PEPPERDINE | GRAZIADIO
BUSINESS SCHOOL

Working Paper Series
Center for Applied Research

Risk Premia in Commodity Price Forecasts and their Impact on Valuation

Dr. Joe Hahn, Dr. James DiLellio, Dr. Jim Dyer

February, 2018

Paper No. GCAR-A004
Category: Research Article

Risk Premia in Commodity Price Forecasts and their Impact on Valuation

Abstract

Commodity price driven valuation models require a stochastic price input if the value of managerial flexibility, such as the option to defer investment until the optimal time and the option to abandon a project, is to be estimated. The risk-neutral version of the stochastic price model is typically used in academic work; however, risk-adjusted models of the expected spot price are often used in practice. These two approaches are connected by a risk premium which is unfortunately often difficult to estimate. In this work, we use natural gas futures prices in a Kalman filter approach with maximum likelihood estimation to parameterize the Schwartz and Smith (2000) stochastic price model, and then apply an asset pricing model to address the large uncertainty of the risk premia parameter estimates. To evaluate the impact of the risk premia and other parameters in the two-factor price model on project valuation, we apply the price model to a prototypical shale gas investment, both for a base reference case as well as for cases where there are real options to optimally time decisions to invest or to abandon the project. Using this approach, we are able to determine the implied risk-adjusted discount rate that would be used with the spot price forecast, given the two-factor model risk premia, and we also discuss the impact of the risk premia on project value relative to other model parameters.

1. Introduction

There are several approaches to developing long-term forecasts for commodity prices, including many types of econometric models, equilibrium models, and expert survey forecasts, as well as methods that utilize empirical data from commodities markets to calibrate commonly-used stochastic process models. Schwartz (1997), Schwartz and Smith (2000), Manoliu and Tompaidis (2002), and others describe how the parameters for these types of process models can be obtained with the Kalman filter and maximum likelihood estimation, and evaluate the performance of these models for capturing the dynamics of futures prices. In this paper, we use this approach to calibrate the Schwartz and Smith (2000) two-factor stochastic process, which involves estimating the values of seven model parameters, including short- and long-term risk premia. The relationship between futures and spot prices in this approach is established within the context of a risk-neutral valuation framework, where futures prices are equal to the expected future spot price under a risk-neutral stochastic process (Duffie, 1992). The implication of this relationship is that futures price data can be used to directly calibrate the price model in its risk-neutral form, and then a risk adjusted (expected future spot price) model can be calculated using the estimated risk premia. Unfortunately, as shown by Cortazar et al. (2015) and others, the risk premia are often difficult to estimate with the Kalman filter approach. To address this problem, Cortazar, et al. (2015) suggest exogenously estimating these parameters, and restricting the Kaman filter approach to the remaining five model parameters. In this work, we are interested in distinguishing between the parameter estimates from these two approaches and investigating how valuation is affected.

Given a commodity price model, discounted cash flow methods (DCF) are commonly used in practice for valuation and decision-making regarding investments in energy producing assets.

In the DCF approach, the deterministic expected cash flows can be estimated using the risk-neutral version of the forecast, in which case the risk-premia are irrelevant, and the cash flows can simply be discounted at the risk-free rate of interest. However, many practitioners prefer to estimate the actual cash flows instead based on the expected commodity spot price, and to apply a risk-adjusted discount rate. In this case, the risk premia are relevant for developing the price forecast, and since the net present value should be unique, the correct risk-adjusted discount rate that accounts for the forecast risk premia can be inferred.

The DCF approach can be augmented by recognizing the analogy between financial options and real options – project decisions that can be made as the uncertainties associated with real assets evolve over time. By this analogy, the methods that were developed by Black and Scholes (1973), Merton (1973), Black (1976) and others to price financial options can also be used to value these real options. This approach has the advantage of including the value of managerial flexibility, which is frequently not captured by the traditional DCF approach. A risk-neutral forecast of the relevant underlying uncertainty (e.g., commodity price) is typically used in this approach, but the forecast risk premia are indirectly involved because estimates of the other parameters in the forecast model can be affected by the method used to estimate the risk premia.

The application of real option valuation (ROV) to oil and gas projects has been used in practice for some time, as these types of projects naturally have attributes that imply potential value associated with optimal management decision making under uncertainty. These attributes include large initial investments, a market-traded underlying commodity, private reservoir uncertainty and strategic flexibility over choices related to timing, drilling and production technologies and capacity. Research in this area of application includes Siegel, Smith and Paddock (1988) who assess investing in offshore petroleum leases and compare ROV with both NPV approaches and observed bid prices; Cortazar and Schwartz (1998) who find the optimal timing of investing in a field with a set oil rate that declines exponentially and with varying but deterministic operating costs; McCardle and Smith (1998) who consider the timing of investment, the option to abandon, and the option to tie in surrounding fields where both prices and production rates are modeled as stochastic processes; Ekern (1988) who values the development of satellite fields and adding incremental capacity; Lund (1999) who considers an offshore field development by using a case from the Heidrun field in the North Sea, taking into account the uncertainty regarding both reservoir size and well rates, in addition to the oil price; Dias, Lazo, Pacheco and Vellasco (2003) who find an optimal development strategy for oil fields when considering three mutually exclusive alternatives; Chorn and Shokhor (2006) who value investment opportunities related to petroleum exploration; and Jafarizadeh and Bratvold (2012) who value an oilfield abandonment option under a two-factor commodity price model. Additional work related to the petroleum industry can be found in a review compiled by Dias (2004).

In this paper, we have selected an unconventional shale gas project as the setting for studying the effect of risk premia on valuation. The renewed and enhanced application of horizontal drilling and hydraulic fracturing technologies to shale reservoirs has dramatically changed the US domestic natural gas production output during the past decade. Prior to that time, it was expected that the US would need to import large volumes of liquefied natural gas (LNG) to make up for an anticipated shortfall in domestic production. The International Energy Agency (IEA) now projects that the United States will be a net exporter of natural gas by as soon as 2020. The gas price environment, through its impact on investment economics, will undoubtedly be a

key factor in determining whether such projections are accurate or not. Figures 1 and 2 show the historical and future projected impact of shale gas production and the resulting realized price, respectively. The fundamental economic relationship between supply and demand seems to hold, with the increase in production coinciding with a decrease in the price index. The relatively new significance of this source of production and its unique characteristics, such as its typical production decline curve, ratio of drilling to completion costs, and risk profile, make it a useful application for continued research. Here, we are interested in exploring this relationship at the project investment level, with the objective of identifying the impact of the current and future forecasted price dynamics.

