DAY-OF-THE-WEEK EFFECTS IN STOCHASTIC-OIL-PRICE MODELS

A Thesis

by

JOHN D. RICHARDSON

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Chair of Committee,	Duane A. McVay
Committee Members,	Ruud Weijermars
	Thomas A. Blasingame
Head of Department,	Jeff Spath

December 2020

Major Subject: Petroleum Engineering

ABSTRACT

Reserves are the technically-recoverable hydrocarbon volumes that can be economically produced given the current economic condition. Uncertainty about future oil prices is experienced in the present, not the future, and is a property of the current economic condition. Average or risked oil price scenarios are deterministic and may not correctly capture the impacts of oil price volatility on field-level reserves or other economic outcomes such as net present value. Stochastic price-volatility models tend to operate on the scale of days to a couple weeks, which is short compared to the life of a hydrocarbon well. The lack of a long-term stochastic model for price volatility motivates a look at how current stochastic models are made, and the creation of a model consistent with the life of hydrocarbon wells. With such a model, the dependence of reserves volume on price volatility can be assessed.

Random-walk models have been used to simulate the behavior of market prices and the uncertainty of future price changes over time, but usefulness is limited when the distribution of observable price changes is not well defined. A new density function is proposed to model returns on oil price. This density function, having a shape that depends on the coefficient of variation of the returns, is formed by the product of two Laplace distributed random variables. Although the new distribution was developed in context of West Texas Intermediate (WTI) spot price, the model was capable of modeling other markets such as Johnson and Johnson stock price.

Traditional methods of calibrating mean and variance behavior of returns for use in random-walk models has been inadequate. Models often err by failing to explicitly include or exclude intra-week behavior. In this thesis, mean and variance metrics are determined in a novel way that defines then removes day-of-the-week effects that may have cumulative bias when estimating project value measured in years or decades.

Using probabilistic decline-curve parameters from the Eagle Ford shale, reserves and profitability were estimated for a synthetic project. The average reserves volume determined under a stochastic scenario was less than the average reserves volume determined using average price as a deterministic input. This means that using average prices when estimating reserves volume does not obtain the true average reserves volume in practice. The magnitude of the difference between stochastic and deterministic price scenarios depends on the variable cost per barrel of oil. For $10/bbl variable cost, average reserves did not differ significantly between deterministic and stochastic price modeling; at $30/bbl variable cost, average reserves were 17% lower using stochastic modeling. Accounting for volatile oil prices is paramount to obtain true average reserves volume, and average-price-in does not always result in true average-reserves-volume out.

DEDICATION

I dedicate this thesis to my family for their unconditional support.

ACKNOWLEDGEMENTS

I would like to thank the support staff at Texas A&M for their assistance in progressing through the Master's program and the professors of the petroleum engineering department for the education they provided me.

I would like to thank my committee chair and advisor, Dr. Duane McVay, for his patient advice and corrections of this work. I would like to thank Dr. Ruud Weijermars who inspired my interest in this topic. I appreciate the teachings of Dr. Thomas Blasingame which were instrumental in breaking down and analyzing data.

I am grateful for the patient ear of my roommates Roman Johnson and Zach Hinsen who helped clarify my explanations.

Finally, I would like to thank my parents, Sandra and Joseph Richardson, and sister, Elizabeth Richardson, for their encouragement of my success.

CONTRIBUTORS AND FUNDING SOURCES

This work was supervised by a thesis committee consisting of Professor Duane A. McVay [advisor] and Professor Ruud Weijermars of the Department of Petroleum Engineering, and Professor Thomas A. Blasingame of the Department of Petroleum Engineering and Geology. All work for the thesis was completed independently by the student.

There are no outside funding contributors to the research, analysis, and compilation of this thesis.

NOMENCLATURE

b	Decline Exponent
C	Coefficient of Variation of R
D	Distribution of Unknown Form
D_0	Initial Decline Rate
D_e	Minimum Decline Rate
$\mathbb{E}(x)$	Expected Value of x
$\mathcal{L}$	Laplace Distribution
m	Months
q	Oil Flow Rate
Q	Cumulative Production, Rate Integral
r	Return, Observed Ratio of Two Spot Prices
R	Observed Set of Returns, r
s	Spot Price
t	Time
$\mathbb{V}(x)$	Variance of x
$\mathscr{y}$	Laplace Distributed Random Variable
$\mathcal{Y}$	Distribution of y variables
z	Z-Factor

α	Exponential Parameter of Log-Logistic Distribution
β	Scale Parameter of Log-Logistic Distribution
γ	Location Parameter of Log-Logistic Distribution
μ	Mean of R
$\tilde{\mu}$	Mean of $\mathcal{Y}$
σ	Standard Deviation of R
$\tilde{\sigma}$	Standard Deviation of $\mathcal{Y}$
τ	Elapsed Time, Holding Period

CF	Cash Flow
FC	Fixed Cost
IRR	Internal Rate of Return
NCF	Present Value of Cumulative Cash Flow
NOP	Present Value of Cumulative Operating Profit
NPV	Net Present Value
OP	Operating Profit
P10	10^{th} Percentile, Cumulative
P50	50^{th} Percentile
P90	90^{th} Percentile, Cumulative
PI	Profitability Index
VC	Variable Cost

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF TABLES

Page

CHAPTER I

INTRODUCTION

According to the Society of Petroleum Engineers (2007), reserves are defined as the portion of the hydrocarbon resource base that is technically and commercially recoverable and are further sub-classified into proved, probable, or possible volumes by the uncertainty in forecasted volume. Proved reserves are defined by United States Securities and Exchange Commission (SEC) regulation (210.4-1 (a) (22)) as a volume that "[…] can be estimated with reasonable certainty to be economically producible – from a given date forward, from known reservoirs, and under existing economic conditions, operating methods, and government regulations […]." A reserves estimate is "proved" if it is the value that is exceeded by 90% of possible outcomes (210.4-1 (a) (24)). Three categories of uncertainty can be extrapolated from the SEC definition. One is uncertainty in government regulation, another is technical uncertainty, the last is economic uncertainty. Technical uncertainty would, for example, account for the uncertainty in being able to successfully drill a three-mile well lateral, when the longest the company has drilled is a 2.5-mile lateral. For the purpose of this thesis, the economic uncertainty is of primary interest.

The oil and gas price that satisfies "existing economic conditions" is standardized by the SEC as the average of the first day's price for each of last 12 months. However, this does not account for uncertainty in what the price *will be* in the future. The 12-month average may be used as a good starting price (and is used that way in this thesis), but where the price goes from there over a well’s lifetime is highly variable and subject to a lot of uncertainty. Uncertainty in future oil prices introduces economic uncertainty. Companies that report reserves according to SEC guidelines may choose to also run internal books for risk assessment and decision making. The average proved, probable, or possible reserves volume reported after accounting for technical uncertainty may be different than the average volume obtained when economic uncertainty is also accounted for. Additionally, investors may find it prudent to understand that risk of oil-price changes may not be entirely captured in the current SEC method of reporting reserves, where average prices such as the SEC 12-month average or New York Mercantile Exchange Strip prices are used to estimate average reserves volume. Use of these averages or strip prices might not accurately represent the true average reserves volume if volatility in oil prices is not accounted for.

Deterministic price timelines akin to “hockey sticks” (Olsen, McVay, and Lee 2005) may be constructed to represent high, average, and low-price scenarios, but this does not account for volatility of price in achieving those high, average, and low-price outcomes. Bootstrap methods (McMichael 1999) that sample from previously observed oil prices can be employed to account for oil-price volatility. Bootstrapping only samples the observed and assumes it is a reasonable approximation of the possible. If a continuous distribution is capable of describing the discrete set of observed price changes, it would, on a technical level, be a better descriptor of all possible values than bootstrapping. To account for all possible outcomes, a stochastic price model needs to be developed that does not rely on bootstrapping changes in oil price.

Categories of models producing long-term forecasts on the scale of years are few, and are primarily limited to National Energy Modeling System (NEMS), neural networks, and futures contract/spot price speculation (Lee and Huh 2017). With NEMS, the Energy Information Administration (EIA) provides long-term forecasts of energy prices (EIA 2014). Relying on finding equilibrium between world energy markets as well as supply-and-demand structures, NEMS is the most comprehensive model. According to documentation on NEMS architecture published by the National Research Council (1992), the process of using a probabilistic model would be “so time-consuming as to be entirely impractical” and “NEMS will not be used routinely to examine probabilistic rational expectations equilibria.” Unfortunately, operators are stuck with static forecasts from the model, because its complexity makes deployment in generating probabilistic-price scenarios for project economics difficult. It may therefore be useful to look at models that are conventionally for short-term price forecasts—which operate on the scale of seconds to weeks—and find a suitable way to extend them to time frames relevant to the life of a well or field.

Neural networks suffer a serious flaw of being a “black box” by making the method by which a prediction was obtained opaque. There is work on the topic of extracting “fuzzy rules” from neural networks (Kolman and Margaliot 2007; Buhrmester, Münch, and Arens 2019). In their current state, neural networks give rise to a lack of transparency and are not easily audited (Bathaee 2018). Lacking fundamental descriptions other than “fuzzy rules” for the system being described prevents using neural networks for the purpose of this thesis. The predictions must be clearly constructed, defined, and auditable, and neural networks fail each aspect.

It may be alluring to use the futures and options markets, where experts speculate on the value of futures contracts. Futures contracts are an agreement today to exchange a good at a specific price at some defined point or period of time in the future. Options contracts give the buyer of the options contract the right, but not the obligation, to buy or sell blocks of futures contracts to the seller of the options contract by or on the expiration date of the options contract at a predetermined strike price. Changes in these markets will impact the underlying spot price. Despite having “future” in the name, the oil futures market is fundamentally *not* predictive of future price, and the price of a futures contract *must* be equal to that of the spot price plus the cost-to-carry, which includes interest rates, storage costs, and opportunity costs (McDonald 2006).

Yanagisawa (2009) claimed futures prices were useful for one-week or one-month predictions of oil price, but not for anything longer term. However, in a review of literature on the topic, Fattouh et al. (2012) concluded that current evidence did not support speculation having an important role in determining spot price of oil. Furthermore, they conclude that oil futures do not improve predictions of future spot price. To the contrary, they found that “co-movement between spot and future prices” instead “reflects common economic fundamentals rather than the financialization of oil futures markets.” These common economic fundamentals are the relationship between spot and futures pricing described in the previous paragraph. Futures price is *not* the price of oil in the future. It is the price of oil now, plus costs associated with the exchange in the future. This is a subtle, but non-trivial nuance to futures pricing. There is a lack of a clear fundamental basis for using speculative markets to predict price and there is a literature support that such markets are not useful at the time scales being used for a hydrocarbon well. The futures and options market cannot be reasonably used for forecasts on the scale of years, which is the need in this thesis.

Simulating changes in price with random-walk models is a common short-term method focused on seconds to days, the most popular of which is the Generalized Auto Regressive Conditional Heteroscedasticity (GARCH) class (Bollerslev 1986). The change from one return to the next, under a GARCH model, has conditional variance and errors. As the GARCH model generates a simulated timeline of prices, the previous values generated will impact the selection and randomness of future values that can be generated. By allowing mean and variance parameters to change at each step of the model, the behavior of an individual GARCH model can be finely tuned to the underlying data.

In an assessment of GARCH class models applied to Brent and West Texas Intermediate (WTI) crude-oil markets, Wei, Wang, and Huang (2010) conclude that model choice will depend on the statistical loss function used to rank them. Minimization of loss functions, such as square-error, are used to maximize model fit. Their work confirms results by Lopez (2001) that the choice of loss function impacts forecast-evaluation results and subsequent ranking of models based on their results, which further complicates finding a clear method of evaluating these models. Hou and Suardi (2011) did a follow-up analysis with a non-parametric method of forming the GARCH model, using only "robust" loss functions. The non-parametric method, which was generated by iteratively solving for the constants, performed the best. Still, there is a complete lack of justification for GARCH model form, other than the ability to reproduce out-of-sample historical prices. This does not mean the model has described the underlying fundamentals; it has merely created correlations. Rather than the model being prediction-capable, the method was capable of making well-fitting models. This is a subtle, but important distinction. If the method is rerun with updated prices or a different price profile, there is nothing that precludes the correlation structure of the model from changing.

The primary objective of this thesis is to develop a better method of handling price data when constructing price-forecast models. In this thesis, a stochastic-price-forecasting model is developed in three steps. First, the density function of returns is described with a fundamental form in Chapter II. Second, the behavior of returns are analyzed in context of correlations and time in Chapter III. Third, a random-walk model is calibrated to WTI spot price history in Chapter IV using the analysis framework developed in Chapters II and III. To demonstrate how the model can be used, the model is employed in Chapter V to estimate reserves for a synthetic field case.

