Exploration Activity, Long-run Decisions, and the Risk Premium in Energy Futures

Alexander David
University of Calgary

Investment by oil firms positively affects the futures basis and negatively predicts excess returns on crude oil futures. I build an equilibrium model of drilling, exploration, and storage to understand these facts. Firms' capital stock lowers extraction costs as firms drill in increasingly expensive fields. Drilled wells produce the resource at a geometrically declining rate; however, by specifying consumers' habit level equaling production from old wells, the futures basis and risk premium are only related to drilling, investment, and inventory. Investment leads to a more elastic drilling response by firms and dampens oil price increases from demand shocks, thus lowering the risk premium. (*JEL* G12, G13, Q31, Q32, Q41, Q43)

Received October 10, 2015; editorial decision March 30, 2018 by Editor Philip Strahan. Author has furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Recent years have seen the development of increasingly sophisticated technologies for the extraction of natural resources from costlier fields.[1] These new technologies, which have been brought to fruition by investments by the resource extraction industry, have changed the current and expected future prices of resources and have important consequences for energy self-sufficiency and stability of growth for North America. In this paper, I ask if investment in the exploration and development (E&D) of resources has an impact or is

I thank the Editor, Phil Strahan; two anonymous referees; Steven Baker (Cavalcade Discussant), Hui Chen (NBER Discussant), Jaime Cassasus (WFA Discussant), Lars Alexander Kuehn, Miguel Palacios, Antonio Diez de los Rios (NFA Discssuant), John Parsons (OU Conference Discussant), Bryan Routledge, José Scheinkman, Chester Spatt, Sheridan Titman, Harrie Vredenburg, and Wei Xiong; and participants at the Alberta Finance Institute Commodities Conference, Carnegie Mellon Finance Area, the EFMA Special Panel on Commodities, the European Finance Assocation Meetings, the HEC-Montreal – McGill Winter Finance Conference, the Northern Finance Association Meetings, the Oklahoma University Energy Finance Conference, the NBER Commodities Group Conference, the Bendheim Centre Seminar at Princeton University, the Western Finance Association Meetings, the SFS Cavalcade, and the Commodity and Energy Markets Association Meetings for their helpful comments. Finally, I wish to thank my late father, Colonel Wilbur David, who encouraged me for several years to work on a paper on oil futures markets. I am grateful to the SSHRC for a research grant. Supplementary data can be found on *The Review of Financial Studies* Web site. Send correspondence to Alexander David,Haskayne School of Business, 2500 University Drive NW, Calgary, AB T2N 1N4, Canada; telephone: (403) 220-6987. E-mail: adavid@ucalgary.ca.

[1] Examples of new technologies include steam-assisted gravity drainage (SAGD) in Alberta's oil sands, and hydraulic fracturing in the shale oil fields in the United States. A recent survey by Maugeri (2012) provides several interesting highlights of the recent revolution in energy production.

doi:10.1093/rfs/hhy067

affected by the keenly watched market statistics of current and future prices of the resource.

One of the most widely watched statistics in the futures market is the slope of the futures curve. I measure it as the weak relative basis, which is the proportional difference between the discounted value of the futures price and the current spot price of the resource (see also Litzenberger and Rabinowitz 1995). When this difference is positive (negative), I say the futures market is in weak contango (backwardation). Of interest to practitioners and researchers is the economic information that determines the relative basis. The theory of storage (Kaldor 1939; Working 1948) implies that the futures relative basis is positively correlated to inventories. I call this the "short-run" information about resource prices in the futures relative basis. In addition, I argue that the futures relative basis contains "long-run" information about resource prices, which has important implications for decisions about exploration and the development of the resource extraction process. In particular, I examine the relationship between the basis and the risk premium on oil futures contracts and energy firms' inventory and investment decisions.

I begin by displaying this paper's main variables of interest in Figure 1. Appendix A provides a detailed description of the construction and sources for the series. The first panel shows the seasonally adjusted futures basis for 1-year crude oil futures.[2] As can been seen, the futures curve has mostly been in contango since 2008, whereas backwardation was more frequently prevalent earlier. The second panel shows the excess returns from holding fully collateralized 1-year futures contracts, which are defined as $S(t)/F(t-11)-1$, where $F(t)$ is the 1-year futures price, and $S(t)$ is the spot price. The third panel shows the ratio of the detrended seasonally adjusted crude oil inventory to the U.S. real gross domestic product (GDP).[3] The fourth panel shows the seasonally adjusted ratio of investment (capital expenditures) of oil and gas firms to the U.S. real GDP. This ratio showed a spectacular increase from around 2000 to 2014, which was the period of rapid deployment of new technologies in the United States, and then fell dramatically after the crash in oil prices. Next, I examine the ability of inventory and investment to explain the variations in the basis and the risk premium.

In Table 1, I report simple linear regressions at a monthly frequency of the futures relative basis on inventory and investment. As can be seen, inventory explains about 5% of the variation in the basis, while investment explains about

[2] Throughout this paper, I look at statistics of the 1-year futures contracts. Although it would be of interest to study longer maturity futures, I am constrained by the lack of long historical time series on these longer term contracts. Two-year contracts started trading actively in mid-1990 and 4-year contracts only in 1997. The correlations between the relative bases of the 2-year and 1-year contracts with the 1-year contract are 99.7 and 98.7 over the subsamples, respectively.

[3] I detrended this series because the ratio of inventories to GDP fell quite steadily in the first half of my sample period. The downward trend likely arose from the increased efficiency of supply-chain operations, a factor that is unrelated to the main interests of this paper.

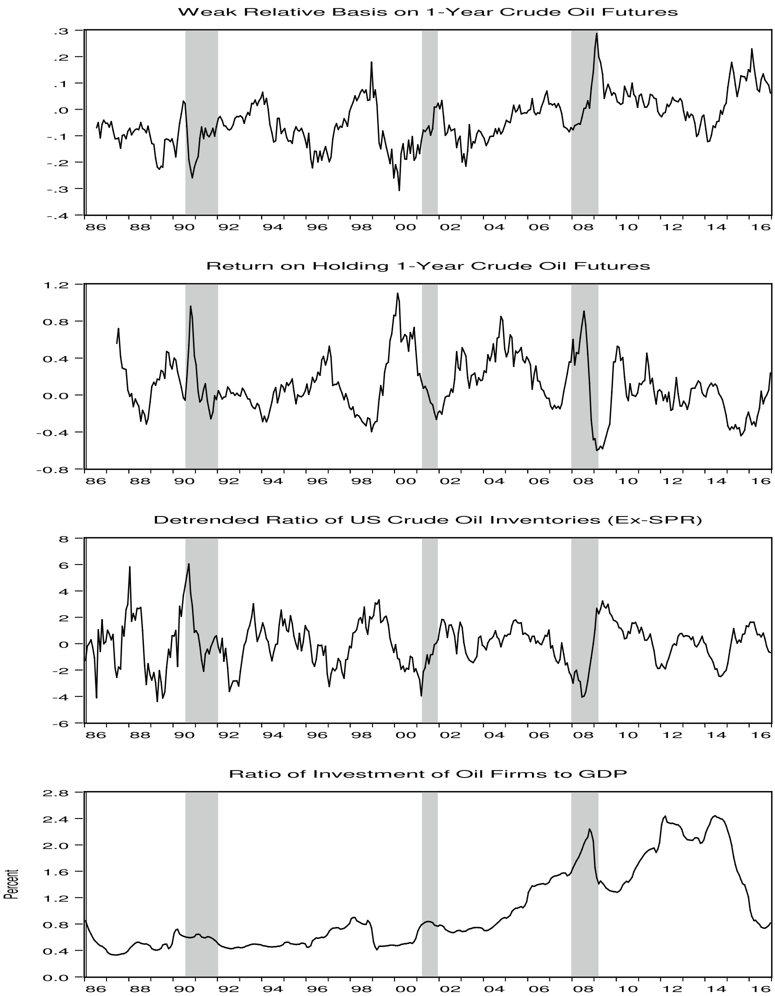

Figure 1
Crude oil futures basis and returns, oil inventory and capital stocks of oil exploration firms (July 1986–December 2016)
The **top panel** shows the seasonally adjusted weak relative basis on 1-year crude oil futures contracts, which in month t is $[e^{-r(t)}F(t)-S(t)]/S(t)$, where $F(t)$ is the 1-year futures prices at the beginning of each quarter and $S(t)$ is the spot price of WTI oil in Cushing, Oklahoma. The **second panel** reports the excess return on a fully collateralized 1-year futures contract defined as $S(t)/F(t-11)-1$, where $F(t)$ is the 1-year futures price, and $S(t)$ is the spot price. The **third panel** shows "Inventory," which stands for the seasonally adjusted total U.S. stock of crude oil and petroleum products (in thousands of barrels) excluding special purpose reserves at the end of each month, normalized by U.S. real GDP and detrended using the Hodrick-Prescott filter. The **fourth panel** shows "Investment," which stands for the seasonally adjusted capital expenditures of oil and gas firms (SIC codes 1311 and 138), normalized by U.S. real GDP. Seasonal adjustment is performed using the X-12 procedure (used by the U.S. Department of Commerce).

Table 1
What explains the futures weak relative basis for crude Oil?

	α	β_1	β_2	R^2
Full sample July 1986–December 2016:				
1	−0.051	0.011		0.048
	[−5.412]	[1.837]		
2	−0.115		0.065	0.236
	[−7.643]		[4.974]	
3	−0.121	0.013	0.069	0.300
	[−8.392]	[2.851]	[5.571]	
Sample from July 1986–December 1999:				
4	−0.080	0.006		0.025
	[−6.987]	[1.013]		
5	−0.185		0.197	0.111
	[−4.423]		[2.473]	
6	−0.182	0.005	0.191	0.118
	[−4.633]	[0.795]	[2.591]	
Sample excluding 2008–2010:				
7	−0.061	0.008		0.032
	[−7.025]	[1.431]		
8	−0.112		0.054	0.196
	[−7.763]		[4.984]	
9	−0.115	0.009	0.056	0.231
	[−8.300]	[1.935]	[5.360]	

I report the coefficients of the fitted monthly regression: Weak relative basis$(t)=$

$$\alpha+\beta_1\,\text{Inventory/Real GDP}(t-1)+\beta_2\,\text{Investment/GDP}(t-1)+\epsilon(t).$$

The seasonally adjusted weak relative basis on 1-year contracts in quarter t is $[e^{-r(t)}F(t)-S(t)]/S(t)$, where $F(t)$ is the 1-year futures prices at the beginning of each quarter and $S(t)$ is the spot price. The explanatory variable "Inventory" stands for the seasonally adjusted total U.S. stock of crude oil and petroleum products (in thousands of barrels) excluding special purpose reserves at the end of each month. I normalize this series by U.S. real GDP and detrend the ratio using the Hodrick-Prescott filter. "Investment" stands for the seasonally adjusted capital expenditures of oil and gas firms (SIC codes 1311 and 138). Seasonal adjustment is done using the X-12 procedure (used by the U.S. Department of Commerce). t-statistics are in brackets and are adjusted for heteroscedasticity and autocorrelation.

24% (lines 1 and 2). Both variables have positive slope coefficients. The two variables jointly explain about 30% of the variation in the relative basis, and each variable remains significant in the multivariate regression (line 3). Lines 4 to 6 show the same regressions for the first half of the sample, which preceded the developments of the new technologies in the current millennium. In this subsample, inventory has an insignificant coefficient while investment has a significant coefficient, and explains about 11% of the variation in the basis, which is smaller than for the full sample. Finally, lines 7 to 9 show the regression fits for the full sample, but excluding the period from 2008 to 2010, and the results are similar to full sample, meaning that the results are not mainly driven by the events during and immediately following the financial crisis.[4] Overall,

[4] Some market analysts have reported that the large contango in crude oil futures in 2009 resulted from the lack of credit for firms to engage in cash-and-carry arbitrage.

Table 2
What explains the risk premium on crude oil futures?

	α	β_1	β_2	R^2
Full sample July 1986–December 2016:				
1	0.096	0.048		0.075
	[2.898]	[2.361]		
2	0.244		−0.146	0.092
	[3.891]		[−3.003]	
3	0.23	0.042	−0.131	0.149
	[3.912]	[2.250]	[−3.077]	
Sample from July 1986–December 1999:				
4	0.063	0.005		0.001
	[1.594]	[0.321]		
5	0.229		−0.306	0.026
	[1.281]		[−0.932]	
6	0.237	0.007	−0.321	0.03
	[1.329]	[.495]	[−.988]	
Sample excluding 2008–2010:				
7	0.095	0.038		0.053
	[2.952]	[1.982]		
8	0.236	−0.142		0.096
	[3.808]	[−3.214]		
9	0.229	0.036	−0.139	0.146
	[3.875]	[1.956]	[−3.441]	

I report the coefficients of the fitted monthly regression: Excess return(t)=

$$\alpha+\beta_1\,\text{Inventory/Real GDP}(t-12)+\beta_2\,\text{Investment/GDP}(t-12)+\epsilon(t).$$

The excess return on a fully collateralized 1-year futures contract is defined as $S(t)/F(t-11)-1$, where $F(t)$ is the 1-year futures price, and $S(t)$ is the spot price. The explanatory variable "Inventory" stands for the seasonally adjusted total U.S. stock of crude oil and petroleum products (in thousands of barrels) excluding special purpose reserves at the end of each month. I normalize this series by U.S. real GDP and detrend the ratio using the Hodrick-Prescott filter. "Investment" stands for the seasonally adjusted capital expenditures of oil and gas firms (SIC codes 1311 and 138). Seasonal adjustment is performed using the X-12 procedure (used by the U.S. Department of Commerce). t-statistics are in brackets and are adjusted for heteroscedasticity and autocorrelation.

the results suggest that both short- and long-run decisions by firms are important determinants of the futures relative basis.