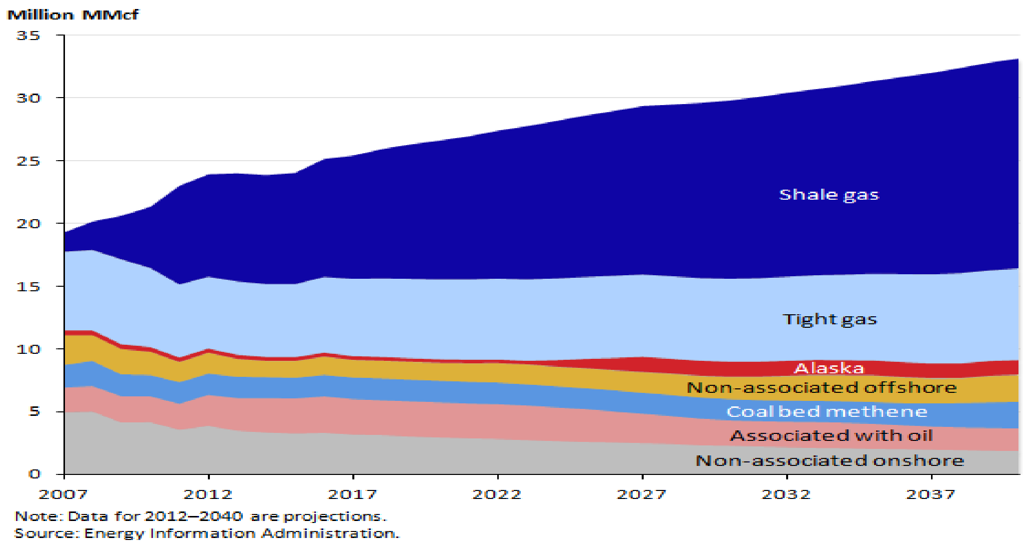

Figure 1 – US Domestic Natural Gas Production 2007-2040

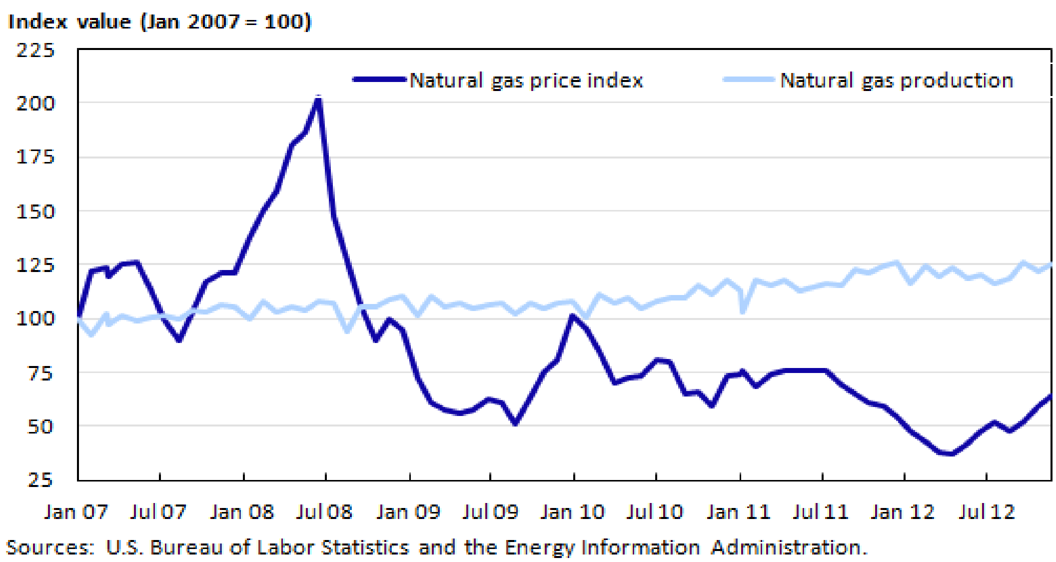

Figure 2 – Natural Gas Price and Production Indices 2007-2012

The prototypical shale gas investment valuation problem we use is a base case model developed by Lake, et al.(2013) for a shale gas well found in the Haynesville shale region of North Central Louisiana (Haynesville #1). The model includes a detailed production decline curve forecast, drilling and completion cost data, annual operating costs data, and a deterministic forecast for the expected price of natural gas. We augment this model with a stochastic forecast for natural gas prices and then account for the decisions about when to drill the well, given an assumed five-year lease, and, after production has commenced, when to divest in the well, given an assumed market value multiple of reserves in two separate analyses.

Our research questions here include: 1) what are the differences in the two-factor natural gas price model parameters when the risk premia are estimated with the Kalman Filter and when the premia are estimated exogenously using an asset pricing model, 2) what are the implied differences in the risks associated with the cash flows (e.g., through implied discount rate) under different estimates of the risk premia, 3) what is the sensitivity of the project value (rather than the price forecast itself) to different parameters in the stochastic price model, and 4) what is the impact on base case project value, as well as on the optimal timing of investment and or abandonment given a market-based price forecast. Our goal is to address each of these questions using both the risk-neutral approach to valuation and the risk-adjusted approach with an implied discount rate. The answers to these questions are potentially interesting for both practitioners assessing shale gas field strategies as well as analysts and researchers forecasting future petroleum supply.

This paper is organized as follows: Section 2 includes a brief overview of the Lake et al. (2013) shale gas well cash flow model and base case assumptions. Section 3 details the results of a study to parameterize the Schwartz and Smith (2000) two-factor model using a recent natural gas futures data set in a Kalman filter – maximum likelihood estimation approach and the development of a forecast over the producing horizon for the shale gas well model. In Section 4, we first present the results of a deterministic base case valuation of the shale gas investment, and then value two associated real options and conduct a sensitivity analysis to input parameters. In Section 5, we discuss the conclusions and implications of the integrated forecasting-option valuation model.

2. Shale Gas Project Cash Flow Model

Our valuation model is developed from an example that was originally discussed in Lake et al. (2013), with a few additional simplifying assumptions. Beginning in year 1, annual cash flows in the t^{th} year, CF_t , are calculated as

$$CF_t = \left(\text{Revenue}_t - \text{Opex}_t - \text{Severance Tax}_t - \text{Depletion}_t\right) \times \left(1 - T_C\right) + \text{Depletion}_t \qquad (1)$$

where

$$\text{Revenue}_t = \text{Gross Production}_t \times \text{Net Revenue Interest} \times \text{Gas Price}_t ,$$

$$\text{Opex}_t = \text{Revenue}_t \times \text{Operating Cost Rate} \ ,$$

$$\text{Severance Tax}_t = \text{Revenue}_t \times \text{Severance Tax Rate} \text{, and}$$

$$\text{Depletion}_t = \text{Gross Production}_t \times \text{Net Revenue Interest} \times \text{Depletion Allowance per MCF} .$$

The decline curve used to model gross production over time corresponds to a specific shale gas well found in North Central Louisiana, and is shown in Figure 3; it decays rapidly over a T = 38 year time period. Gross annual gas production was reported in Lake et al (2013) and it can be modeled as a harmonic decline ($q_t = q_i/(1 + Dt)$) with q_i = 2.405 Bcf and D = 2.0395, which is very steep to properly model a well that employs hydraulic fracturing. While this model depicts the production forecast for a specific shale gas well, the rapidly decaying decline curve is representative of a typical shale gas well. The net gas production is assumed to be a constant fraction of gross production; in this case a net revenue interest of 15.75% applies.

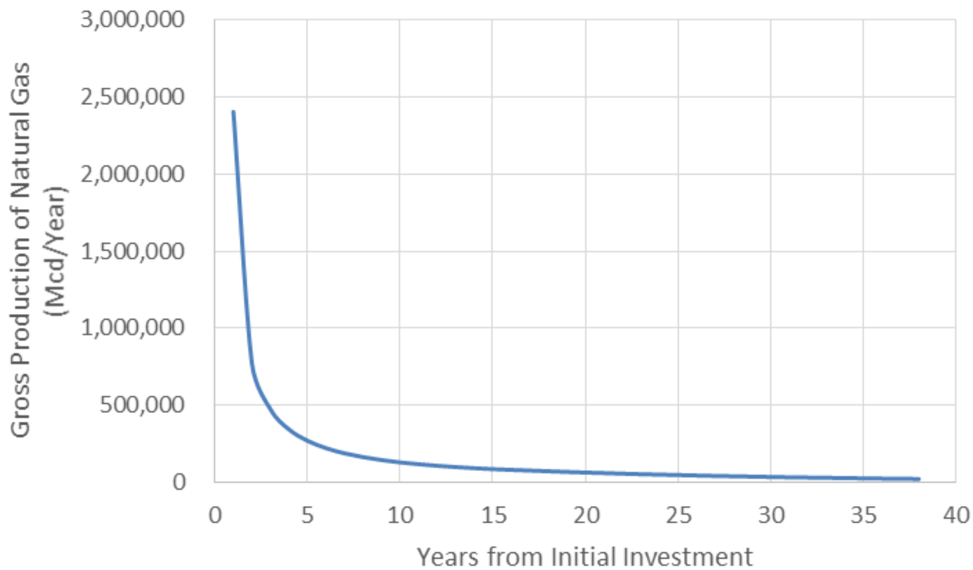

Figure 3 - Gross Production of a Prototypical Shale Gas Well in North Central Louisiana

The expenses that appear in Eq. (1) fall into three categories. Total net operating expenses correspond to a 2.9% rate applied annually to total net revenue. Similarly, severance and Ad Valorem taxes are each assumed to be 8.4% of total net revenue. Estimating these expenses as a fraction of revenue is appropriate because they typically will scale with the increases/decreases in the gross profit from the well. The depletion allowance in Eq. (1) is a non-cash expense similar to depreciation, and is calculated using:

$$\text{Depletion Allowance per MCF} = \frac{\text{Net Revenue Interest} \times \text{Total Capital Expenditures}}{\sum_{t=1}^{T} \text{Gross Production}_t \times \text{Net Revenue Interest}} \qquad (2)$$

In Eq. (2), the total capital expenditures consist of drilling and completion costs, assumed to be \$5.1M and \$3.4M, respectively in this example. Lastly, the initial investment in Year 0 is estimated to be \$1,912.5M, based on a 22.5% share of total capital expenditures.

Assuming a fixed \$4.50/Mcf natural gas price forecast and using a 30% corporate tax rate and a 10% cost of capital, we obtain an NPV of just over \$0.309M, as reported in Lake et al (2013). Although common in practice, these assumptions imply that the cost of capital appropriately discounts the future cash flows associated with the fixed price forecast. Unfortunately, this

approach may not yield correct market-based valuations. In Section 5, for the purpose of real option valuation, we will first use a risk-neutral pricing framework in which we use the risk-free rate to discount cash flows resulting from a risk-neutral price forecast, followed by a risk-adjusted expected spot price framework in which we will use a risk-adjusted discount rate. Accordingly, the following section applies the Schwartz and Smith (2000) modeling framework to developing risk neutral and expected spot natural gas price forecasts to be used in place of the simple $4.50 flat forecast.

3. Two-factor stochastic process model parameterization and price forecast

For the Schwartz and Smith (2000) two factor model, the natural log of the natural gas price S_t at any point in time *t* is expressed as the sum of short-term deviations χ_t and a long-term equilibrium level ξ_t, such that $\ln(S_t) = \chi_t + \xi_t$. Assuming the short-term deviations follow a mean-reverting process and the long-term equilibrium follows geometric Brownian motion, the two factor process can be summarized as:

$$d\chi_t = \kappa(0 - \chi_t)dt + \sigma_\chi dz_\chi \tag{3}$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \tag{4}$$

where κ is the reversion rate of the deviation to a mean value of zero, σ_χ is the volatility of the short term deviation, μ_ξ is the drift rate of the long term equilibrium level, σ_ξ is the volatility of the long term equilibrium, and $\rho_{\chi\xi}$ is the correlation between the random increments of the two processes, dz_χ and dz_ξ.

