-month futures contract is a good proxy for the average production

[4]In this framework the process describing the stochastic behavior of the crude price assumes that returns are normally distributed.

[5]Figure B.4 is an estimate of the annualized standard deviations of futures returns for each contract maturity.

price of a well. Additionally, we can use the 24-month option contract's implied volatility to model the stochastic process driving the 24-month expected price dynamics. These mean-reverting expectations allow us to use the 24-month contract instead of all the maturities between the front-month and 24-month maturities, without loss of critical information and with improved tractability. Long-term, less liquid maturities represent production cash flow of less than one quarter of profits. Therefore, when modeling the expected price used by producers, we draw on an estimated time series of 24-month futures prices.

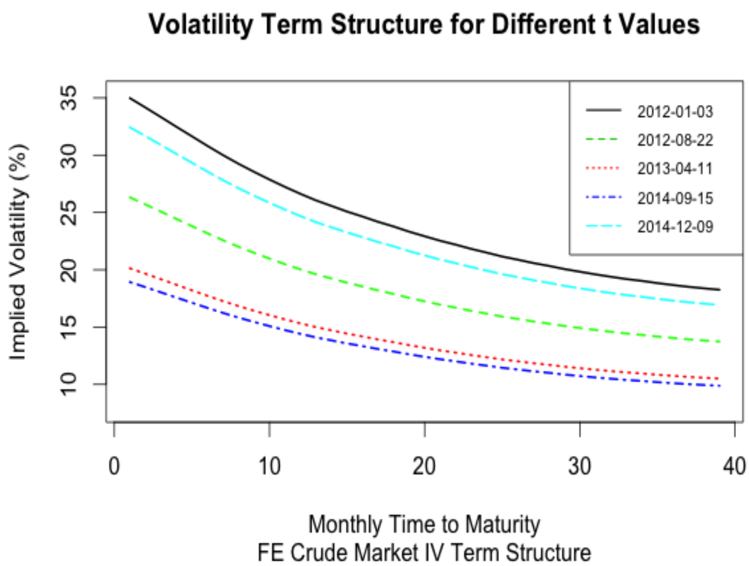

Figure A.3: WTI Fixed Effect Implied Volatility Term Structure

We briefly discuss the liquidity concerns of futures contracts; however expected volatility derived from options is the constraining factor as it is even less traded for further out maturities. Futures are widely traded up to 24 months out but this isn't the case for option contracts. For crude oil, the widely traded future contracts are short-term—from one to 6 months—with an average of 125,548 contracts traded daily. Beyond the short-term, the high volume contracts are always June and December: the 24-month out June and December maturities average 3,543 trades, while the other 24-month out contracts average less than 40 trades a day. Meanwhile, 24-month June and December options are traded less than 100 times in a day. Therefore, we are limited by the traded volume and must construct an implied volatility term-structure (parametric estimation) or surface (non-parametric measure).[6]

We introduce both approaches and show that the non-parametric implied volatility surface doesn't make an assumption regarding the shape of the term-structure for the sample and yields a better fit of firm response to volatility. We start by constructing an implied term structure of volatility by combining short-term option volatilities with an estimated realized term structure. The standard deviation of realized futures price from the crude curve and short-run futures option volatilities—1-, 3-, and 6-month implied volatilities—are combined to estimate the 24-month implied volatilities. Using short-term implied volatility lets the market do the work of forecasting expected future volatility by extrapolating 24-month implied volatility time series over the realized volatility term structure.

We obtain expected futures price volatility from the CME, given that market-implied volatility (Black (1976)) is explicitly traded on the exchange using volatility-capturing option strategies such as calendar and butterfly spreads. However, since the 24-month option contract is illiquid we first look at the annualized standard deviations of futures returns for each contract maturity. We then obtain an estimate of the overall volatility structure of the crude curve—a graphical depiction of volatility for a calendar date across different maturities. Second, applying a principal component analysis (PCA), we estimate the term structure of volatility at each date t in the sample. In the appendix, section A.4.1 goes into greater detail on the derivation of the 24-month implied volatility time series. Finally, as the front month option contracts[7] are liquid, we use them to control for the level of volatility and estimate the 24-month contract.

[6]Estimating implied time-varying non-parametric volatility uses all market information. The implied volatility grid is measured for a range of time-to-maturities and strikes—known as a volatility surface. For an environment with shifting volatility—term structure structural shifts—a non-parametric estimation of a a volatility surface is an improvement over parametric estimation improvement. In fact, it yields a better fit of firm behavior.

[7]The 1-month, 3-month and 6-month are widely traded and are used to control for the level of volatility across the estimated curve structure of future prices.

A.4.1 Parametric Implied Volatility

The PCA technique estimates the covariance matrix between every pair of futures price returns from the crude curve. The eigenvector decomposition yields a set of independent factors, which drives the evolution of the variables underlying the covariance matrix. The associated eigenvalues from the decomposition represent the variances of the independent "factors" which drive the maturities, τ, in proportions determined by the eigenvectors. Therefore, PCA helps identify a number of driving factors per maturity, thereby functionalizing volatility, beyond the fixed effect t. We focus on decomposing futures price returns into three primary driving factors, known as principal components (PCs) or eigenvalues (λ). By selecting the first three PCs, we capture the primary causes of variance across the term structure, ignoring the noisier and less informative factors. Hence, PCA gives us the following functional form, where returns $r_{t,\tau}$ are equal to the first three eigenvalues (λ) times the first three eigenvectors (γ) containing each maturity τ at time t; the result is shown in Figure A.4:

$$r_{t,\tau} \approx \gamma_{t,\tau} * \lambda_{1,\tau} + \gamma_{t,\tau} * \lambda_{2,\tau} + \gamma_{t,\tau} * \lambda_{3,\tau} \tag{A.3}$$

$r_{t,\tau}$ is log returns of each maturity τ over time t

$\lambda_{1,\tau}$ (or PC_1) is the level control of each maturity τ over t

$\lambda_{2,\tau}$ (or PC_2) is the slope control of each maturity τ over t

$\lambda_{3,\tau}$ (or PC_3) is the curvature control of each maturity τ over t

We divide the crude curve into its first three risk drivers (or PCs or eigenvalues), which explain the majority of the curve variance. As depicted in Figure A.4, the eigenvalues (or PCs) resulting from the PCA eigenvector decomposition indicate the importance of each eigenvector and hence the number of components we should include in a crude curve analysis. Figure A.4 plots the eigenvalues for the CME crude oil futures in descending order of size. The figure depicts the first eigenvector as explaining 92.2% of the total variation in the evolution of the crude curve. If we add the two factors it explains 96.2% of the total variation and, together with the 3^{rd} factor, they explain 97.2%.

In this case, the first three factors explain a sufficient amount of the crude curve

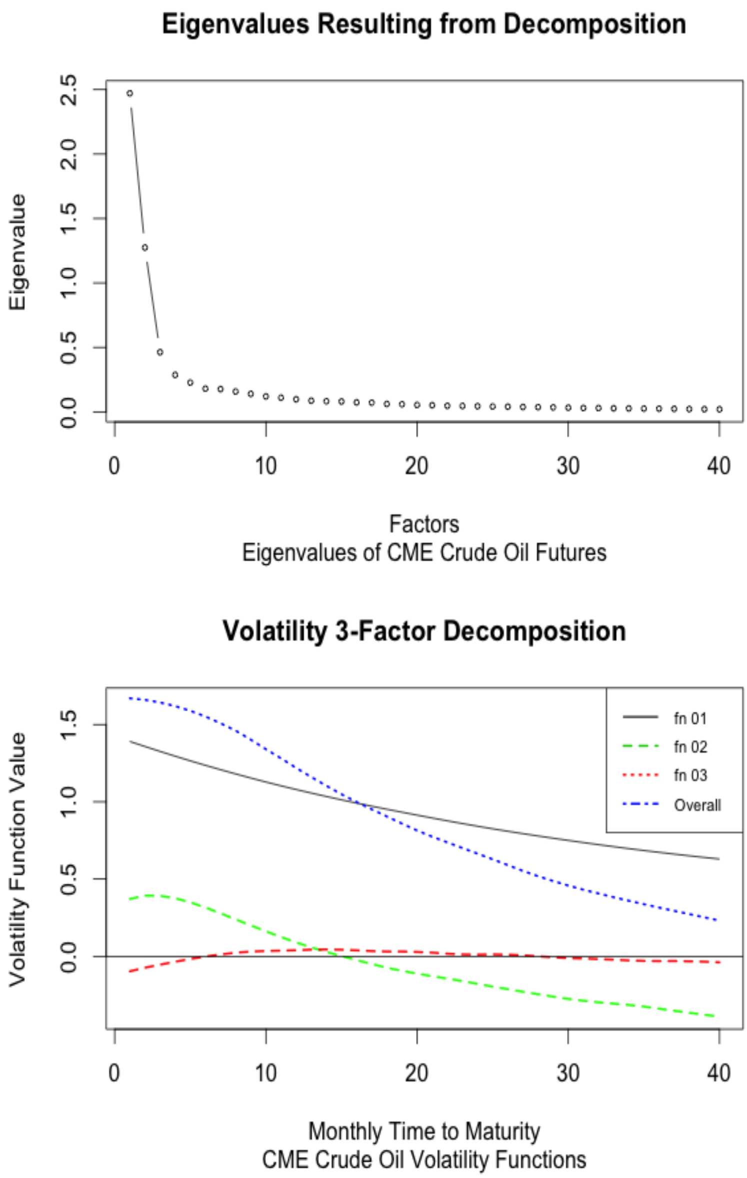

Figure A.4: Forward Curve Volatility Functions

evolution over the 2009-2015 period. The most important factor, PC1, is positive for all maturities, implying that a positive shock to the crude oil market causes all prices to shift up by a marginally different amount. PC2 is the slope factor, which causes the short and long maturity ends of the curve to move in opposite directions. Finally, PC3 is the curving factor, which causes the short and long ends to move counter to the middle section; for the crude curve, these factors are commonly referred to as as level, slope, and curvature, respectively, as illustrated in Figure A.4. Hence, the total volatility is approximated by the combined effect of all three factors, while the remaining factors are viewed as white noise (Clewlow and Strickland (2000)).

A.4.2 Non-Parametric Implied Volatility

Financial options give the holder the right (but not the obligation) to engage in a financial transaction at a predetermined price (strike price) and date (time-to-maturity). For example, an oil call option gives a trader the right to buy barrels of oil at a preset strike price, date, and delivery point.