Table 2 examines the role of inventory and investment in explaining excess returns on 1-year crude oil futures. For the full sample in lines 1 to 3, inventory positively predicts excess returns, while investment negatively predicts them. Both variables are statistically significant, and jointly explain nearly 15% of the variation. However, lines 4 to 6 show that their explanatory power is much lower at only about 3% in the sample until 1999, and the inventory coefficient is insignificant in this subsample. Finally, the results for the sample excluding 2008–2010 are quite similar to those of the full sample. Therefore, both variables are again important determinants of the variation in excess returns, although the explanatory power of investment is greater.

Figure 2 shows some additional data series relevant to modeling of oil drilling in increasingly costly wells in recent years. The top panel shows that the real spot price of crude oil (in 1984 dollars) has been on average higher in the sample

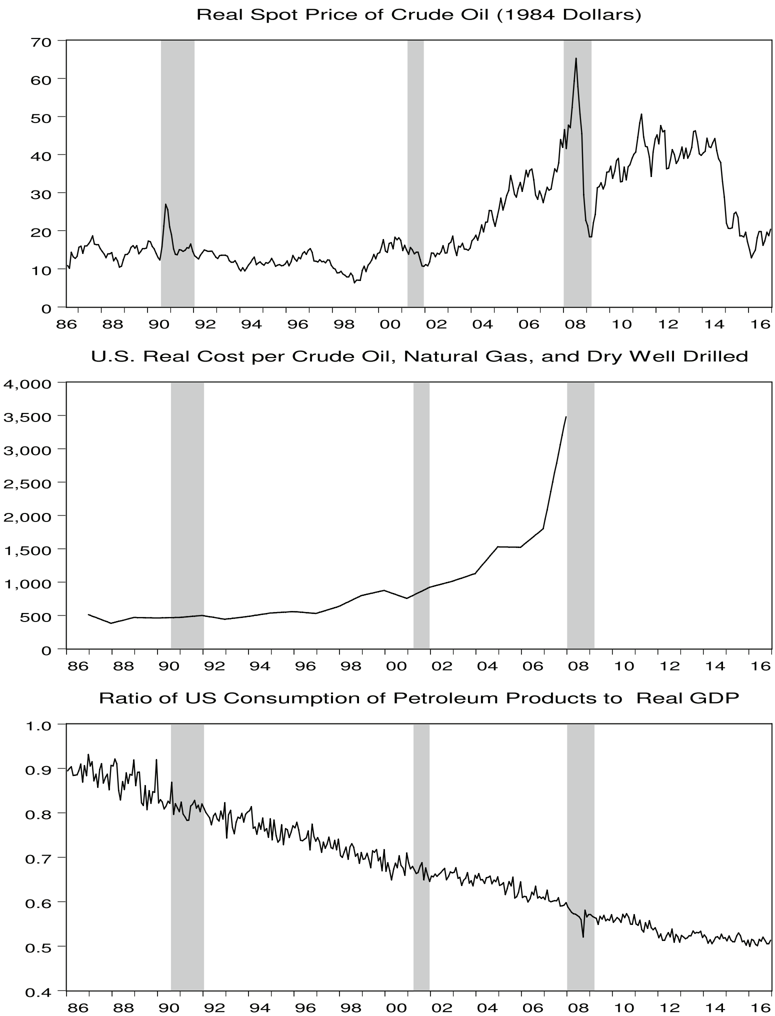

Figure 2
Real spot price of crude oil (July 1986–December 2016), consumption of petroleum products and real drilling costs per well (1996–2007)
The **top panel** shows the monthly U.S. real spot price of crude oil. The spot price is approximated as the WTI futures price with less than one month to maturity reported at the beginning of the month. The real spot price is the spot price divided by core CPI (excluding food and energy). The **middle panel** shows the real cost per crude oil, natural gas, and dry well dug in the U.S. from 1986 to 2007 as provided by the EIA. The **bottom panel** shows the monthly U.S. consumption of petroleum products as reported by the EIA normalized by U.S. real GDP.

post-2000 despite some large fluctuations. The middle panel shows that the real cost of new drilling increased steadily until 2000 and then far more rapidly until 2007 (this data series is not available after that). Finally, the bottom panel shows that the ratio of U.S. consumption of petroleum products to real GDP has been steadily declining in my sample of 30 years.

In this paper, I build a model of the short- and long-run decision-making of resource producing firms that in equilibrium will generate relations between inventory, investment, the futures basis, and the risk premium similar to that in the data. I assume that there is a known and constant total amount of the resource. The process of making some of the resource economically recoverable through investment is explicit in my model. I model demand shocks that drive the business-cycle fluctuations in oil prices, but I also build in the implications of increasing drilling costs and firms' attempts to manage these costs by investing in exploration and technology development. In addition, firms choose inventories to smooth fluctuations in demand and extraction. The simultaneous modeling of inventory and investment is essential to not only understanding the positive relation between the basis and each of these variables, but also the role each of these variables plays in enabling firms to transfer resources over time. In a nutshell, a steeper futures slope increases the attractiveness of accumulating inventory in some periods and investment in others. As will be clear in the model analysis, the short-run (inventory) and long-run (capital accumulation) decisions are tied: if capital is currently high but demand is moderate so that further investment is not optimal, then extraction costs are expected to increase as capital depreciates, and increasing current extraction rates to carry forward inventory is optimal. Consistent with this model implication, inventory and investment are negatively correlated in the 30-year sample.

I start with a two-period version of the model with a linear demand function and for which I can provide closed-form expressions for the value of extraction options, and can characterize the optimal investment, extraction, and inventory policies quite tractably. In particular, it shows how the futures basis and risk premiums in my model directly depend on investment, since the latter affects the future supply of the resource. Both also indirectly depend on investment through the effect of investment on the firm's extraction and inventory decisions. Also notably, I show that the risk premium on a fully collateralized futures position in the model is not simply the risk premium for bearing demand shocks, but it also depends on the firm's policies that affect the future supply of the resource, and hence either amplify or dampen these shocks. I return to this point later when discussing the infinite-horizon model, since it is the key insight of this paper.[5]

To study the dynamics of these variables and their relationship to the data, I next build an infinite horizon model with three added features relative to the two-period model. First, following Anderson, Kellogg and Salant (2018), my model makes a distinction between drilling in new fields and production from all existing fields.[6] These authors point out that drilling activity is quite

5 In the two-period model that is solved, the sign of the risk premium is constant because of the linearity of the demand function, an assumption that I change in the infinite-horizon model.

6 Hamilton (2012) shows that oil production in some regions show peaks as new technologies are developed and installed, and subsequent and slow declines as oil wells get depleted. Chen and Linn (2017) provide further

volatile while production is much smoother. This is because, once drilled, oil wells produce at close to zero marginal cost at geometrically declining rates as well pressure declines over time. Second, I model external habit formation in oil consumption. I assume that the level of habit equals the production from all previously drilled wells, which is predictable and smooth. In effect, I assume that consumers in the economy make plans to consume the predictable part of oil production so that only deviations from it affect the marginal utility of oil consumption. This assumption ensures that the futures basis and risk premium in my model only depend on incremental decisions on drilling and not all production. Finally, as is standard in the investment literature in macroeconomics, I assume that there are adjustment costs in investment. This assumption is made to smooth out the investment process to match the volatility of the investment-to-capital ratio in the data.

In dynamic models with inventory and investment restricted to be positive, it is not possible to obtain closed-form solutions as has been pointed out by several authors. So I solve the model using projection methods that have been popularized by the work of Judd (1999). Simulations of the model show that both inventory and investment affect the basis and the risk premium with the same signs as in the data. In particular, I show that both capital and investment are negatively related to the futures risk premium. Similar to the two-period model, a higher investment lowers extraction costs and increases the supply of the resource in the future. This effect notably weakens the correlation between demand shocks in the economy and resource prices and can even reverse it in periods when the capital stock is very high, thus leading to a conditionally negative risk premium in some periods. I show that the relationship between investment and the risk premium is nonlinear in the model so that a linear regression of excess returns on futures on investment leads to a relationship similar to that in the data, but is a misspecified regression in the context of my model. The correctly specified regression of excess returns on the model risk premium (which is negatively related to investment) provides an economically significant regression coefficient. In addition, in my model simulations, I have a negative relationship between the risk premium and the weak basis. This finding is consistent with the findings of Fama and French (1987) and several other papers.[7]

evidence that drilling in different countries is positively related to futures prices. However, none of these papers study the determinants of the futures basis or the risk premium, which is the main focus on this paper. In fact, I find that drilling has a very weak relationship with each of these variables.

[7] See, for example, Gorton and Rouwenhourst (2006), Erb and Harvey (2006), and Koijen et al. (2013). Baker and Routledge (2012) show that such a relationship can arise from the speculative strategies of different agents like in the equity premium literature on heterogeneous beliefs (see David 2008). Ready (2018) attributes the increase in the futures slope and the decline in the risk premium in 2005–2012 to an increase in the uncertainty about the long-run supply of oil.

My model contributes to the literature on resource extraction, storage, and investment. Most existing models have either one of these features. Models of storage assume exogenous extraction decisions (e.g. Gustafson 1958; Williams and Wright 1991; Deaton and Laroque 1992; Routledge, Seppi, and Spatt 2000; Pirrong 2012, and Gorton, Hayashi, and Rouwenhourst 2013). Starting with the seminal work of Hotelling (1931), models of exhaustible resource extraction, on the other hand, allow no storage (e.g. Pindyck 1980, Campbell 1980; Litzenberger and Rabinowitz 1995; Carlson, Khokher, and Titman 2007; Ghoddusi 2010).[8] The same is true of models that allow production of commodities (see, e.g., Casassus, Collin-Dufresne and Routledge 2008; Kogan, Livdan, and Yaron 2009; Yang 2013). Pindyck (1994) allows for production and storage but uses exogenous prices as opposed to equilibrium prices, which equate demand to supply for the resource in my model.[9]

My model has some features in common with production-based models of commodity prices (see, e.g., Casassus, Collin-Dufresne and Routledge 2008; Kogan, Livdan, and Yaron 2009; Yang 2013) but it also has several important differences. Like in these models, the elasticity of supply of commodities positively depends on investment: higher investment implies a more elastic future supply. This feature helps us obtain a lower risk premium in periods of higher investment, as discussed above. Besides not modeling storage, these models do not impose exhaustibility of the resource but instead assume unlimited oil production with an increase in the capital stock of oil firms, which is the only factor of production. However, such models are inconsistent with the reality of oil production, which is restricted to geographical areas endowed with oil wells. In addition, in the data, oil production does not always increase with capital installed by oil companies. In the model of Yang (2013), the risk premium differences across commodities is not explained by the investment rates but instead by investment shocks, which are proxied empirically by changes in the prices of investment goods, which I hold constant in my model. My model in contrast does not have productivity shocks in oil production (like in Kogan, Livdan, and Yaron 2009) or in capital accumulation (like in Yang 2013) but instead has endogenously generated changes in extraction costs of the resource determined by firms' capital accumulation decisions. As such, my technological progress specification is novel to the literature on investment and real options.

[8] Campbell (1980) and Ghoddusi (2010) consider investment to increase production capacity in an exhaustible resource setting, but besides having zero extraction costs, they do not study the relationship between investment and either the futures basis or its risk premium.

[9] In the context of agricultural commodities, there is a literature that has production and storage in equilibrium, but the analysis in such models does not apply to exhaustible resources, where equilibrium profits are compatible with competitive equilibrium due to limitations in supply (e.g., Scheinkman and Schechtman 1983). With exhaustible resources, as pointed out in Litzenberger and Rabinowitz (1995), non-zero profits at the time of extraction are consistent with competition that are optimized by the extraction timing decision of resource firms, a feature that I model explicitly.

1. A Two-Period Model of Resource Extraction, Storage, and Exploration

I build on the two-period version of the model of Litzenberger and Rabinowitz (1995) with several generalizations. The most significant addition is of an E&D (investment) decision that reduces the costs of future extraction of the resources. To tractably analyze the investment decision with technology spillover, I introduce multiplant firms. Assume a continuum of price-taking identical resource production multiplant firms, each of which owns an equal share of reserves. I focus my analysis on the representative firm.

I start with a description of the demand side of the model. The demand function for the resource at time t is given by simple function $q_t = f(S_t;\epsilon_t)$, where S_t is the price of the resource at t, and ϵ_t is a demand shock realization for the resource at date t. Without loss of generality, I set $\epsilon_0 = 0$, and $\epsilon_1 = \epsilon$. Conditional on a realization of ϵ, the inverse demand function is $s = f^{-1}(q_t;\epsilon_t)$.