The calibration of the two factor process utilizes futures prices, and because futures prices are equal to the expected future spot price under a risk-neutral stochastic process (Duffie, 1992), it is necessary to develop a risk neutral version of the two factor process. This version can also then be used to value investments that derive their value from the underlying commodity (i.e., natural gas in our case), including real options, without the requirement of estimating a risk-adjusted discount rate. The modifications to Equations (3) and (4) for a risk neutral process involve relatively straightforward adjustments to the drift rates of the processes for χ_t and ξ_t using risk premiums λ_χ and λ_ξ, respectively:

$$d\chi_t = (0 - \kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi^* \tag{5}$$

$$d\xi_t = \mu_\xi^* dt + \sigma_\xi dz_\xi^* \tag{6}$$

where the parameters are defined the same as in Equations (3) and (4), except μ_ξ^* is the risk neutral drift rate of the long term equilibrium level, calculated as $\mu_\xi^* = \mu_\xi - \lambda_\xi$ and the random increments of the two processes, dz_χ^* and dz_ξ^* are increments of the risk neutral process.

The expectation and variance of the risk neutral process, as shown in Schwartz and Smith (2000) are:

$$E[\ln(S_t)] = e^{-\kappa t}\chi_0 + \xi_0 - (1 - e^{-\kappa t})\frac{\lambda_\chi}{\kappa} + \mu_\xi^* t \tag{7}$$

$$Var[\ln(S_t)] = (1 - e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 t + 2(1 - e^{-\kappa t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \tag{8}$$

Here, S_t is lognormally distributed, so that $\ln(S_t/S_0)$, which are the returns from period 0 to period t, is normally distributed. Under risk-neutral valuation, the futures prices will equal the expected spot prices (Black, 1976). Therefore the expectation and variance in (7) and (8) can be used to derive the following expression for the futures prices:

$$\ln(F_{T,0}) = e^{-\kappa t}\chi_0 + \xi_0 + A(T),$$

where $F_{T,0}$ is the current (time 0) market price for a futures contract with maturity at time T, and

$$A(T) = \mu_\xi^* T - (1 - e^{-\kappa T})\frac{\lambda_\chi}{\kappa} + \frac{1}{2}\left[(1 - e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2(1 - e^{-\kappa T})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right].$$

Thus, there are seven parameters required to specify the model in its basic and risk neutral versions. We estimate the values of these parameters using a Kalman filter with maximum likelihood estimation. The Kalman filter is a recursive procedure for optimally estimating unobserved state variables based on observations that influence these state variables (Kalman, 1960). In this case, the Kalman filter can be applied to estimate the unobservable state variables χ_t and ξ_t, making it possible to calculate the likelihood of a set of observations given a particular set of parameter values. By varying the parameter values and re-running the Kalman filter, the value of the parameters that maximize the likelihood function can be identified. A detailed description of this technique can be found in Harvey (1989).

For the Kalman filter, the stochastic process must be represented in a state space formulation. This representation consists of a transition equation to describe the evolution of the state variables over time and a measurement equation to relate the state variables to the observable data. Schwartz and Smith (2000) specify the *transition equation* for the two factor model as:

$$x_t = c + Gx_{t-1} + \omega_t, \qquad t = 1, \dots, n_T$$

Where n_T is the number of time periods,

$x_t = \begin{bmatrix}\chi_t \\ \xi_t\end{bmatrix}$ is a 2x1 vector of state variables;

$c = \begin{bmatrix}0 \\ \mu_\xi \Delta t\end{bmatrix}$ is a 2x1 vector and Δt is the length of time steps;

$G = \begin{bmatrix}e^{-\kappa t} & 0 \\ 0 & 1\end{bmatrix}$ is a 2x2 transition matrix;

ω_t is a 2x1 vector of serially uncorrelated normally-distributed disturbances with

$E[\omega_t] = 0$, and

$$Var[\omega_t] = Cov[\chi_{\Delta t}, \xi_{\Delta t}] = \begin{bmatrix} (1-e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} & (1-e^{-2\kappa t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1-e^{-2\kappa t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2 t \end{bmatrix} .$$

The corresponding *measurement equation* is:

$$y_t = d_t + F_t' x_t + v_t, \quad t = 1, \dots, n_T ,$$

where $y_t = \begin{bmatrix} \ln(F_{T,1}) \\ \vdots \\ \ln(F_{T,n}) \end{bmatrix}$ is a n x 1 vector of observed (log) futures prices for the n maturities $T_1, T_2, \dots, T_n$, $d_t = \begin{bmatrix} A(T_1) \\ \vdots \\ A(T_n) \end{bmatrix}$ is a n x 1 vector, $F_t' = \begin{bmatrix} e^{-\kappa T_1} & 1 \\ \vdots & \vdots \\ e^{-\kappa T_n} & 1 \end{bmatrix}$ is a n x 2 matrix, and v_t is a n x 1 vector of serially uncorrelated normally-distributed disturbances (measurement errors) with $E[v_t] = 0$ and $Cov[v_t] = V$.

With a state space formulation and a set of historically observed futures prices for different maturities, the Kalman filter runs recursively beginning with a prior distribution of the initial values of the state variables (χ_0, ξ_0). In addition, the terms in the covariance matrix (V) for the measurement errors for each of the futures contract maturities in the data must also be estimated. The measurement errors can be simplified by making the common assumption that they are uncorrelated with each other, so that V is a diagonal matrix as in Schwartz (1997) and Schwartz and Smith (2000). The objective is to maximize the log-likelihood function for a joint normal distribution.

Our sample of futures data consisted of 969 weekly observations of futures prices at maturities of 1, 3, 6, 12, 18, 24 and 36 months, as shown in Figure 4. The time period for these observations is from the week of June 6, 1997, the date from which all seven contracts above were first continuously traded, to the week of January 1, 2016. We also worked with a subset of this data, which included 366 weekly observations, from the week of January 2, 2009 through the same end date as the full data set. This subset was selected to correspond with the estimated date of the effective beginning of the shale gas era, because it coincides with the approximate date when natural gas produced by hydraulic fracturing started to significantly influence market prices.

The data from both time horizons (June 1997 start and January 2009 start) was analyzed using the Kalman filter to maximize the likelihood function and to produce the parameter estimates reported in Table 1. To characterize the uncertainty in the parameter estimates, the standard errors are also shown in Table 1.

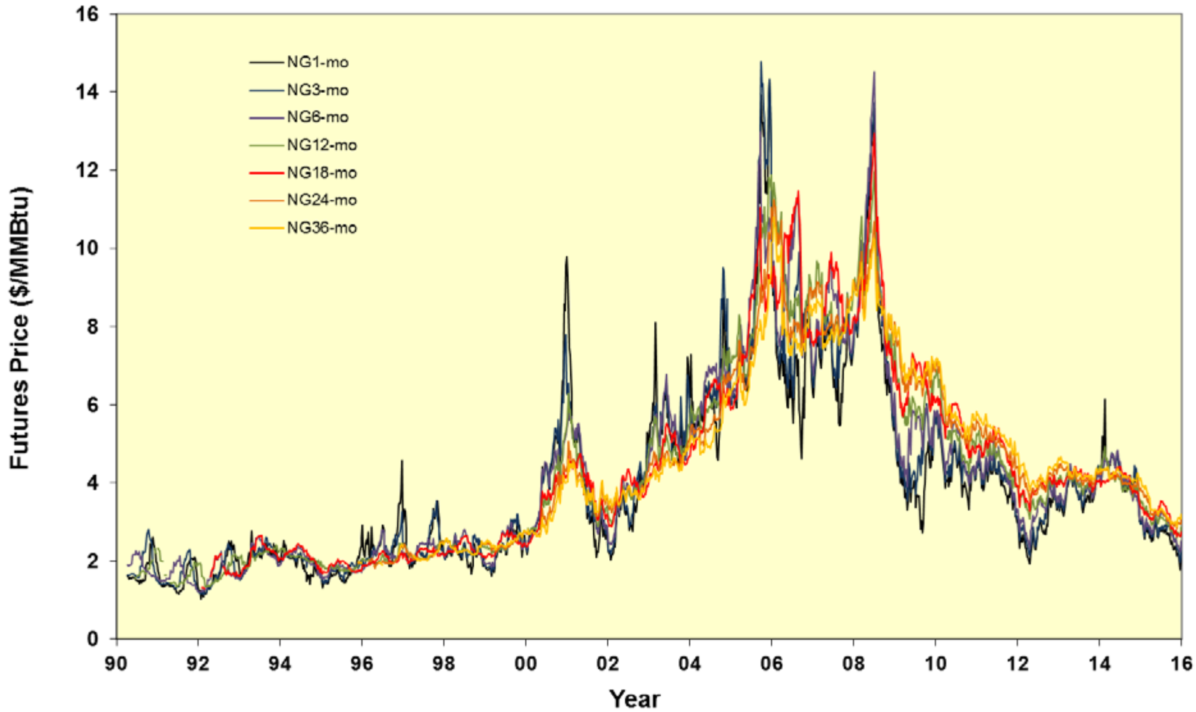

Figure 4 – Natural gas futures price data

Several observations regarding the risk premia in the two-factor model can be made from reviewing Table 1. First, the long-term equilibrium drift rate term μ_ξ is negative for both data sets, especially for the shale gas era data set. As noted by Schwartz and Smith (2000) and others, a negative estimate for this drift rate may not be realistic, given that even a small inflation rate over the forecast period should generate a slightly upward-trending forecast. In addition, the standard errors of these estimates are relatively large, indicating that these terms were not estimated very precisely.