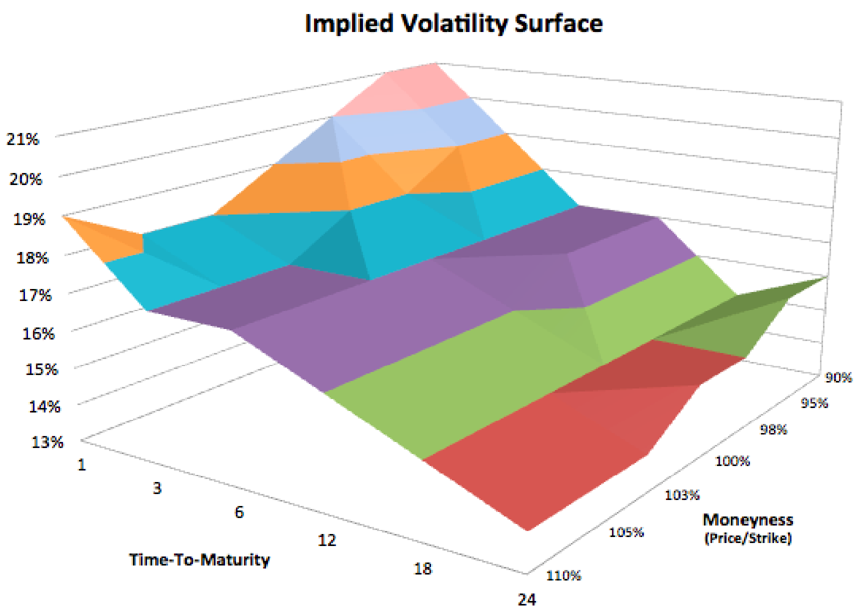

Figure A.5: Crude Oil Implied Volatility Surface for September 2014

We first estimate all missing option prices applying the Heston (1993) and Chen (2011) method and then use the Black equation to measure implied volatility at every available (or traded and estimated) strike and maturity. Because not all possible (call and puts) options are necessarily traded at every strike price and maturity, we need to extrapolate/interpolate.[8] First, for each expiry month (traded time-to-maturity), we

[8]Most of the time, implied volatility is quoted for at-the-money options and is often measured based on the average of an at-the-money straddle, i.e. a call and put with the same strike price.

estimate the missing strike data by non-parametrically interpolating implied volatilities using a quadratic interpolator with Gaussian kernel across expiries.[9] Second, we use cubic-spline interpolation across expiries on this grid to get the implied volatilities for each of the fixed expiry grid points (e.g. 1 month, 2 months, 3 months etc).[10] To estimate the implied volatilities across maturities, we apply the Heston (1993) and Chen (2011) method; the improvement on the Heston model assures that longer-term volatilities are stable. For example, as illustrated in Figure A.5, we can now estimate the volatility surface for September 2014 for a select few time-to-maturities and strikes (or moneyness). Note a 100% moneyness (oil over strike price) specifies that the underlying oil price is equal to the specified (buying/selling) option contract strike price. The recovered implied volatility for 24-month TTM and 100% is our implied volatility $\sigma_P(t)$ deduced from markets.
Under the Heston model, crude oil price can be made to follow the process:

$$dS_t = \mu S_t + \sqrt{\nu_t} S_t dW_t^S \tag{A.4}$$

where ν_t, the instantaneous variance, is a Cox-Ingersoll-Ross (CIR) process

$$\begin{aligned} d\nu_t &= \kappa(\theta - \nu_t)dt + \xi\sqrt{\nu_t} dW_t^\nu \\ dW_t^S dW_t^\nu &= \rho dt \end{aligned} \tag{A.5}$$

and dW_t^S, dW_t^ν are Wiener processes with correlation ρ

μ: s the rate of return of the asset

θ: is long-run variance

κ: is speed of mean reversion of ν_t to θ

ξ: is vol of vol

ν_0: is initial variance[11]

Although a numerical integration can be used to evaluate the model, we solve it using a closed form Fourier Fast Transform. As stated above, the model allows for non-Gaussian

[9] This is similar to a curve fitting model, which will return a smooth curve that runs very close to all the traded option (input) data points. Hence, options prices in this model are given simply by a weighted sum of Black prices. Implied volatilities are found by numerically inverting the Black formula.

[10] If we wanted, we could also linearly interpolate along the fixed expiry grid points to get the implied volatilities for the corresponding maturity τ at time t. See Chen (2011) "Mixed-lognormal Volatility Extrapolation Adjustment" to see the steps for extrapolating call/put prices when missing. Beyond the last market maturity, he calculates parameters for mixed-lognormal method.

[11] All parameters (μ, κ, θ, σ, and ρ) are state and time homogeneous.

specification of oil returns.[12] In particular, ρ, the correlation between the returns and volatility of oil, impacts kurtosis and skewness, while ξ influence the kurtosis of the distribution. When option price data is unobserved, the calibration of these parameters influences the interpolation and extrapolation of implied volatility.[13] Furthermore, κ is a mean-reverting parameter, which captures the levels of volatility clustering. These κ, θ, σ, ν_0, ρ parameters are calibrated by minimizing the error, using a non-linear least-square optimization, between our model specification and the market price using N options.[14]

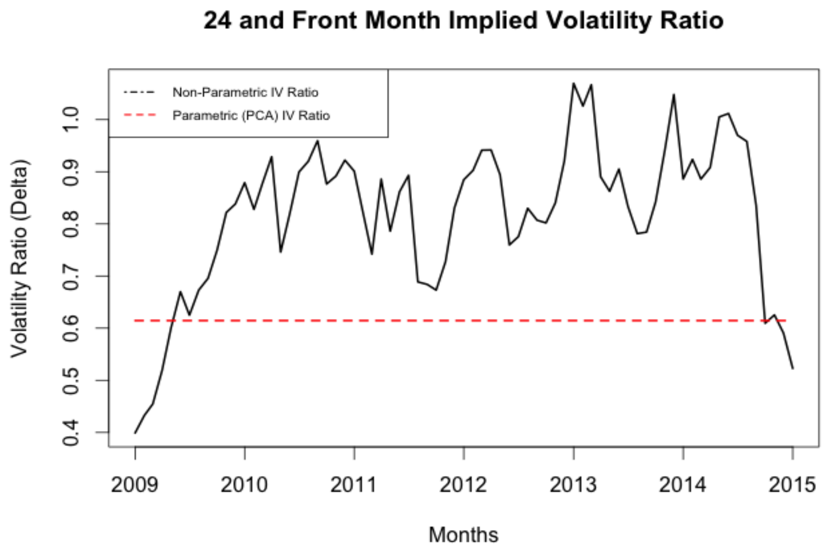

Figure A.6: Crude Oil Parametric and Non-Parametric Volatility Ratio

We compare the parametric and non-parametric estimation of front month and 24 month time-to-maturity volatility ratio $\frac{\sigma_P(t)^{24}}{\sigma_P(t)^{Front}}$. The parametric approach estimates the 24-month implied volatility ($\sigma_P(t)^{24}$) by extrapolating front-month ($\sigma_P(t)^{Front}$) applying an average volatility ratio ($\frac{\sigma_P(t)^{24}}{\sigma_P(t)^{Front}}$) over the sample.[15] The non-parametric measure of

[12] Circumvents the shortcoming of Black's formula constant volatility and application of Gaussian distribution.

[13] For example, ξ enhances the distribution skew (or volatility smile). Specifically, an increasing ξ implies higher volatility of volatility.

[14] We use the average between the bid-offer spread for the option price.

[15] Using fixed effect for time-to-maturity to control for the level of volatility on each date t still only

$\sigma_P(t)^{24}$ is based on the re-estimation of the implied volatility surface (Figure A.5) and the volatility ratio every period; the volatility ratio then becomes: $\frac{\sigma_P(t)^{24}}{\sigma_P(t)^{Front}}$. Hence, when measuring expected price volatility, it is more accurate not to constrain the estimation by assuming that the volatility term-structure is stable, by measuring a parametric form, for the sample.[16]

A.5 Total Cost per Well

For our real option modeling, of primary importance is estimating the total cost level when data is missing as it impacts the empirical log-likelihood fit. Of secondary importance is the accuracy of our total cost per well estimates as it affects the production (bbl) threshold—the production level that leaves a producer indifferent between investing now or waiting—required to invest in a well. Figure A.7 shows our estimate of monthly total cost per well. To estimate monthly total costs per well, we combine that year's total cost from a recent Energy Information Administration (EIA)-commissioned IHS Global Inc. (IHS) study with the corresponding monthly drilling and completion ratio using *Spears & Associates DCS Report* and *Day Rate Report* surveys. The IHS report measures total cost per well as a unit cost with constant characteristics, which is then also estimated with changing factor costs (e.g. pumping, drilling, proppant, and fluid casing). The primary cost item in drilling is rig rental and in completion is fracking costs. The *Spears & Associates DCS Report* provides us with quarterly fracking cost data for Eagle Ford as of 2012 and for the Permian starting in 2015, while the *Day Rate Report* has rig rental data from 2009-2012. Rig rental data is quarterly from 2009-2012 and monthly from 2012-2015.[17]

To estimate monthly total costs per well, we combine that year's total cost per well

---

gives us the average 24-month tenure (τ) dynamics. That is, every month, a contract (e.g. a 2 month time-to-maturity) rolls to the adjacent contract (e.g. 1 month time-to-maturity) and our analysis constrains us to using months as well investment data is monthly.

[16]For Eagle Ford, we compare the model fit and estimate a log-likelihood fit of −6,365.13 for parametric and −6,362.99 for non-parametric estimation of 24-month implied volatility.

[17]Note that for late 2014 and early 2015, our estimate are in-line with the EIA's IHS commissioned report when we apply drilling and completion ratios from the *Spears & Associates DCS Report* where costs increase by a little less than half a million dollars and, later in 2015, start going downwards.

with the corresponding monthly drilling and completion ratio. When only drilling rental data is available, we normalize the year's total cost per well by drilling rental changes. Lastly, we interpolate missing monthly data early in the sample using a polynomial regression fit (or spline) between quarters. From Figure A.7, we note that total cost per well went up from 2009 to 2012 because of increased well investment activity. Note that the productivity of inputs increased, that is, according to the IHS report as of 2012, total cost per well for a given design and overall well investment activity decreased—vertical, oil, and natural gas wells.