Let R_0 be the total resource available at date 0. Supply of the resource is optimally determined by the firm. The resource can be extracted from wells of varying quality, which is parsimoniously captured by a variable x. Wells x are uniformly distributed, $x \in [0,\bar{x}]$, in period 0, and each well is endowed with the same infinitesimal amount of resource $R_0/\bar{x}\,dx$. Well x is operated by plant x, which is owned by the firm, and has access to technology with extraction cost at date t of $x\,g(K_t)$, where K_t is the amount of capital installed by the firm. Therefore, total extraction costs of any given well depend on both the quality of the well, x, and on the amount of capital deployed. I specify $g(K_t) = \gamma_1/K_t^2$. At date 0, the plant-level decisions determine the cutoff reserve quality (extensive margin), x_0^e.

As discussed in the introduction, there are interesting relationships between the slope of the futures curve, real decisions, and expected returns on futures strategies. To address these, my model must build in the price of risk in energy commodities. Following a long literature in asset pricing, I specify an exogenous pricing kernel with a constant price of risk of the form:

$$M_{t+1} = M_t \cdot \exp(-r - \sigma_M \epsilon). \tag{1}$$

To keep things simple, I have specified that the kernel depends on the shock to energy demand so that marginal utility is high in periods of low energy demand. While oil and total consumption are not perfectly correlated in real data, I could generalize this assumption with an increase in computational complexity by adding a second not perfectly correlated shock to the kernel. Using the kernel, I can compute all expectations under the risk-neutral measure, which I will denote as $E^Q[\cdot]$.

I assume that investment and inventory decisions are made at the firm level, while extraction decisions are made at the plant level. Essentially, in the model with storage, the firm has two substitutable ways of providing the resource to consumers at date 1: it can either defer date 0 extraction and extract in date 1 or it can extract in date 0 and carry inventory to date 1. Which strategy is more profitable? Each has its own advantages, and the tradeoff is, to a large part, determined by storage costs and the expected change in extraction costs. If the latter are expected to increase rapidly, for example, it might be worthwhile for the firm to extract at date 0 and carry inventory. In addition, the price protection offered by holding the resource in the ground (like in the case of no storage) implies that an increase in uncertainty will make the delayed extraction choice more profitable.

I first discuss the firm's investment decision. To tractably capture the social learning-by-doing aspect of technological innovations in energy production (see, e.g., Covert 2015), I model capital expenditures by firms that lower the extraction costs of multiple wells owned by the firm.[10] The timing of capital installation is as follows: at date 0, the firm inherits capital of K_0 from past decisions. The firm can augment this capital stock by incurring E&D expenses, which I call investing. The new capital will follow the standard process,

$$K_{t+1} = e^{-\delta} K_t + I_t. \tag{2}$$

For the two-period model therefore, investment at date 0, I_0, affects extraction costs at date 1, and no investment at date 1 is required.

1.1 The resource extraction decision at the plant level

Conditional on the inventory and investment choices at the firm level, each plant chooses its extraction amount, Q_0^x, to maximize the profits of the plant. Conditional on the firm-level investment, the plant-level maximization can be written as:

$$\pi_0^x = \max_{0 \le Q_0^x \le \frac{R_0}{\bar{x}} dx} S_0 Q_0^x - Q_0^x x\, g(K_0) + e^{-r} E^Q[(S_1 - x\, g(K_1))^+](\frac{R_0}{\bar{x}} dx - Q_0^x), \tag{3}$$

where S_0 is the price of the resource at date 0, which the plant assumes is fixed. I will show how it is determined after the description of the firm's problem. In particular, for a firm with positive and interior production

$$S_0 - x\, g(K_0) = C(x\, g(K_1)), \tag{4}$$

where $C(x\, g(K_1))$ is the value of a one-period call option with exercise price of $x\, g(K_1)$. The left-hand side is the net gain to current extraction, while the right-hand side is the value of delaying extraction. It is useful to note at this point that the call option valuation in Equation (4) is quite similar to a regular

[10] My assumption of investment leading to social learning is consistent with Covert (2015), who documents that in the Bakken shale area, firms learn from experimenting with various combinations of inputs at actual on-site production locations, and doing so requires physical investment. Further, the results of each driller's input choices and results are made public by the regulator. Firms therefore learn about the technology and gain efficiency from the drilling experience of other drillers. Covert cites a broader literature starting with Arrow (1962) that documents such social learning-by-doing in agriculture and manufacturing. Besides the sharing of information by the regulator, in oil drilling there are additional mechanisms that facilitate social learning, including the spread of technology by the movement of specialists between regions, and firms like Schlumberger, which assist various drilling firms in different regions.

American option, with the only difference being that the price at each date of the resource is determined by the aggregate optimal extraction decision of all producers using the inverse demand function. The Kuhn-Tucker optimality condition at the boundaries for the extraction choice of plant x satisfies

$$[S_0 - x\,g(K_0) - C(x\,g(K_1))]\,Q_0^x = 0 \text{ or}$$

$$[S_0 - x\,g(K_0) - C(x\,g(K_1))]\,(\frac{R_0}{\bar{x}}dx - Q_0^x) = 0. \tag{5}$$

I complete the analysis of the model by determining the investment choice at date 0 in the following subsection.

1.2 Firm-level decisions

At date 1, since there are no further options, all plants with available resource and extraction costs smaller than the price ($x\,g(K_1) < S_1$) will extract. Given installed capital of K_1, therefore, aggregate production at date 1 will be,

$$Q_1(x^e, \epsilon) = \left(\int_{x_0^e}^{S_1/g(K_1)} \frac{1}{\bar{x}} dx \right) R_0 = \frac{S_1/g(K_1) - x_0^e}{\bar{x}} R_0. \tag{6}$$

Hence, the date 1 price is $\tilde{S}_1 = s(Q_1(x^e, \epsilon) + Z_1 e^{-u}; \epsilon)$, where Z_1 is the inventory carried by the firm from date 0 to date 1, and u is a proportional storage cost parameter. Similarly, the date 0 price is given by $S_0 = s(\frac{x_0^e}{\bar{x}} R_0 + Z_0 - Z_1; 0)$, where Z_0 is the starting inventory assumed at date 0. The firm behaving competitively takes the price at each date as given, and in equilibrium, the realized price equals the price consistent with the real choices.

Given the plant level extraction decisions, the profit at the firm level is

$$\pi_0 = \max_{I_0 > 0} \max_{x_0^e \in [0, \bar{x}]} \max_{Z_1 \in [0, \frac{x_0^e}{\bar{x}} R_0 + Z_0]} S_0 \left[\frac{x_0^e}{\bar{x}} R_0 + Z_0 - Z_1 \right] - 0.5 \frac{(x_0^e)^2}{\bar{x}} g(K_0) R_0 - p_I I_0$$

$$+ E^Q \left[e^{-(r+u)} \tilde{S}_1 Z_1 \right] + \left(\int_{x_0^e}^{\bar{x}} C(x\,g(K_1) | Y_0) dx \right) \frac{R_0}{\bar{x}}, \tag{7}$$

where $Y_0 = (x_0^e, I_0, Z_1)$ denotes the vector of policy variables. The optimality conditions for the extraction policy of the firm are

$$S_0 - x_0^e g(K_0) = C(x_0^e g(K_1) | Y_0), \text{ if } 0 < x_0^e < \bar{x}, \tag{8}$$

$$x_0^e = 0 \text{ if } s(0) < C(0 | Y_0), \tag{9}$$

$$x_0^e = \bar{x} \text{ if } s(\bar{x}) - \bar{x}\,g(K_0) > C(\bar{x} | Y_0), \tag{10}$$

where for parsimony, I have written the date 0 price, $s(x_0^e)$, only as a function of the choice of the extensive margin, x_0^e, even though it depends on the entire vector Y_0.

I can similarly formulate the firm's optimal storage policy conditional on the investment and extraction decisions. Continuing to assume that the firm is a price taker, the first order condition with respect to inventory, Z_1, is

$$-S_0+e^{-(r+u)}E^{\mathcal{Q}}[S_1]=0 \text{ if } 0<Z_1<\frac{x^e}{\bar{x}}R_0+Z_0, \tag{11}$$

$$<0 \text{ if } Z_1=0, \tag{12}$$

$$>0 \text{ if } Z_1=\frac{x^e}{\bar{x}}R_0+Z_0. \tag{13}$$

The interior case in Equation (11) determines the regular textbook equation for the value of a forward contract, $F_0=E^{\mathcal{Q}}[S_1]=S_0e^{-(r+u)}$, whereas Equation (12) occurs in "stockouts," when all available resource is consumed, and hence no inventory is carried. Finally, (13) occurs in periods when nothing in consumed at date 0, and all available resource is stored for future consumption.[11]

To compute the expected profit in (7), I can calculate the maximal investment choice numerically by choosing over a grid of values, while varying the extensive margin and inventory choices for each level of investment using the first-order conditions in Equations (8)–(10) and (11)–(13), respectively.

1.3 The futures basis and risk premium

For the linear demand case, $q_0=a-bS_0$, and $q_1=a\cdot e^{\mu+\sigma\epsilon}-bS_1$, enables us to solve for resource prices and extraction options in closed form. Equilibrium at date 1 now requires:

$$\frac{1}{\bar{x}}\left(S_1/g(K_1)-x_0^e\right)R_0+Z_1e^{-u}=ae^{\mu+\sigma\epsilon}-bS_1.$$

Solving for prices, I now have

$$S_0=1/b\left(a+Z_1-Z_0-\frac{x_0^e}{\bar{x}}R_0\right), \tag{15}$$

$$S_1=\frac{ae^{\mu+\sigma\epsilon}+\frac{x_0^e}{\bar{x}}R_0-Z_1e^{-u}}{b+\frac{R_0}{\bar{x}g(K_1)}}. \tag{16}$$

I now solve for the extraction option value in closed-form conditional on all firm-level decisions.

[11] For the interior case, by setting $S_0=e^{-(r+u)}F_0$, I have the inventory in closed form as

$$Z_1=\frac{e^{r+u}(-a+\frac{x^e}{\bar{x}}R_0+Z_0)(bg(K_1)\bar{x}+R_0)+bg(K_1)\bar{x}(ae^{\mu-\sigma_M\sigma+0.5\sigma^2}+\frac{x_0^e}{\bar{x}}R_0)}{e^{r+u}(bg(K_1)\bar{x}+R_0)+e^{-u}bg(K_1)\bar{x}}. \tag{14}$$

This equation explicitly shows how the inventory choice depends on the drilling and the investment choices.

Proposition 1. The value of the extraction call option at date 0 in the presence of a storage technology with proportional storage costs of u, given installed capital K_1, cutoff resource quality $x_0^e \in [0,\bar{x}]$, and carried inventory of Z_1 for a resource with a current extraction cost of x is:

$$C(x\,g(K_1)|Y_0)=\frac{a e^{-r}}{D}\left[e^{(\mu-\sigma_M\sigma+0.5\sigma^2)}N(-d_1)-kN(-d_2)\right],$$

$$d_1=\frac{\log(k)-m-\sigma_M\sigma-\sigma^2}{\sigma};\qquad d_2=\frac{\log(k)-m-\sigma_M\sigma}{\sigma};$$

$$k=\frac{1}{a}\left(D x\,g(K_1)-\frac{x_0^e}{\bar{x}}R_0+Z_1 e^{-u}\right);$$

$$D=b+\frac{R_0}{\bar{x}\,g(K_1)}.$$

Appendix B contains the proof.

The call option value formula resembles the call option value formula as formulated by Black and Scholes, but, importantly, it separately includes terms related to demand shocks for the resource, as well as the supply response of firms from existing capital, new investment, extraction, and inventory accumulation.

Using the stock price in Equation (16) implies that the forward price for the linear demand case satisfies,

$$F_0=E^{\mathcal{Q}}[S_1]=\frac{a e^{\mu-\sigma_M\sigma+0.5\sigma^2}+\frac{x_0^e(I_0)}{\bar{x}}R_0-Z_1(I_0)e^{-u}}{b+\frac{R_0}{\bar{x}\,g(K_1)}},\tag{17}$$

where, for the sake of exposition, I have written the extensive margin and inventory as functions of the investment decisions of the firm. What does this simple two-period model imply about the static relationship between investment and the futures basis? It is difficult to sign this relationship in general since the extensive margin and inventory choices depend on incoming capital and investment choices at date 0. Therefore, to evaluate the effects of an increase in the capital stock, K_1, on the futures price, I must consider a direct effect (in the denominator), and indirect effects through $x_0^e(I_0)$ and $Z_1(I_0)$ in the numerator. First, let's consider the direct effect by holding $x_0^e(I_0)$ and $Z_1(I_0)$ fixed. In this case, the futures price increases in the technological component of extraction costs, $g(K_1)$, and the spot price as calculated by (15) does not change. Under the assumption that $g'(K_1)<0$, higher investment implies higher future capital, which implies lower extraction costs, resulting in an increase in future supply so that the futures basis declines. The indirect effect, on the other hand, implies that if $x_0^e(I_0)$ is higher in periods of high capital (the firm supply decision is more elastic due to lower extraction costs), the future prices will he higher due to a decline in the future availability of the resource, while date 0 prices will be lower, resulting in an increase in the futures basis. In addition, investment can also affect inventory, Z_1, and have further indirect effects on the futures basis. If firms choose lower inventory with higher capital, this will increase the futures price and lower the spot price and will again raise the basis.[12]

Combining the futures and expected spot price, I have

$$\frac{E[S_T]-F_0}{F_0}=\frac{a e^{\mu+0.5\sigma^2}(1-e^{-\sigma_m\sigma})}{a e^{\mu-\sigma_m\sigma+0.5\sigma^2}+\frac{x_0^e(I_0)}{\bar{x}}R_0-Z_1(I_0)e^{-u}} \tag{18}$$

which is the expected excess return, or the risk premium, from a fully collateralized long futures position. Keynes (1930) observed that speculators mostly take the short side of futures contracts and therefore require a risk premium for holding the commodity risk. Therefore, the futures price that the producers (hedgers) would sell at should be lower than the expected spot price that they could obtain by holding the commodity and selling in the future. The use of the left-hand side as a measure of the risk premium is now standard in the literature (see, e.g., Gorton and Rouwenhourst 2006). However, here I notice that the risk premium for the commodity is not simply $\sigma_M\sigma$, which is the premium for bearing the risk of the demand shocks, but it also depends on the firm's investment decision through its effect on the extensive margin and inventory, each of which is endogenous. These variables affect the future supply of the resource, which then either amplifies or dampens the demand shocks. I note that these endogenous variables affect the denominator of (18), which is the cost of entering into the collateralized futures and bond position. Finally, the risk premium is positively related to inventory and negatively related to the extensive margin, and I will discuss intuition for the latter in Section 3.4.