	1997-2016 Data Set		2009-2016 Data Set	
	Parameter Estimate	SE(Est.)	Parameter Estimate	SE(Est.)
μ_ξ	-0.0465	0.0519	-0.2060	0.0605
κ	0.9451	0.0197	1.3998	0.0270
λ_χ	0.4794	0.0480	-0.7452	0.0489
σ_χ	0.4961	0.0192	0.3770	0.0201
σ_ξ	0.2223	0.0058	0.1628	0.0062
$\rho_{\xi\chi}$	-0.2534	0.0480	0.3320	0.0672
μ_ξ^*	-0.0380	0.0019	0.0159	0.0014
$\lambda_\xi = \mu_\xi - \mu_\xi^*$	-0.0085		-0.2219	

Table 1 - Parameter estimates and standard errors – unrestricted case

This limited observability in the drift rate of the long term equilibrium level has also been documented by others, including Schwartz and Smith (2000). Although the risk neutral long-term drift rate term μ_ξ^* is more accurately estimated in both data sets, the calculated values for λ_ξ, shown in the bottom row of Table 1 cannot be precisely estimated due to the uncertainty in μ_ξ. The short term risk premium λ_χ estimates from both data sets also had relatively high standard errors.

However, the larger issue for this parameter was the variability of the estimates from the two data sets, and apparent sensitivity to the data set time horizon. Other researchers have noted similar issues with estimating this risk premium parameter (see for example Cortazar, et al., 2015). The implication of these issues is that, while it would be possible to parameterize a risk neutral price model which does not depend on the estimated risk premia, it would not be possible to specify a reliable model of the expected spot price without better estimates of λ_χ and λ_ξ.

It may appear that the risk premia, λ_χ and λ_ξ, are required to parameterize the risk neutral version of the two factor model because Equation 8a of Schwartz and Smith (2000) includes both the risk neutral drift rate of the equilibrium level, which is connected to the true drift rate of the equilibrium level by the long-term risk premium ($\mu_\xi^* = \mu_\xi - \lambda_\xi$), as well as the risk premium for the short term deviation, λ_χ. However, neither of the risk premia has any effect in the risk-neutral model. The risk-neutral drift rate for the long term equilibrium level μ_ξ^* is directly estimated by the Kalman Filter, therefore λ_ξ does not directly or indirectly factor into their Eq. 8a. While λ_χ does appear in that equation, any changes to it must be accompanied by adjustments to the values of the two factors (χ_0 and ξ_0) and there is no net effect on the risk-neutral distribution of prices. This relationship is discussed in detail in Section 6.1, p. 906 of Schwartz and Smith (2000).

To improve the estimates of λ_χ and λ_ξ, Cortazar et al. (2015) suggest the use of one or more of three different approaches to exogenously estimate these risk premia. The three approaches are 1) using a capital asset pricing model (CAPM)-like asset pricing model to estimate the risk premia, 2) setting the risk premia to zero and 3) using expert opinion of the values of the risk premia. In this work, we used the asset pricing model approach, as it is most likely to align with our goal of obtaining market valuations.

The asset pricing model approach is analogous to estimation of β for the CAPM risk premium for a stock. In this case, however, there are seven assets, which are the seven different maturity futures contracts. It is the different maturities that facilitate estimation of the short and long term risk premia. To calculate the seven beta coefficients, we regressed the weekly returns for each futures contract against a market proxy, the Vanguard 500 Index (VFINX). The beta coefficients were then multiplied by the average market premium since 2000 (4%, from Graham and Harvey, 2012) to produce the expected return for each contract, with the results for both of our data sets shown below in Table 2.

	1997-2016 Data Set			2009-2016 Data Set		
Contract Maturity (days)	β	t-statistic	$E[R_T]$	β	t-statistic	$E[R_T]$
30	0.1277	1.11	0.51%	0.4165	2.32	1.67%
90	0.1167	1.21	0.47%	0.3696	2.49	1.48%
180	0.1564	2.14	0.63%	0.2841	2.26	1.14%
360	0.1272	2.26	0.51%	0.2769	3.07	1.11%
540	0.1397	2.79	0.56%	0.2149	2.85	0.86%
720	0.1225	2.61	0.49%	0.2207	3.39	0.88%
1080	0.1241	2.72	0.50%	0.2234	3.74	0.89%

Table 2 – Beta coefficients and asset pricing model expected returns

The vectors of expected futures contract returns from Table 2 can be equated with the expected contract returns from the Schwartz and Smith (2000) model, which was shown by Cortazar et al. (2015) to be:

$$E_0\left[r_{T,\Delta t}\right] = exp\left(\lambda_\xi \Delta t - e^{-\kappa T}(1 - e^{\kappa \Delta t})\frac{\lambda_\chi}{\kappa}\right) - 1 \qquad (9)$$

This yields a system of seven equations with three unknowns; κ, λ_χ and λ_ξ. We used the κ estimates from Table 1 for each data set and solved for the risk premia numerically with Excel® Solver by minimizing the sum of the squared differences between the expected futures return for each contract in Table 2 and the calculated expected returns from Equation (9). The results of this analysis are shown below in Table 3, again, for both data sets. Relative to the estimates in Table 1, these estimates show more reasonable values for the risk premium for the long term equilibrium level λ_ξ, and that the risk premium for the short term deviation λ_χ is less sensitive to the data time horizon.

With the asset pricing model estimates for the risk premia determined, a restricted case can then be run where only the remaining five parameters are estimated by the Kalman filter, and the two risk premia are entered as deterministic model inputs. Our results for the restricted case for both data sets, which are summarized in Table 4, can then be compared against the unrestricted case results in Table 1.

	1997-2016 Data Set			2009-2016 Data Set		
Contract Maturity (days)	APM $E[R_T]$	Eq. (9) Calculated $E[R_T]$	SSE	APM $E[R_T]$	Eq. (9) Calculated $E[R_T]$	SSE
30	0.0051	0.0053	2.72E-08	0.0167	0.0162	1.93E-07
90	0.0047	0.0053	3.48E-07	0.0148	0.0145	5.59E-08
180	0.0063	0.0052	1.03E-06	0.0114	0.0127	1.71E-06
360	0.0051	0.0052	1.64E-08	0.0111	0.0104	4.30E-07
540	0.0056	0.0052	1.48E-07	0.0086	0.0093	4.98E-07
720	0.0049	0.0052	8.72E-08	0.0088	0.0087	6.45E-09
1080	0.0050	0.0052	4.84E-08	0.0089	0.0083	3.60E-07
		Sum =	1.70E-06		Sum =	3.25E-06
	λ_χ	**0.01%**		λ_χ	**0.41%**	
	λ_ξ	**0.52%**		λ_ξ	**0.82%**	

Table 3 – Risk premia estimates from the asset pricing model approach

	1997-2016 Data Set		2009-2016 Data Set	
	Parameter Estimate	**SE(Est.)**	**Parameter Estimate**	**SE(Est.)**
μ_ξ	-0.0372	n/a	0.0228	n/a
κ	0.9481	0.0199	1.4010	0.0269
λ_χ	0.0001	n/a	0.0041	n/a
σ_χ	0.4556	0.0170	0.4139	0.0226
σ_ξ	0.2318	0.0061	0.1768	0.0075
$\rho_{\xi\chi}$	-0.2957	0.0421	0.0890	0.0639
μ_ξ^*	-0.0424	0.0020	0.0146	0.0017
$\lambda_\xi = \mu_\xi - \mu_\xi^*$	0.0052		0.0082	

Table 4 - Parameter estimates and standard errors for restricted case

Using the parameter estimates in Table 4, we can finally develop forecasts and confidence envelopes for both the risk neutral price and the expected spot price, which are shown in Figure 5 for our 1996-2016 data set and in Figure 6 for the 2009-2016 data set. In both cases, we can see the effect of the risk premia through the lower risk neutral forecasts. In Figure 5, it is apparent that the long term equilibrium level drift in the two factor model is influenced by the longer term (i.e., back to circa 2000) historical trend from a period of high, and occasionally very high, prices (circa 2000 to 2008) to a period of low and declining prices (circa 2009 to present). As a result, the expected spot price is forecasted to rise to just above \$3.00/million Btu during the next two years and then stay nearly flat at this level out to the end of the forecast horizon. This may be interpreted as a somewhat conservative forecast, although it roughly aligns with the High Oil and Gas Resource scenario projections from the EIA 2015 Energy Outlook (Figure 7).