CHAPTER II

STANDARD PROBABILITY DENSITY FUNCTION

Methodology for generating a density curve for returns and basic functional relations is presented in this chapter. Some markets appear to have Gaussian mixture, lognormal, or generalized log-F distributed returns, while others appear Laplace distributed (Behr and Pötter 2007). Use of a hyperbolic density function of four parameters appears to model the various density functions (Küchler and Neumann 1999). It may be possible that the observed behavior of returns on oil spot price can be represented using a simpler function with fewer parameters.

In this thesis, all analysis is performed on WTI spot prices between January 2, 1986, and December 31, 2015. No data were excluded. These data were obtained online from the EIA's website (EIA 2020). The EIA uses the unweighted average closing spot price and reports values Monday to Friday. The timeline of these prices is presented in **Fig. 1**.

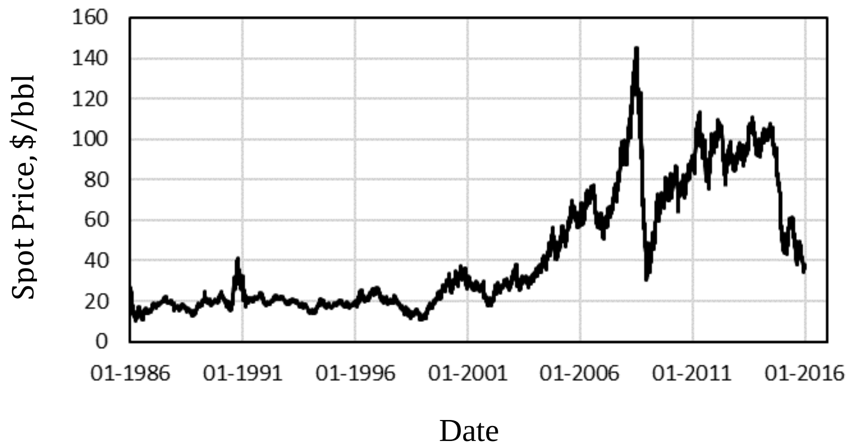

Fig. 1—WTI spot price between January 2, 1986, to December 31, 2015.

Density Function

I begin by establishing relations between common terminology and mathematical representation used in this thesis. A holding period is the time between buying and selling a specific asset, in this case, oil. A return is defined as the relative change in price with respect to time. Return (r_τ) on an asset held for a period of time before selling is the ratio of the price τ days ahead, $s_{t+\tau}$, to the price at time t, s_t:

$$r_\tau = \frac{s_{t+\tau}}{s_t} \tag{1}$$

The holding period (τ) over which this return occurs defines the step width of the change. When the asset is purchased is defined by t. A specific and singular value of r_τ is obtained by specifying a value of t for a particular value of τ. The possible values of r_τ are described by set R_τ. A set (R_τ) of possible returns for a given value of τ is described by a probability density function (D) having a mean and variance of μ_τ and σ_τ^2.

$$r_\tau \in R_\tau \sim D(\mu_\tau, \sigma_\tau^2) \tag{2}$$

The proposed density function of R_τ is the product of two elements from the Laplace distribution $\mathcal{Y}$:

$$y_\tau \in \mathcal{Y}_\tau \sim \mathcal{L}(\tilde{\mu}_\tau, \tilde{\sigma}_\tau^2) \tag{3}$$

$$r_\tau = y_\tau y_\tau \tag{4}$$

Mean value of $\mathcal{Y}_\tau$ is calculated by:

$$\tilde{\mu}_\tau^2 = \mu_\tau \tag{5}$$

Using a general form of the notation, variance of $\mathcal{Y} \times \mathcal{Y}$ with independent sampling is:

$$\mathbb{V}(R) = \mathbb{V}(\mathcal{Y}\mathcal{Y}) = \mathbb{V}(\mathcal{Y})\mathbb{V}(\mathcal{Y}) + \mathbb{V}(\mathcal{Y})\mathbb{E}(\mathcal{Y})^2 + \mathbb{V}(\mathcal{Y})\mathbb{E}(\mathcal{Y})^2 \tag{6}$$

Substituting in symbolic notation and simplifying yields:

$$\sigma_\tau^2 = \tilde{\sigma}_\tau^2\tilde{\sigma}_\tau^2 + \tilde{\sigma}_\tau^2\tilde{\mu}_\tau^2 + \tilde{\sigma}_\tau^2\tilde{\mu}_\tau^2 \tag{7}$$

$$0 = \tilde{\sigma}_\tau^4 + 2\,\tilde{\mu}_\tau^2\tilde{\sigma}_\tau^2 - \sigma_\tau^2 \tag{8}$$

Applying Eq. 5:

$$0 = (\tilde{\sigma}_\tau^2)^2 + 2\mu_\tau(\tilde{\sigma}_\tau^2) - \sigma_\tau^2 \tag{9}$$

Where $\tilde{\sigma}_\tau^2$ can be solved for via the quadratic formula:

$$\tilde{\sigma}_\tau^2 = -\mu_\tau + \sqrt{\mu_\tau^2 + \sigma_\tau^2} \tag{10}$$

Elements of $\mathscr{y}_\tau$ are thus defined as:

$$\mathscr{y}_\tau \in \mathcal{Y}_\tau \sim \mathcal{L}\left(\sqrt{\mu_\tau}, -\mu_\tau + \sqrt{\mu_\tau^2 + \sigma_\tau^2}\right) \tag{11}$$

If r_τ is formed by the product of two Laplace distributed elements, let its distribution be called the double Laplace Distribution ($\mathcal{L}\mathcal{L}$):

$$r_\tau \in R_\tau \sim \mathcal{L}\mathcal{L}(\mu_\tau, \sigma_\tau^2) \tag{12}$$

Z-Table

A probability density function can be changed to a standard scale by converting the described data to z-scores. In this section, it is shown that the standard density function of $\mathcal{Y}_\tau\mathcal{Y}_\tau$ describing R_τ is uniquely determined by the coefficient of variation (*c.v.*) of R_τ. To get Eq. 12 into a more generic form, the distribution is normalized by the mean value:

$$\frac{r_\tau}{\mu_\tau} \in \frac{R_\tau}{\mu_\tau} \sim \mathcal{LL}\left(1, \frac{\sigma_\tau^2}{\mu_\tau^2}\right) \tag{13}$$

Letting the *c.v.* equal C:

$$\frac{r_\tau}{\mu_\tau} \in \frac{R_\tau}{\mu_\tau} \sim \mathcal{LL}(1, C^2) \tag{14}$$

The standard score, z, of r_τ is then:

$$z = \frac{r_\tau - \mu_\tau}{\sigma_\tau} = \frac{\mathcal{LL}(\mu_\tau, \sigma_\tau^2) - \mu_\tau}{\sigma_\tau} \tag{15}$$

The far-right side of this equation can be rearranged in terms of C by multiplying by $\frac{1/\mu_\tau}{1/\mu_\tau}$:

$$z = \frac{r_\tau - \mu_\tau}{\sigma_\tau} = \frac{\mathcal{LL}(\mu_\tau, \sigma_\tau^2) - \mu_\tau}{\sigma_\tau} \left(\frac{1/\mu_\tau}{1/\mu_\tau}\right) \tag{16}$$

$$z = \frac{r_\tau - \mu_\tau}{\sigma_\tau} = \frac{\mathcal{LL}(1, C^2) - 1}{C} \tag{17}$$

Thus, the standard density function of R_τ depends entirely on the coefficient of variation (C) and admits no explicit dependence on time (t) or holding period (τ). This allows for a table of standard density curves to be constructed. The C value of R_τ selects the curve and the mean and variance metrics shift and scale the curve:

$$r_\tau = \mu_\tau + \sigma_\tau \times z|C \tag{18}$$

The tables of $z|\mathrm{C}$ form curves for the standard density of returns. Although there is a solution for the product of Laplace-distributed variables of zero mean (Nadarajah 2007), there is no known solution for the product of Laplace-distributed random variables with a non-zero mean. Lacking a known analytical solution, the product distribution needs to be simulated. Simulation method and a results table are detailed in Appendix A.

In order to highlight the time-independent stability of density matching, the observed z-scores (left side of Eq. 19) for returns on WTI spot price are plotted against the simulated z-scores of the density curve (right side of Eq. 19) in **Fig. 2**. A more detailed view of the density matching is presented in Appendix B.

$$\frac{r_\tau - \mu_\tau}{\sigma_\tau} = z|C \tag{19}$$

As τ increases, there is no trend in density-matching capability for the model, even though the mean and variance of the dataset change with τ. It is immediately evident that given the mean and variance of the dataset, the density function can be reliability reproduced, and that the holding period (τ) associated with those mean and variance metrics is ultimately unimportant. Applying this mythology to returns on Johnson and Johnson (JNJ) stock showed similar performance (**Fig. 3**).

Along (-2,-2) to (2,2) in Fig. 2, the function is approximately linear; the slight curve at the ends imply *slightly* heavier tails than the fitted distribution. One could adjust variance on the fitted distribution to bring the approximately linear region to a slope of one, while causing greater error in the tails. This effect is small and was not present on the fit to JNJ stock price in Fig. 3. It would be interesting to know if exclusion of the 2008 price rush and subsequent crash resulted in the slightly heavier tales, but this was not investigated.

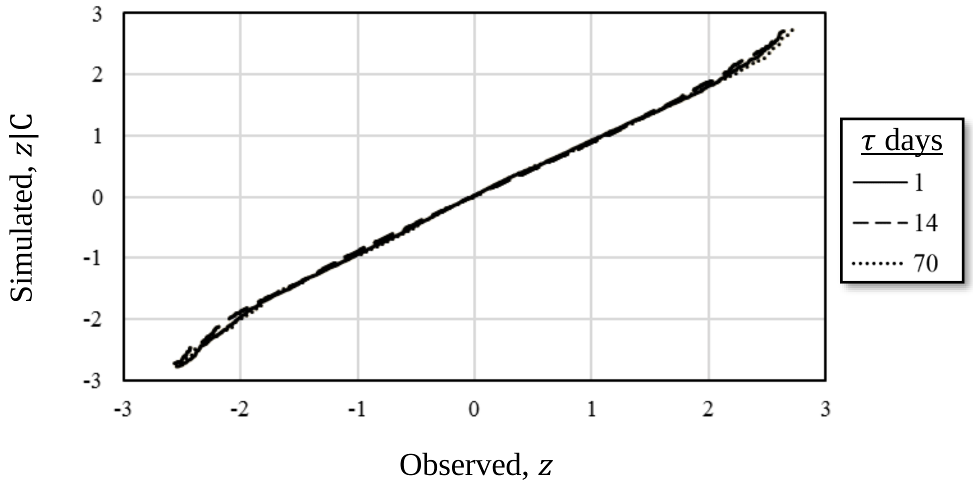

Fig. 2—ZZ plot of WTI spot-price returns using observed mean and variance. Inner 98% of data shown. January 2, 1986, to December 31, 2015.

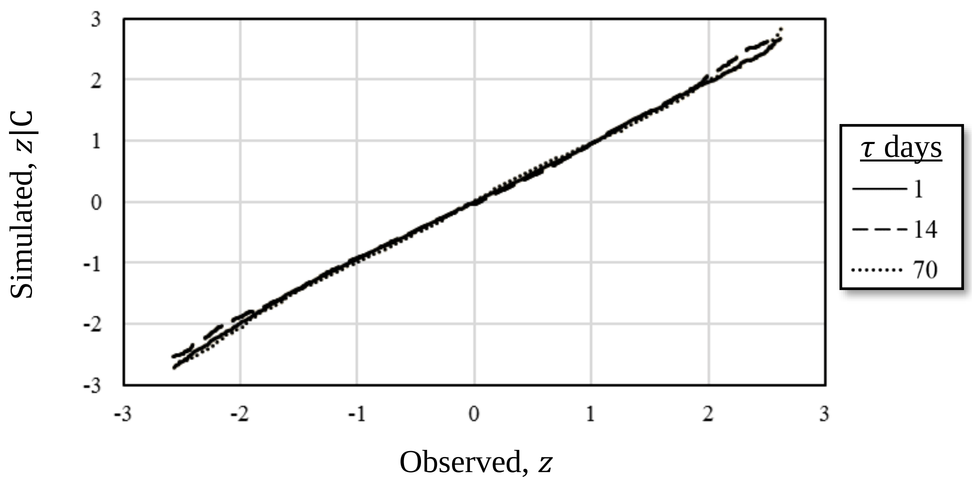

Fig. 3—ZZ plot of Johnson and Johnson stock returns using observed mean and variance. Inner 98% of data shown. January 2, 1986, to December 31, 2015.