2. The Infinite Horizon Model with Production, Exploration, and Storage

I preserve much of the structure of the two-period model. However, I introduce four additional assumptions: first, wells produce for multiple periods, albeit at declining rates; second, the inverse demand function for oil exhibits habit formation; third, there are quadratic adjustment costs in capital accumulation; and fourth, I specify a mean-reverting process for demand shocks. The first assumption is made to differentiate the processes of drilling and resource extraction, since the former is fairly irregular, whereas the latter is fairly smooth, as is observed in the data. The second assumption, on habit formation, implies that oil prices respond to deviations of consumption from the habit level leading to volatile prices despite relatively smooth consumption. The third assumption is made to capture the difficulties firms have in changing the scale of their investments quickly, and as a result, to smooth out the investment process to match the volatility of the investment-to-capital ratio in the data. Finally, the

12 Some comparative static results are evaluated numerically and illustrate these effects. These results are available from the author on request.

last assumption is made so that prices have a mean-reverting component, as is commonly modeled in the literature on pricing commodities.

The demand function for the resource at time t is given given $q_t = f(S_t, H_t, \epsilon_t)$, where H_t is the level of the external oil consumption habit. The assumption that habit is external means that agents take the habit level as exogenously given and behave as if their current decisions do not affect its future evolution (see Abel 1990; Campbell and Cochrane 1999). I will specify the process that determines H_t when I discuss the oil production process.

The demand shock realization for the resource at date t is ϵ_t. Conditional on a realization of the habit and demand shock, the inverse demand function is $s_t = f^{-1}(q_t; H_t, \epsilon_t)$. For testing the predictions of my model, I use an inverse demand function of the form:

$$s(q_t; H_t, \epsilon_t) = e^{b+\epsilon_t} / (\omega + q_t - H_t)^{\alpha}, \tag{19}$$

where $\alpha > 1$. The quantity $\omega > 0$ is some exogenous consumption of a substitute to oil. I use $\alpha > 1$ since most empirical estimates of the elasticity of oil consumption are smaller than 1. I note that this inverse demand function does not satisfy the Inada condition at zero oil consumption. In addition, I assume the same form for the pricing kernel like in the two-period model as specified in (1) with a constant interest rate and price or risk.[13]

I assume that the demand shock follows a mean-reverting Ornstein-Uhlenbeck (OU) process:

$$\epsilon_{t+1} - \epsilon_t = -k_\epsilon \epsilon_t + \sigma_\epsilon (1 + |\epsilon_t|) e_t, \tag{20}$$

where $e_t \sim N(0,1)$. The use of mean-reverting demand shocks is standard in the commodity pricing literature (e.g., Carlson, Khokher, and Titman 2007; Pirrong 2012). The process exhibits time-varying volatility, which has a V-shaped relation to the demand shock. Therefore, volatility is high when demand is extremely low or extremely high. This feature captures the essence of time-varying uncertainty of fundamentals that is now standard in macroeconomics and finance (see, e.g., Bloom 2009; David and Veronesi 2013). I model this feature to slow the process of extraction in my model, since, even in periods of high demand there is value to waiting because of higher uncertainty.[14]

Let x_t^e be the extensive margin at the start of time t. I assume that drilling decisions are made at the plant level, as in the two-period model in Section 1.

---

[13] A constant interest rate and price of risk arises in general equilibrium models where aggregate consumption follows a homoscedastic process and investors have standard constant relative risk aversion preferences over aggregate consumption (see, e.g., Chapter 1 of Campbell and Cochrane 2001). Here I do not model habit formation in aggregate consumption, but only in oil consumption (a part of aggregate consumption), and I will further motivate this next. Ravn, Schmitt-Grohe, and Uribe (2006) assume habits for different goods. Relatedly, Chetty and Szeidl (2007) find that U.S. consumers get committed to consuming different goods, and doing so affects their preferences on future consumption changes.

[14] In the Online Appendix I report results for the model with constant volatility of demand shocks. I obtain the same signs on the basis and risk premium regressions on investment as for the above model, but the R^2s are smaller.

An increase in the extensive margin by i_t raises the extensive margin to $x^e_{t+1} = x^e_t + i_t$. Then, at date t, the resource production from newly drilled wells equals

$$Q^n_t = R_0 \cdot \int_{x^e_t}^{x^e_t + i_t} \frac{1}{\bar{x}} dx = R_0 \frac{i_t}{\bar{x}}. \tag{21}$$

The drilling costs incurred by the firm at date t are:

$$C_t = g(K_t) \cdot R_0 \cdot \int_{x^e_t}^{i_t + x^e_t} \frac{x}{\bar{x}} dx = \frac{1}{2} g(K_t) R_0 \frac{(x^e_t + i_t)^2 - (x^e_t)^2}{\bar{x}}. \tag{22}$$

It is useful to note that drilling costs are not simply proportional to i_t^2, but are instead proportional to $i_t^2 + 2x_t i_t$. This is because an increase in the extensive margin leads to higher resource extraction costs as lower quality wells are accessed. An interesting implication is that the industry will have to maintain a higher level of capital stock over time to maintain a constant level of extraction costs.

Following Anderson, Kellogg and Salant (2018), I assume that drilled wells continue to produce at close to zero marginal extraction costs, albeit at geometrically declining rates due to a steady loss of well pressure. I further assume that production at all wells declines at the same rate, $0 < \beta < 1$. This implies that total resource production at date t equals

$$Q_t = \beta Q_{t-1} + Q^n_t. \tag{23}$$

Therefore, existing oil production implies a deterministic amount of production at future dates, which will be augmented by production from new drilling.

I now return to the specification of the habit process in the demand function. In particular, I assume that the habit at t satisfies $H_t = \beta Q_{t-1}$; that is the level of habit at time t equals the continued production from wells drilled in previous periods. As discussed in the previous paragraph, production from existing wells declines deterministically at a geometric rate. Since consumers in the economy would rationally expect a level of production of at least H_t in period t, it is reasonable that they would make plans for this level of consumption so that deviations from it, rather than the level of consumption, would lead to fluctuations in marginal utility (inverse demand). My assumption on a persistent habit process is similar to most papers in the literature (see, e.g., Campbell and Cochrane 1999), but these papers assume that the habit process is purely exogenous. My specification here, of habit in oil consumption, ties it specifically to the predictable part of production that arises from the natural decline rate of production of drilled wells.

The final assumption I introduce in the dynamic model is the quadratic adjustment cost of capital. In particular, adjustment costs satisfy

$$\text{ac}(I, K) = \gamma_2 \frac{I^2}{K}, \tag{24}$$

where γ_2 is a constant. Such an assumption is standard in macroeconomic models since the work of Hayashi (1982) for controlling the volatility of the investment-to-capital ratio.

As for the two-period model, the firm has a costly storage technology. It is able to place a nonnegative quantity, Z_{t+1}, in storage at date t. Storage costs are a proportion, u, of the quantity stored, so that an amount, $Z_{t+1}e^{-u}$, is available for consumption at date $t+1$. The firm behaves competitively in production markets, and I assume here that its storage decision has no price impact either. For the competitive case, I alternatively could assume that inventory decisions are made by speculators. However, with complete markets, the equilibrium will be identical with storage by either the firm or the speculators. Combining production like in (21) and inventory, the total amount available for consumption in period t is

$$q_t = Q_t + Z_t - Z_{t+1}. \tag{25}$$

If there is a stockout, then $Z_{t+1}=0$; that is, all available resource is consumed in date t.

To solve for equilibrium prices and quantities, I solve the related problem of a social planner who maximizes the discounted expected consumer plus producer surplus (see, e.g., Weinstein and Zeckhauser 1975; Carlson, Khokher, and Titman 2007). The social surplus at time t is therefore,

$$\mathrm{SS}_t = \int_0^{q_t} s(q; H_t, \epsilon_t)dq - C_t - p_I I_t - \mathrm{ac}(I_t, K_t), \tag{26}$$

where total production from new wells drilled and existing wells, drilling costs, consumption, and adjustment costs are given in Equations (23), (22), (25), and (24), respectively, and p_I is the constant price of capital goods in units of consumption goods.

The social-planning problem can be solved using standard dynamic programming methods. The Hamilton-Jacobi-Bellman equation is:

$$\begin{aligned} &J(x_t^e, Z_t, K_t, \epsilon_t, Q_{t-1}) \\ &= \max_{i_t \in [0, \bar{x}-x_t^e],\, Z_{t+1} \in [0, Z_t + Q_t^n],\, 0 \le I_t} SS_t + e^{-r} E^{\mathcal{Q}}[J(x_t^e + i_t, e^{-u} Z_{t+1}, e^{-\delta} K_t + I_t, \epsilon_{t+1}, Q_t)]. \end{aligned} \tag{27}$$

Because the level of habit in my specification equals lagged total output, I do not need an additional state variable. To economize on notation below, I suppress the arguments of the J function and write $J_t = J(x_t^e, Z_t, K_t, \epsilon_t, Q_{t-1})$ and $J_{t+1} = J(x_t^e + i_t, e^{-u} Z_{t+1}, e^{-\delta} K_t + I_t, \epsilon_{t+1}, Q_t)$.

The first-order conditions for this problem are:

$$\frac{R_0\left(s(q_t; Q_{t-1}, \epsilon_t) - (x_t^e + i_t) g(K_t)\right)}{\bar{x}} + e^{-r} E^{\mathcal{Q}}[J_{x,t+1}] \le 0; = 0 \text{ if } i_t > 0, \tag{28}$$

$$\frac{R_0\left(s(q_t;Q_{t-1},\epsilon_t)-(x_t^e+i_t)g(K_t)\right)}{\bar{x}}+e^{-r}E^{Q}[J_{x,t+1}]\geq 0 \text{ if } i_t=\bar{x}-x_t^e, \quad (29)$$

$$-s(q_t;Q_{t-1},\epsilon_t)+e^{-(r+u)}E^{Q}[J_{Z,t+1}]\leq 0;=0 \text{ if } Z_{t+1}>0, \quad (30)$$

$$-s(q_t;Q_{t-1},\epsilon_t)+e^{-(r+u)}E^{Q}[J_{Z,t+1}]\geq 0 \text{ if } Z_{t+1}=Z_t+Q_t^n, \quad (31)$$

$$-p_I-\frac{2\gamma_2 I}{K}+e^{-r}E^{Q}[J_{K,t+1}]\leq 0;=0 \text{ if } 0<I_t. \quad (32)$$

A few points are worth noting. First, the optimality of the extensive margin and inventory must be checked at both the lower and upper boundaries. Second, my assumption that the level of habit equals lagged production implies that the spot price in Equations (28)–(31), $s(q_t;Q_{t-1},\epsilon_t)=e^{b+\epsilon_t}/(\omega+\frac{i_t}{\bar{x}}R_0+Z_{t+1}-Z_t)^{\alpha}$, depends on new production and current inventory accumulation, but not on lagged production. Finally, this observation implies that lagged production does not appear in the social surplus at any time t, and hence it affects neither the first order conditions nor the value function. I solve the HJB equation using projection methods as described in Judd (1999). Appendix C provides details of the approximation method.