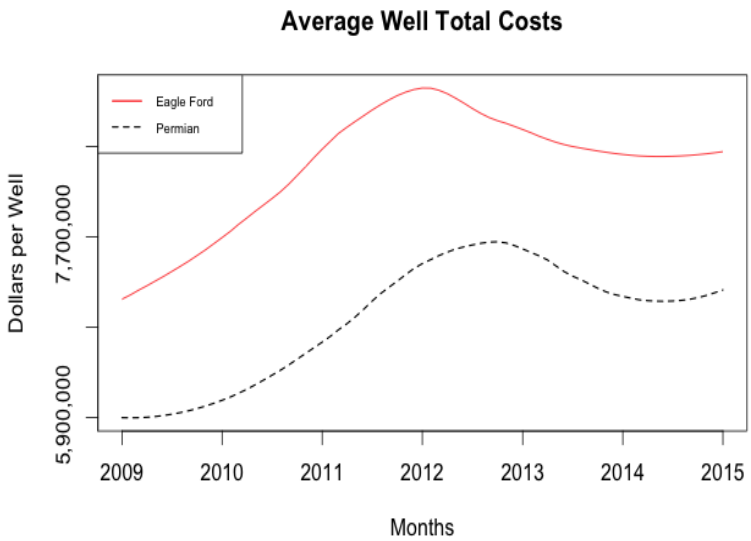

Figure A.7: Eagle Ford and Permian Total Cost

Previous mentioned, shale oil fixed cost is divided into two categories: drilling and completion. The remaining costs—such as setup, fluid, transportation and fuel, and service rental equipment—are less variable and can therefore be modeled as a function of rig costs and steel prices.[18] Rig rental (or dayrates) information is published in the

[18]From Federal Reserve Economic Data (FRED) and Quandl, we know that steel was stable over the period, with a low of $1,517(real 2015 US$) and high of $2,920 per gross tons.

Day Rate Report firm survey.[19,20] The reports have quarterly rig cost data from 2008 through the third quarter of 2012, beyond which the data is monthly. To convert the initial quarterly data to months we first equate the months at the center of the quarter to the quarterly dayrates and then use polynomial regression fit to interpolate the dayrates for the intervening months.[21]

As shown in Figure A.7, total cost per well differs from basins, these differences are captured in rig rental rates. Rig dayrates differ from basin to basin: Eagle Ford uses rigs that drill on average 10,000-13,000 feet deep, running on about 1,000-1,500 horse power (hp), and averaging $19,026 (real 2015 US$) a day, while the Permian digs on average 6,000-10,000 feet in depth, running on 500-1,000 hp, and averaging $14,463.[22] As depicted in Figure A.8, in 2009—when horizontal drilling was primarily used in natural gas and the commodity's price plummeted from oversupply—horizontal drilling demand went bust but, as drilling and fracking technology moved from shale-gas to shale oil plays, rig demand and therefore dayrates picked up.[23]

A.5.1 Drill Rental Days

In Figure A.8 and A.9, we observe that the reduction in number of days a rig is rented contributed to the shale oil revolution from 2009 to 2015, undoing any upward cost trend in rig rental. In drilling, innovation translated primarily into a decrease in the average drilling time, which went from 32 to 15, and 33 to 23 in Eagle Ford and the Permian,

[19] Surveyed drillers for the period report that contract terms are generally mixed between term, multi-well, and spot contracts; towards the end of the period, more term contracts were signed. The mix of contract specifications and the varied start and end dates of term contracts allow the surveyed prices to approximate spot price dayrates.

[20] Rig rental data is drawn from Rigdata's 2Q2010, 1Q2012, and JAN2015 *Day Rate Report*.

[21] Alternative methods, such as equating each month within a quarter to the quarterly rate and smoothing, do not statistically impact the results.

[22] Eagle Ford rigs drilling from 10,000-12,000 feet and 13,000-15,000 feet cost $16,966 and $21,087, respectively, while Permian rigs cost $13,182 and $15,543 for 6,000-9,000 and 10,000-12,000 feet, respectively. Even when we control for depth, Permian rigs trade at a discount to Eagle Ford. Eagle Ford's higher rig rental costs are the result of higher demand for shale oil services and the Permian—which was a conventional play prior to being a shale oil one—being closer to drilling service hub.

[23] From 2012-2013, 35% of surveyed producers reported they had to wait a month or two before getting a rig. However, there isn't any data to control for waitlists systematically, aside from adding a lag between the decision to drill and commencement of drilling.

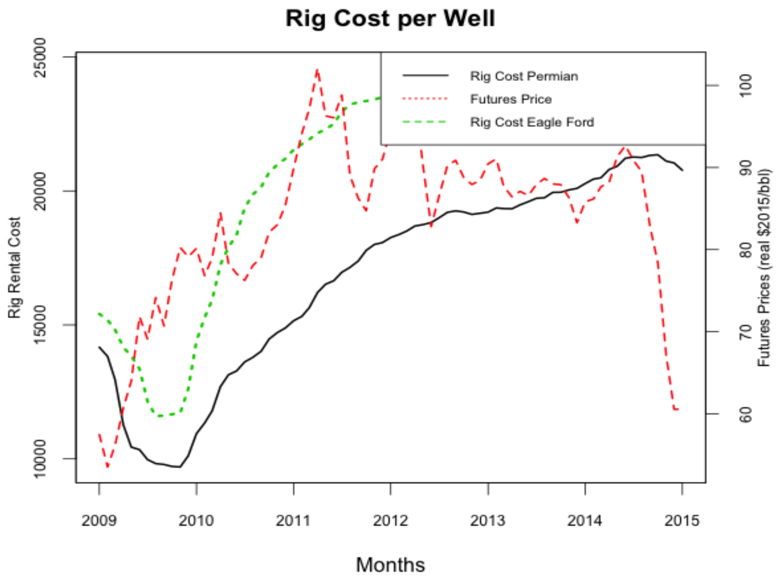

Figure A.8: Drilling Cost and Price

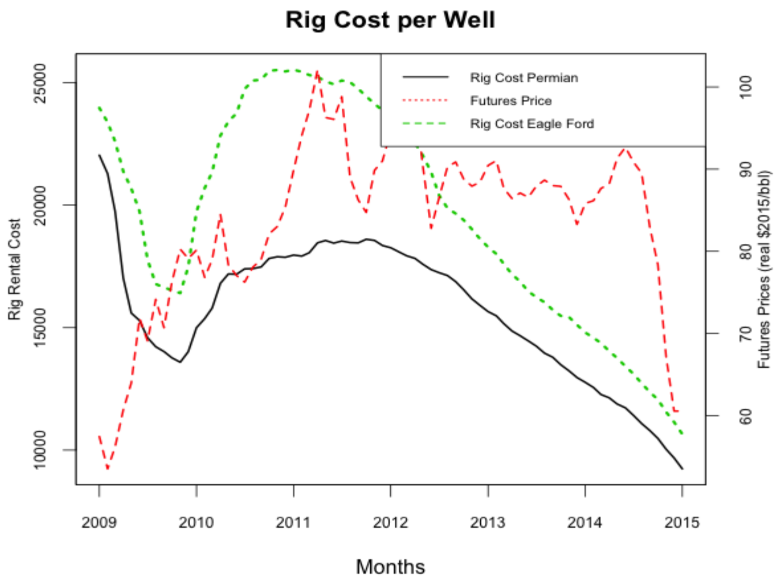

Figure A.9: Day Count Adjusted Drilling Cost and Price

respectively.[24] For illustrative purposes, we linearly normalize drilling rental cost by the number of days using a 22-day sample average to estimate rental rates for the period. For

[24]The decrease in number of days to drill is due to improved setup time and drilling techniques, such as shorter transport distances between wells and platforms that can handle up to ten more wells.

example, the resulting mean day-adjusted rig rental cost for the sample was $21,664 in Eagle Ford and $20,146 in the Permian, as illustrated in Figure A.9. In 2009, when shale gas equipment shifted towards shale oil, there was an initial drop in rental costs; however, as demand for shale oil picked up, costs initially increased. Finally, with increased rig efficiency, the number of days a rig was rented fell and rig costs adjusted for days did as well. When measuring total cost per well, we control for the reduction in the number of days a rig was rented using estimates from the Energy Information Administration (EIA) commissioned IHS Global Inc. (IHS) study.

A.5.2 Well Productivity

To estimate productivity, we look at the production gains for the first month of production (bbl) post completion. To get a stable estimate of productivity, we average the first three months of a producing well and refer to that as first month production. For our estimate, we select wells on a single lease because wells are not flow metered but leases are. The estimate, shown in Figure A.10, gives us a direct measure of well productivity over time. A well's (consistent) production hyperbolic decline and the fact that the first two years of production represent half to thirds of a well's production makes our productivity measure not only an appropriate estimate of productivity, but also of profitability.

To empirically fit the model and control for increases in expected productivity ($\mathrm{E}(x_{it})$), we normalize total cost per well by improved production per well. This allows us to measure a "true" producer's costs per well.[25] We illustrate these normalized total cost per well estimates in Figure A.11 and A.12 for Eagle Ford and the Permian, respectively. In late 2012-2013, productivity growth in Eagle Ford and the Permian came on the heels of increased drill and completion costs. Bigger rigs and increased fracking stages were used, which explains the rise in costs and continued productivity gains. Pre-2013[26] cost declines in Eagle Ford are in-line with logistical productivity gains from innovation, as

[25]The normalization captures the idea that deeper and longer lateral wells increase cost but also improve well productivity.

[26]Production efficiency is in line with the Energy Information Administration (EIA) June 2015 "Drilling Productivity Reports." Also, for the period we document a similar per well production efficiency ratio of 1.25 with respect to both Eagle Ford and the Permian.