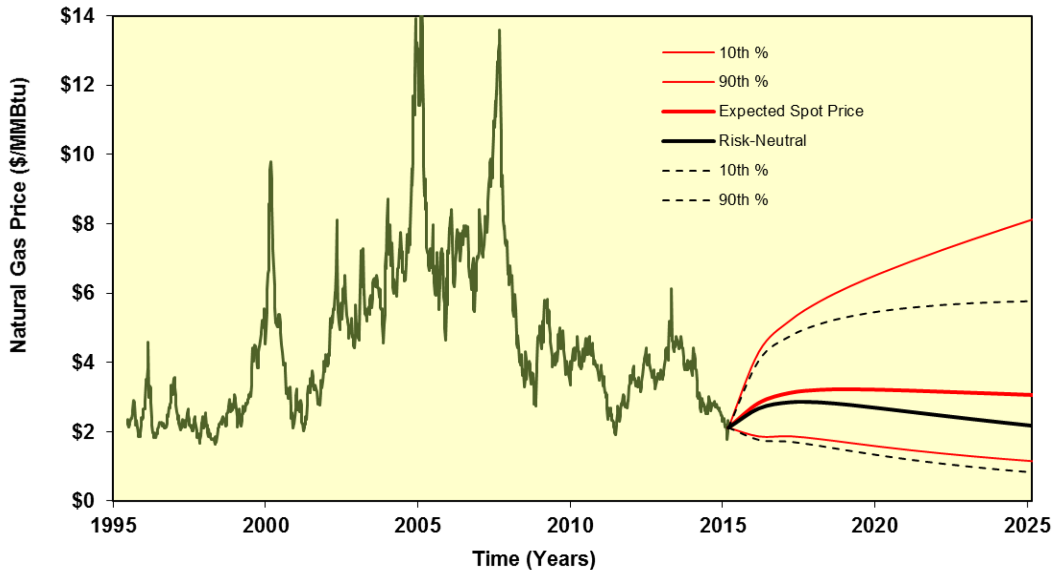

Figure 5 – Natural gas historical and forecasted prices (1996 – 2016 data set)

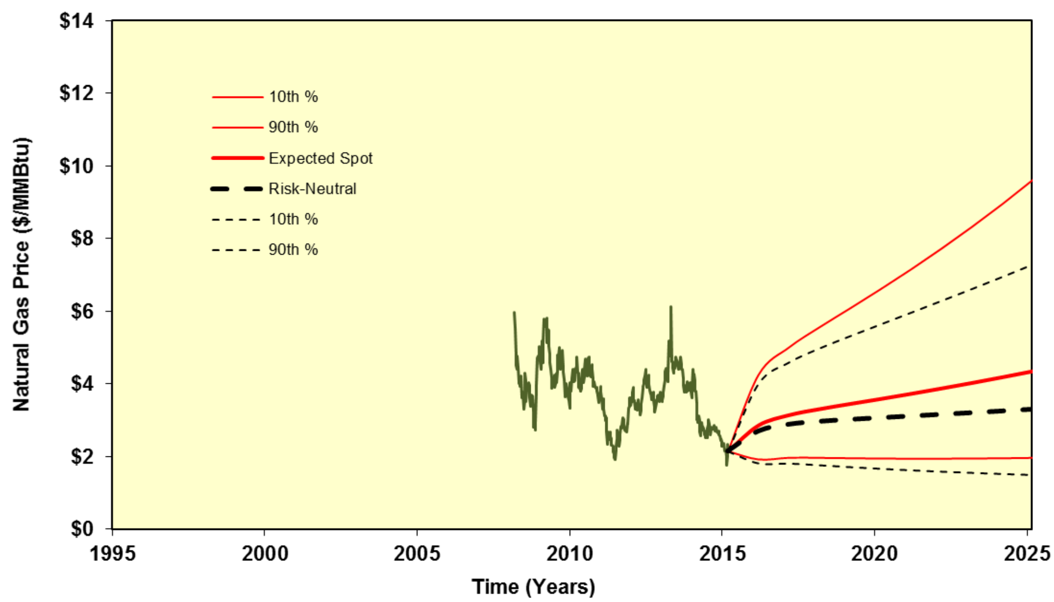

Figure 6 – Natural gas historical and forecasted prices (2009 – 2016 data set)

The forecast in Figure 6, by comparison, is not influenced by the memory of high longer term historical price levels and downward trend. During the period from 2009 through 2014, spot prices oscillated around a mean price just under $4/million Btu, and then prices dropped significantly and rapidly in 2015. As a result, the expected spot price is forecasted to recover to about $3.00/million Btu over the next two years, but then grow at a moderate rate to a price of $4.35/million Btu by the end of the forecast horizon. This forecast calls for higher gas prices than the forecast in Figure 5, although it is still below the Reference case projections from the EIA 2015 Energy Outlook in Figure 7 below.

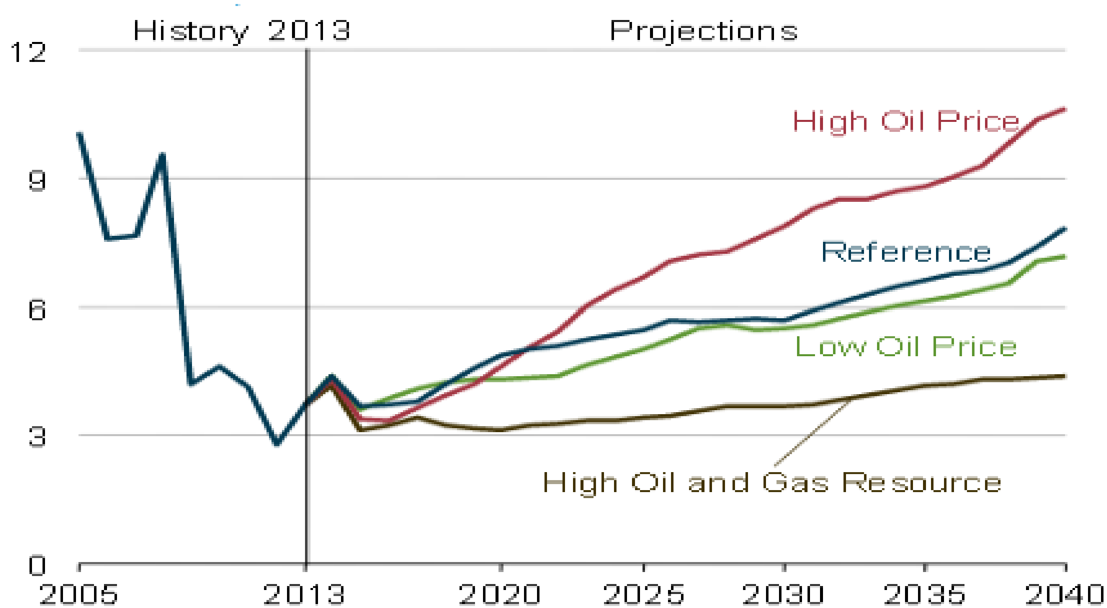

Source: EIA Annual Energy Outlook 2015

Figure 7 – EIA projections for Henry Hub natural gas spot prices under different oil price and resource scenarios

4. Valuation Results

In this section, we use the forecast information from the previous section to first estimate the market value of a prototypical shale gas well in the deterministic base case, and then value two basic managerial flexibilities: 1) the option to decide when to drill the well during a five year lease of the prospective acreage and 2) the option to sell the well at any time after production commences up until reserves are exhausted, for a multiple of the remaining reserves.

If we simply apply the expected value for natural gas price using the 2009-2016 data in Table 4 to the base case project economic model of a shale gas well that was outlined in Section 2, we see a fairly significant change in value from Lake et al. (2013), as indicated below in Table 5. This change reflects the shift from a subjectively chosen $4.50 flat natural gas price forecast and 10% discount rate, to an NPV based upon a risk neutral version of the forecast shown in Figure 6 and a risk free rate of 3%. The latter value is explicitly based on the market's view of future natural gas prices, and it also includes an estimate of the uncertainty associated with future prices, as reflected in the risk premia that would be used to shift to the expected values for the future spot

prices. As discussed earlier, we could also use the expected spot price forecast, if we knew the correct risk-adjusted discount rate. This rate can be found by changing the prices in the project value model to the expected spot prices and solving for the rate that produces the correct NPV of -$91,999, which is 5.2% as indicated in Table 5. Therefore, the risk premia in the natural gas price forecast using the 2009-2016 data imply a 2.2% increase in discount rate for this investment. Similarly, if we were to use the original $4.50 flat price forecast, we could solve to obtain the correct risk-adjusted discount rate, which is 16.3%, rather than the assumed 10%. We note that these results underscore the deterioration of the value of shale gas investments due to the lower price environment since the Lake et al. (2013) paper was written.

Assumptions	$4.50/Mcf flat forecast 10% discount rate	Risk-neutral price forecast 3% discount rate	Expected price forecast 5.2% discount rate
NPV	**$317,990**	**-$91,999**	**-$91,999**

Table 5 – Project NPV for risk-adjusted versus risk-neutral cases

To value real options, the base case deterministic model must first be enhanced so that the option decisions be modeled and recursively solved under the price uncertainty to yield the project value with options. One such method of enhancement is to construct a discrete lattice model of the stochastic natural gas price forecast. If the discrete model is calibrated to market futures prices, as in the procedure in the previous section, it can be considered a risk-neutral forecast of future gas prices, and the future cash flows resulting from the project's options can then also be discounted at the risk-free rate (forecasted to be 3% over the project life).