Next, I turn to the computation of futures prices and the risk premium for the infinite horizon model. The nth-period-ahead futures prices is the nth-period-ahead expected spot price under the risk-neutral measure, which can be written as $F_{t,n}=E^{Q}(S_{t+n})$. Using the policy functions written in polynomial form, I can calculate production at each future date and state, and, using the inverse demand function and the Markovian shocks, I can compute the futures prices. In particular, I recursively calculate the nth-period-ahead futures price as:

$$F(x_t^e,Z_t,K_t,\epsilon_t;n)=E^{Q}[F(x_t^e+i_t,e^{-u}Z_{t+1},e^{-\delta}K_t+I_t,\epsilon_{t+1};n-1)], \quad (33)$$

$$F(x_t^e,Z_t,K_t,\epsilon_t;0)=\frac{e^{b+\epsilon_t}}{(\omega+\frac{i_t}{\bar{x}}R_0+Z_t-Z_{t+1})^{\alpha}}. \quad (34)$$

It is again important to observe that, since firms' optimal policy does not depend on lagged production, Q_t, it is not a state variable for the futures price recursion. Analogously, the expected spot price can be recursively calculated at each horizon n as:

$$\mathcal{E}(x_t^e,Z_t,K_t,\epsilon_t;n)=E[\mathcal{E}(x_t^e+i_t,e^{-u}Z_{t+1},e^{-\delta}K_t+I_t,\epsilon_{t+1};n-1)], \quad (35)$$

$$\mathcal{E}(x_t^e,Z_t,K_t,\epsilon_t;0)=\frac{e^{b+\epsilon_t}}{(\omega+\frac{i_t}{\bar{x}}R_0+Z_t-Z_{t+1})^{\alpha}}. \quad (36)$$

where now the expectation is under the objective measure. All expectations are calculated using Gaussian Quadrature.[15] Finally, the risk premium is

$$\Omega(x_t^e, Z_t, K_t, \epsilon_t; n) = \frac{\mathcal{E}(x_t^e, Z_t, K_t, \epsilon_t; n)}{F(x_t^e, Z_t, K_t, \epsilon_t; n)} - 1. \tag{37}$$

3. Explaining the Stylized Facts

In this section, I provide simulation results from the infinite horizon model in Section 2 and then provide some intuition as to why the risk premium on futures is inversely related to investment.

3.1 Model calibration

I start with the calibration of the preference parameters. For the inverse demand function, I choose $\alpha = 2$, and $b = 0$. I set the persistence of the external habit level, $\beta = 0.9$, at an annual rate, which is the value used in Campbell and Cochrane (1999). I set $\omega = 0.075$ so that alternative substitutes to oil are about one-third of oil consumption of about 0.21 in the first 100 years of the simulation. I assume that the price of risk is $\sigma_m = 0.3$. This is around the standard level used in asset-pricing models to justify an aggregate Sharpe ratio of 30% on stocks, which close to its historical average.

Next, I calibrate the technology parameters. The extraction cost function that I continue to use is $g(K_t) = \frac{\gamma_1}{K_t^2}$. This implies that extraction costs decline as capital accumulates, but explode as capital tends to zero, so that positive capital is required to ensure the supply of the resource. The technology parameter I use is $\gamma_1 = 150$. I set the adjustment cost parameter to $\gamma_2 = 1$, which is the same as the level set in several recent papers on investment (see e.g., Belo, Lin, and Bazdresch 2014). In my model, the persistence in production from wells equals that of the habit level, which for $\beta = 0.9$, implies an average half-life across wells of 6.57 years. While wells from fracking are estimated to have shorter half-lives than the average, conventional wells' half-lives are longer, some of which produce for decades. Using data from the Federal Reserve Bank of St. Louis, I calculate the average relative price of investment goods to consumption goods in the U.S. from 1986 to 2015 of 1.27, and hence set the price to $p_I = 1.27$.

The parameters for the inverse demand function determine the level of prices, while the technology parameters determine extraction costs. Together, the parameters chosen determine the gross margin of resource production. The parameters are chosen jointly to match the historical average gross margin of oil-producing firms. Using data from Compustat, I find that the historical average of the gross margin of oil firms is about 20%. I set the rate of capital depreciation at 10% a year, a standard rate assumed in the real business-cycle literature. I set

15 See, e.g., Tauchen and Hussey (1991). For a textbook description of this method of integration, see Judd (1999).

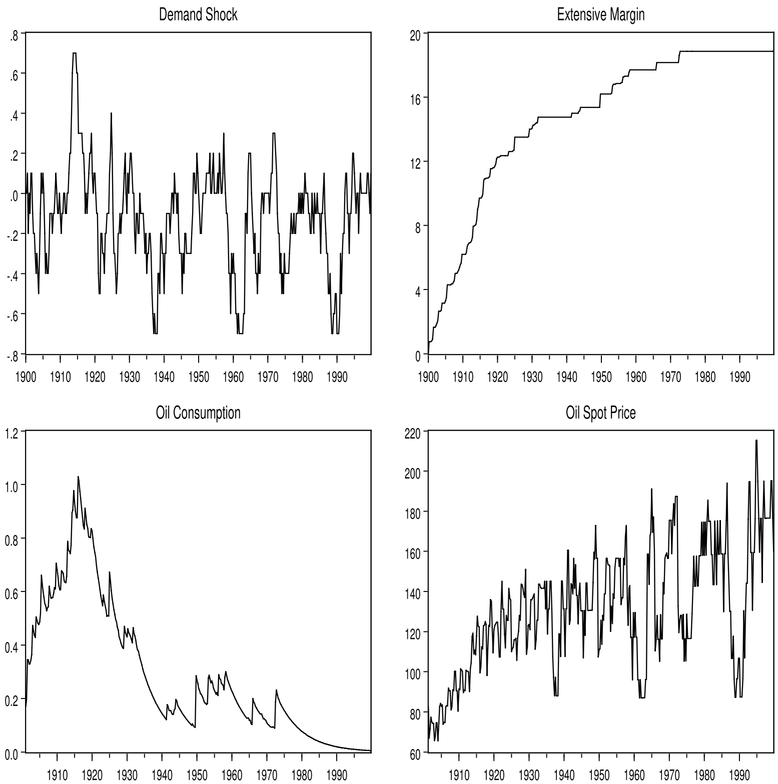

Figure 3
Optimal firm's decisions from a single simulation of infinite horizon model
I show the simulated time series of the model in Section 2 using the model's parameters in Section 3.1.

proportional storage costs of 5% a year, similar to that in Routledge, Seppi, and Spatt (2000) and Pirrong (2012).

The parameters for the demand shock process in Equation (20) that I use are $k_\epsilon = -0.3$, and $\sigma_\epsilon = 0.2$. The drift parameter governing the speed of mean reversion is the same as that in Carlson, Khokher, and Titman (2007), whereas I choose a lower volatility, since I scale up the volatility by the amount $(1+|\epsilon_t|)$.

I also make some choices on the scale of the problem. I assume that $\bar{x} = 50$, and the total reserves of the resource, $R = 10$. I believe these do not affect the results of the paper.

3.2 Results from a single simulation

I highlight in this section that this model displays inventory and investment cycles that help explain the variation in the relative basis and the risk premium. I start by showing in Figures 3–6 the first 100 years of a 300-year simulation of the model's real and financial variables.

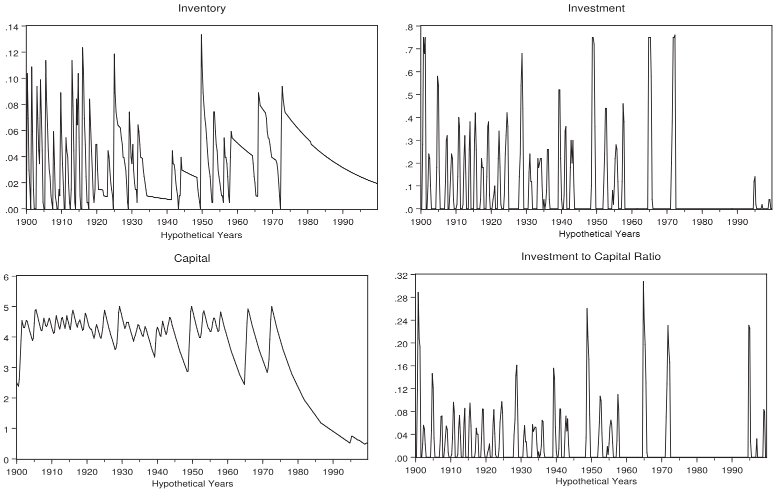

Figure 4
Optimal firm's decisions from a single simulation of infinite horizon model
I show the simulated time series of the model in Section 2 using the model's parameters in Section 3.1.

The top-left panel of Figure 3 shows the simulated demand shock process, which fluctuates around 0. The demand process exhibits some persistent booms and busts that can be relatively short or as long as 10 years, but there are shorter episodes as well. The top-right panel shows how the extensive margin expands over time. The extensive model is determined by firms' drilling decisions, as was discussed in the previous section. Indeed, new drilling can be calculated as the first difference of the extensive margin series. As can be seen, after 70 years, resources with an extraction cost of 17 are extracted, while the maximum is 50; that is 34% of the resource is extracted. The bottom left panel shows the consumption of the resource, which results from production from new as well as old wells, and inventory decisions as shown in (25). In the simulation shown, consumption peaks at about 15 years, and continues to fluctuate until about year 70. It then declines to a trickle after year 75 and remains so for several decades. Consumption does resume after year 100 (data not shown). In the next subsection, I will see that the point at which extraction of the resource stops is highly variable, and, in some simulations, it is as high as 250-300 years. The bottom-right panel shows the spot price in the model has a positive trend despite some large fluctuations.

Figure 4 shows the firm's optimal decisions. The top-left panel shows the level of inventory. Inventory fluctuates more so in the first 30 years of the simulation than in later years. In the first 30 years, resource extraction is more frequent; hence inventory plays a role in smoothing its fluctuations at a relatively high frequency. In later years, extraction slows, and now inventory is carried longer as consumption drops. Finally, after year 70, inventory is held for even

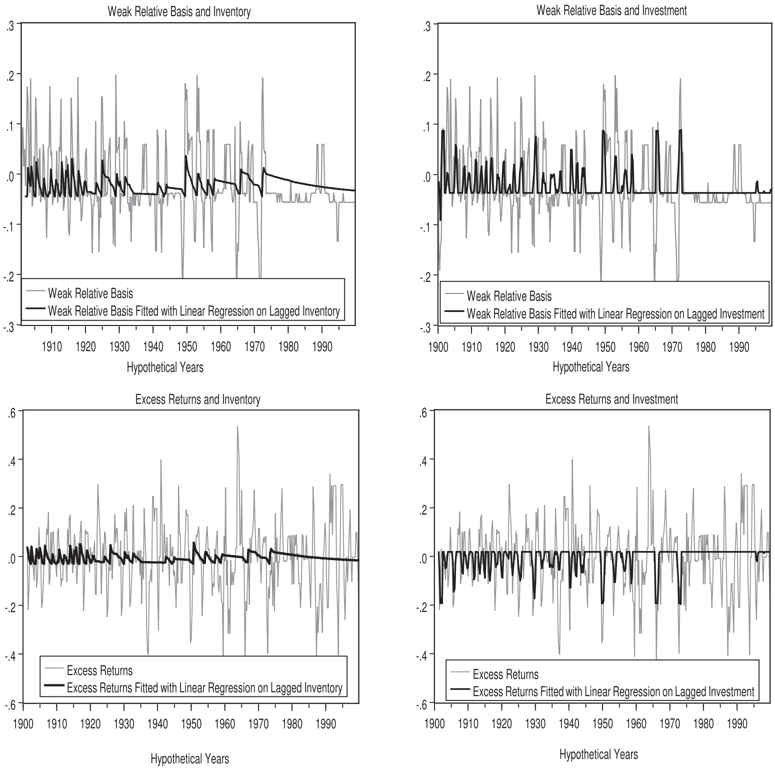

Figure 5
Weak relative basis and excess returns from a single simulation of infinite horizon model
I show the simulated time series of the model in Section 2 using the model's parameters in Section 3.1. Futures prices in the model at each date are calculated using equation (33), while the spot price is calculated using equation (34). Excess returns are calculated as $S_t/F_{t-11,12}-1$. The fitted values of the weak basis shown in the top two panels are from the estimated regression coefficients in lines 1 and 2 of Table 3, respectively. The fitted values of the excess returns shown in the bottom two panels are from the estimated regression coefficients in lines 1 and 2 of Table 4, respectively.

longer periods until extraction restarts (after year 100). The top-right panel shows that investment also fluctuates at a higher frequency until about year 30, as firms maintain their capital stock levels to resume extraction quickly in periods of demand spikes. In later years, as increasingly more costly resource is extracted, investment stops in periods of low demand and only resumes when demand returns to high levels. The capital stock of oil firms (shown in the bottom-left panel) fluctuates around a level of 4 until year 70, as the firm keeps investing albeit at a declining frequency. After year 70, when extraction costs rise and investment becomes even more irregular, the capital stock declines slowly and is only reaccumulated once demand is high again (after year 100). Inventory and investment have a correlation of about −0.4 in the shown sample path, while it is about −0.2% in the data. The bottom-right panel shows the

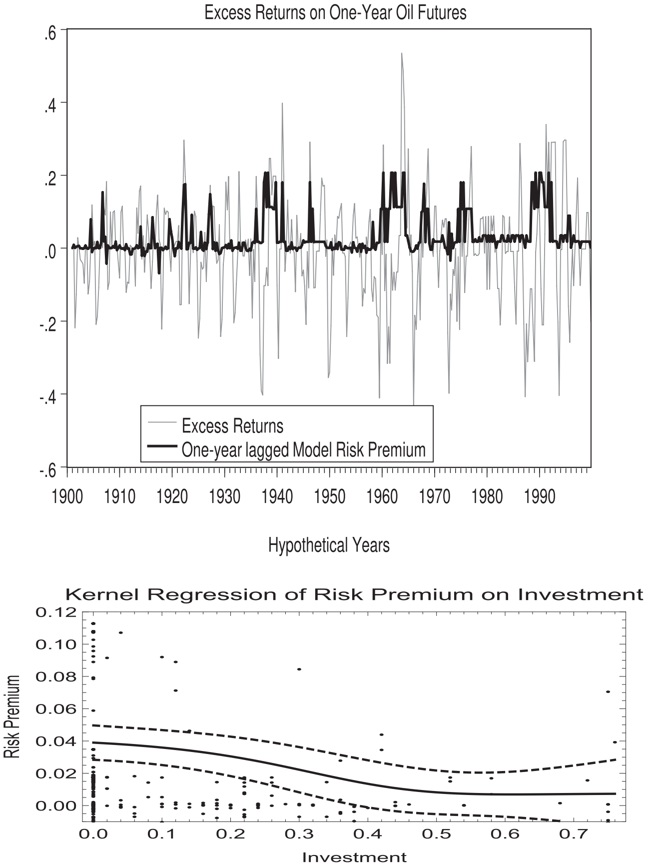

Figure 6
Model's risk premium and excess returns on one-year futures
In the top panel I the simulated time series of the model in Section 2. The parameter choices are in Section 3.1 . The model risk premium at each date is calculated using equation (37). The bottom panel shows the fit from a nonlinear kernel regression using the Nadarya-Watson kernel (see, e.g., Campbell, Lo, and MacKinlay 1995) of the model's simulated risk premium (shown in the top panel) on the model's investment (shown in the top-right panel of Figure 4). 95% confidence bands are shown in dashed lines.

investment-capital ratio. The volatility of this ratio is around 8% at an annual rate, which is similar to the data.