Analysis

In the early exploration of the data, prior to choosing the double-Laplace distribution as a descriptor, the set R_1 (collection of all observed one-day returns on WTI spot price) was fit with multiple distributions using the Excel add-in @Risk (v. 7.5). Based on AIC rank criterion, Hyperbolic Secant (HS) distributions fit the best, with the Laplace distribution second best out of 33 tested distributions. However, the HS distribution does not have skew, and the data set does (skewness = 0.2 at $\tau = 7$ days, and 0.6 at $\tau = 70$ days). In this chapter, the Laplace distribution was used to generate the double-Laplace distribution, which produces skew. With nothing but a substitution of notation, the equations in this chapter are valid for a double-HS distribution. In both the double-Laplace and double-HS distributions, two values are sampled and multiplied together.

To see how the double-Laplace and double-HS distributions compare in performance to a Laplace and HS distribution, a goodness-of-fit test was performed on the inner 98% of data using the coefficient of determination at a τ of 1, 7, 14, 35 and 70 days. In all cases tested on both WTI spot price (**Fig. 4**) and Johnson and Johnson closing price (**Fig. 5**), the double-Laplace distribution outperformed the Laplace distribution. The double-Laplace distribution outperformed the double-HS distribution for low values of τ, but as τ increases the difference in performance was minimal. Interestingly, the double Laplace was *extremely* similar in goodness-of-fit to the HS distribution. The HS distribution does not have skew, but the double Laplace does, and thus the double-Laplace distribution was chosen due to underlying data for WTI (and JNJ) having skew.

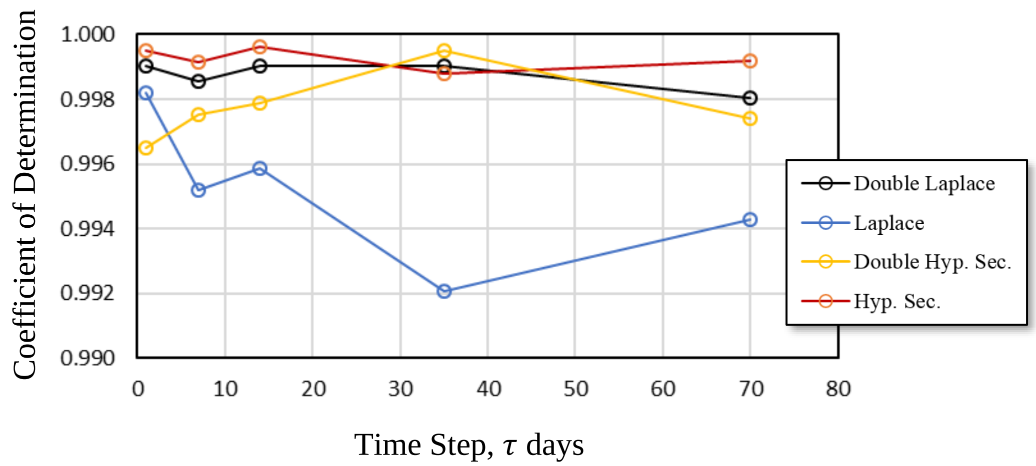

Fig. 4—Coefficient of determination for the match of several distributions fit to the observed returns, R_τ, on WTI spot price with τ values of 1, 7, 35 and 70 days.

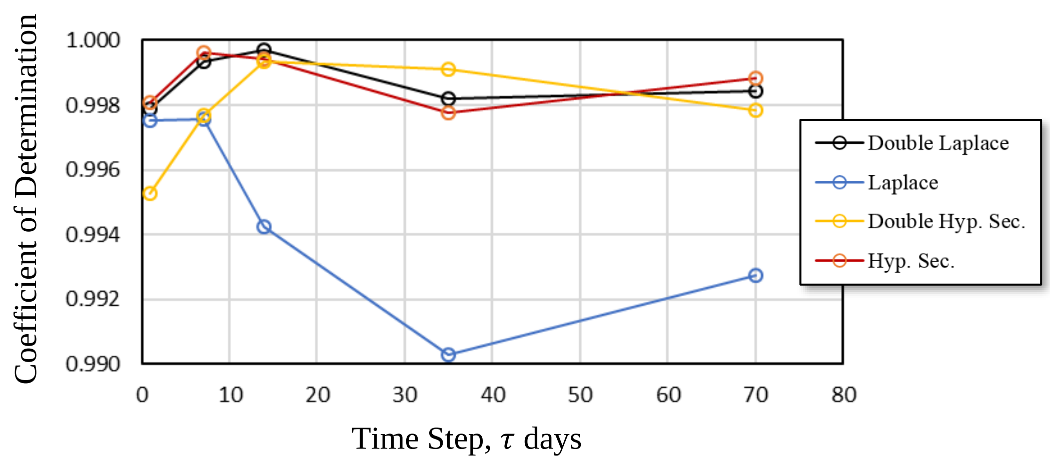

Fig. 5—Coefficient of determination for the match of several distributions fit to the observed returns, R_τ, on Johnson and Johnson closing price with τ values of 1, 7, 35 and 70 days.

CHAPTER III

TIME SCALE AND INTRA-WEEK BEHAVIOR

In the previous chapter, the standard density function of returns was developed as a function of the coefficient of variation. The next step is to examine the definition of time (t) on which returns are analyzed. In the previous chapter, t was used to define a particular date (i.e., 1/1/2000). For convenience in the equations, t can also be used as an elapsed time from a particular date. If the earliest date in the set is defined as $t = 0$, then t can represent the number of days after the earliest date in the set. Doing this does not change the math, results, or conclusions in the previous chapter, but it does make t easier to work with. Time t relates clearly to τ, the length of a holding period in days. If one were to hold between $t = 1$ and $t = 5$, then $\tau = 4$ days. The defined time scale used in this thesis on which τ is calculated departs from methodology common in literature. This departure is essential to separate day-to-day price volatility from long-term uncertainty in what oil is worth. This chapter explores that departure, and its implications for analysis of returns.

It is seductive to collect all of the changes that occur between adjacent spot prices and fit the data with a distribution. However, if one spot price is on a Friday and the other on the following Monday, there may be more uncertainty than the change between Monday and Tuesday. Recent papers on modeling WTI spot prices have not accounted for this effect, even though intra-week behavior has been known to exist for a while (Maberly 1995). For instance, Sadorsky (2006) reports 3,910 observations for 1-day returns on WTI spot price between February 5, 1988, and January 31, 2003, which would require direct comparison of the 3,911 prices. Kang, Kang, and Yoon (2009) reported 3,413 observations, when there are only 3,128 days not followed by a weekend (January 6, 1992, to December 29, 2006). Wei, Wang, and Huang (2010) report 4,474 observations when there are only 3,756 days that are not followed by a weekend (January 6, 1992, to December 31, 2009). None of these papers accounted for intra-week behavior in the handling of their data, which may impact their results. These papers were selected as examples due to their analysis of WTI spot prices, but this type of analysis is prolific in literature on other markets. It is common in literature to define t as the indexing of trading days (rather than calendar days), and this results in an improper calculation of τ.

Consider this: if you were to throw a dart at a calendar, the odds of landing on a Tuesday are lower than the combined odds of landing on a Saturday, Sunday, and Monday. Similarly, oil may have perceived value changes on Saturday, Sunday, and Monday. Any events that occurred over the weekend get priced on Monday in WTI spot price. This means that Mondays do not see the same set of possible changes as the rest of the weekdays. An event changing traders' opinions of asset evaluation can manifest any day of the week, even the weekend. A plane crash, a war, an embargo, or the death of a highperforming executive can all occur on weekends and materially impact the perceived value of an asset.

When trying to define the set of returns over one *trading* day of holding, the return from Friday to Monday is not describing a return for $\tau = 1$ day, but rather $\tau = 3$ days. The return from Monday to Tuesday is for a holding period of $\tau = 1$ day. They are not equivalent elements in the same set. The issue is, holding from Monday to Tuesday is not the same as holding from Friday to Monday, or in the event of a holiday, Friday to Tuesday. While the number of trading days may be the same, the amount of time elapsed is not. All have different risks and must be considered elements of separate sets. Problematically, this means the set $r_{\tau=1}$ cannot contain returns over the weekend and a solution will be developed in this chapter.

Analysis methods that use holding periods (τ) larger than or equal to one week might not experience this issue. An example is the analysis of Akilu, McVay, and Lee (2006) in which they forecasted prices using bootstrapped data. The authors perform analysis with a monthly holding period, avoiding the common pitfall in relating relative position with calendar position. However, this method is based on the observed price timeline, which has discrete and finite variability of returns. *Possible* returns are continuous in possible value (not discrete) and can take on an infinite number of values (not finite).

Day-of-the-week effects are not new to literature; others have noted that they do exist (Charles 2010; Yu, Chiou, and Jordan-Wagner 2008). In spite of that, if returns are considered constant across the calendar like interest, the total return from Friday to Monday is still not the same as Monday to Tuesday. Hence, spot prices are placed on the basis of time, t, rather than relative position in the list, i. Plotting the mean and variance of returns against calendar-day holding period elapsed (τ days) shows clear oscillatory patterns (**Figs. 6** and **7**) corresponding to one week.

Holding periods that do not correspond to increments of one week force the exclusion of data points from analysis. When a holding period of one day is applied, Friday is excluded because Saturday is an empty element; the return from Friday to Saturday cannot be explicitly measured. A holding period of two days excludes Thursday and Friday because Saturday and Sunday are empty elements. Until a holding period of 7 days is used, the empty elements on the weekend force exclusion of data (like Friday being excluded in $r_{\tau=1}$). If all trading days were the same, the exclusion of some portions of the week would not impact the data. However, if they are not all the same, exclusion will cause perturbations in the trend of mean and variance metrics as seen in Figs. 6 and 7. This chapter details how to handle these effects and how to extract useful data for use in a price model.

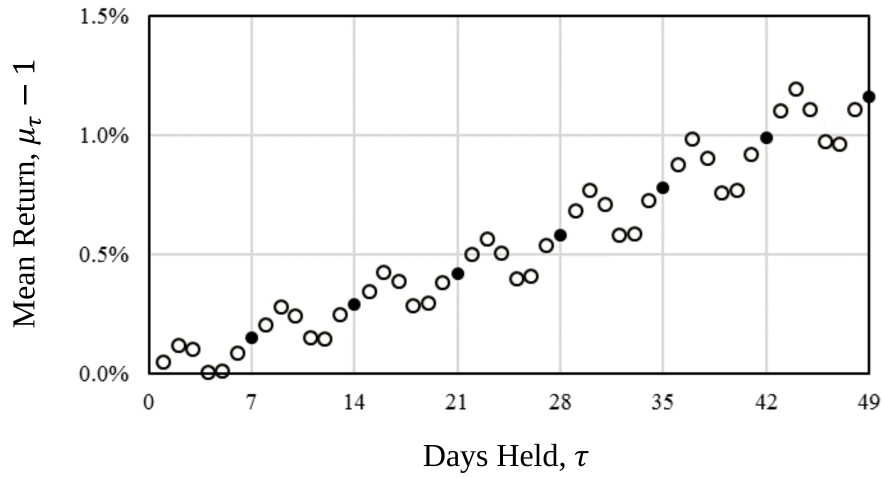

Fig. 6—Mean calendar-day returns for a given holding period, $\boldsymbol{\tau}$. Every 7 days is marked with a solid circle.

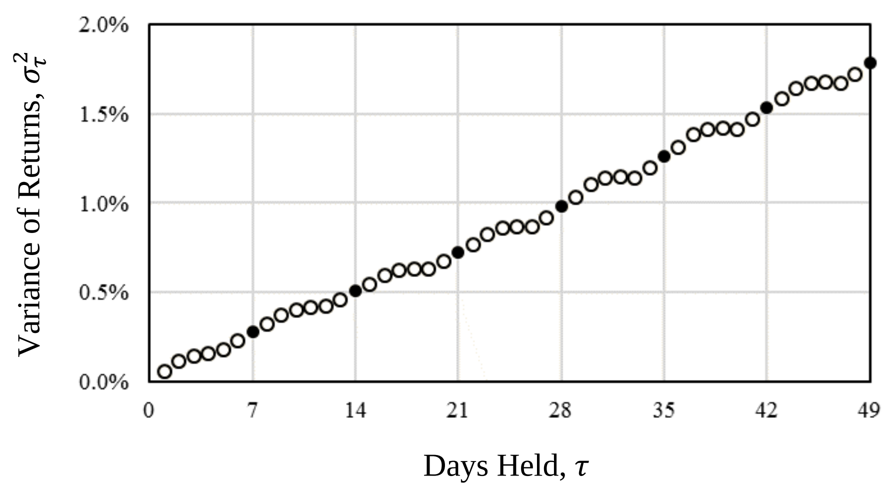

Fig. 7—Variance of calendar-day returns for a given holding period. Every 7 days is marked with a solid circle.