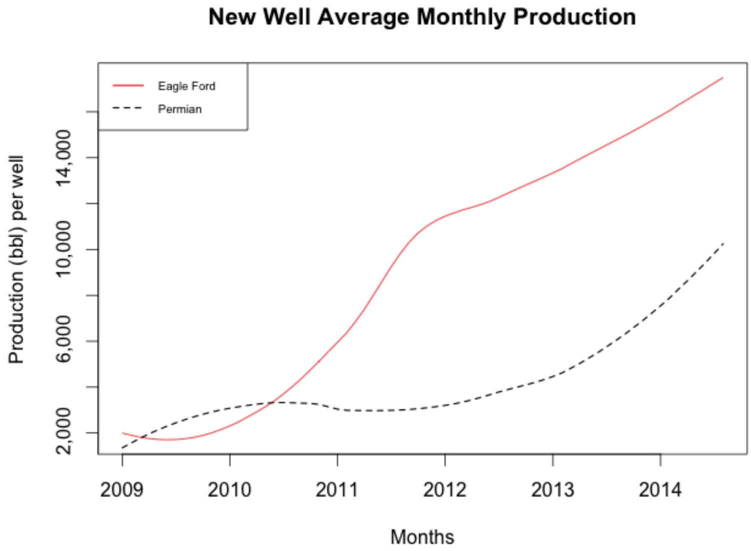

Figure A.10: First Month Productivity (bbl) per Well

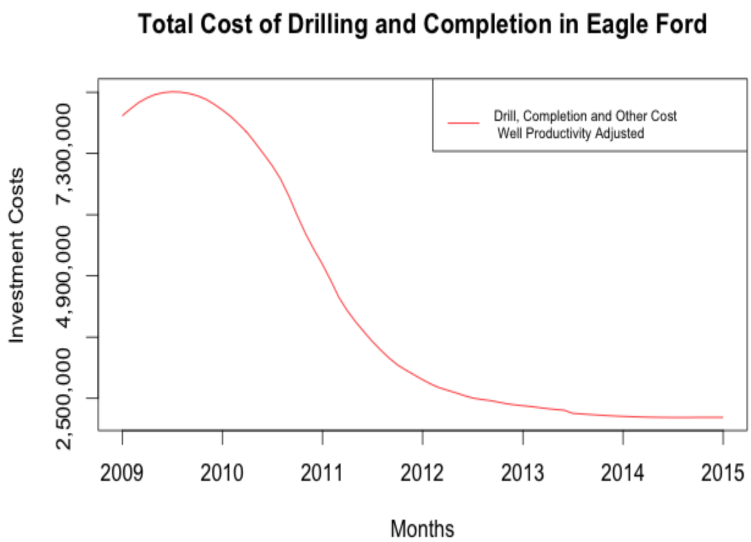

Figure A.11: Eagle Ford Cost Adjusted by Productivity per Well

shown in Figure A.10, while for the same period Permian, total costs were lower and productivity took some time to pick up. Permian low initial productivity gains can be explained by well investments being more spread out and fewer horizontal shale oil wells

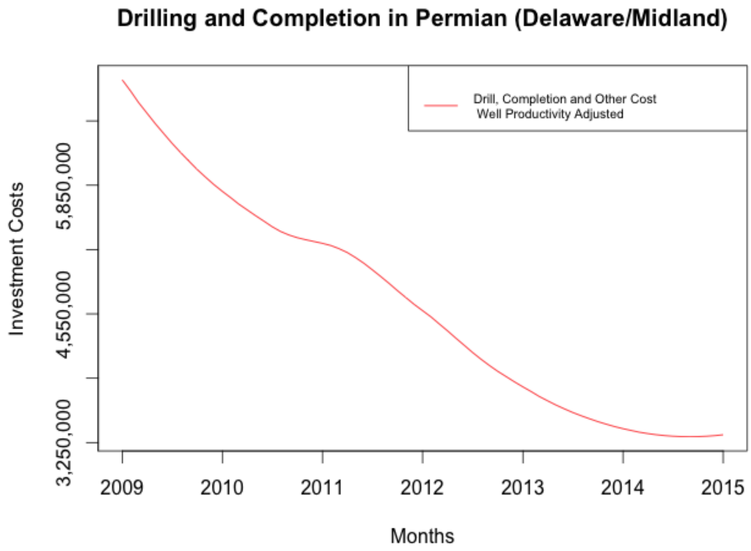

Figure A.12: Permian Cost Adjusted by Productivity per Well

being drilled on average for the sample—141 in Eagle Ford versus 50 in the Permian.

A.6 Descriptive Decision-Making

We first motivate the relationship of completion activity in a volatile price environment and later discuss firm responsiveness to market-based expectations of volatility. To examine the relationship between well drilling and completing activity, and crude price and volatility, we look at small firm completed well activity in two leading Texas shale-oil basins—Permian and Eagle Ford—for the period 2009-2015. Encouraged by an initially high-price and low-volatility environment, shale-oil technological innovation cascaded in across the basin as firms cooperated and shared information—incentivized to increase lease values to be sold in a secondary market. Shale-oil growth took off in the Permian and Eagle Ford, which are the biggest and most productive shale oil producing basin. The relationship between independent price and volatility variables, and the dependent well activity variable is shown in Figure B.5.

We find that volatility and price are negatively correlated throughout the period and fluctuate quite a lot: they start off high and low, respectively, and then quickly switch—with price high and volatility low for most of the remaining period. At the tail end of the period, prices drop and volatility spikes. In fact, at the beginning of the period, dynamics are dictated by the expected logistical drop in costs,[27] while for the remainder of the period price and volatility are the primary drivers.

According to Dunne and Mu (2010), the hazards model is widely used in the literature to measure the effect of uncertainty on the timing of energy investment projects; they use the approach to analyze lumpy refinery capacity expansion under uncertainty. Hence, using a simple regression analysis we explore the negative relationship between drilling costs and expected price volatility and the positive relationship with price, as depicted in Figure B.5. The unit of analysis to infer these dynamics is an individual drilling prospect, where we control for fields. As discussed above, we estimate a lower bound of around 18,102 prospects for Eagle Ford in 2009-2015. During the period, out of the feasible prospect set, 10,353 oil wells are drilled and completed in Eagle Ford.

Shale oil sediment rock formations are geologically mapped out early in the sample,[28] suffer minimal common-pool infield concerns, and are effectively homogeneous over several mile ranges.[29] Furthermore, the data confirm that within field drilling and completion does not often occur serially and thus, within field drilling prospects are considered independent from each other.[30] The independence of drilling and completion decisions within the same field is also due to capacity constraints, such as gathering equipment construction lags and capital and labor prior project commitments. Hence, drilling activity within the same field is modeled independently, given that they are mainly driven by exogenous drivers such as price, volatility, and cost (we do not account for common factors that might drive the simultaneous decision to drill and complete more than one well within a

[27] The same rigs are used to drill and frack multiple wells using the same processes in similar locations. The factory drilling and fracking approach, as with many repeated manufacturing processes outlined by Grübler (1990), Benkard (1999), and Gruber (1992) generates strong productivity gains.

[28] Shale oil plays were discovered and mapped centuries ago; recent technological innovation made them feasible.

[29] For the sample period, industry experts contend that "factory drilling" applied similar technology and practices to all prospects within a basin, avoiding experimental risk; returns were so high with current techniques that there was little incentive for high-grading in a high-price environment.

[30] Throughout the sample period, Eagle Ford and the Permian combined drill and complete only 7.5% of within-field wells in quick succession.

field serially, finding the effect to be negligible).

Like Dunne and Mu (2010) and Kellogg (2014), we test whether price increases and volatility and cost decrease the likelihood of drilling and completing by running a hazard regression and then construct a structural model to capture causality.[31] The binary LHS variable is whether a firm drilled and completed a prospect or not and is represented by κ:

$$\kappa(t) = e^{\beta_0 + \beta_p Price_{t-2} + \beta_v Vol_{t-2}} \tag{A.6}$$

Taking logs, we obtain the additive log-linear model:

$$\ln \kappa(t) = \beta_0' + \beta_p Price_{t-2} + \beta_v Vol_{t-2} \tag{A.7}$$

$Price_{t-2}$ & Vol_{t-2} covariates are lagged by 2 because there is a 2-month wedge between initial decision to drill and final completion of a well.[32]

All covariates are lagged by 2 months as permitting, rig contracting, and drilling and fracking create a wedge between the initial decision to drill and completion. Hence, statistical inference with clustered data is misleading if we do not account for serial correlation and field-specific spatial correlation problems that emerges with clustered panel data. We apply a robust covariance matrix estimator to correct measured standard errors for within field correlation of the likelihood scores.

For Eagle Ford, the relationship between independent price and volatility variables and the dependent well activity variable is shown in Figure B.5.

A.7 Real Option Model Intuition

After numerically simulating Equation (2.5) with state space price (P), price volatility (σ), and cost (TC), we notice that a firm's failure to heuristically account for volatility using a framework that mimics that of a real option leads to different completion state variable trigger rules.

---

[31]For the same base specification of price, volatility and cost covariates we run an OLS specification. The OLS results confirm our findings where we regress log of wells drilled each month on the oil price, volatility, and drilling cost to estimate coefficients of 0.063, -0.031, and -0.158, respectively (and -0.308, in \$1M, using total cost indexed by productivity).

[32]Other lags were tried: 2 months provided the best fit and is in line with market dynamics.

Table A.2: Hazard Function: Eagle Ford Likelihood of Drilling & Completion

	Base Model (1)	Total Costs (2)	Prospect FE (3)	Total Cost & Time Trend (4)
Oil futures price ($/bbl)	1.065***	1.042***	1.045***	1.046***
	(0.018)	(0.016)	(0.021)	(0.021)
Implied volatility (%)	0.976**	0.973***	0.972**	0.972**
	(0.014)	(0.012)	(0.012)	(0.012)
Total Cost ($100,000)	–	0.853***	0.897***	0.823***
	–	(0.014)	(0.014)	(0.015)
Linear Time Trend (Months)	–	–	–	0.971*
	–	–	–	(0.020)
Field FE (Field Heterogeneity)	NO	NO	YES	NO
Log Likelihood	−11,524.7	−11,517.7	−11,517.1	−11,472.2

Notes: The exponential coefficients are interpreted as multiplicative effects on the hazard rate of a one unit increase in the covariates. All estimates use 24 month futures prices and option volatilities. The models are estimated using hazard functions. Covariates are lagged by 2 months. Field fixed effects and time trend dummies are applied. Also, standard errors are clustered at the field level. *p<0.1; **p<0.05; ***p<0.01