Figure 5 shows the main variables of interest of the paper. The top-left panel shows the path of the model's weak relative basis and its fitted value from a linear regression on lagged inventory. Line 1 of Table 3 shows the regression fit. As seen, lagged inventory has a positive coefficient, and it explains 4.7% of the variation in the basis, similar to my data sample in Table 1. The panel conveys much of the intuition behind the relation between inventory and the basis from the storage literature. The two series are positively related, especially when the basis is positive. In such periods, the positive basis provides incentive to carry

Table 3
What explains the futures weak relative basis for crude oil in 100-year model simulation?

	α	β_1	β_2	R^2
1	−0.044	0.605		0.047
	[−5.825]	[3.661]		
2	−0.036		0.165	0.153
	[−7.833]		[6.015]	
3	−0.056	0.565	0.162	0.199
	[−7.077]	[3.544]	[6.326]	

I report the coefficients of the fitted regression for the model simulation:

$$\text{Weak relative basis}(t) = \alpha + \beta_1 \text{Inventory}(t-1) + \beta_2 \text{Investment}(t-1) + \epsilon(t).$$

The weak relative basis on 1-year contracts in quarter t is $[e^{-r(t)}F(t) - S(t)]/S(t)$, where $F(t)$ is the 1-year futures prices at the beginning of each quarter and $S(t)$ is the spot price. t-statistics are in brackets and are adjusted for heteroscedasticity and autocorrelation.

the resource forward into time. The inventory has limited explanatory power when the basis is negative, since it is constrained to be nonnegative. Indeed there are large downward spikes (backwardation) in the panel when inventory becomes zero. In such periods, the agents would like to transfer resources from the future to the current, but since they are unable to, spot prices spike and the futures become backwardated. This impact of "stock outs" on backwardation is similar to that in Routledge, Seppi, and Spatt (2000) where production decisions are exogenous.

The top-right panel of Figure 5 shows the basis and its fitted value from a linear regression on lagged investment whose fit is shown in line 2 of Table 3. As seen, the spikes in lagged investment (also evident in Figure 4) are able to explain a number of spikes in the basis. Overall for the 100-year simulation, lagged investment has a positive coefficient and it explains about 15% of the variation in the basis, which is smaller than in my data sample (Table 1). As seen in the panel, the basis and investment are more strongly positively related until year 70 (as in the data) when investment is more frequent. The positive correlation shows that, in the model during periods of a positive basis, firms also have the incentive to accumulate capital to lower their future extraction costs. Using both inventory and investment as explanatory variables for the basis, I obtain an R^2 of nearly 20% compared to about 30% in my data sample.

The bottom panels of Figure 5 show the 1-year excess returns on the fully collateralized 1-year futures contract, which is calculated using (37), along with its fitted values from the univariate regressions on inventory and investment, respectively. Lines 1 and 2 of Table 4 show the fits of the regressions. While the left panel shows a weak relation between inventory and excess returns, the right panel shows a stronger relationship: between investment and excess returns the R^2 on inventory is only about 1.5%, whereas that on investment has an R^2 of over 11%. It is useful to note that investment mainly predicts negative returns in periods when it is high, but being constrained to be positive, it has little explanatory power in periods when it is close to zero. Line 3 shows

Table 4
What explains the risk premium on crude oil futures in 100-year model simulation?

	α	β_1	β_2	R^2
1	−0.029	0.644		0.014
	[−1.868]	[2.361]		
2	0.018		−0.281	0.113
	[1.676]		[−6.377]	
3	0.032	−0.335	−0.305	0.116
	[1.697]	[−1.082]	[−5.858]	

I report the coefficients of the fitted regression for the model simulation: Excess return(t)=

$$\alpha+\beta_1 \text{Inventory}(t-12)+\beta_2 \text{Investment}(t-12)+\epsilon(t).$$

The excess return on a fully collateralized 1-year futures contract is defined as $S(t)/F(t-11)-1$, where $F(t)$ is the 1-year futures price, and $S(t)$ is the spot price. t-statistics are in brackets and are adjusted for heteroscedasticity and autocorrelation.

the fit when both regressors are used, and in it only investment is significant. The fit of inventory in the data is better; however, the lower R^2 of my model arises mostly from the smoothness of inventory after year 70 when consumption declines steadily.

In the final exhibit of the simulation, in the top panel of Figure 6, I plot the 1-year excess returns and the 1-year lagged model implied risk premium as calculated in (37). The model risk premium aggregates the information in both inventory and investment, which I examined above. The coefficient for the univariate regression of excess returns on the 1-year lagged premium is 0.7, while the R^2 is about 7%. Notably, the positive slope coefficient is not statistically different from 1. As with the other regressions, the R^2 is higher, at nearly 11%, for the first 70 years of the simulation. The lower R^2 relative to the multivariate regression in line 3 of Table 4 results from the nonlinear relationship between investment and the risk premium. I make a couple of important observations. First, in contrast to the bottom-right panel of Figure 5, the model's risk premium does a better job of predicting high excess returns than low excess returns. I shed further light on this issue by plotting the nonlinear fit from a nonparametric kernel regression of the model's risk premium on the model's investment in the bottom panel. As seen, the relationship between these variables is mainly negative but flattens out considerably for higher levels of investment. The risk premium is positive (it averages 3.3% for the full sample), but is significantly lower for higher levels of investment. Second, the model's risk premium turns slightly negative in about 20% of the observations. I will provide an explanation for these features in Section 3.4 below.

3.3 Results from multiple simulations

While the results from the single simulation shed light on the observed relations between real and financial variables, the exact sequence of demand shocks determines the decisions of firms and potentially the statistical significance of my results. In this section, I simulate 100 sample paths of 300 years and examine the distribution of several statistics for the first 100 years of each sample.

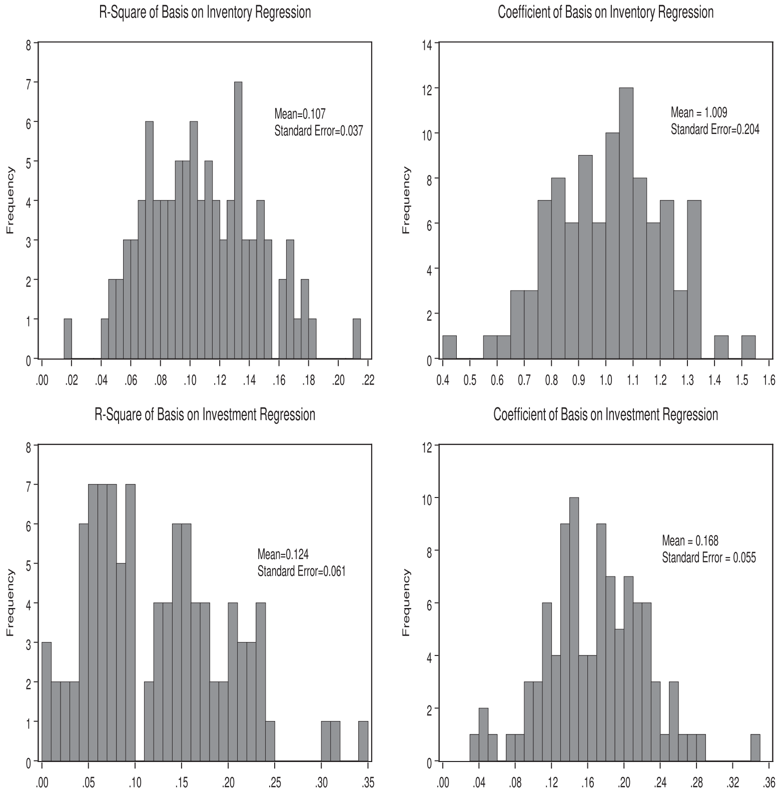

Figure 7
Distribution of statistics for basis regressions from multiple simulations
I simulate time series of the model in Section 2 using the model's parameters in Section 3.1. I report the average R^2 and coefficients of the univariate regressions in Table 3 across simulations. For the basis calculations, futures prices in the model at each date are calculated using Equation (33), while the spot price is calculated using Equation (34).

I show the distribution of the regression statistics for the basis in Figure 7. The top panels show the distribution of the R^2 and the slope coefficient of the univariate regression of the basis on lagged inventory, respectively. The basis coefficient averages about 1 across samples and the R^2 averages 10.7%, so that my model captures well the relationship between the basis and inventory. The bottom panels show the analagous statistics for the univariate regression on lagged investment. The slope coefficient is positive (0.168), and the R^2 average is 0.124, which is higher than that for inventory, like in the data.

Similarly, I show the distribution of the regression statistics for excess returns in Figure 8. As for the single sample path in the previous subsection, and for the data, the coefficient on inventory is positive while that of investment is negative. I will provide intuition for the sign of the investment slope coefficient in the next subsection. The average R^2 for inventory of 2.1% is smaller than that for

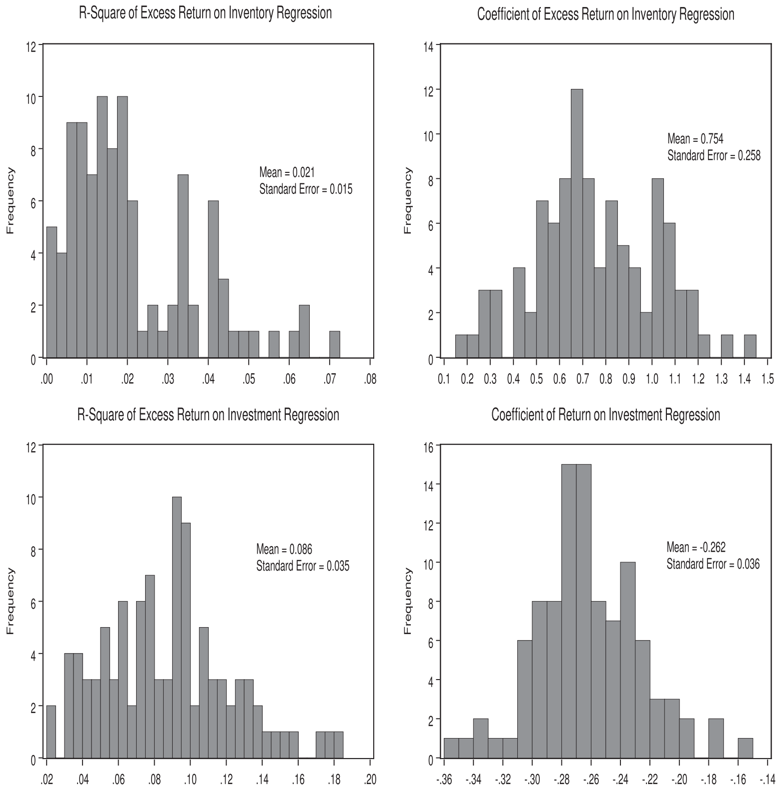

Figure 8
Distribution of statistics for excess return regressions from multiple simulations
I simulate time series of the model in Section 2 using the model's parameters in Section 3.1. I report the average R^2 and coefficients of the univariate regressions in Table 4 across simulations. Futures prices in the model at each date are calculated using Equation (33) and the spot price is calculated using Equation (34). Excess returns are calculated as $S_t/F_{t-11,12}-1$.

investment of 8.6%, once again consistent with the data. As discussed in the previous subsection, the R^2 is higher in the first 70 years of the subsamples, since investment tends to drop off around this time.