Incremental Returns

A single observed element r_τ is formed by N sequential returns with holding periods less than τ. Incremental returns having a holding period of $\Delta\tau$ days form a return having a holding period of τ days by the following relations:

$$r_\tau = \prod_{n=1}^{N=\tau/\Delta\tau} r_{\Delta\tau,n} \tag{20}$$

If holding period of incremental returns is one day ($\Delta\tau = 1$ day), and the holding period of the total return is three days ($\tau = 3$ days), then the incremental returns would be from R_1. For example:

$$r_3 = \frac{s_1}{s_0}\frac{s_2}{s_1}\frac{s_3}{s_2} = r_{\Delta\tau,1}r_{\Delta\tau,2}r_{\Delta\tau,3} \tag{21}$$

When looking at a specific element in R_τ, the values $r_{\Delta\tau,n}$ have specific values as well, they are *not* randomly sampled from $R_{\Delta\tau}$. In other words, any observed return with a holding period τ is formed by one or several observed "incremental" returns with holding period $\Delta\tau$ that are not randomly sampled. The set $R_{\Delta\tau=1}|\tau$ can be viewed as R_τ "unitized." $R_{\Delta\tau=1}|\tau$ is the definition of the smallest incremental change that can be used to define R_τ. This unitized set of returns will be useful in analyzing day-of-the-week effects present in the data.

Intra-Week Trends

Mean values of the average daily variance, $R_{\Delta\tau=1}|\tau$, are certainly geometric (multiplicative); thus, the average daily mean is the τ^{th} root of μ_τ. Variance is less obvious. There will be correlated short-term behavior on the scale of days to weeks that needs to be separated from long-term behavior. To determine when this occurs, Eq. 6 was sequentially chained and the uncorrelated daily variance needed to produce the variance after a holding period τ was calculated iteratively. This is not implying that the variance is uncorrelated; it is determining the variance *implied* if no correlation is *assumed*. These normalized metrics are the average values for daily returns over the course of the holding period:

Avg. daily mean of R_τ: $$\sqrt[\tau]{\mu_\tau} \tag{22}$$

Avg. daily variance of R_τ: $$\bar{\sigma}_\tau^2 \tag{23}$$

These values would be the mean and variance of $R_{\Delta\tau=1}|\tau$. Average daily mean is easily contextualized as the average gain each day required to reach the total gain after holding the asset for τ days. Average daily variance is best understood by drawing a line between the origin and a data point in Fig. 7. The slope of this line is similar to the average daily variance and can be used as an "almost right" conceptual tool to understand the math. However, the slope of this line is not *exactly* the value of $\bar{\sigma}_\tau^2$. Per Eq. 6, we know that the mean impacts variance. To account for deviations in mean between time steps seen in Fig.

7 and obtain the exact value of $\bar{\sigma}_\tau^2$, the mean value of each step needs to be accounted for. For the first time step, $\bar{\sigma}_1^2$ is the same as σ_1^2. For the second time step, $\bar{\sigma}_2^2$ is the value that satisfies:

$$\sigma_2^2 = \bar{\sigma}_2^2\bar{\sigma}_2^2 + \bar{\sigma}_2^2\mu_1^2 + \bar{\sigma}_1^2\frac{\mu_2^2}{\mu_1^2} \tag{24}$$

For the third time step, let x be a temporary variable:

$$x = \bar{\sigma}_3^2\bar{\sigma}_3^2 + \bar{\sigma}_3^2\mu_1^2 + \bar{\sigma}_3^2\frac{\mu_2^2}{\mu_1^2} \tag{25}$$

In terms of x and the average daily variance, $\bar{\sigma}_3^2$, the observed variance σ_3^2 is:

$$\sigma_3^2 = \bar{\sigma}_3^2 x + \bar{\sigma}_3^2\mu_2^2 + x\frac{\mu_3^2}{\mu_2^2} \tag{26}$$

Minor differences between the value of the quotients (i.e., μ_2^2/μ_1^2 vs μ_3^2/μ_2^2) prevents the nested form of these equations at higher tier steps from collapsing to a simpler form, so the value of $\bar{\sigma}_3^2$ that satisfies the equation is solved for using Excel's Solver function. The value of the quotients could be assumed to be the same, but this would not be *technically* correct. By accounting for the mean values observed, the noise contribution of the mean to the variance has been removed from $\bar{\sigma}_\tau^2$. What little noise remains is actual variability of observed average daily variance, and any trending with time will be informative of how variance behaves as holding period increases.

In **Figs. 8** and **9**, the mean and variance values of $R_{\Delta\tau=1}|\tau$ are presented as a function of weeks where the horizontal axis is in increments of $\tau/7$, and data shown for every τ. The choice of axis in weeks was intentional to highlight the periodicity of the data, as every whole increment on the axis corresponds to one period. Understanding the trends in these figures is improved when the average daily mean and variance are shown in context of the 7-day holding period. Reframing the time in terms of k weeks and j days, $\tau = 7k + j$. The value of j specifies the position of the data point within the 7-day oscillations seen in Figs. 6 and 7. The value of j then specifies which days are force-excluded from analysis due to empty elements in the set (**Table 1**). In the example, a j value of one excludes Friday from the dataset, as there is no observed spot price on a Saturday that Friday can be compared to.

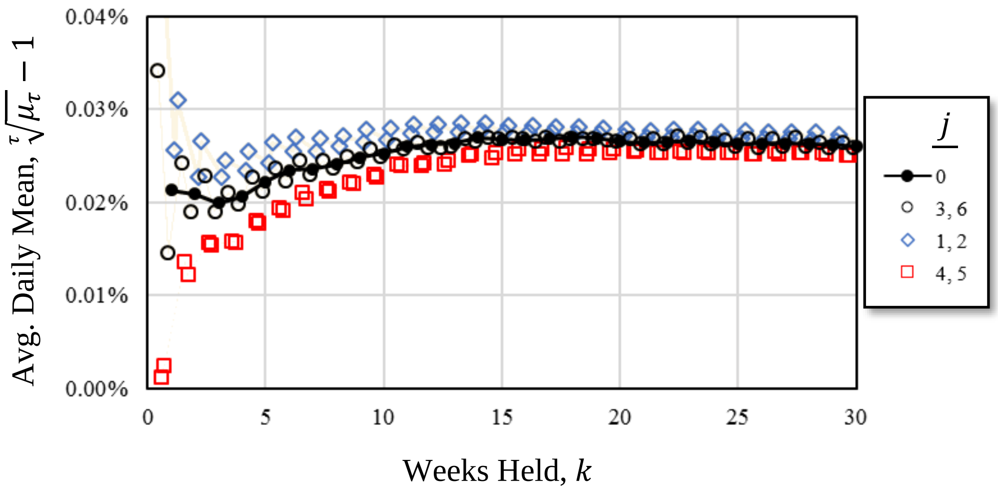

Fig. 8—Time-normalized mean ($\sqrt[\tau]{\mu_\tau}$) of R_τ, where $\tau = 7k + j$. Data grouped based on value relative to center line defined by $j = 0$ (black line).

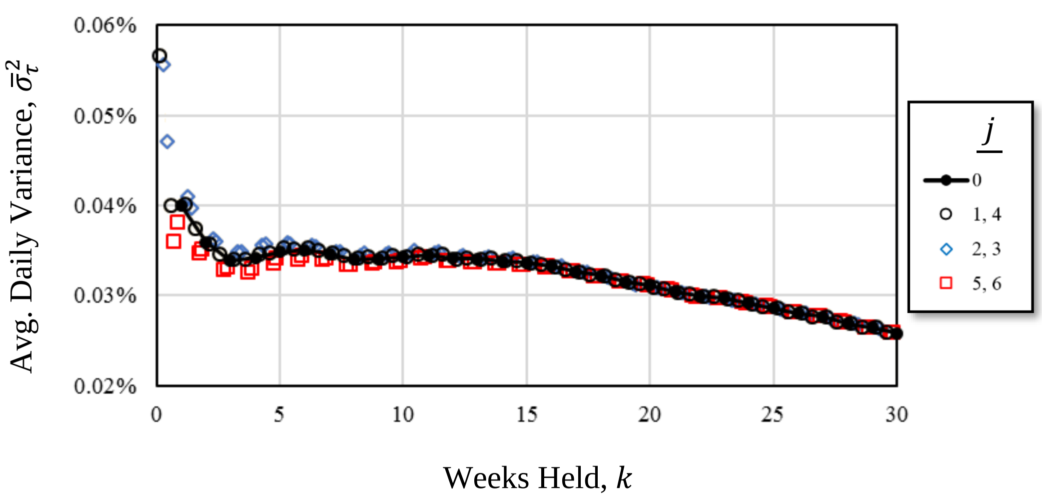

Fig. 9—Time-normalized variance ($\bar{\sigma}_\tau^2$) of R_τ, where $\tau = 7k + j$. Data grouped based on value relative to center line defined by $j = 0$ (black line).

$\boldsymbol{R_{7k+j} \equiv R_\tau}$			**Day of the Week**						
$\boldsymbol{j}$	$\boldsymbol{\sqrt[\tau]{\mu_\tau}}$	$\boldsymbol{\bar{\sigma}_\tau^2}$	**Su**	**M**	**T**	**W**	**Th**	**F**	**S**
0	$\mid$	$\mid$	$\times$	$*$	$*$	$*$	$*$	$*$	$\times$
1	$\mid + +$	$\mid$	$\times$	$*$	$*$	$*$	$*$	$\times$	$\times$
2	$\mid + +$	$\mid + +$	$\times$	$*$	$*$	$*$	$\times$	$\times$	$\times$
3	$\mid +$	$\mid + +$	$\times$	$*$	$*$	$\times$	$\times$	$*$	$\times$
4	$- - \mid$	$\mid$	$\times$	$*$	$\times$	$\times$	$*$	$*$	$\times$
5	$- - \mid$	$- \mid$	$\times$	$\times$	$\times$	$*$	$*$	$*$	$\times$
6	$- \mid$	$- - \mid$	$\times$	$\times$	$*$	$*$	$*$	$*$	$\times$

Table 1—Root-mean and normalized-variance for R_{7k+j} relative to the average behavior by R_{7k}, qualitatively assessed. Plus signs ($+$) indicate values above R_{7k}, dashes ($-$) indicate values below R_{7k}. Day of the week from which predictions are made are indicated by an asterisk ($*$), the rest of the days ($\times$) were force-excluded due to prediction to/from the weekend.

In context of Eq. 1 the formulation for $7k + j$ subscripts is:

$$r_{7k+j} = \frac{s_{t+7k+j}}{s_t} \tag{27}$$

$$r_{7k+1} = \frac{s_{t+7k+1}}{s_t} \tag{28}$$

When the value of t places the observation of s_t on a Friday, s_{t+7k+1} will always be a value from Saturday. Spot prices are not reported on Saturday and Sunday, and therefore r_{7k+1} can only be calculated on Monday, Tuesday, Wednesday, and Thursday. All values of j from one to six force exclusion of data due to the weekend. However, when j is equal to zero, the weekend does not force exclusion of any data:

$$r_{7k+7} = \frac{s_{t+7k+7}}{s_t} \quad (29)$$

In the case of Eq. 29, s_{t+7k+7} will be the same day of the week, and k weeks ahead of s_t. Analyzing in context of j value highlights intra-week behavior (Figs. 8 and 9) and produces a clear trend; when Thursday and/or Friday were force excluded from the dataset ($j = 1$ or 2) the mean daily returns were higher than when Monday and/or Tuesday ($j = 5$ or 6) were force excluded from the data set (Table 1). Returns on an asset purchased at the beginning of the week tended to be greater and had more variance than assets purchased at the end of the week! As τ increases, intra-week trends described collapse and converge on the center line of the periodicity $j = 0$. The magnitude of this trend collapses and values of $j \neq 0$ converge on values where $j = 0$.