Development of a discrete model of the natural gas price process requires an extension of the standard Cox, Ross and Rubinstein (1979) approach. Because the Schwartz and Smith (2000) model includes two factors, a two-dimensional binomial approximation must be developed for this process, which results in a four-branch node for the joint process each discrete period. This can alternatively be modeled as the product of the marginal binomial process for long-term deviation ξ and the conditional binomial process for short-term deviation χ:

$$p(\xi_t, \chi_t) = p(\chi_t | \xi_t) p(\xi_t).$$

To define the joint probabilities for the nodes in the binomial lattice in this approach, the increments for each factor value are first set to:

$$\Delta_\xi = \sigma_\xi \sqrt{\Delta t}$$

$$\Delta_\chi = \sigma_\chi \sqrt{\Delta t}\ .$$

Denoting the drifts of the Brownian motion process for ξ and the Ornstein-Uhlenbeck process for χ as:

$$\nu_\xi = \mu_\xi$$

$$\nu_\chi = \kappa(\overline{\chi} - \chi_t),$$

the marginal probabilities of up and down moves in equilibrium level are, respectively:

$$p_u = \frac{1}{2} + \frac{1}{2}\frac{\nu_\xi \Delta t}{\Delta_\xi}, \quad \text{and}$$

$$p_d = 1 - p_u.$$

Expressions for the conditional probabilities for χ are:

$$p_{u|u} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi + \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|u} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi - \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

and

$$p_{u|d} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\rho\sigma_\xi\sigma_\chi - \Delta_\chi \nu_\xi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|d} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) - \Delta t\left(\Delta_\chi \nu_\xi + \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}.$$

Detailed discussions of this approach and example implementations can be found in Hahn and Dyer (2008) and Hahn and Dyer (2011). The implementation is relatively straightforward in a decision tree format using software that will allow node probabilities to be conditioned on prior nodes in a tree. However, since a decision tree with this configuration would have 4^N endpoints, where N is the total number of discrete periods, for most real problems it is desirable to implement this model in more computationally efficient recombining two-dimensional lattices. This alternative implementation requires a few additional steps, but is also relatively straightforward. An example of such an algorithm is presented in Clewlow and Strickland (2000) for the case of two GBMs and based on four branch nodes at each step. In this case, one of the GBMs is simply replaced with a one-factor Ornstein-Uhlenbeck process and the four-branch node has also been replaced with a two node marginal-conditional sequence.

In what follows, we summarize the results from applying our real option model to optimally time investment during a typical five year lease and show the sensitivity of the lease value to the value of the parameters of the stochastic price model. We again use the restricted case with the 2009-2016 data set as the base case set of parameters (i.e., right-hand column of Table 4). The option to invest was modeled over five discrete periods (T=5 years, N=5 periods, so n=N/T=1 period per year). The optimization algorithm begins in the last period by selecting the maximum of zero (abandoning the lease) or the value of investing at the price given at each of the different terminal nodes, and then recursively working backwards to time zero, maximizing the value of investing at each node versus waiting. For a five period model, the number of nodes in the terminal period would be $(N+1)^2 = 36$. It is well known that binomial models more accurately approximate true values as the number of periods for a given time horizon increases (thereby increasing the number n of the discrete periods per unit of time). Figure 8 below shows this convergence behavior

for our model, assuming a 5-year option period. Convergence appears to occur for $n > 4$, corresponding to quarterly time periods, and the value of the shale gas project with the option stabilizes at $483k. Thus, the project's NPV from Table 5 changes from negative to positive, and the value of the option to optimally time investment is worth $575k = 483k – (-92k).

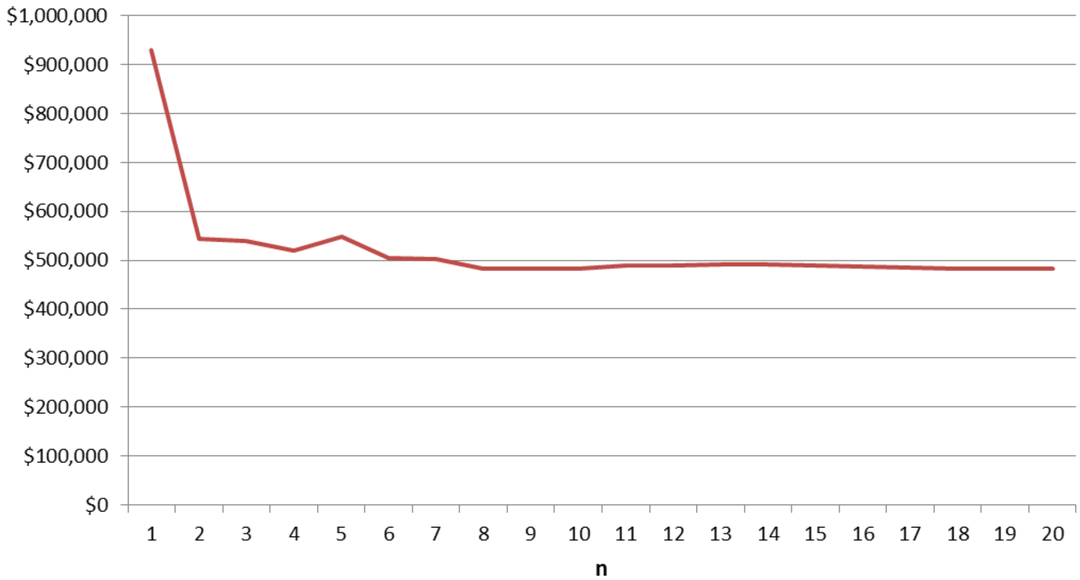

Figure 8 – Convergence of value for investment timing option

We also checked the assumption of a five year lease period, to see whether additional option value could be obtained from extending this investment horizon up to T=20 years. The results shown in Figure 9 indicated that a significant part of the option value is derived from the first 5-7 years, perhaps supporting the intuition behind this industry practice.

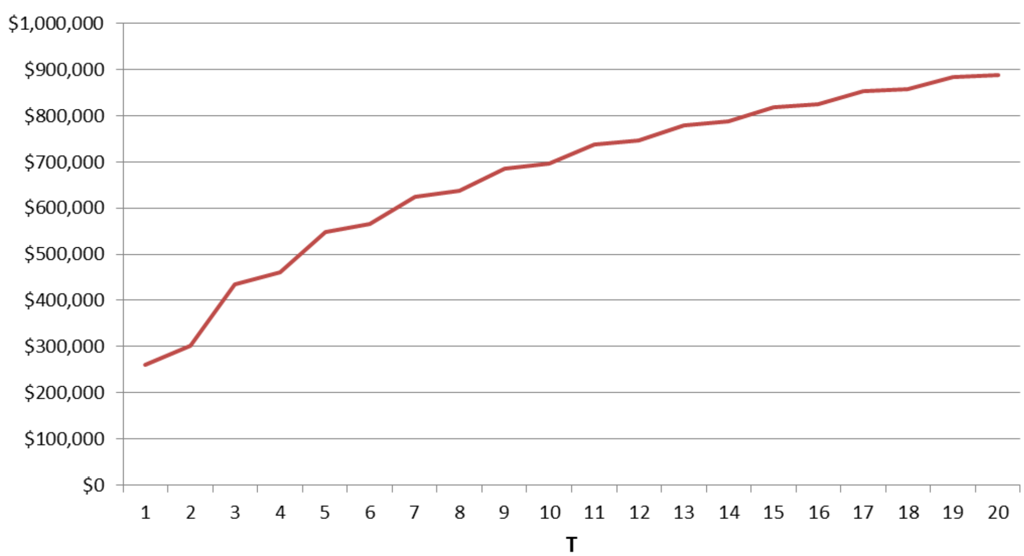

Figure 9 – Sensitivity of investment timing option value to lease expiration time (T)

We further conducted a sensitivity analysis of the two-factor model parameters for the 2009-2016 data set in Table 4 as the base case, and with the ranges of variation of the inputs shown

in Table 6. The ranges for variation of inputs were selected based on ranges of estimates seen in different studies, such as Manoliu and Tompaidis (2002) and Carvalho (2010), as well as our own experience with different natural gas futures data sets. The results are summarized in Figure 10. The individual sensitivity curves were all increasing functions in the inputs except for the mean reversion coefficient and the risk-free rate. In those case, higher values for the inputs resulted in lower option values relative to the base case. These relationships confirm intuition in most cases. It is expected that the value of the option to invest decreases as the mean reversion coefficient κ increases, because a larger mean reversion coefficient reduces the time to revert to the long-term equilibrium level, and has the effect of reducing the project's overall volatility. As expected, the value of the option increases with both the short- term volatility and long-term volatility, although the effect is stronger for the latter. Also as expected, the value of the project with the option increases with the long-term drift, since this increases the likelihood of realizing higher prices and project values earlier during the investment time horizon. Although not technically a model parameter, we did confirm that the value of the option decreases with an increased risk free rate. This simply reflects the effect discounting (at the risk-free rate) has on future cash flow valuations.