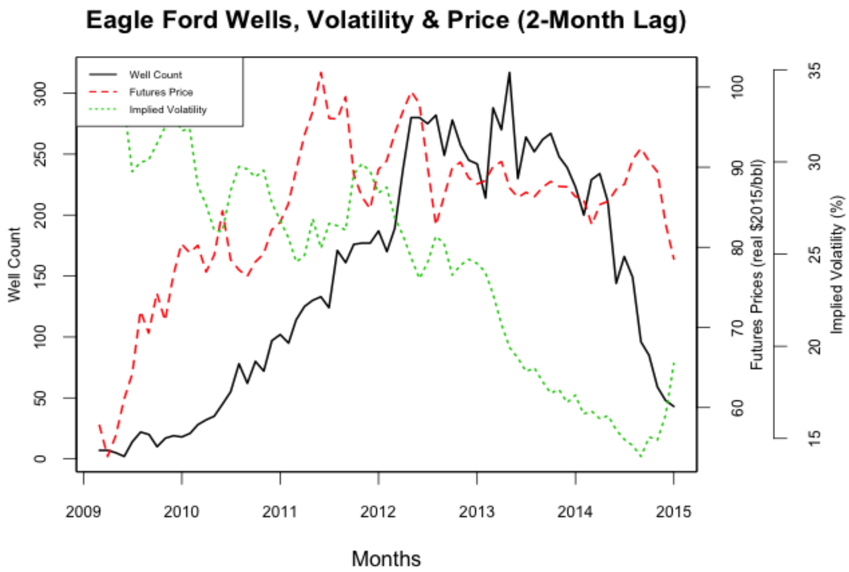

Figure A.13: Monthly Well Count, Futures Price & Implied Volatility

As illustrated by the real option payoff depicted in Figure A.14, for a given expected production of 40,000 barrels (bbl) and an average cost of $23/bbl, if a producing firm

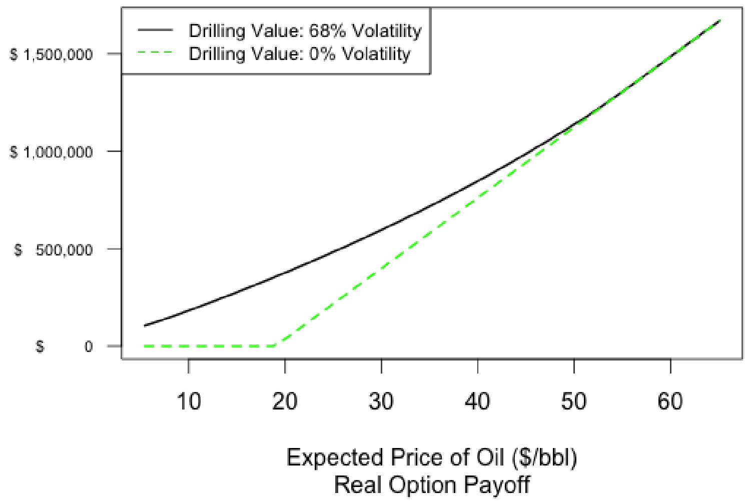

Figure A.14: Drilling Expected Payoff

assumes a volatility belief of 0% when, in actuality, volatility is 68%, the producer would make a $1 million mistake—completing a well at a price of $20/bbl with 0% versus $50/bbl for a volatility of 68%. Hence, failure to account for volatility in a dynamic real option model is costly and may lead to bankruptcy in the long run. We go on to show that producer behavior incorporates time-varying market-based volatility, setting β responsiveness equal to one.

A.8 Log-Likelihood

From Table A.2, we first run the base model with only price and volatility covariates. The results in column 1 are statistically significant at the 1% level and confirm that a $1 increase in expected price increases the likelihood of drilling and completion by 6.5%, while a 1% increase in expected price volatility decreases the probability of drilling and completing by 2.3%. In column 2, we add total cost per well and find a statistically significant, at the 1% level, negative effect in line with theory: a thousand dollar increase in cost decreases the likelihood of drilling by 15%. Finally, we progressively include within-field heterogeneity and a time trend and show that signs and magnitude are robust correlation structures with expected price positively, and volatility and cost negatively, impacting drilling and completion activity.

The descriptive statistics show that the correlation structure of volatility is robust to model specifications, but offers no causal inference. To test whether high stakes shale oil firms' stochastic decision-making is in line with the theoretical optimum, we develop a structural model. We can then draw causal inferences as well as estimate policy impacts on the decision to drill and complete. Thus, the real option structural model allows us to measure the importance of volatility on high stakes investing, whether firms respond to market expected price and volatility state variables, and the distortionary impact and costliness of policies such as the export ban on well activity.

To motivate the state variables price, volatility, and cost used in the real option model, we conduct a descriptive statistical analysis using a hazard function and report the result in appendix A.6. Our findings show that all three state variables are statistically significant when explaining well count.

In Equation (2.5), current profitability π_{it} implicitly accounts for current volatility $\sigma_P(t)$ used to initiate future expected price, price volatility, and cost evolutions. Current profitability is not modeled using any explicit price or cost probabilities of switching to new price and cost states. Current price and cost are estimated using 24-month futures contracts and current cost respectively. Hence, we write current profit as a function of only the state variables P_t and TC_t, omitting $\sigma_P(t)$ in $\pi_i(P_t, TC_t)$. The value function (2.5) shows that $\sigma_P(t+1)$ impacts the expected distribution of future payoffs from drilling in the future. As $\sigma_P(t+1)$ increases, the variance of P_{t+1} increases with respect to P_t, thereby heightening the value of waiting to drill and complete a well. By juxtaposing current and future evolution of expected profits, we solve for the trigger rule that gives us a unique level of expected productivity (x_{it}) that leaves a producer indifferent between investing now or waiting—e.g. when $(x_{it} \geq x_t^*)$, a producer will drill now and not wait.[33]

[33]Producer investment threshold (x_{it}) depictions and consequences of uneconomical investing from a failure to incorporate volatility are illustrated in Appendix A.7. When the option to delay is valuable, failure to properly integrate volatility leads to uneconomical well investments that don't account for the real option value. Our state space stochastic process and a Dixit and Pindyck (1994) fixed point contraction mapping argument are sufficient for a trigger rule x^*(P, D, σ) to exist.

To measure producer responsiveness to volatility we designate a free parameter β. Applying the dynamic programming algorithm, we first fit the β firm volatility behavioral parameter and the normal distribution parameters μ and ζ for the log of expected productivity $log(x_{it})$. A β equal to one indicates that firm behavior is in line with time-varying market-based volatility, while a β of zero signifies constant volatility. Volatility beliefs that differ from market indicate either a forecast that is orthogonal to the market's volatility or a firm's failure to correctly specify its investment decision. The results and robust standard errors are displayed in Table 2.1.

The reported log-likelihood (LL) fit[34] in Table 2.1 measures the probability that x_{it} exceeds the trigger productivity given time-varying realized state values (P,TC,σ), probabilities of transitioning between these realized price and cost states, and behavioral fitting parameters (β, μ, ζ).

$$\log \mathcal{L}(\theta \mid x) = \log \mathrm{Prob}(x \mid \theta) \tag{A.8}$$

$$\begin{aligned} \log \mathcal{L}(\beta, \mu, \zeta \mid P, TC, \sigma) = \sum_{t=1}^{T} \Big[N_t \log \mathrm{Prob}(I_{it} = 1 \mid P, TC, \sigma; \beta, \mu, \zeta) \Big] \\ + N_T \log \mathrm{Prob}(I_{it} = 0 \; \forall t \mid P, TC, \sigma; \beta, \mu, \zeta) \end{aligned} \tag{A.9}$$

I_{it} equal to 1 if prospect i is drilled in month t and is 0 otherwise

N_t is the number of wells actually drilled at t

N_T is the number of prospects not drilled for the whole period T

In equation (A.9) we write the probability that a prospect is drilled at $t = 1$ as the conditional probability that it is drilled at $t = 1$ multiplied by the probability that it was not drilled at $t = 0$. The objective is to fit an estimate of β, μ, and ζ that maximizes LL(β) by locating the optimal $\hat{\beta}$. Numerically, the maximum can be found at the point where there is no further increase for a predefined tolerance.[35] For a more complete

[34]The estimation equation A.9 follows the same procedure outlined in Kellogg (2014).

[35]Each iteration estimates a new $\hat{\beta}$ and log-likelihood value, which is higher than at the previous value until the desired tolerance between steps is achieved. Additionally, we apply the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method to approximate an arc-Hessian (Train (2009)), as the log-likelihood function can be non-quadratic. The method uses information at more than one point on the likelihood function to estimate the curvature.

visualizing of the algorithm's estimation procedure used to solve and fit the model see appendix section A.10.

A.9 Diffusion Process Drift Term

Although we find a random walk drift, the price and cost drift terms are functionalized to test whether there is any attenuated mean in the long-term—beyond the 24-month maturity. Price and total cost drift terms are functions of both level and volatility. This allows us to test and capture the possible attenuated mean-reverting nature of forward prices beyond 24 months. First, as discussed by Clewlow and Strickland (2000), heightened spot volatility affects the convenience yield of holding an additional unit of crude inventory—marginal value of storage. Expected storage value—the difference between injected spot and expected withdrawal future price—increases with spot price volatility; that is, higher volatility increases the probability of spot and future spreads expanding. In fact, as price volatility increases, there is higher inventory demand used to smooth crude supply. Therefore, price volatility leads to inventory buildup and raises expected prices. Second, as we show in this thesis, increased volatility raises a shale-producer's investment option—opportunity cost of completing now versus postponing. Similarly, the volatility drift term is a function of volatility because of the mean-reverting nature of the term-structure of volatility.

A.10 Dynamic Programming Algorithm Depiction

Rust's (1987) dynamic programming hierarchical structure with an inner and outer loop is depicted in Figure A.15. The transition matrix is used by the dynamic programming algorithm to estimate future payoffs from postponing drilling and completion. The 3D matrix's state space estimation is depicted in Figure A.16.

The transition matrix is a 3D matrix of state transition probabilities, where we estimate a price and cost transition matrix for some volatility level. The structural value function draws from the matrix when numerically solving the future payoffs from postponing drilling and completion.