In Figure 9, I show the distribution of the final year of extraction across simulations. In the model, the quality of wells accessed declines over time and extraction requires increasingly larger investments for the firm to be able to extract profitably. Eventually, extraction stops when the new investment required is too large relative to the value of the resource produced. As can be seen, the most frequent time of stopping extraction is 120 years, although there is considerable variation, and there is significant probability that extraction continues, albeit at a declining rate, to years 250–300. As time proceeds, increasingly costly resource is extracted, and extraction occurs mainly in periods of high demand. It is finally noteworthy, that the regression results in the

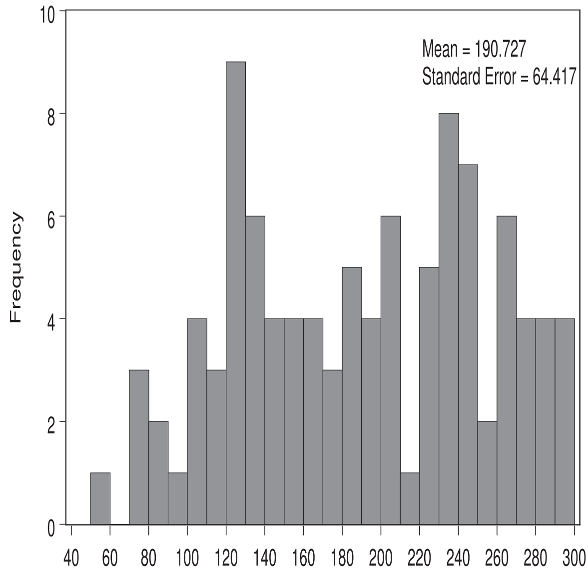

Figure 9
Distribution of final year of extraction
I simulate time series of the model in Section 2 using the model's parameters in Section 3.1.

model are similar to those in the data despite explicitly modeling exhaustibility, which adds an element of nonstationarity to my model. The declining rate of consumption in the model is not at odds with the declining share of oil consumption in the past 30 years that I showed in the middle panel of Figure 2.

3.4 Why is investment negatively related with the risk premium?

In this subsection, I provide intuition on why the risk premium on futures contracts is negatively related to the capital stock in the oil industry. Since investment and capital are positively related in my simulations using the calibrated parameters, this also means a negative relation between investment and the risk premium. The risk premium on a futures contract with n periods to maturity is calculated using Equation(37), which can also be expressed as

$$\Omega_t = \frac{\mathcal{E}_t}{F_t} - 1 = \frac{E^{\mathcal{P}}(S_{t+n})}{E^{\mathcal{Q}}(S_{t+n})} - 1,$$

where $E^{\mathcal{P}}(\cdot)$ and $E^{\mathcal{Q}}(\cdot)$ are the expectations under the objective $\mathcal{P}$-measure, and the risk-neutral $\mathcal{Q}$-measure, respectively.[16]

16 For notational convenience, whenever I skip a superscript on the expectations operator, I mean expectations under the $\mathcal{P}$-measure.

The demand shock process, $\{\epsilon_t\}$, under the objective measure is like that in Equation (20), whereas under the $\mathcal{Q}$-measure using Girsanov's theorem it is

$$\epsilon^*_{t+1} - \epsilon^*_t = -(k_\epsilon + \sigma_M)\epsilon^*_t + \sigma_\epsilon(1 + |\epsilon^*_t|)e^*_t. \tag{38}$$

With a negative adjustment to the drift of the shock, therefore, expectations under the $\mathcal{Q}$-measure put greater weight on low demand shocks than the $\mathcal{P}$-measure. I now examine how prices in the model respond to demand shocks. The spot price at t is

$$s(q_t; Q_{t-1}, \epsilon_t) = e^{b+\epsilon_t} / (\omega + \frac{i_t}{\bar{x}} R_0 + Z_{t+1} - Z_t)^\alpha. \tag{39}$$

There are two effects of demand shocks on oil prices. First, a higher shock implies a higher price from the numerator. Second, a higher demand shock could imply a lower or higher quantity extracted (and consumed) in the denominator. For relatively low levels of the capital stock, extraction costs are high and supply is fairly inelastic, and hence the quantity extracted will only increase slightly, exerting only a small downward pressure on prices, and hence future spot prices are strongly positively correlated with demand shocks. For higher levels of the capital stock, extraction costs are low, supply is more elastic, and hence the quantity extracted can increase rapidly with a higher demand shock. In such situations, the second effect becomes larger, puts downward pressure on spot prices, and lowers the correlation of demand shocks and spot prices.

It is now evident why the risk premium is negatively related to the capital stock. When capital is relatively low, the first effect in the previous paragraph dominates, so that demand shocks and future spot prices are positively correlated. Since the $\mathcal{Q}$-measure puts greater weight than the $\mathcal{P}$-measure on low demand shocks, the futures price then becomes lower than the expected (under the $\mathcal{P}$-measure) spot price, which means a high risk premium. When capital is very high, the second effect in the previous paragraph dominates, so that demand shocks and future spot prices are negatively correlated. Now, the $\mathcal{Q}$-measure puts greater weight on low demand shocks, or higher future spot prices, so that they are higher than expected spot prices under the $\mathcal{P}$-measure, which means a low (and possibly negative) risk premium.

The relationship between the capital stock and the risk premium also can be explained from the covariance of the pricing kernel with spot prices. I have $F_t = E^{\mathcal{Q}}_t(S_{t+n}) = E(\frac{M_{t+n}}{M_t} S_{t+n})$, which also can be written as

$$0 = E(\frac{M_{t+n}}{M_t} \frac{S_{t+n} - F_t}{F_t}) \tag{40}$$

$$= E(\frac{M_{t+n}}{M_t}) E(\frac{S_{t+n} - F_t}{F_t}) + \mathrm{Cov}(\frac{M_{t+n}}{M_t}, \frac{S_{t+n} - F_t}{F_t}). \tag{41}$$

Therefore, the risk premium can be written as

$$E(\frac{S_{t+n} - F_t}{F_t}) = -\mathrm{Cov}(\frac{M_{t+n}}{M_t}, \frac{S_{t+n} - F_t}{F_t}) \cdot 1 / E(\frac{M_{t+n}}{M_t}). \tag{42}$$

Since the second term of (42) is a positive constant (see (1)), and I have assumed that the pricing kernel is perfectly negatively correlated with demand shocks, the risk premium is proportional to the negative of the covariance of the future spot price and demand shocks. However, I have argued above that the covariance between demand shocks and spot prices is generally positive, decreases in the level of oil capital in the economy, and turns negative for very high levels of capital. Therefore, the oil risk premium is generally positive, is negatively related to capital, and turns negative for very high levels of capital. The top panel of Figure 6 shows that periods of negative risk premiums are quite infrequent in my model (only about 20% of the observations), and that they are not too negative. Indeed, the bottom panel of Figure 6, where I plot the non-parametric relationship between capital and the risk premium, the risk premium is as high as 4% for low levels of capital, but near zero at high levels of capital (negative risk premiums are within the 95% confidence bands only at high levels of capital).

The change in the sign of the risk premium in the model, which is driven by a single shock (to oil demand) and a constant price of risk, is a bit surprising. Generally, a single-factor model with a constant price of risk would imply a risk premium of a constant sign; in my case the risk premium for bearing this demand shock would be $\sigma_M \sigma_\epsilon (1+|\epsilon_t|)$. Despite there being only one shock, this is not really a single-factor model because as discussed above, the demand shock affects prices both directly, and indirectly through an effect on the supply response of firms. The relative magnitudes of the two effects therefore change the sign of the effect of shocks on oil prices, leading to possible changes in sign of the risk premium. I believe this a novel feature of my model.

I end this section with some further evidence that investment by oil firms reduces the effect of demand shocks on future oil price changes. Determining the relative importance of supply and demand shocks on oil prices has been controversial in the literature (see, e.g., Hamilton 2003; Killian 2009). Here, I examine the coefficient of the effect of oil consumption demand shocks on oil prices as identified in a recent paper by Baumeister and Hamilton (2017).[17] This paper uses a Bayesian framework to avoid imposing strong views on the identification of the vector autoregression (VAR) system, which has been the standard methodology in previous research to determine the relative impact of the different shocks. The time series of the moving average of the coefficient along with the moving average of the investment-to-GDP ratio, are in Figure 10. The two variables are negatively correlated. Table 5 shows some simple regressions for these variables. For the full sample, the coefficient of the fit of the regression of the demand shock coefficient on investment is close to -1, and is statistically significant. The R^2 of the regression is about 6%, that is, the variables have a negative correlation of about -0.25. For different

[17] I thank Jim Hamilton for providing the time series of this coefficient.

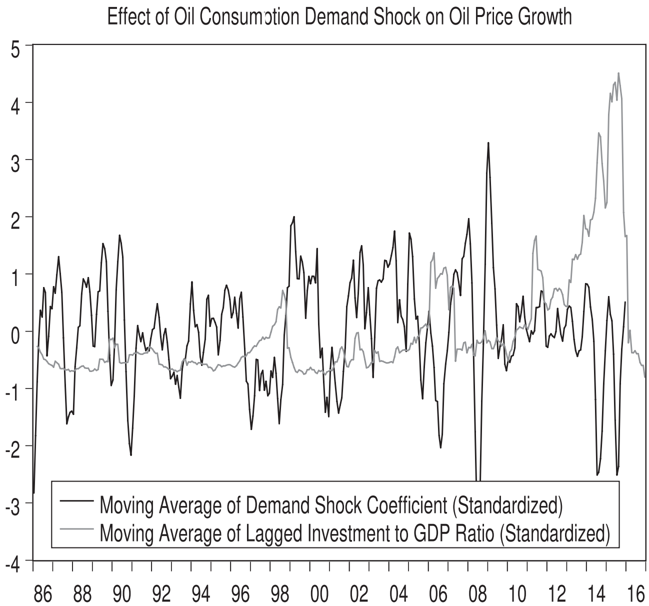

Figure 10
Investment and the effect of oil consumption demand shocks on oil price growth
I report the time series of the standardized variable $\sum_{i=1}^{6}\Gamma(t+i)$, where $\Gamma(t)$ is the VAR coefficient of the effect of oil consumption demand shocks on real oil price growth at date t as shown in figure 10 of Baumeister and Hamilton (2017), and the time series of the standardized variable $\sum_{i=1}^{2}$ Investment/GDP$(t-i)$, where "Investment" stands for the seasonally adjusted capital expenditures of oil and gas firms (SIC codes 1311 and 138).

Table 5
Investment and the effect of oil consumption demand shocks on oil prices?

	α	β	R^2
Full Sample July 1986–December 2016:			
1	1.981 [−0.986]	−0.986 [−2.874]	0.062
Sample from July 1986–December 1999:			
2	5.798 [1.654]	−4.329 [−2.973]	0.1054
Sample excluding 2008–2010:			
3	1.973 [1.008]	−0.998 [−2.909]	0.082
Sample from 2000–2016:			
4	5.052 [1.614]	−1.192 [−2.811]	0.112

I report the coefficients of the fitted regression:

$$\sum_{i=1}^{6}\Gamma(t+i)=\alpha+\beta\sum_{i=1}^{2}\text{Investment/GDP}(t-i)+\epsilon(t),$$

in which $\Gamma(t)$ is the VAR coefficient of the effect of oil consumption demand shocks on real oil price growth at date t as shown in figure 10 of Baumeister and Hamilton (2017), and "Investment" stands for the seasonally adjusted capital expenditures of oil and gas firms (SIC codes 1311, 138). t-statistics are in brackets and are adjusted for heteroscedasticity and autocorrelation.

subsamples of my data, the coefficient remains significant, while the R^2s are higher. The regressions provide further evidence that investment by oil firms reduces the sensitivity of oil price changes to shocks in oil demand, which is a key part of the mechanism on why investment lowers the futures risk premium for oil in my model.

4. Conclusion

In the past 30 years, investment by oil firms and the futures basis have been positively related, while investment has negatively predicted excess returns on oil futures. Oil inventory has also been positively related to the basis, but had had little predictive power for futures' excess returns. In addition, inventory and investment have been negatively related. In this paper, I provide a new equilibrium model of exhaustible resource extraction, inventory accumulation, and investment in exploration and development to understand these facts. I model demand shocks that drive the business-cycle fluctuations in oil prices, and also build in the implications of the declining quality of the total resource base, and firms' decisions on capital accumulation that manage their extraction costs as they extract from increasingly costly fields. Drilled wells produce the resource at a geometrically declining rate; however, by specifying consumers' habit level equaling production from old wells, the futures basis and risk premium are only related to drilling, investment, and inventory. In my model, a steeper futures curve increases the attractiveness of of accumulating inventory in some periods and investment in others. The main reason why investment is negatively related to the risk premium is that it leads to more aggressive drilling by firms and dampens oil price increases from demand shocks, thus lowering the covariance of the oil prices with the pricing kernel in the economy.

My results are not driven by productivity shocks in oil production or in capital accumulation, but instead by endogenously generated changes in extraction costs of oil determined by firms' drilling and capital accumulation decisions. As such, my technological progress specification is novel to the literature on investment and real options.

Appendix A

I obtain historical crude oil futures contracts prices from July 1986 to December 2016, from the *Chicago Mercantile Exchange* (CME). The data series provided summarize the prices from all publicly traded exchanges. I filter the series and use only prices for contracts with positive volume. I obtain the series of constant maturity Treasury yields from the *Federal Reserve Board*. The weak relative basis on 1-year contracts in quarter t is $[e^{-r(t)}F(t)-S(t)]/S(t)$, where $F(t)$ is the 1-year futures prices at the beginning of each quarter and $S(t)$ is the spot price of WTI oil in Cushing, Oklahoma. "Return" is the excess return on a fully collateralized 1-year futures contract defined as $S(t)/F(t-11)-1$, where $F(t)$ is the 1-year futures prices at the beginning of each quarter and $S(t)$ is the spot price (see, e.g., Gorton, Hayashi, and Rouwenhourst 2013) for this definition of excess-returns). The spot price is proxied by the nearby futures price at the beginning of each month as is standard in the finance literature (see, e.g., Fama and French 1987). Full collateralization means that, at the time of entering a futures contract, the buyer has to post collateral equal to the present value of the futures price, which earns the riskless rate.