Parameter	Low	Base	High
Mean-reversion speed, κ	2.000	1.401	0.500
Short-term volatility, σ_χ	0.300	0.414	0.500
Long-term volatility, σ_ξ	0.100	0.167	0.300
Correlation, $\rho_{\chi\xi}$	(0.500)	0.089	0.500
Long-term drift, μ_ξ*	0.000	0.015	0.050
Risk-free rate	0.050	0.030	0.010

Table 6 – Input ranges for sensitivity analysis of price model parameters

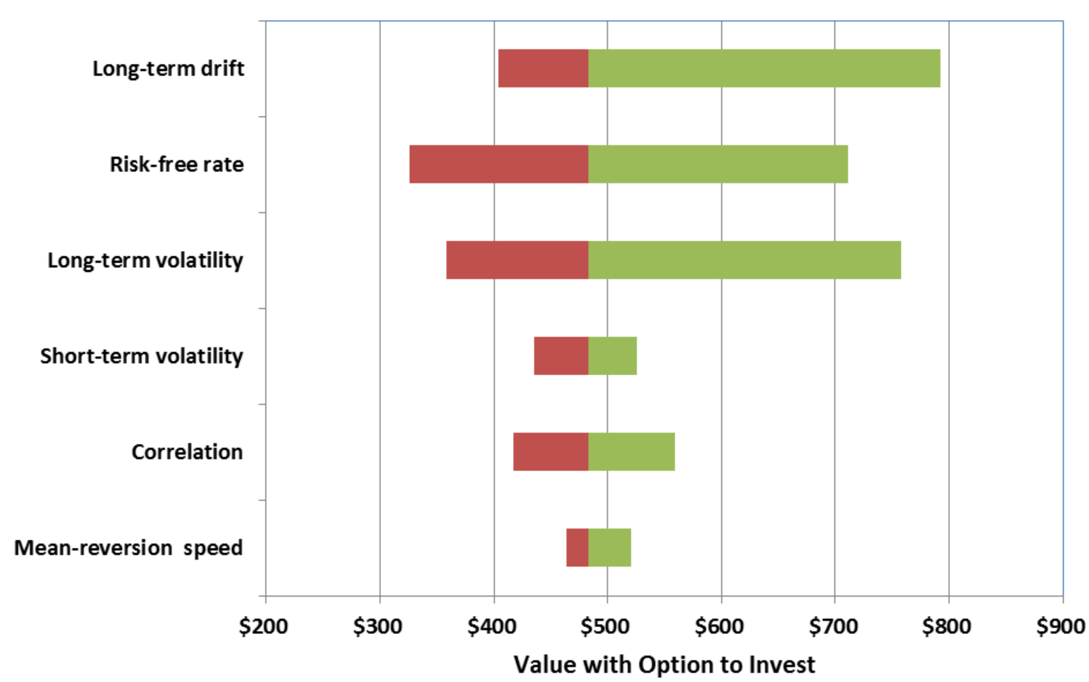

Figure 10 – Sensitivity of investment timing option value to price model parameters

By comparing the parameter estimates used for the base values in Table 6, which were obtained in the restricted case (as summarized in Table 4) with the parameter estimates from Table 1, and referencing Figure 10, we can evaluate the change in value that would result from using the unrestricted parameter estimates. The most sensitive model parameters, the long-term risk neutral drift rate and the long-term volatility would be lower and higher respectively; however, as discussed above, the two effects would tend to cancel each other out.

Last, we discuss the results from our option pricing model for the option to divest at any time after investing in the shale gas well at time zero, again with the objective of showing the sensitivity of the option value to the different parameters in the stochastic price model. The option to divest was modeled over 152 discrete steps (T=38 years, n=4 periods/year), with the option pricing model algorithm beginning in the last period by selecting the maximum of last cash flow at the price given at each of the different terminal nodes or the value of divesting at a given multiple ($1.60/Mcf) of the remaining unproduced reserves, and then recursively working backwards to time zero, maximizing the value of continuing to produce at each node plus the expected value of the four possible values in the subsequent period versus selling immediately. This model is much larger than the investment timing model due to the long time horizon, and the resulting number of nodes in the terminal period is $(N+1)^2 = 23{,}104$ nodes. This makes it less practical to decrease the size of the discrete time steps in the model, but we observed accurate convergence to values with quarterly time periods, and even fairly good convergence with semi-annual time periods.

We used a base case valuation multiple of $1.60/Mcf for developed but unproduced gas reserves. With this assumption, the value of the shale gas project with the divestment option is now positive at $47.2k, which is a $139.3k premium over the base case value. But, the option to divest has a much lower option value compared to the investment timing option. This option does rapidly start to become more valuable as the valuation multiple increases above $1.60/Mcf as shown in Figure 11, although it is unlikely that higher multiples will be realized unless gas prices also see sustained increases. Furthermore, the valuation multiple might be expected to decrease, rather than remain fixed, if prices decrease. In that case, the option to divest would be worth less because it is more likely to be exercised when prices, and the remaining cash flows, have low values. Therefore it appears that, under the current price forecast, waiting to invest in a well over a typical five year lease is a more valuable strategy than investing in a well now and factoring in the value of the option to divest later.

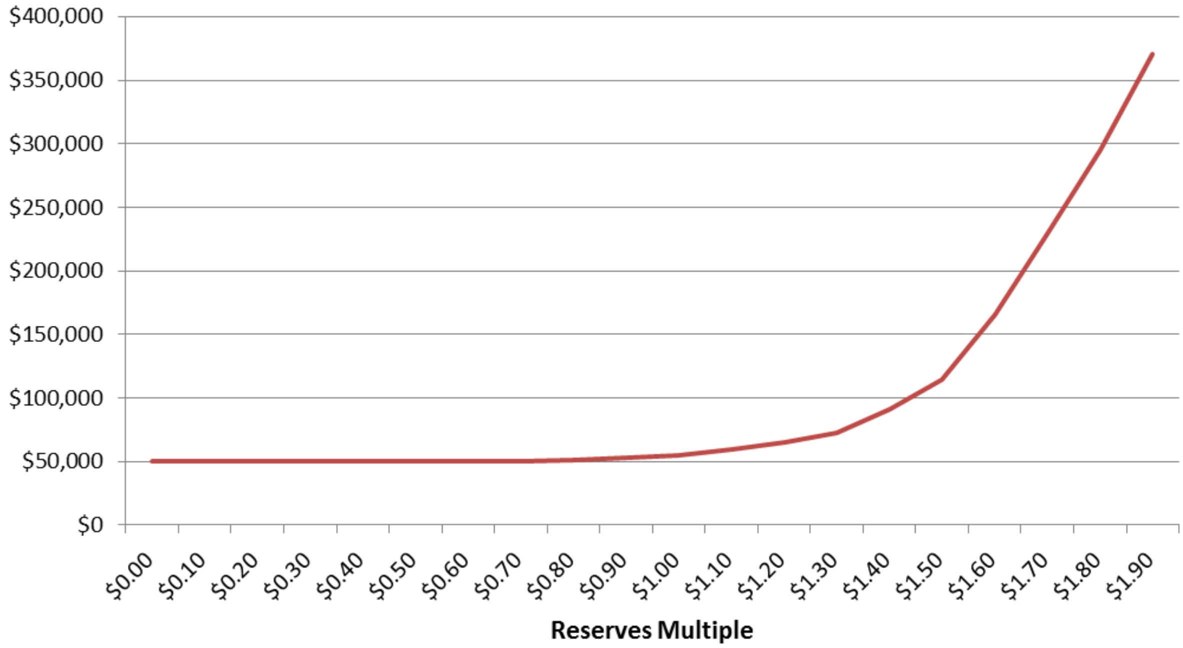

Figure 11 – Sensitivity of divestment option value to reserves multiple

As with the investment timing option, we conducted a sensitivity analysis of the two-factor model parameters, with the same ranges of variation of the inputs as shown in Table 6. Figure 12 summarizes the results (using the same scale and order, for comparison to Figure 10). In this case, the long-term drift rate has very little effect on the value of the option to divest. This situation occurs because the instances in which the option is likely to be exercised occur late during the option period, when the long-term drift apparently has little effect on dropping values below divestment levels.

The individual sensitivity curves were again mostly all increasing functions of the inputs; however, there were a couple of changes in the exceptions. The mean reversion coefficient in this case becomes an increasing function, because in addition to counteracting volatility κ, increasing it also has the effect of keeping future prices lower. This time, the value of the project with the option *decreases* as the long-term drift increases, since this increases the likelihood of realizing higher prices and project values earlier during the investment time horizon and therefore reduces the value of the option to divest. The effect of the other parameters continued to be as expected, with the value of the option increasing with both the short- term volatility and long-term volatility, this time due to higher likelihood of lower future price and value realizations.

By again comparing the parameter estimates in Tables 1 and 4 (unrestricted and restricted cases), we can see the change in value that would result for this option. In this case, higher long-term risk neutral drift rate and lower long-term volatility estimates in Table 1 would not only tend to cancel each other out, but the abandonment option value is also not very sensitive to changes in either of those directions, as shown in Figure 12.