Three distinct regions of average daily variance are present in Fig. 9. The first region is defined by the presence of a non-zero slope of data in the range $0 \leq k < 4$, which is indicative of correlation and collapse of the intra-week trends. The second region $4 \leq k < 13$ is approximately linear and shows no clear dependence on τ (recalling $\tau = 7k$). The final region $13 \leq k$ shows breakdown of the dataset. It is in this final region that there is not enough data to describe a continuous distribution. For an extreme example, if τ spans the entire timeline in the analysis, the variance of R_τ would be zero. This is because $R_{\tau=\min(t)-\max(t)}$ would have a single element and would no longer be a function of the path price took, and simply the start and end values of the path price took. When the value of τ places R_τ in the third region of Fig. 9, the sets no longer describe the possible paths, and begin describing the start and end of the price path. In the first region where $0 \leq k < 4$, the data describe parts of the path (intra-week effects) that do not really impact the final price for τ because they also contain information about the intra-week noise in price. This intra-week noise is unimportant, but at the same time the set R_τ used when building a model still needs to be descriptive of possible paths, and not just the start and end of the observed path.

Discussion

This chapter addresses day-of-the-week effects present in the data when constructing the definition of returns. When grouping data into R_τ, the elements must be from the same set, therefore, it is inappropriate to group Monday-to-Tuesday returns with Thursday-Friday returns. The returns have different behaviors; these elements come from fundamentally different sets. By correctly modeling time scale when grouping data into R_τ, day-of-the-week trends become apparent (Fig. 6 and 7). However, grouping Monday-to-Monday returns with Thursday-to-Thursday returns *is* appropriate because the returns have the same value for τ.

In order to better analyze the data in context of data being excluded due to the weekend, the metrics for R_τ were "unitized" to the sets $R_{\Delta\tau=1}|\tau$. Based on the variance values for $R_{\Delta\tau=1}|\tau$ shown in Fig. 9, if someone were to purchase oil and hold for one week their average daily variance ($\mathbb{V}(R_{\Delta\tau=1}|7) \cong 0.0004$) is less than if they were purchase and sell the asset seven random times during the year ($\mathbb{V}(R_{\Delta\tau=1}|1) \cong 0.00055$). The value of τ clearly has an impact on the average daily variance experienced while holding the asset for τ days. As holding period increases in length, this effect dampened and loses influence relative to the total change in price. When τ is between 28 and 91 ($4 \leq k < 13$), the set's $R_{\Delta\tau=1}|\tau$ are close in variance. Variance collapses to a constant value in this region. If observed returns were not correlated as assumed in this chapter, the variance of holding for seven days would be the same as holding for seven random days during the year. This is not correct; therefore, the incremental returns *are* correlated, and this correlation is related to intra-week trends as seen in the periodic behavior centered on $\tau = 7k$. Other correlation may exist, but the intra-week effect has been specifically observed in this chapter.

Although interesting, the region of constant variance will not be useful for determining metrics used in the model. Despite variance collapsing to a constant value in Fig. 9, the mean values in Fig. 8 still have periodic behavior. No matter how great τ becomes, a j value of two will always exclude Thursdays and Fridays from R_τ, biasing the mean to a higher value.

Days of the week excluded from analysis matter, and sets $R_{\tau=7k+j}$ cannot be compared to each other if they have different values of j. The sets are different because they do not contain the same underlying data. Simply increasing τ cannot eliminate this. Furthermore, increasing τ causes the set R_τ to provide less information on the path the price took, and more information on the relationship between the starting and ending prices. This is understood by taking the extreme value of τ that spans the entire timeframe observed, resulting in a single element in R_τ. That single element would be calculated by the most recent price divided by the first price. There would be no other values in this list, as no other prices could be compared for this maximum value of τ. The set would contain no information on *how* the price got from the first value to the last value. This would be equivalent to using the slope of a straight line between endpoints of a curve to define the derivative, rather than tangent lines. The closer the analysis can get to a $\tau = 1$, the better the metrics will describe the price path, and not the price outcome. Yet, as τ goes to one, the intra-week trends become important. This presents a conundrum, where neither the minimum value of τ nor the maximum value of τ are informative for the desired long-term-forecasting model.

During the analysis of this chapter, it was not assumed that $R_{\Delta\tau=1}|\tau$ had a double Laplace distribution. Multiplication or division of randomly sampled elements of two distributions will not necessarily reproduce the same standard density (curve shape) as the parent distributions. This is a property reserved for stable distributions (Mandelbrot 1960), and the double Laplace distribution is not. This chapter exploited the fact that child distribution $R_{\Delta\tau=1}|\tau$ will have predictable mean and variance. The density function of $R_{\Delta\tau=1}|\tau$ is unimportant for this chapter. Whatever mean and variance is chosen for modeling changes in price, the distribution will ultimately need to have a double-Laplace distribution. The next chapter addresses this issue.

What is useful about the analysis in this chapter is the definition of $R_{\tau=7k+j}$, and understanding that the periodic behavior of the data is centered on holding periods that are multiples of seven days. Assumptions of non-correlation in the math allowed the correlated behavior to be demonstrated and analyzed. It allowed for the conclusion that sets of $R_{\tau=7k+j}$ for which j are the same *are* comparable. When j is zero, $R_{\tau=7k}$ does not exclude weekends. All of this makes $R_{\tau=7k}$ the most informative group of sets, and the only sets that are not biased by intra-week trends. Because lives of hydrocarbon wells are on the scale of years, the intra-week behavior is unlikely to be of any importance. It would be better to model weekly holding periods where the intra-week problem is avoided. Analysis of the sets that satisfy $R_{\tau=7k}$ is performed in the following chapter to obtain mean and variance metrics suitable for a random-walk model.

CHAPTER IV

INTER-WEEK TRENDS AND A RANDOM WALK MODEL

In the previous chapter, the time scale was defined by calendar-days elapsed rather than relative juxtaposition of prices. A justification for calibrating functions to inter-week behavior and assuming them to be approximate for intra-week behavior was also proposed. The analysis assumed that sequential returns were not correlated in order to show that they were correlated. The current chapter will accept that correlation exists and will not assume that it is absent in the data. The definitions in the previous chapters used τ. In the current chapter, analysis is performed in context of weeks held, k. Recalling Eq. 27:

$$r_{7k+j} = \frac{s_{t+7k+j}}{s_t}$$

using sets of R_{7k+j} that have the $j = 0$ are the most informative because they do not exclude any day of the week due to the presence of the weekend. This definition will allow for the analysis of inter-week trends and calculation of mean and variance metrics for use in a random-walk model.

Calibration of Inter-Week Trends

When analyzing the behavior of R_τ, empty elements on τ cannot be predicted from or to. When j is equal to zero, the holding period will never end on the weekend, and the problem weekends present is averted. An additional problem arises where recent price values have no prices to predict to. Although it would be tempting to use the mean and variance of R_7 if the holding period of the model is one week, it would be prudent to consider the impact of recent spot prices on the metrics of R_7. For example, when comparing a 7-day change, the most recent 6 prices have no subsequent value with which to compare them; when looking at a 14-day change, 13 of the most recent prices have no subsequent value with which to compare. This causes variations in the datasets as holding period is increased. Although the effect is slight, it may obfuscate analysis of the data. Even though the problem of weekends is resolved by using holding periods of k weeks and $j = 0$, the data sets are still not *exactly* comparable. Larger values of k still cause data to be excluded at the start and end of the series of spot price.

To establish the contribution of these effects to the observed trend in mean between sets R_{7k}, the data must be broken down into sets that differ only by the excluded data points. In the previous chapter, the mean and variance of $R_{\Delta\tau=1}|\tau$ were calculated as inferred average-daily metrics for R_τ. In this chapter, the observed returns that created every element in $R_{7(k+1)}$ from R_{7k} need to be obtained. To do this, Eq. 21 is reformulated in terms of τ:

$$r_{\tau} = \prod_{n=1}^{N=\tau} r_{1,n} \tag{30}$$

Next, τ is put in terms of k weeks:

$$r_{7k} = \prod_{M=1}^{M=k} r_{7,7M-6} \tag{31}$$

Appendix C presents detailed position of elements $r_{\tau,n}$ relative to the calendar. Then for a return after 14 days ($k = 2$):

$$r_{14} = r_{7,1} r_{7,8} \tag{32}$$

Each observed value of r_{14} would have two known values in R_7 that created it. This means an element of R_{21} would be composed of three elements from R_{21}, or one element of R_{14} and one element of R_7:

$$r_{21} = r_{7,1} r_{7,8} r_{7,15} \tag{33}$$

$$r_{21} = r_{14,1} r_{7,15} \tag{34}$$

More generally:

$$r_{7(k+1)} = r_{7k,1} r_{7,7k-6} \; \forall \; k > 0 \tag{35}$$

The question is, how much does the subset of elements from R_7 used to transform R_{14} to R_{21} differ from the subset used to go from R_{21} to R_{28}? More generally, how much do the metrics of the subset of elements from R_7 used to go from R_{7k} to $R_{7(k+1)}$ depend on k? The only difference between these sets is due to data-exclusion issues from holidays and the most recent spot prices. This cannot be done by comparing the metrics between sets of R_{7k} if a subset of R_7 is going to be obtained. The analysis must be performed on the elements of each set $R_{7(k-1)}$ and R_{7k}. Let δr_{7k} represent the element from R_7 that changes r_{7k} to $r_{7(k+1)}$:

$$\delta r_{7k} \; = \frac{r_{7(k+1)}}{r_{7k}} \tag{36}$$

$$\delta r_{7k} \in \delta R_{7k} \sim D(\delta\mu_{7k}, \delta\sigma_{7k}^2) \tag{37}$$

The set δR_{7k} will be a subset of R_7 for all values of $k > 1$, and will identical to the set R_7 for $k = 1$. The "δ" notation on δR_{7k} is useful for denoting that this is the subset that defines the change from R_{7k} to $R_{7(k+1)}$. The set δR_{7k} is the observed set of incremental returns for the week between R_{7k} to $R_{7(k+1)}$. Importantly, the set δR_{7k} *will* be correlated to R_{7k}. It will be the exact set of elements that changed R_{7k} to $R_{7(k+1)}$, and will be a subset of R_7.

In Eqs. 33 and 35 it was shown that $r_{7(k+1)}$ is formed by one element from R_{7k} and one element from R_7. Notably, the element $r_{7,k}$ is not used to generate the element $r_{7k,1}$. Thus, the ratio of observed elements from $r_{7(k+1)}$ and r_{7k} recreate a unique element observed in R_7. If the second subscript continues to be used to designate position in the set, this redundancy is easy to identify. As an example, the first element of δR_{14} is defined by the ratio of the first element of R_{14} and R_7, and it is the same as the 8th element of R_7 (i.e. $\delta r_{14,1} = r_{14,1}/r_{7,1} = r_{7,8}$). Hence, the set R_7 is the same as δR_7, and contains all the other sets δR_{7k}. Means of the δR_{7k} sets are plotted in **Fig. 10** and the *c.v.* is plotted in **Fig. 11**. In Figs. 10 and 11, each data point represents the mean or *c.v.* of the set of seven-day-holding-period returns (δR_{7k}) that transformed R_{7k} into $R_{7(k+1)}$.

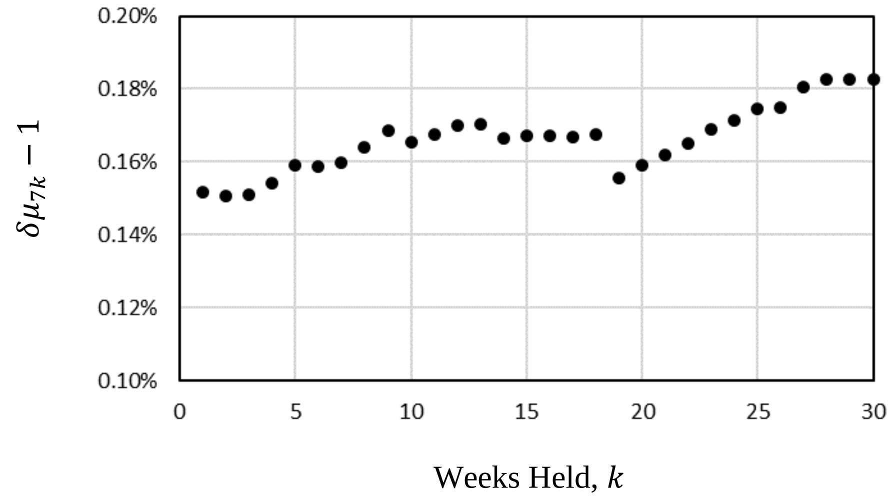

Fig. 10—Mean value of δR_{7k}, the set of returns that generates $R_{7(k+1)}$ from R_{7k}.