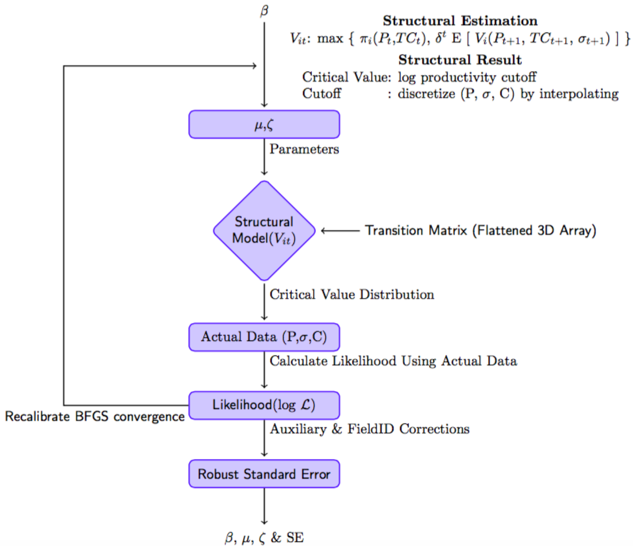

Figure A.15: Outer(Parameter Values)/Inner(Structural and LL) Loop

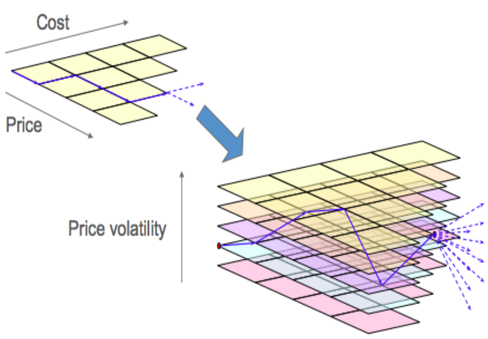

Figure A.16: Price, Cost, and Volatility Transition Array

Appendix B
TECHNICAL DETAILS, CHAPTER 3

B.1 Transportation Basis

B.1.1 Motivating Transportation Basis

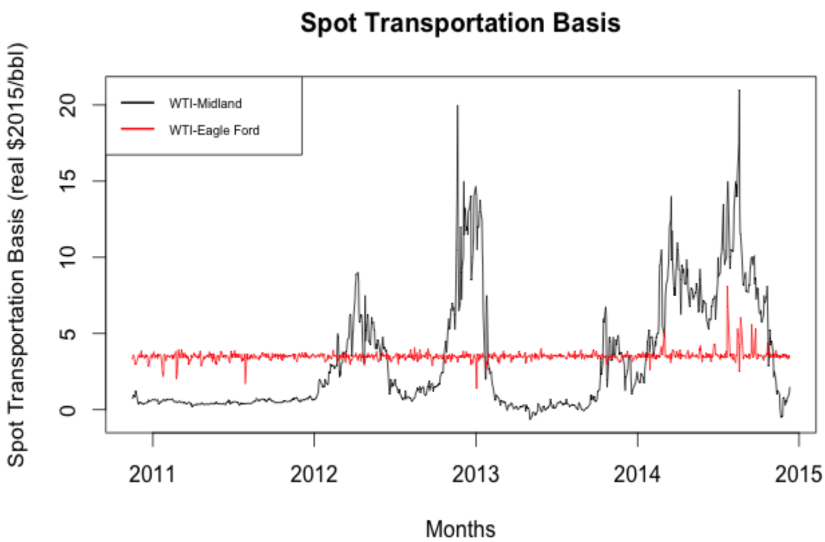

Figure B.1: Eagle Ford and Permian Spot Transportation Basis

Figure B.1 illustrates Eagle Ford and the Permian's respective stable and unstable transportation bases. The Permian's spot transportation basis shows signs of reversion to a long-term mean with the occasional jump from congestion—due to demand for transportation outstripping supply. Without pipeline congestion between the Permian and Cushing, transportation basis equals around \$1/bbl. In 2014, when drilling activity took off in the Permian, production doubled from 2011 levels. Growth in production was expected to match increases in pipeline capacity. However, with midstream capacity pipeline expansion projects not underway until early 2013, there were unexpected delays coupled with larger than expected shale oil volumes, which lead to congestion. This led transportation basis to go as high as $15/bbl. Signifying that producers made well investment decisions using a lower Permian Midland price.

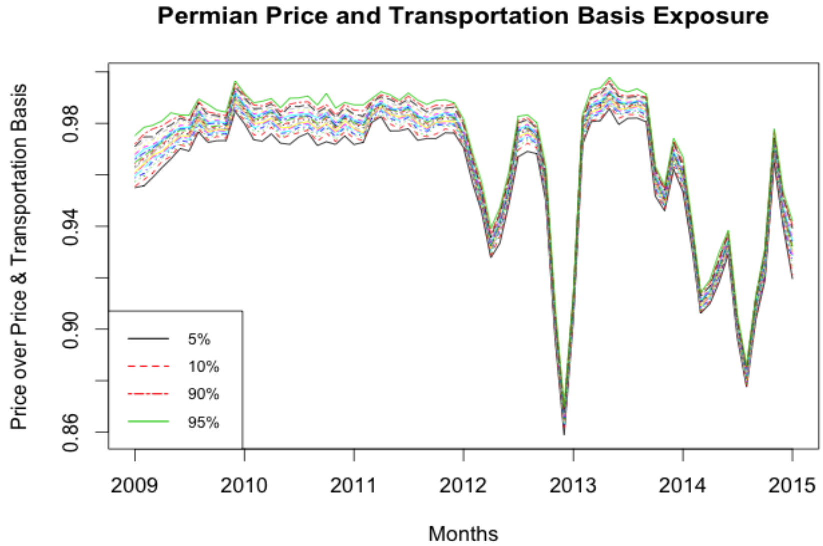

Figure B.2: One Month Ahead Price and Transportation Risk Identification

Figure B.2 simulates a month ahead quartile analysis of price and transportation and its percent contribution to producer revenue. In the beginning of the sample, price contributed 97% to revenue risk. However, towards the end of the sample congestion in pipeline capacity increased the importance of transportation basis, which then contributed as much as 14% to revenue risk.

Figure B.3 illustrates a 30-month out of sample simulation of transportation basis exposure—a quartile distribution of 24-month out contract—using our fitted mean-reverting jump diffusion (MRJD) from Appendix B.1.3. Note that the 2009 simulation period is relatively free of congestion as shale oil production was just starting to pick up. For 2010, a spike in transportation basis was correctly simulated with a skewed confidence interval between $1 & $6/bbl and an expected transportation basis of $2.05/bbl.

To illustrate the impact of including transportation basis exposure, we simulate forward transportation basis with crude curve prices. Although the 2009-2011 period was

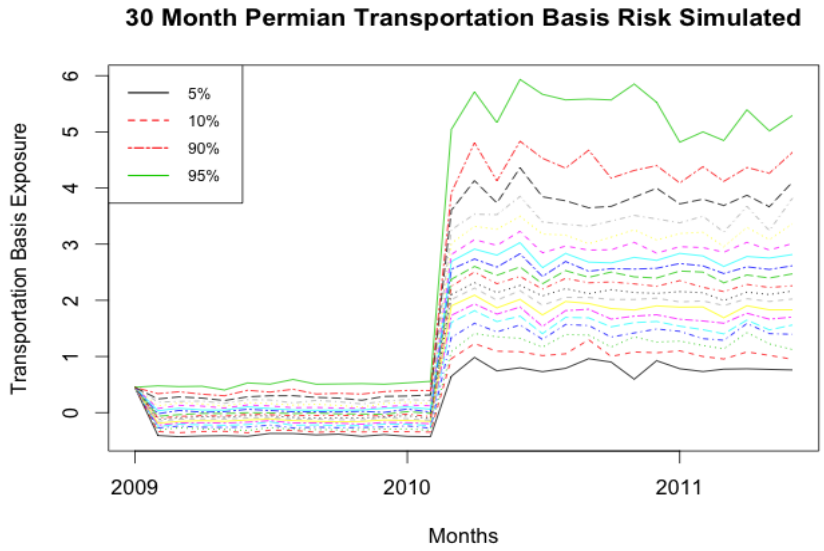

Figure B.3: 30-Month Forward Simulation

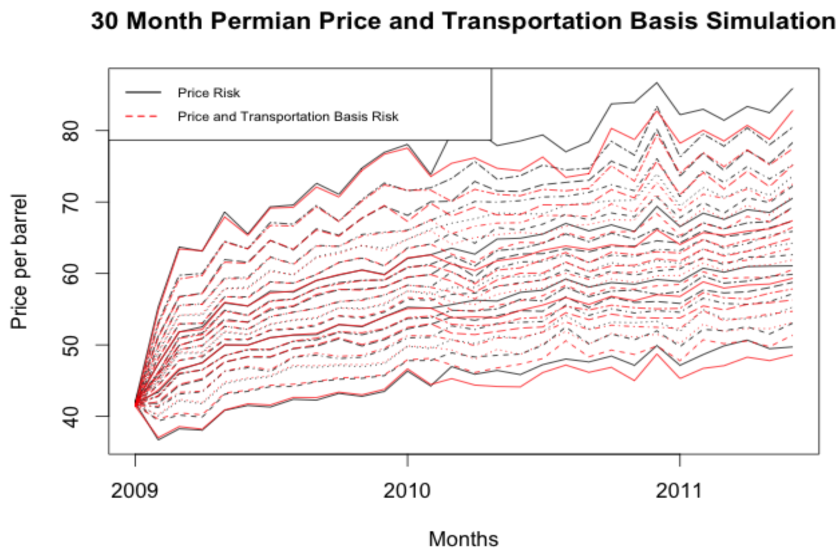

Figure B.4: 30-Month Dual Risk Simulation

relatively free of congestion, Figure B.4 shows that integrating transportation basis lowers the expected Permian price relative to Eagle Ford—a region without transportation basis exposure.

B.1.2 General Mean-Reversion Jump Diffusion Model

The mathematical description of mean-reverting jump dynamics is given by the stochastic differential equation ((B.1). To estimate mean-reversion and jump characteristics in transportation basis, we combine models by Schwartz (1997), and Clewlow and Strickland (2000) into the same standard stochastic differential equation (B.1), where the jump process dq is assumed independent of dz, the Wiener increment from the continuous process—the Poisson process is independent of the mean-reverting process:

$$dTB_t = \alpha(\mu_t - \ln TB_t)dt + \sigma_t TB_t dz_t + \kappa TB_t dq_t \tag{B.1}$$

α: speed given by the mean reversion rate

μ_t: instantaneous drift, used to estimate the long-term transportation basis level

We use $\mu_t(P_t > \bar{P})$ as a dummy to control for a higher long-term transportation basis mean (or high congestion), which results in a higher expected transportation basis

κ: proportion of random jump size

σ_t: transportation basis volatility

dq_t: a discrete time process and takes value 1 when random jump occurs.

dTB_t represents the increment in transportation basis process for a time interval dt and dz is the underlying uncertainty driving the model and represents an increment in a Weiner process during dt.

Using equation (B.1), we simulate forward Permian spot transportation basis and volatility. First, we fit the mean-reversion rate α and long-term mean μ_t using a linear regression. Second, we use a recursive filter to fit the jump-diffusion parameters, mainly the standard deviation of jump returns, frequency of jumps, and standard deviation of non-jump returns.[1] Lastly, the volatility diffusion process is normalized for mean-reverting volatility term structure dynamics.