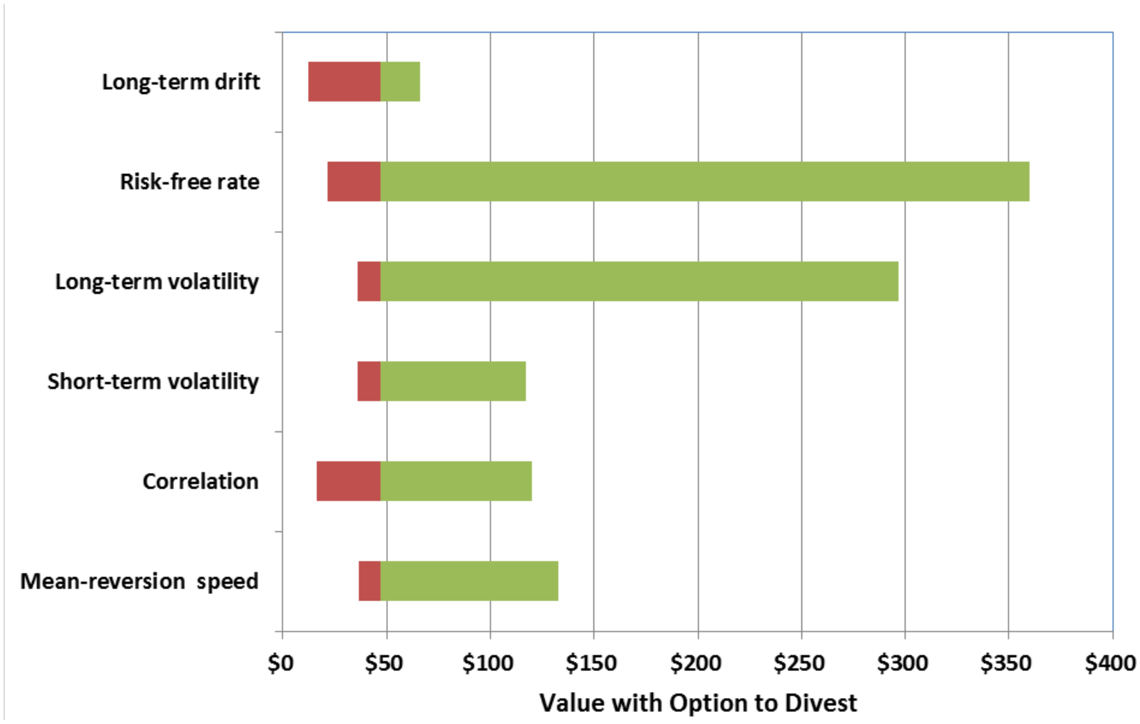

Figure 12 – Sensitivity of divestment option value to price model parameters

5. Conclusions

In this paper, we used natural gas futures prices in a Kalman filter – maximum likelihood estimation approach to parameterize the Schwartz and Smith (2000) stochastic process model and then applied it to forecast prices and value a prototypical shale gas well. We find, as have other researchers in other applications, that the short and long term risk premia are not well estimated by the Kalman filter approach. We also find that an asset pricing model approach can be used to obtain more realistic estimates of risk premia parameters, and that the Kalman filter can be used on a restricted basis to estimate the remaining two-factor model parameters. The estimates of these remaining parameters are only marginally (range from 8-10%) changed in the restricted case, except for the correlation parameter, which decreased by 73%, and their standard errors are not significantly changed. The improved estimation of the risk premia allows the development of reasonable forecasts for both the risk neutral price and the expected spot price.

We used two data sets in this work; a large set of natural gas futures prices, beginning in 1996 when all of our contracts of interest began to be continuously traded, as well as a more abbreviated data set that is intended to represent the shale gas production era, which began to impact production volumes and prices in the 2009 timeframe. We found that the choice of data set has some effect on the two factor model parameter estimates and the resulting forecast, with the longer term data set resulting in a slightly more conservative forecast due to the long term downward trend from the high prices realized in the mid- to latter part of the 2000-2010 decade. With either data set, however, we obtain forecasts that roughly align with the High Oil and Gas Resource and Low Oil Price scenarios from the 2015 EIA Energy Outlook, two outcomes that seem increasingly likely as judged by current market sentiment.

Finally, we illustrated an application of the parameterized two factor model by valuing two basic operating options associated with a shale gas investment – the option to optimally time development and the option to divest (abandon for salvage) an asset. We find that the option to divest is less valuable, even with optimistic expectations for the divestment multiple. Our integrated price forecasting – option valuation model also yielded insights about the relative effects of each of the parameters in the two-factor price model on the project value via sensitivity analysis. Our findings indicate that the most impactful model parameter in the case of the option to develop is the drift rate of the long-term equilibrium level, while in the case of the option to abandon, the volatility of the long-term equilibrium level predominates. Therefore, in terms of the processes for the two factors in the Schwartz and Smith (2000) model, we find that even with the rapid decline in production and cash flows from a shale gas well, the parameters for the long-term component of the two-factor model are the most significant drivers of project option value.

References:

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

Black, F. and Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* **81(3)**, 637-654.

Boyle, P., 1988. A lattice framework for option pricing with two state variables. *Journal of Financial and Quantitative Analysis* **23(1)**, 1-12.

Carvalho, A., 2010. Calibration of the Schwartz-Smith model for commodity prices. *Instituto de Matemática Pura e Aplicada* (Brasil).

Chorn, L. and Shokhor, S., 2006. Real options for risk management in petroleum development investments. *Energy Economics* **28(4)** 489 – 505.

Clewlow, L. and Stickland, C., 2000. *Implementing derivatives models*. Wiley, West Sussex, U.K..

Cortazar, G. and Schwartz, E., 1998. Monte Carlo evaluation model of an undeveloped oil field. *Journal of Energy Finance & Development* **3(1),** 73–84.

Cortazar, G., Kovacevik, I., and Schwartz, E., 2015. Expected commodity price returns and pricing models. *Energy Economics* **49,** 60–71.

Cox, J., Ross, S. and Rubinstein, M., 1979. Option pricing: A simplified approach," *Journal of Financial Economics* **7(3),** 229-263.

Dias, M., 2004. Valuation of exploration and production assets: An overview of real options models. *Journal of Petroleum Science and Engineering* **44(1-2)** 93 – 114.

Dias, M., Lazo, J., Pacheco, M., and Vellasco, M., 2003. Real option decision rules for oil field development under market uncertainty using genetic algorithms and Monte Carlo simulation. 7th Annual Real Options Conference, Washington DC.

Duffie, D., 1992. Dynamic asset pricing theory. Princeton University Press, Princeton, NJ.

Ekern, S., 1988. An option pricing approach to evaluating petroleum projects. *Energy Economics* **10(2)** 91–99.

Hahn, W., and Dyer, J., 2008. Discrete time modeling of mean-reverting stochastic processes for real option valuation. *European Journal of Operational Research* **184**(2) 534-548.

Hahn, W. and Dyer, J. 2011. A discrete time approach for modeling two-factor mean-reverting stochastic processes. *Decision Analysis* **8**(3) 220-232.

Harvey, A., 1989. Forecasting, structural time series models and the Kalman filter. Cambridge University Press, Cambridge, U.K.

Hull, J., and White, A., 1993. One-factor interest rate models and the valuation of interest rate derivative securities. *Journal of Financial and Quantitative Analysis* **28(2),** 235-254.

Hull, J., and White, A., 1994. Numerical procedures for implementing term structure models II: Two-factor models. *Journal of Derivatives* **2(2),** 37-49.

Jafarizadeh, B., and Bratvold, R., 2012. Two-factor oil-price model and real option valuation: An example of oilfield abandonment. *SPE Economics and Management* **4(3)**, 158-170.

Kalman, R., 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME – Journal of Basic Engineering* **82(1)**, 35-45.

Lake, L., Martin, J., Ramsey, J and Titman, S., 2013. A primer on the economics of shale gas production, just how cheap is shale gas?”, *Journal of Applied Corporate Finance*, 25(4),87-96

Longstaff, F. and Schwartz, E., 2001. Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies* **14(1)**, 113-147.

Lund, M., 1999. Real options in offshore oil field development projects. 3rd Annual Real Options Conference, Leiden.

Manoliu, M., Tompaidis, S., 2002. Energy futures prices: term structure models with Kalman filter estimation. *Applied Mathematical Finance* **9(1)**, 21-43.

McCardle, K., and Smith, J., 1998. Valuing oil properties: Integrating option pricing and decision analysis approaches. *Operations Research* **46(2),** 198 – 217.

Merton, R., 1973. The theory of rational option pricing. *Bell Journal of Economics and Management Science* **4(1),** 141-183.

Nelson, D. and Ramaswamy, K., 1990. Simple binomial processes as diffusion models in financial models. *The Review of Financial Studies* **3(3),** 393-430.

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

Schwartz, E. and Smith, J., 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46(7),** 893-911.

Siegel, D., Smith, J., and Paddock, J., 1988. Valuing offshore oil with option pricing models. *Midland Corporate Finance Journal* **5(1),** 22–30.