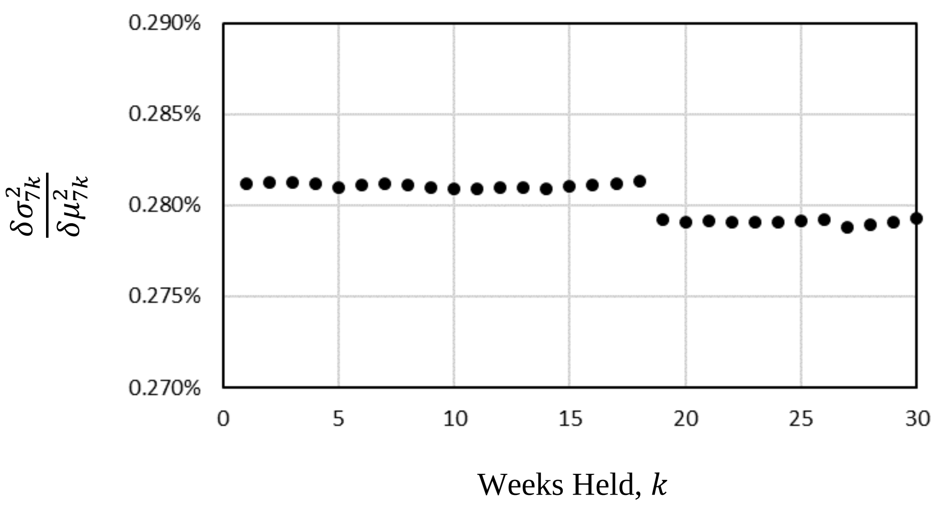

Fig. 11—Square coefficient of variation of δR_{7k}, the set of returns that generates $R_{7(k+1)}$ from R_{7k}.

Fig. 10 is best interpreted as the gain for each sequential week of holding. The first point represents the gain over the first week, the second point represents the gain from the end of week one to the end of week two. Fig. 11 can be understood similarly, where each point is that week's contribution to the coefficient of variation. For both figures, there is no expectation of dependence on k. Any variation on k is a result of recent prices being excluded from the dataset. For example, the change in variance (Fig. 11) between $k = 18$ and 19, is a result of a sharp change in spot price in August 2015 (18 to 19 weeks before December 31, 2015).

Recalling that it could not be assumed that $R_{\Delta\tau}|\tau$ had a double-Laplace distribution like R_τ, it can now be concluded that δR_{7k} *does* have the same density function as R_7. It was proposed in this thesis that the dataset R_7 comes from a double-Laplace distribution. It is reasonable to expect that the most recent spot prices have experienced random returns selected from the distribution that describes R_7. Therefore, if these random points are removed from R_7, the double-Laplace distribution is still the descriptor of the new subset. δR_{7k} is only a smaller observation of potential outcomes from the double-Laplace distribution. By looking at the dependence of δR_{7k} on k, the impact on R_7 of the most recent spot prices can be observed. Analysis of the dependence of δR_{7k} on k provides a quantitative view of how sensitive those metrics are to recent changes in price (Figs. 10 and 11).

As k increases, more data points are excluded from the most recent spot prices. For every increase in k, the exclusion of recent prices causes $\delta\mu_{7k}$ to change in value. By this method the impact of recent prices on $\delta\mu_{7k}$ can be observed. The most recent spot prices have a significant impact on the mean value of $\delta\mu_{7k}$. As k is increased and more of the recent prices are excluded, the value of $\delta\mu_{7k}$ increases (Fig. 10). This is due to the recent prices in the list trending downward. If the average of $\delta\mu_{7k}$ over the analyzed range of k is used, additional spot prices being observed will have less of an impact on the mean value used in the model. Using the average value of $\delta\mu_{7k}$ over $1 \leq k \leq 30$ as modeling the set that represents a holding period of one week:

Weekly Holding $$\overline{\delta\mu_{7k}} = 1.001661482 \tag{38}$$

The average value of $\delta\sigma_{7k}^2/\delta\mu_{7k}^2$ over $1 \leq k \leq 30$ is:

Weekly Holding $$\overline{\frac{\delta\sigma_{7k}^2}{\delta\mu_{7k}^2}} = 0.002802865 \tag{39}$$

Monthly Holding Periods

Because the financial model to which WTI price data are applied is in time steps of one month, it is beneficial to obtain the values of the mean and variance for a holding period of equal length ($\tau = 30.4375$ days). Recalling that the definition of holding period used in this chapter is seven days, mean and variance metrics will be approximate if τ is not an increment of seven. All weekly trends will oscillate as noise around this model, and on the scale of years and decades for which this model will be used, these trends are not important. To change the mean value in Eq. 38 to a monthly (rather than weekly) holding period, it must be scaled geometrically. Mean values for returns combine geometrically, such that:

$$\mu_\tau = \mu_{\Delta\tau}^{\tau/\Delta\tau} \tag{40}$$

Extending $\overline{\delta\mu_{7k}}$ to a holding period of one month, the mean value becomes:

Monthly Holding $$\delta\mu_{30.4375m} = \overline{\delta\mu_{7k}}^{30.4375/7} = 1.007244601 \tag{41}$$

Scaling the weekly variance to 1, 2, 3, 4, and 5 weeks using the product variance formula similar in form to Eq. 6 produces a linear trend (not shown, coeff. of det. = 0.999997). Interpolating a value for 30.4375 days, or k = 4.34821 weeks yields:

Monthly Holding $$\frac{\delta\sigma_{30.4375m}^2}{\delta\mu_{30.4375m}^2} = 0.012245498 \tag{42}$$

Values for mean and *c.v.* for the model are summarized in **Table 2.**

Metrics by Holding Period			
Period	$\Delta t =$	$\mu_{\Delta t} - 1$	$\sigma_{\Delta t}^2/\mu_{\Delta t}^2$
Week	k	1.6615 E-3	2.8029 E-3
Month	m	7.2446 E-3	1.2245 E-2

Table 2—Mean and variance metrics of returns, based on the chosen holding period. Data based on inter-week behavior.

Price Model

Up until this point, time (t) has been indexed to the earliest date in the data set. To make future predictions, the starting (current) date is indexed as $t = 0$. With this indexing, t represents how far into the future the forecasted price is. This doesn't change the math, it simply changes the initial point in time for the model. To get to some time t in the future, sequential holding periods of length Δt are used. Previously, days (τ), and weeks (k) were used for Δt. At this point, the difference between t and holding period disappears. Holding one year into the future ($t = 1$ year), can be modeled as being made of 52 sequential holding periods of one week ($\Delta t = 1$ week). Similarly, holding one week into the future (t = 1 week), can could be modeled as 7 sequential holding periods of one day ($\Delta t = 1$). In terms of t:

$$r_t = r_{t-\Delta t} r_{\Delta t} \tag{43}$$

$$r_t = r_{t-\Delta t}(\mu_{\Delta t} + \sigma_{\Delta t} z) \tag{44}$$

An expanded form of Eq. 44 for $t = 3$ weeks, and $\Delta t = 1$ week would be:

$$r_{3k} = (\mu_k + \sigma_k z_1)(\mu_k + \sigma_k z_2)(\mu_k + \sigma_k z_3) \tag{45}$$

Each value of z would be independently sampled from the double-Laplace distribution (subscript added on z to denote each sample). In this case, three returns would be generated, and this could be used to obtain a price at the end of each week:

Week 1 $$\frac{s_1}{s_0} = r_{1k} = (\mu_k + \sigma_k z_1) \tag{46}$$

Week 2 $$\frac{s_2}{s_0} = r_{2k} = r_{1k}(\mu_k + \sigma_k z_2) \tag{47}$$

Week 3 $$\frac{s_3}{s_0} = r_{3k} = r_{2k}(\mu_k + \sigma_k z_3) \tag{48}$$

Multiplying r_t by some initial price, s_0, would produce a price at some time, t, in the future. The path price takes between s_0 and s_t would be made up of sequential holding periods of length Δt. These sequential holding periods of length Δt are the "time steps" on which the model forecasts price.

Reforming Eq. 44 while preserving the *c.v.* used for the standard density distribution Z (of which z is an element):

$$r_t = r_{t-\Delta t}\mu_{\Delta t}\left(1 + \frac{\sigma_{\Delta t}}{\mu_{\Delta t}} z\right) \tag{49}$$

For sequential holding periods of one week ($\varDelta t$ = k) or one month ($\varDelta t = m$), the metrics used for $\mu_{\varDelta t}$ and $\sigma_{\varDelta t}^2/\mu_{\varDelta t}^2$ are summarized in Table 2. From Chapter II, the standard variate z equals:

$$z = \frac{\mathcal{L}\mathcal{L}\left(1, \frac{\sigma_{\varDelta t}^2}{\mu_{\varDelta t}^2}\right) - 1}{\frac{\sigma_{\varDelta t}^2}{\mu_{\varDelta t}^2}} \qquad (50)$$

The variable z is a random element generated by the distribution function defined in Chapter II. It adds randomness to each change in price. Outcomes of the model defined by the density function z in Eq. 50, and the propagation sequence defined by Eq. 49 were simulated with the Excel add-in @Risk (v. 7.5). Percentile outcomes of price changes for 40 years (**Fig. 12**) were simulated with time steps ($\varDelta t$) of one month (values from Eqs. 41 and 42). One million iterations of Latin-hypercube sampling were used to populate the distribution of possible price paths.

When a single run of the model is made (rather than one million), there will be a price for each timestep $\varDelta t$ between 0 and t. This price timeline can be used as a single predicted price timeline. Each price outcome of the model (examples in **Fig. 13**) functions as a price prediction where the next price (s_t) is related to the current price ($s_{t-\varDelta t}$) by:

$$s_t = s_{t-\varDelta t}(\mu_{\varDelta t} + \sigma_{\varDelta t} z) \qquad (51)$$

In Fig. 13, five price timelines predicted as possible by the model are graphed alongside the observed prices for WTI spot price. In this case, the prices were forecasted using the starting price for the WTI dataset ($25.56/bbl). To construct price timelines predicting into the future, the initial $s_{t-\Delta t}$ in Eq. 49 would be the current price of oil.

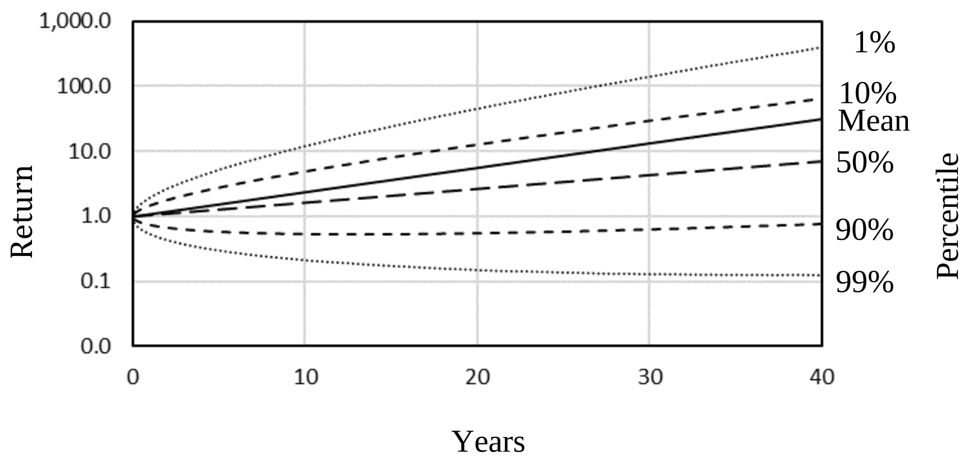

Fig. 12—Forty-year forecast of possible returns on WTI spot price. $\Delta t = 1$ month.

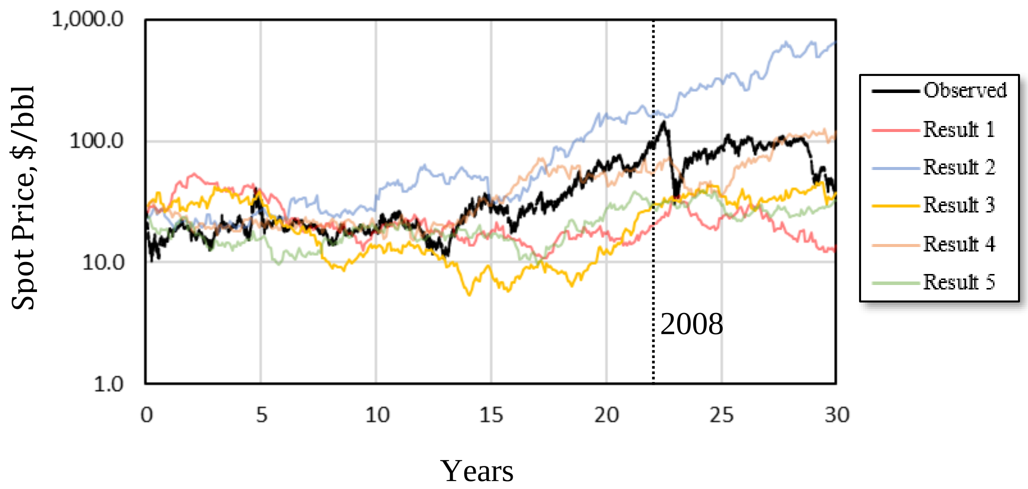

Fig. 13—Observed WTI spot prices on t (black line) and random model outputs (colored lines). A time step of $\Delta t = 1$ month was used for simulation. Time span shown starts January 2, 1986, and ends December 31, 2015.