The explanatory variable "Inventory" stands for the seasonally adjusted total U.S. stock of crude oil and petroleum products (in thousands of barrels) excluding special purpose reserves at the end of each month as provided by the U.S. Energy Information Association (EIA). I normalize this series by U.S. real GDP and detrend the ratio using the Hodrick-Prescott filter. The nondetrended series declines steadily in the first half of my sample for efficiency reasons likely unrelated to my model. "Investment" stands for the seasonally adjusted capital expenditures of oil and gas (SIC codes 1311 and 138) from Compustat. Code 1311 includes firms in "Crude and Natural Gas," and code 138 includes firms in "Oil and Gas Field Services". Firms in code 1311 had 88% of the total capital expenditures over my full sample. In an earlier version of this paper, I reported results with firms only in code 138, and I find similar results here. Indeed, the correlation of capital expenditures of firms in codes 1311 and 138 is 98% in my sample. I also measure the "Capital Stock" of these firms as the "Property, Plant, and Equipment" variable in Compustat. Seasonal adjustment of all variables is performed using the X-12 procedure (used by the U.S. Department of Commerce).

Appendix B

Proof of Proposition 1.

Using the equilibrium stock price at date 1 in (16), I have that the call option value is simply

$$
\begin{aligned}
&C(xg(K_1)|Y_0)\\
&=e^{-r}E^{\mathcal{Q}}\left[\left(\frac{a\,e^{\mu+\sigma\epsilon}+\frac{x_0^e}{\bar{x}}R_0-Z_1e^{-u}}{b+\frac{R_0}{\bar{x}g(K_1)}}-x\,g(K_1)\right)^+\right]\\
&=\frac{e^{-r}}{D}E^{\mathcal{Q}}\left[\left(a\,e^{\mu+\sigma\epsilon}-(Dx\,g(K_1)-\frac{x_0^e}{\bar{x}}R_0+Z_1e^{-u})\right)^+\right]\\
&=\frac{a\,e^{-r}}{D}\left[E[e^{\mu-\sigma_M\sigma+\sigma\epsilon^*}|\mu-\sigma_M\sigma+\sigma\epsilon^*>\log(k)]-k\text{Prob}[\mu-\sigma_M\sigma+\sigma\epsilon^*>\log(k)]\right]\\
&=\frac{a\,e^{-r}}{D}\left[e^{(\mu-\sigma_M\sigma+0.5\sigma^2)}N(-d_1)-kN(-d_2)\right],
\end{aligned}
$$

as stated. I note that in the third line I use the definition of the 'risk-neutral shock' $\epsilon^*=\epsilon+\sigma_M$, while in the fourth line I use the conditional expectation for lognormal variables (see, e.g., proposition 2.29 in Nielsen 1999). ■

Appendix C

I proceed by formulating an "approximate" solution to the Hamilton-Jacobi-Bellman equation in Equation (27) using projection methods (Judd 1999, chapter 11). The value function is denoted as $J(x_t^e, Z_t, K_t, \epsilon_t)$. Because lagged production does not affect the value function as discussed in Section 2, I do not include it as a state variable.

STEP 1. Choice of individual basis functions. I choose the Chebyshev polynomials in each of the four dimensions: the Chebysev polynomials on $[-1,1]$ for the basis for each dimension are given by

$$q_m(x)=\cos(m\cos^{-1}x),$$

for $m=1,2,\cdots$, which satisfy the recursive scheme

$$q_{m+1}(x)=2xq_m(x)-q_{m-1}(x). \tag{C1}$$

These polynomials are restricted for the interval $[a,b]$ using the transformation

$$p_m,(x) = \frac{q_m(\frac{2x-a-b}{b-a})}{||q_m(\frac{2x-a-b}{b-a})||}.$$

I solve the value function on bounded spaces in each dimension: $[0,\bar{x}]\times[0,\bar{Z}]\times[0,\bar{K}]\times[-\bar{\epsilon},\bar{\epsilon}]$. The family $\{p_m(x)\}_{m=1,2,\cdots}$ are orthonormal polynomials over the chosen intervals.

STEP 2. Choose a basis of "complete" polynomials over the space

The basis of degree M over the four dimensions is given by

$$\mathcal{P}_M = \{p_{1,i_1}(X)\cdot p_{2,i_2}(Z)\cdot p_{3,i_3}(K)\cdot p_{4,i_4}(\epsilon)|\sum_{n=1}^{4} i_n \le M, 0\le i_1,\cdots,i_3\}$$

I write the generic element of $\mathcal{P}_M$ as $\phi_m(X,Z,K,\epsilon)$, $m=1,2,\cdots M^c$, where M^c is the length of the complete polynomial basis. The set of complete polynomials for a 4 dimensional problem grows polynomially in 4, as opposed to the tensor product basis that would use every possible product of the Mth degree individual basis functions, and hence would grow at the rate of M^4 (see, e.g., Judd 1999, p. 239). The complete polynomials asymptotically, as M becomes large, provide as good an approximation as the tensor product, but with far fewer elements. Extending the L^2 norm over the four-dimensional space as the four-fold integral, it can be verified that the basis of complete polynomials is orthonormal on the bounded Cartesian product space.

STEP 3 Let $b^{(n)}$ be the nth guess on the coefficients of the polynomial, that is, $J^{(n)}(x_t^e, Z_t, K_t, \epsilon_t) = \sum_{m=1}^{M^c} b^{(n)}\cdot\phi_m(X,Z,K,\epsilon)$. Then I solve (27) for the $n+1$th guess as $J^{(n+1)}(x_t^e, Z_t, K_t, \epsilon_t)$, using the first-order conditions in Equations (28)–(32). Note that I are able to take partial derivatives of the nth guess value function, which is just a polynomial sum. The first-order conditions are solved on a discrete grid of values for inventory, investment, and the extensive margin.

STEP 4 I now appeal to the Chebyshev interpolation theorem (see Judd 1999) to find an approximate solution to the Bellman equation. Denote $W=(x^e,Z,K,\epsilon)$. The approximation is made by evaluating the $J^{(n+1)}(W)$ at the Chebyshev zeros in the Cartesian space, given the coefficients $b^{(n)}$. Each interpolation point therefore provides us a linear equation in the coefficients $(b_m)^{(n+1)}$. With M^I interpolation points, I have an overidentified system of equations in M^c unknown coefficients, and I solve for $(b_m)^{(n+1)}$ using linear regression. I then repeat until convergence.

References

Abel, A. B. 1990. Asset prices under habit formation and catching up with the Joneses. *American Economic Review, Papers and Proceedings* 80:38–42.

Anderson, S. T., R. Kellogg, and S. W. Salant. 2018. Hotelling under pressure. *Journal of Political Economy* Advance access published January 26, 2018, 10.1086/697203.

Arrow, K. J. 1962. The economic implications of learning by doing. *Review of Economic Studies* 29:155–73.

Baker, S. D., and B. R. Routledge. 2012. The price of oil risk. Working Paper, Carnegie Mellon University.

Baumeister, C. J., and J. D. Hamilton. 2017. Structural interpretation of vector autoregressions with incomplete identification: revisiting the role of oil supply and demand shocks. Discussion Paper National Bureau of Economic Research.

Belo, F., X. Lin, and S. Bazdresch. 2014. Labor hiring, investment, and stock return predictability in the cross section. *Journal of Political Economy* 122:623–89.

Bloom, N. 2009. The impact of uncertainty shocks. *Econometrica* 77:623–89.

Campbell, H. F. 1980. The effect of capital intensity on the optimal rate of extraction of a mineral deposit. *Canadian Journal of Economics,* 13:349–56.

Campbell, J.Y., and J. H. Cochrane. 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107:205–51.

Campbell, J.Y., A. W. Lo, and C. A. MacKinlay. 1995. *The econometrics of financial markets*. Princeton NJ: Princeton University Press.

Carlson, M., Z. Khokher, and S. Titman. 2007. Equilibrium exhaustible resource price dynamics. *Journal of Finance,* 62:1663–703.

Casassus, J. P., P. Collin-Dufresne, and B. R. Routledge. 2008. Equilibrium commodity prices with irreversible investment and non-linear technologies. Working Paper, Carnegie Mellon University.

Chen, F. and S. C. Linn. 2017. Investment and operating choice: oil and natural gas futures prices and drilling activity. *Energy Economics* 66:54–68.

Chetty, R. and A. Szeidl. 2007. Consumption commitments and risk preferences. *Quarterly Journal of Economics* 122:831–77.

Cochrane, J. H. 2001. *Asset pricing*. Princeton, NJ: Princeton University Press.

Covert, T. R. 2015. Experiential and social learning in firms: the case of hydraulic fracturing in the Bakken shale. Working paper, University of Chicago.

David, A. 2008. Heterogeneous beliefs, speculation, and the equity premium. *Journal of Finance* 63:41–84.

David, A., and P. Veronesi. 2013. What ties return volatilities to price valuations and fundamentals. *Journal of Political Economy* 121:643–81.

Deaton, A., and G. Laroque. 1992. On the behavior of commodity prices. *Review of Economic Studies* 59:1–23.

Erb, C. B. and C. R. Harvey. 2006. The strategic and tactical value of commodity futures. *Financial Analysts Journal* 62:47–68.

Fama, E. F. and K. R. French. 1987. Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage. *Journal of Business* 60:55–73.

Ghoddusi, H. 2010. Dynamic investment in extraction capacity of exhaustible resources. *Scottish Journal of Political Economy* 57:359–73.

Gorton, G. B., and K. G. Rouwenhourst. 2006. Facts and fantasies about commodities futures. *Financial Analysts Journal* 62:47–68.

Gorton, G. B., F. Hayashi, and K. G. Rouwenhourst. 2013. The fundamental of commodities futures returns. *Review of Finance* 17:17–35.

Gustafson, R. 1958. Carryover levels of grains: A method of determining amounts that are optimal under specified conditions. Discussion Paper, USDA Technical Bulletin.

Hamilton, J. D. 2003. What is an oil shock. *Journal of Econometrics* 113:363–98.

Hamilton, J. D. 2012. Oil prices, exhaustible resources and economic growth. In *Handbook of energy and climate change,* pp. 1–44. Ed. Roger Fouquet. Cheltenham, UK: Edward Elgar Publishing.

Hayashi, F. 1982. Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica* 50:213–24.

Hotelling, H. 1931. The economics of exhaustible resources. *Journal of Political Economy* 37:137–75.

Judd, K. L. 1999. *Numerical methods in economics.* Cambridge, Massachusetts: MIT Press.

Kaldor, L. 1939. Speculation and economic stability. *Review of Economic Studies* 7:1–27.

Keynes, J. M. 1930. *A treatise on money, Volume II* London, UK: Macmillan and Co., Limited.

Killian, L. 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99:1053–69.

Kogan, L., D. Livdan, and A. Yaron. 2009. Oil futures prices in a production economy with investment constraints. *Journal of Finance* 64:1345–75.

Koijen, R. S. J., T. Moskowitz, L. H. Pederson, and E. B. Vrugt. 2013. Carry. Discussion Paper, National Bureau of Economic Research.

Litzenberger, R. H., and N. Rabinowitz. 1995. Backwardation in oil futures markets: theory and empirical evidence. *Journal of Finance* 50:1517–45.

Maugeri, L. 2012. Oil: the next revolution. The unprecedented upsurge of oil production capacity and what it means for the world. Geopolitics of Energy Project 2012–10, Harvard Kennedy School: Belfer Center for Science and International Affairs

Nielsen, L. T. 1999. *Pricing and hedging of derivative securities.* Oxford, UK: Oxford University Press.

Pindyck, R. S. 1980. Uncertainty and exhaustible resource markets. *Journal of Political Economy* 88:1203–25.

———. 1994. Inventories and short-run dynamics of commodity prices. *Rand Journal of Economics* 25:141–59.

Pirrong, C. 2012. *Commodity price dynamics.* Cambridge, UK: Cambridge University Press.

Ravn, M., S. Schmitt-Grohe, and M. Uribe. 2006. Deep habits. *Review of Economic Studies* 73:195–218.

Ready, R. 2018. Oil consumption, economic growth, and oil futures: The impact of long-run oil supply uncertainty on asset prices. *Journal of Monetary Economics* 94:1–26.

Routledge, B., D. Seppi, and C. Spatt. 2000. Equilibrium forward curves. *Journal of Finance* 55:1297–338.

Scheinkman, J., and J. Schechtman, 1983. A simple competitive model with production and storage. *Review of Economic Studies* 50:427–41.

Tauchen, G., and R. Hussey. 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59:371–96.

Weinstein, M. C. and R. J. Zeckhauser. 1975. The optimal consumption of depletable natural resources. *Quarterly Journal of Economics* 92:371–92.

Williams, J., and B. Wright. 1991. Storage and commodity markets. Cambridge, Massachusetts. Cambridge University Press.

Working, H. 1948. Theory of inverse carrying charge in futures markets. *Journal of Farm Economics* 30:1–28.

Yang, F. 2013. Investment shocks and the commodity basis spread. *Journal of Financial Economics* 110:164–84.