[1] Clewlow and Strickland (2000) note that the recursive filter approach is better at modeling lower frequency high volatility jumps, while a maximum-likelihood approach is better at estimating high frequency low volatility jumps.

B.1.3 Estimating Spot Mean-Reversion and Jumps

The mean reversion rate of spot transportation basis can be estimated using a linear regression. We discretize Equation B.1, which gives us:

$$\Delta TB_t = \alpha_0 + \alpha_1 TB_t + \sigma\epsilon_t \tag{B.2}$$

where from Equation B.1 $\alpha_0 = \alpha\overline{TB}\Delta t$ and $\alpha_1 = -\alpha\Delta_t$.[2] We estimate an $\alpha = 19.71$, $\mu_t(P_t > \bar{P}) = 4.93$ and $\mu_t(P_t < \bar{P}) = 2.52$. We use $\mu_t(P_t > \bar{P})$ as a dummy to control for a higher long-term transportation basis mean (or high congestion), which results in a lower forecasted Permian Midland price.

Estimating transportation basis jump parameters is complicated since the jumps are only observed as part of a time series which includes the normal non-jump behavior. Clewlow and Strickland (2000) apply what they call a recursive filter. Large transportation basis spikes are attributed to jumps, given that a normal Brownian motion has a low probability of generating these jumps. If jumps are infrequent with low amplitudes, we can estimate volatility using the sample standard deviation of returns. We then assume that our sample is absent any congestion jumps and identify returns that are three standard deviations. Finally, we enter an iterative process whereby we exclude the previously identified jumps (storing the number of jumps) and recompute volatility by reestimating the sample standard deviation. This process is repeated until volatility estimates converge within a tolerance requirement. It took four iterations before convergence was achieved. Once recursively estimated, we report $\phi = 0.50$[3], $\sigma = 0.98$, $\bar{\kappa} = -0.497$,[4] and $\gamma = 1.41$[5] and three jumps on average. When we simulate the 24-month out transportation basis, we normalize volatility for the mean-reverting characteristics and control for the correlation structure between the spot and 24-month contract of the WTI price and volatility term-structure. Forecasting transportation basis using simulations allows us to measure the expected 24-month Permian Midland price (24-month WTI Cushing price

[2] The half-life of the mean reverting process can be estimated as $t_{\frac{1}{2}} = \frac{ln(2)}{\alpha}$.

[3] The average number of jumps per year or annualized frequency of jumps.

[4] Mean jump size.

[5] Jump volatility.

minus transportation basis).

B.1.4 Expected 24-Month: Simulating Transportation Basis

Once the mean-reverting jump diffusion model is fitted,[6] we can simulate transportation basis, as illustrated in Figure B.5. That is, with Equation B.1 estimated, we recursively simulate N=500 paths for 24 monthly time-steps for every spot transportation basis for $\tau \in \{1, \ldots, 24\}$ over t $\in \{1, \ldots, 73\}$—i.e. for 2009 to 2015 monthly time steps—to estimate the $\text{E}(TB_{t,\tau=24})$ time series. In fact, at t $= 1$ we start by estimating spot transportation basis $\tau = 1$, dynamics to then simulate forward $\tau = 24$ transportation basis (or, more specifically, at t estimate transportation for months-τ such that we have: $\{TB_{t,1}, \ldots, TB_{t,24}\}$). Equation B.3 depicts the process for estimating 24-month expected transportation basis:

$$TB_{t,\tau+1} = TB_{t,\tau}e^{-\alpha\Delta_t} + \ln(\mu_t)(1-e^{-\alpha\Delta_t})\sigma_\tau\sqrt{\frac{(1-e^{-2\alpha\Delta_t})}{2\alpha}}N(0,1) + \text{jumps}(P_t > \bar{P}) \quad \text{(B.3)}$$

to reduce the notational complexity, when presenting the results, we drop the τ-index and present $\text{E}(TB_{t,\tau=24})$ as current expected transportation basis $\text{E}(TB_t)$; the resulting time-series is illustrated in Figure B.5.

We see that over the sample, (spot) transportation basis ranges from $1/bbl-$12/bbl and that the expected long-term price without congestion is $2.52/bbl and $4.93 otherwise. There is a higher likelihood of jumps towards the end of the sample: high prices increase the probability of congestion, as well as the likelihood that new pipeline projects are undertaken.

When modeling basis in the Permian, we need to integrate expected pipeline completion dates (e.g. 2010-13 and 2014-15). The opening of new pipelines reverts transportation basis back to the long-term mean. However, once pipelines were finished in 2013, they were expected to be quickly overwhelmed by new production supply, which led once again to an increase in expected congestions. Essentially, as long as expected prices were higher than the front month prices, congestion probabilities remained high.

---

[6] See Appendix B.1.3 for an explanation on how to fit the model.

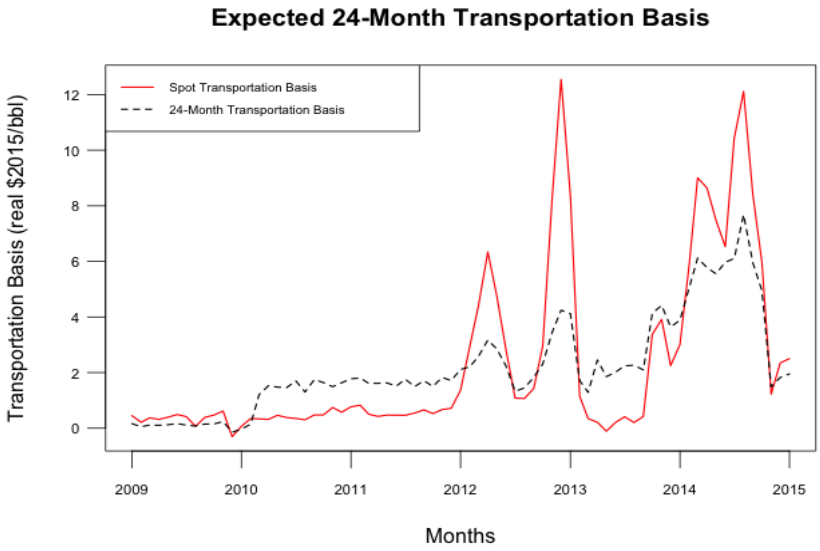

Figure B.5: Expected Transportation Basis

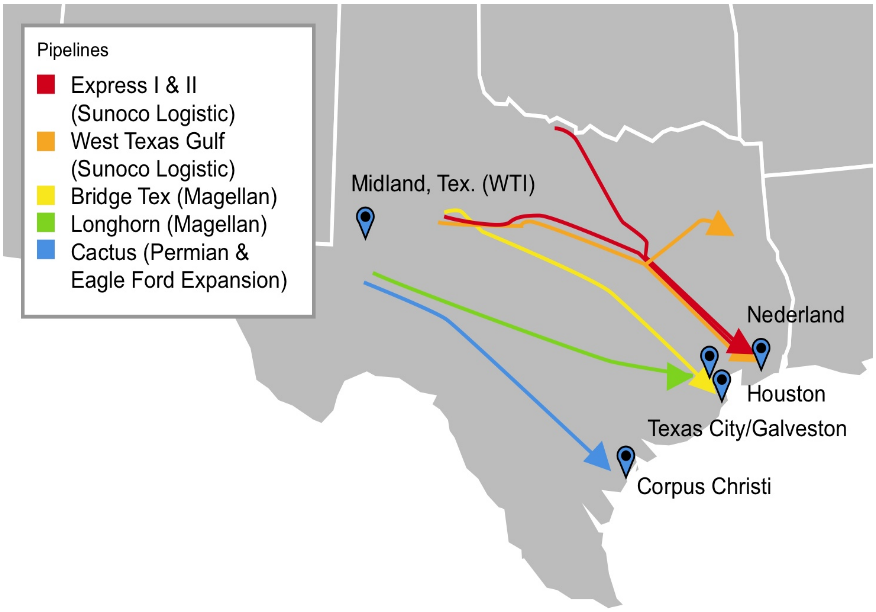

Figure B.6: Permian Pipeline Projects

As observed in 2013-2014, by the time a project is realized, production may outstrip transportation capacity. For example, as shown in Figure B.6, in 2013, basis exposure decreased when Magellan Midstream Partners reversed a 225,000 bbl/day Longhorn

Pipeline to move crude from the Permian to Houston. Also, Sunoco Logistics Partners finished their 150,000 bbl/day Permian Express I & II pipelines. However, by 2014, crude production growth outstripped capacity and transportation risk increased once again. Projects to expand pipeline transportation capacity and connect producing wells to the market are overseen by the Texas Rail-Road Commission (TRRC) for intrastate and Federal Regulatory Commission (FERC) for interstate traffic. The regulatory approval process, or a midstream company's inability to gather enough pre-commitments in a timely manner—a threshold of pre-commitments must be locked in before banks finance construction—causes pipeline construction delays.

B.2 Permian Graphical Fitting

As we discuss extensively in the thesis and Appendix B.1.2 and B.1.1, the Permian is exposed to both crude oil price and transportation basis uncertainty. We define price risk as time-varying uncertainty and complex time-varying uncertainty as the integration of both price and transportation risk. Both of these time-varying uncertainties are a better fit than myopic constant volatility. As shown in Figure B.7, constant uncertainty—equal to the average of first year volatility[7]—does not yield as good a fit. Time-varying complex volatility is an improvement over time-varying crude oil price volatility.[8]

As illustrated in Figure B.7, expected price volatility decreased in mid-2013 and late-2014. Therefore, the higher myopic expected constant volatility decreased well investment activity. This lower expected price volatility reduced the 24-month and front month volatility ratio and, because we model expected transportation basis (or expected Permian Price) term-structure dynamics using the crude curve, the same holds for the transportation basis volatility ratio. At the same time, unsurprisingly, as transportation basis contributed more to revenue, the correlation relationship between expected Permian and transportation basis returns was more negative ($\rho = -0.35$), indicating that expected 24-month complex volatility was lower than price time-varying alone. This explains why

[7] See Appendix A.2 for myopic versus time-varying identification structure.

[8] When we run the reduced form of the complex specification as just time-varying volatility, we get a statistical improvement at the 10% level.