Discussion

When intra-week trends were explicitly removed by a change in definition of returns from a time step of one day to a step of seven days, the oscillations having an interval of one week disappear (Figs. 10 and 11). Specifying the day of the week the model is predicting from and to may serve to clarify the fact that returns are expected to be higher or lower depending on the day of the week (Table 1). However, using time steps of seven days to obtain mean and variance metrics, as done in this chapter, avoids this issue altogether. As shown in the previous chapter, the periodicity of the intra-week behavior is centered on the inter-week trend. Thus, if the inter-week behavior is defined, it will be approximate for intra-week behavior.

No price-shock events (such as the one in 2008) that saw rapid correlated changes in price were excluded from this analysis. If these shocks are considered to be black-swan events then it may be more appropriate to identify and exclude them. Taleb (2010) characterizes black-swan events by “rarity, extreme impact, and retrospective (though not prospective) predictability.” By this definition, black swans cannot be predicted, therefore their impact on a model should not be included. If the data for a black-swan event were included in a model, it would suggest that the model is able to predict the possibility of an event that is, by definition, unpredictable. What constitutes a black swan in oil-price history, and the dates that define such an event, are beyond the scope of this thesis. For this reason, observed price shocks (such as in 2008) are not considered to be black-swan events and are not excluded from the data set.

If shock events are included in the static variance used in this thesis, it would imply the impact of these shock events on the price of oil is permanent. Li and Thompson (2010) found that a deterministic trend in oil price is likely to exist and that shock events may be transitory. The authors argue that price controls imposed by governments may not be necessary if oil prices were to follow a deterministic trend and shock events are only transitory. Still, volatility in oil price is not necessarily static over time. This is where GARCH models (Bollerslev 1986) may do a better job than the simple geometric-random-walk model when fitting the observed behavior of oil price. GARCH models allow for variance to change over time. It may therefore be more descriptive to calibrate GARCH models to seven-day returns, something that is not currently being done. Previous works (Sadorsky 2006; Kang, Kang, and Yoon 2009; Wei, Wang, and Huang 2010) using GARCH models for oil returns are hampered by not defining returns in terms of calendar days, and by not explicitly accounting for day-of-the-week effects presented in this thesis. The material in this thesis for handling day-of-the-week effects should serve as a correction to how data are handled when constructing price models.

Model development and data treatment in this thesis are a return to the basics to correct a technical deficit in handling intra-week behavior of returns in the literature. Conditional-variance models such as those in the GARCH class would better constructed by explicitly accounting for what day of the week they are forecasting on. The previous chapter showed that day of the week has a direct impact on the mean and variance. In this thesis, intra-week behavior was ignored and variance was assumed to be static. This work could be improved upon by constructing conditional-variance models for inter-week behavior and ignoring intra-week behavior, or by explicitly defining what day of the week the model is predicting from and to.

Consider an analysis that correctly calculated returns based on calendar days rather than trading days, but did not account for intra-week effects. The analysis would not remove noise added from intra-week effects, and would calculate a variance of approximately

0.00057 for a simple geometric-random-walk model with $\tau = 1$ day. By comparison, this thesis eliminated the noise of intra-week effects and calculated a variance of approximately 0.00037. Day-of-the-week effects add more than 50% to the variance if they are not explicitly accounted for. Although this thesis removes them, they could be added back into the model.

What if the variance for day-of-the-week effects was added to the model? The model could, for example, predict a value of $r_{t=1year}$ and land on a Monday. Depending on the day it started on, it could alternatively land on any day of the week. The contribution of day-of-the-week uncertainty was 0.0002 (0.00057-0.00037). At one year of holding, it would only increase uncertainty of the return by 0.11%. By two years of holding, intra-week effects increase uncertainty by 0.04%. At ten years, it only adds 0.002% to the uncertainty in price. This is because what day of the week it is does not contribute nearly as much uncertainty to price as ten years of holding. Modeling intra-week effects was assumed to be an unnecessary complication with insignificant impact when modeling prices over 40 years. The uncertainty added by what day of the week the model ends on is small compared to the uncertainty of what the price is after a year of holding. Removal or description of intra-week behavior can only be done if returns are first defined by calendar days, on which GARCH-model literature is silent. It took a breakdown of the count of price inputs and time frames at the beginning of Chapter III to determine that the data were not being defined correctly in literature. One *could* argue that a model having mean and variance conditional on previous returns is capable of describing intra-week behavior without explicitly knowing it exists, but such a feat would be impossible if the dataset is not correctly defined, which it has not been. The price model in this thesis is a return to basics to correct a technical deficit present in literature, and only pursues the ramifications in the most elementary form of a geometric-random-walk model that has no conditional mean or variance. Once constructed and justified with a dataset defined on calendar days, GARCH models may be better descriptors than a simple geometric-random-walk model.

The base of the model developed in this chapter is a conventional geometric random walk. The data used to determine the metrics for this model were treated in a novel way by removing intra-week effects, which were observed when returns were defined on the basis of calendar days. In the following chapter this model is used in assessing the impact of price variability on reserves volume and other economic metrics.

CHAPTER V

WELL AND FIELD MODEL

Production Model

A synthetic field case was created to determine how modeling uncertainty impacts project economics and to demonstrate the use of the price model constructed in this thesis. Two sources of uncertainty were included in this test case—uncertain production behavior from wells and uncertain oil price. Cumulative well production at time t is defined as a function of the initial rate (q_0), initial decline rate (D_0), and the decline exponent (b) (Arps 1945). The choice of function depends on the value of b:

$$b \neq 0, b \neq 1 \qquad Q = Q_0 + \frac{q_0}{D_0(1-b)}\left[1 - \left(1 + bD_0(t - t_0)\right)^{1-\frac{1}{b}}\right] \tag{52}$$

$$b = 0 \qquad Q = Q_0 + \frac{q_0}{D_0}\left(1 - e^{-D_0(t-t_0)}\right) \tag{53}$$

$$b = 1 \qquad Q = Q_0 + \frac{q_0}{D_0} ln\left(1 + D_0(t - t_0)\right) \tag{54}$$

Hyperbolic decline (Eq. 52) must have a b value between, but not equal to, zero and one. A b value greater than one creates a non-zero flow rate at infinite producing time, resulting in an infinite volume in the reservoir. This must correspond to transient behavior of the well, as a reservoir will not actually be infinite in volume. For this reason, a minimum 10% nominal annual decline was imposed (EIA 2014), at which point the behavior shifts to exponential decline (Eq. 53). This causes flow rate to approach zero at infinite time.

Production behavior was modeled after horizontal tight-oil wells in the Eagle Ford shale. Wachtmeister et al. (2017) published distributions for DCA parameters (q_0, D_0, b) based on 294 wells of this type. They used a log-logistic distribution for all three parameters, which has the functional form shown in Eqs. 55 and 56. Notation for variables in these functions were kept consistent with the notation used by the Excel add-in @Risk (v. 7.5) (**Table 3**). Wachtmeister et al. (2017) did not state truncation limits or γ values, so the density curves were observationally matched to those published, and truncation limits imposed (**Table 4**).

PDF	$f(x) = \frac{\alpha}{\beta}\left(\frac{x-\gamma}{\beta}\right)^{\alpha-1}\left[1+\left(\frac{x-\gamma}{\beta}\right)^{\alpha}\right]^{-2}$	(55)
CDF	$F(x) = \frac{1}{1+\left(\frac{\beta}{x-\gamma}\right)^{\alpha}}$	(56)

Table 3—Log-logistic distribution function.

	q_i	D_i	b
	bbl/D	*/mo (nominal)*	___
γ	-300	0	-8.314
α	762.27	0.369	9.314
β	5.339	2.866	32.061
Min	20	0.048	0
Max	3,000	4.5	3.5
Mean	512	0.483	1.047
Std. Dev.	284.4	0.357	0.488

Table 4—Distribution parameters and truncation limits for decline curve variables based on Wachtmeister et al. (2017). Mean and standard deviation metrics listed are estimated by @Risk for the truncated distributions.

Project Economics

Revenue. I assumed this forecast is being performed in December of 2016. Using this starting month, price ($s_0 = \$50/\text{bbl}$) is calculated as the average of the first day's price for the last 12 months, per SEC regulation. Simulated future price (S_m) is obtained each month (m) based on the return on the previous price (r_m) determined by the random-walk model or deterministic-price scenario. Mean and variance metrics for a time step of one month are selected from Table 2 for the variable-price case, and Eqs. 49 and 50 are used to simulate changes in price. The deterministic-price model uses the same mean as the variable-price case, but the deterministic-price model uses a *c.v.* of zero. The prices over time for the deterministic-price scenario are graphed along with several outputs from the variable-price scenario in **Fig. 14**.

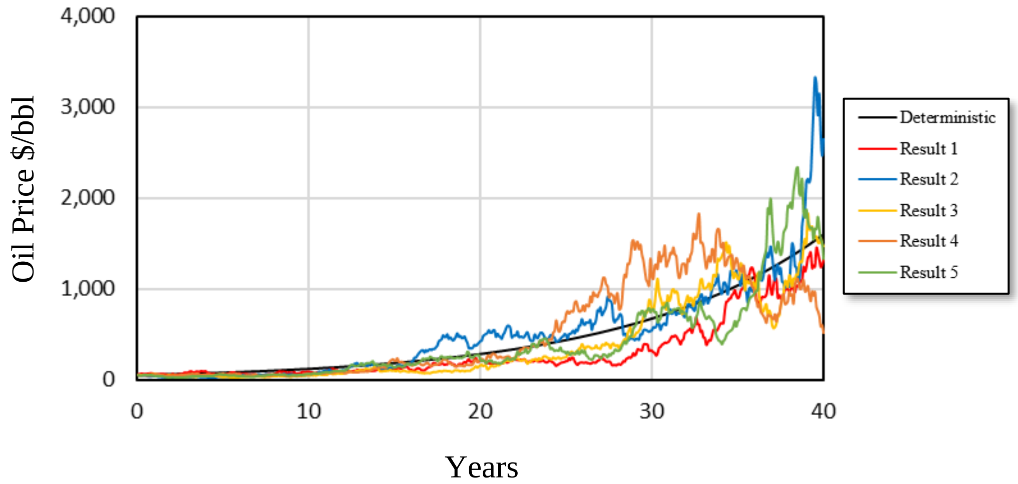

Fig. 14—Deterministic-price profile alongside several outcomes (Results 1-5) from the variable-price scenario.

Monthly production from a well ($Q_{\Delta m}$) was assumed to be shipped to market at the end of each month (m) and was multiplied by the modeled price path (s_m) to generate a timeline of revenue. One well was drilled each month for the first two years. A 5% chance of a dry hole as well as uncertain production (per uncertain decline parameters in Table 4) was imposed for each well.

Finance. Using Texas tax structure, a severance tax of 4.6% was applied to gross revenue. An initial capital expense (CAPEX) of $6.5 million per producing well and $2.5 million per dry hole was incurred, based on 2015 Eagle Ford values published by the EIA (EIA 2016). Initial CAPEX was experienced in the first producing month of each well. An abandonment CAPEX of $20 thousand was scaled with inflation and applied at the end of a well's life or at 40 years if a well was still economical. Variable operating costs (VC) from $0 to $40/bbl were tested in increments of $10/bbl, and fixed operating costs (FC) from $0 to $4,000/mo were simulated in increments of $1,000/mo. An annual inflation rate of 2% was used to discount oil price with time, as no consideration for inflation was made in the development of the model, and no inflation was accounted for in the costs associated with the wells (Eq. 58).

Internal rate of return (IRR), defined as the discount rate at which the net cash flow (NCF) becomes zero, was determined using Excel's built in function. A weighted average cost of capital of 10% was used as the discount factor for determining the net present value (NPV) of the field. Profitability index (PI) was used to contextualize NPV with the required capital investment. Financial formulas are based on models presented in Mian (2011), and are summarized in **Table 5**.

Wells were individually abandoned when operating profit first fell to zero. Furthermore, no production beyond 40 years was considered, and abandonment cost was applied at the end of 40 years if the well had not yet been abandoned.