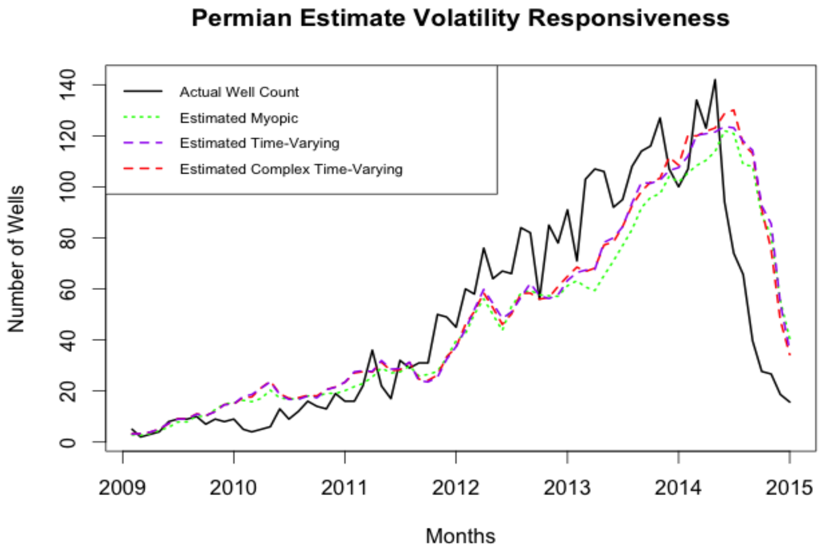

Figure B.7: Permian Myopic $\beta = 0$ & Time-Varying $\beta = 1.04$

the complex volatility model is a better fit and predicts higher investment levels towards the end of the sample.

Appendix C

TECHNICAL DETAILS, CHAPTER 4

C.1 Medium-Run Global Supply

From 2009 to 2015, there was rapid growth in on-shore production in the US, using hydraulic fracturing in tight rocks.[1] Although the shale cost structure varied for different basins, an industry report estimated that the majority of US shale oil lies in the middle of the medium-run cost curve, as illustrated in the stylized Figure C.1.

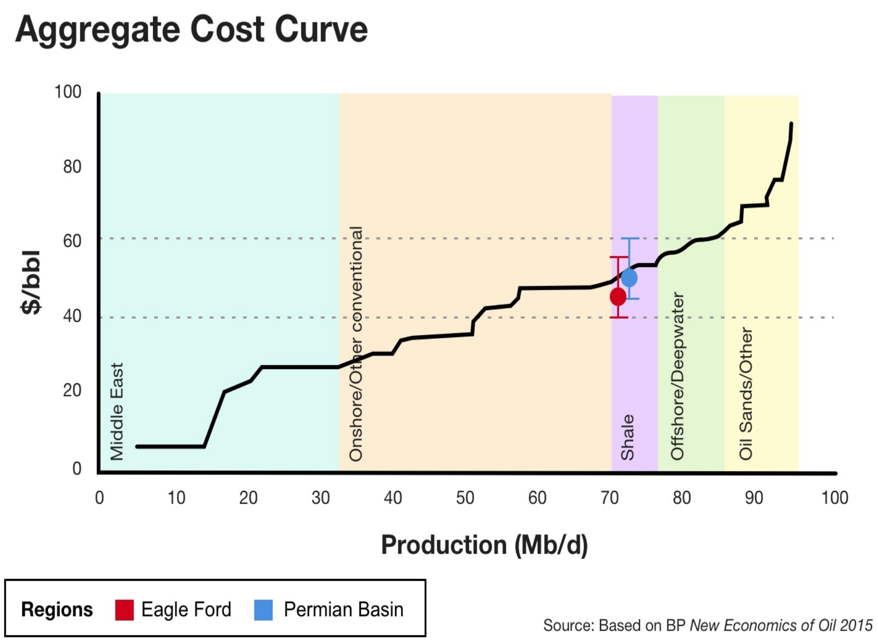

Figure C.1: Small Shale Oil 2014 Production & Breakeven

In Figure C.1, for 2014 we superimpose our breakeven measures for shale oil producers in the Permian and Eagle Ford—they fit in the middle of the cost curve between con-

[1]Note that even with its growth, shale oil production accounted for less than 5% of the global oil market.

ventional and offshore oil production. Hence, over time, these producers quickly became marginal price setters. Their exposure to short investment lead times, sharp well production declines curve, and high technological innovation (or cost reduction) means that in the short- and medium-run production became more price elastic. This also meant that there was a closer relationship between investment and production than previously with conventional/offshore oil production. For the macroeconomy, for example, this signifies that as prices fall, well investments will decline with production following shortly thereafter. This new short-run level and volatility responsiveness has ramifications for energy policy and managing shocks.

C.2 Counterfactual Cost Adjustment

For illustrative purposes, we focus on rig cost for Eagle Ford because there is a theoretically strong relationship between rig cost and oil prices. Note that completion services is a more concentrated market and therefore, fracking stimulation is less responsive to price changes. Figure C.2 shows a graphical fit for rig rental rates. We focus on a sample after 2010, when the impact from the transition of substitutable rig equipment from shale gas isn't a confounding factor. We conduct the analysis in first differences because we cannot reject (using the GLS procedure of Elliott, Rothenberg and Stock 1996) that both log(cost) and log(24 month price) are unit root processes.[2] Using an AIC[3] selection procedure, we regress cost against lagged 24-month price (P) and rig cost (RC) in equation C.1:

$$\Delta\log(\mathrm{RC}_t) = \Delta\log(\mathrm{P}_{t-1}) + \Delta\log(\mathrm{RC}_{t-1}) \tag{C.1}$$

Lagged price, with coefficient 0.106, and rig cost, with coefficient 0.716, are significant at the 5% and 1% level, respectively[4] with $R^2 = 0.4912$ and DF = 56.

We find that when futures prices go from domestic WTI to counterfactual Brent rig rental rates increase by 2-8.3%, depending on the rig size. Rig rental represents 15% of

[2] There is an endogeneity concern that might bias the inference results.

[3] Additional lags do not qualitatively change the results.

[4] Where the constant is equal to 0.003124.

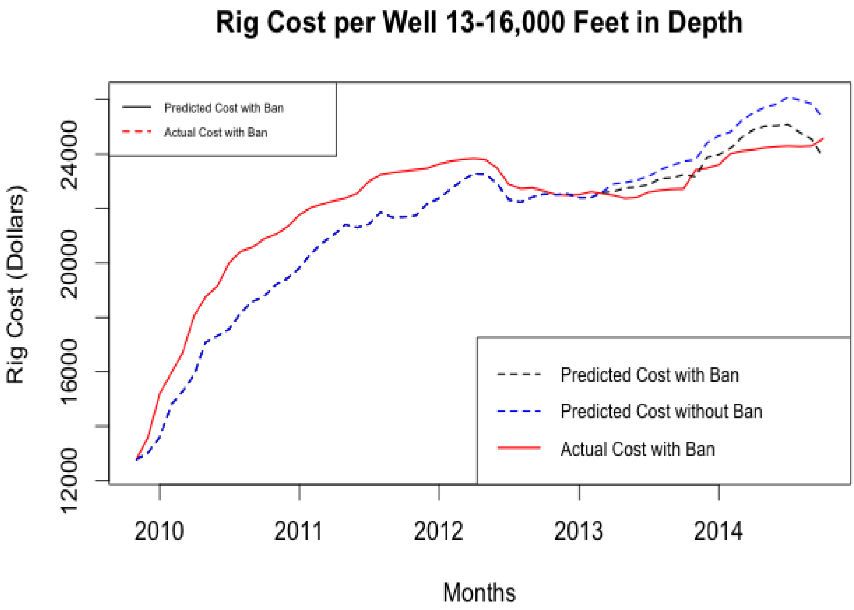

Figure C.2: Rig Cost with and without the Export Ban

total cost—with drilling representing 42% of total cost. Furthermore, the number of well drilled went from 33 to 15; hence the impact of rig rental increase on real option well investment decision-making is insignificant.[5]

C.3 Ineffective Export Ban as of Dec 2014

The drop in prices and increase in volatility towards the end of the sample is attributed by Baumeister and Kilian (2016) to an unanticipated growth in supply from shale oil and refusal by OPEC to reduce production quotas. As shown in Figure A.1, the sudden drop in prices led to a convergence in the WTI-Brent spread, and at lower prices, the binding export ban (as of August 2013) became ineffective. Even though the export was repealed in December 2015, its ineffectiveness as of December 2014 allows us to state that

[5]If we were to model total cost per well, we'd also have to account for productivity—learning-by-doing and engineering productivity, which would theoretically cancel any cost increase. Also, high grading wasn't significant when prices were above $80. Finally, waitlisting isn't an issue as of August 2013, when rental days had decreased because of improved productivity. Furthermore, rig wait lists and work backlogs were surveyed to be on the decline in Q4 2014 and Q1 2015.

we capture the full distortionary impact of the export ban on shale oil well investment activity.

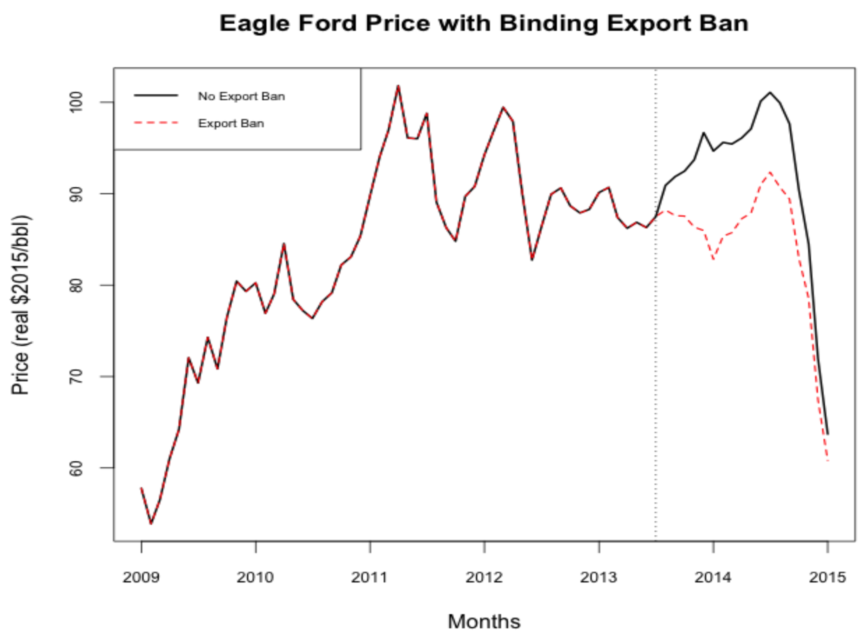

Figure C.3: Price Level with and without the Export Ban