<u>Price</u>	$s_m = r_m s_0$	(57)
<u>Operating Profit</u>	$OP_m = s_m (1.02)^{-m/12} \times Q_{\Delta m} (1 - 4.6\%) - (VC \times Q_{\Delta m} + FC)$	(58)
<u>Cash Flow</u>	$CF_m = OP_m - CAPEX_m$	(59)
Total	$NCF = \sum CF_m$	(60)
NPV	$NPV = \sum CF_m (1 - 0.1)^{m/12}$	(61)
PI	$PI = 1 + \frac{NPV}{CAPEX}$	(62)

Table 5—Financial formulas and metrics used when assessing field profitability and risk. Based on relations in Mian (2011).

<u>Simulation</u>. Ten thousand iterations of the price model were performed to determine the distribution of possible outcomes. During a simulation, project economics were determined for 24 wells with the variable-price model developed in this thesis. For each iteration, the production from those 24 wells was also run through economic calculations using the average price gain as a deterministic-price scenario ($\sigma_{\Delta t}/\mu_{\Delta t} = 0$ in Eq. 49).

<u>Results</u>. Mean, P90, P50 and P10 values for reserves, NPV, PI, and well life for each simulation are detailed in Appendix D. It is important to remember that in the variable-price scenario as well as the deterministic-price scenario, the well production was variable according the distribution of DCA parameters in Table 4. Variable production is what

produces variable results in the deterministic-price scenario. On each iteration of the simulation the project metrics for the variable-price case are compared to the deterministic-price to see how using a deterministic-price scenario causes bias in the expectations and risk profile for the project. The ratio of the metrics is informative of how much ignoring price volatility biases project economics. Reserves volume, well life, and profitability index (PI) for the project that used the variable-price model are compared to the values under the deterministic-price model. This is *not* a direct comparison for mean and variance of each set of outcomes; it is the mean and variance of the ratio of individual outcomes. Using reserves as an example:

$$\mathbb{E}\left(\frac{Reserves_{var}}{Reserves_{det}}\right) \tag{63}$$

$$\mathbb{V}\left(\frac{Reserves_{var}}{Reserves_{det}}\right) \tag{64}$$

During a single iteration of the simulation, the reserves volume for 24 wells is calculated with a variable-price timeline ($Reserves_{var}$), and the exact same production profile is used to calculate reserves volume using the deterministic-price timeline ($Reserves_{det}$). That profile may however be truncated differently depending on the economic limit in each scenario. This is a direct comparison of outcomes to show how ignoring price volatility impacts project economics. This was done for reserves volume, well life, and profitability index (**Figs. 15** to **20**).

In the case of reserves (Fig. 15), the average result is that reserves volume calculated using a variable-price model will be less than using a deterministic-price model. This means that on average, using a deterministic-price model when forecasting reserves will result in the reported reserves volume being overstated. When there were no fixed costs and the variable cost was \$40/bbl, reserves volume calculated under the variable-price scenario was on average 86.5% of the reserves volume calculated under the deterministic-price scenario. Comparing the *c.v.*'s of the data (Fig. 16) shows that there was more uncertainty in the stated volume if variable prices were used. Of course, it *is* expected that adding a new source of uncertainty will add uncertainty to the results. The results were, however, highly dependent on the VC for the project. For low VC (\$0 to \$10/bbl), the deterministic- and variable-price cases would, on average, return the same reserves volume for the project. The ratio of the outcomes under each FC scenario did not vary much ($c.v. \approx 0$), meaning there was little difference when using variable prices to calculate reserves in a low-VC scenario. At \$30/bbl VC, the outcomes of the FC scenarios begin to diverge. This is seen by an increase in the *c.v.* Divergence in the reserves volume between the two price cases is due to differences in well life (Figs. 15 and 17). For a given production outcome for the 24 wells, the variable-price case would, on average, result in a shorter well life. The shorter well life is a result of wells becoming uneconomic earlier in their lifetime under a variable-price case. At \$40/bbl VC and \$0 FC, the volatility in prices resulted in a 20% lower well life than the well life the deterministic-price scenario would determine.

Reductions to well life and, consequently, reserves volumes did not have a significant impact on expected profitability index (Fig. 19), except for \$40/bbl VC scenarios.

As variable cost approaches \$40/bbl, the PI decreases relative to the deterministic-price case. To understand the mechanics of this, consider the difference between initial price and variable cost to be ΔP:

$$\Delta P = s_0 - VC \tag{65}$$

Greater values of ΔP give a greater allowance for prices to change and the project to still remain economical. The smaller the value of ΔP, the greater the chance price will cross over into becoming uneconomical. The moment price drops below the variable cost the well is no longer economical and it is shut in (ignoring the marginal effects fixed costs impose). There is some bias in this evaluation, as wells were assumed to be permanently abandoned when they first become uneconomical. If the well was shut in while having low production volumes, it is possible the company would abandon it permanently. However, wells that become subeconomic due to low oil price may also be temporarily abandoned and reopened when the oil price becomes more favorable. Temporarily abandoning a well is still a negative toward the economics of a project. Consider a thought experiment where a well is shut in for six months then turned back on when oil price is more favorable. Even if the well were to produce the missing six months of production in the first month it comes back online, that production would still be worth less due to the time value of money. It is *still* a negative impact on the project economics. These events are more frequent than a deterministic-price scenario would predict because prices dip and jump under a volatile-price scenario. This may be a result of imposing a maximum well life restriction, which is a common practice in project evaluation. In the cases where wells are shut in during low-price environments and turned back on in a more favorable economic climate, the project economics might improve if the well is allowed to produce beyond the 40-year cutoff imposed by this thesis. It is possible that temporarily abandoning wells may result in more favorable well economics by taking advantage of oil price volatility. This however was not explored and would be an interesting point for further research.

Even “hockey stick” lines used to model a low-price scenario would not be sufficient to capture percentile outcomes for project metrics produced by modeling volatile oil prices. They are *still* deterministic scenarios. There is no volatility to the price. Deterministic scenarios, as shown in this thesis, are not able to capture the true average reserves volume. High and low-price scenarios via a hockey stick method are no different. They *do not* model volatility. Hockey sticks, whether high, average, or low, *are* deterministic. If the deterministic-scenario (which modeled average price) used in this thesis was not capable of producing true average reserves volume, then it is unreasonable to expect a lower 10% price scenario to represent the impact of low prices on reserves (or other project metrics). Simulation of volatile-price behavior is needed to determine how prices limit well life and ultimately reserves volumes and profitability. It is thus recommended to simulate price volatility when determining the economics, as it has a tangible impact on project metrics.

While in this thesis these effects showed at a high VC that might not be realized by a company, the impact of price volatility may be show up at a different VC depending on the anticipated production from a field. Furthermore, this is likely to be project specific. While not specifically examined in this thesis, having a set of wells that produce worse than what was presented in this thesis may result in the impact of price volatility showing up at a lower VC. The ranges for VC chosen were intended to bracket most economic scenarios, and a \$0/bbl VC or a \$40/bbl VC might not actually be realized. In a report by the EIA (2016), lease operating expenses (including water disposal) were between \$2/boe (barrel of oil equivalent) and \$14.5/boe over the life of a well. Transportation expenses were between \$0.25/bbl and \$13/bbl. They also state general and administrative costs between \$1/boe and \$8/boe. Taking the extreme of each cost category produces a VC of \$35.5/bbl, which the chosen extreme of \$40/bbl safely brackets. A company would need to use their own cost structure to determine how price volatility impacts them.

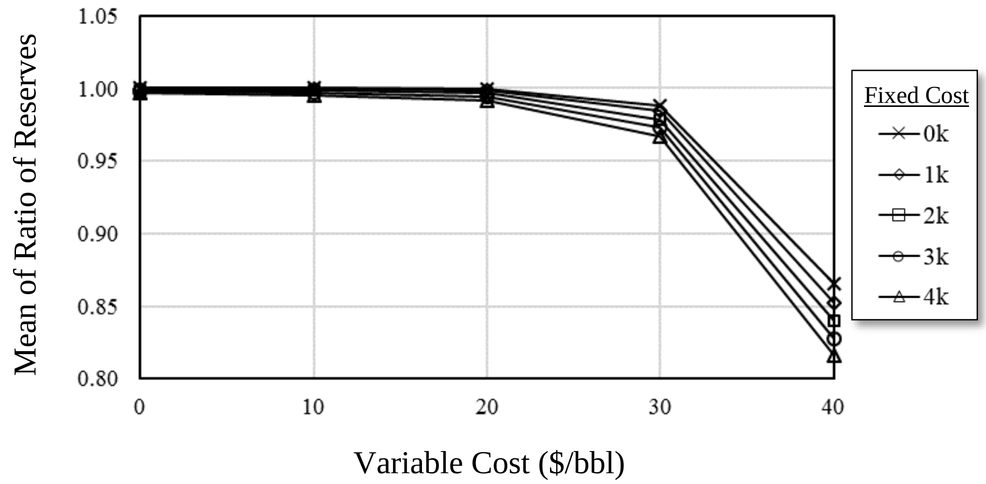

Fig. 15—Average of the ratio of reserves determined using variable prices to value returned using deterministic prices.

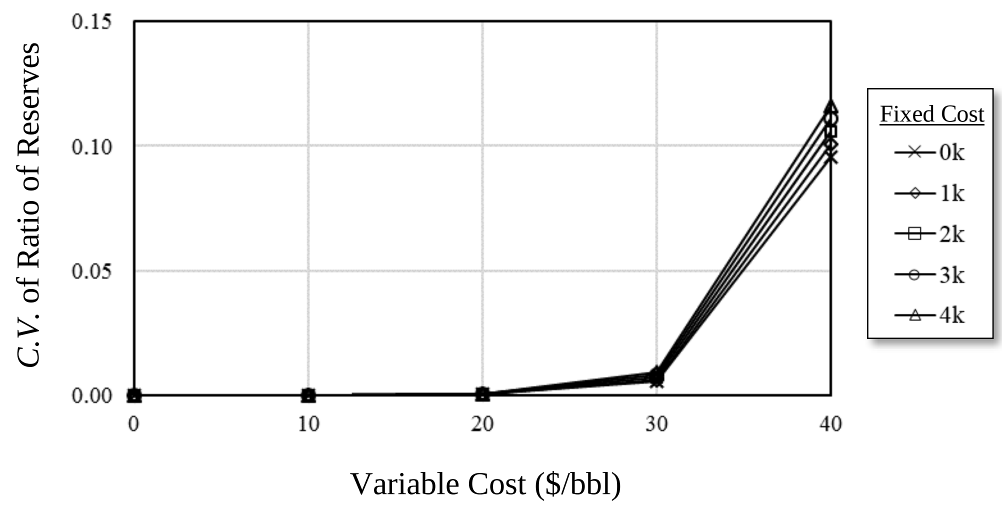

Fig. 16—Coefficient of variation of the ratio of reserves determined using variable prices to value returned using deterministic prices.

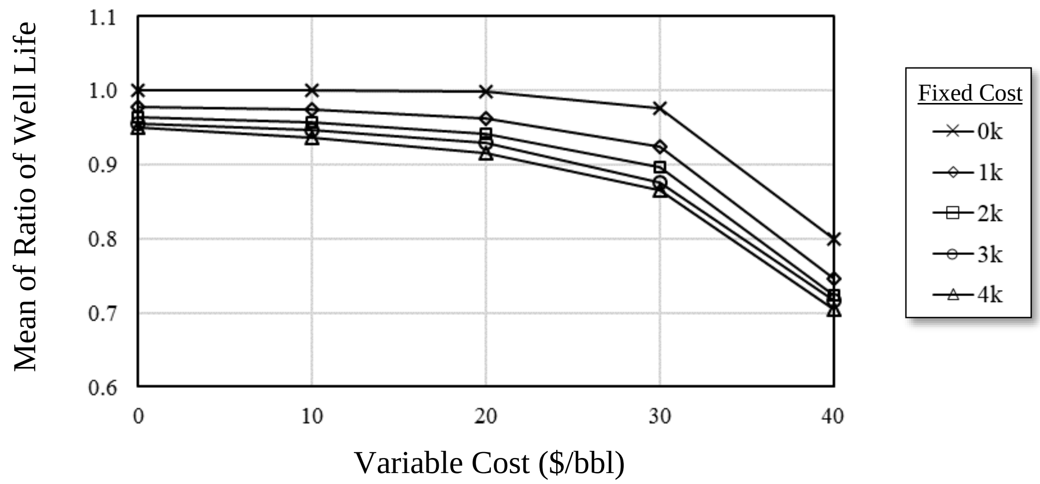

Fig. 17—Average of the ratio of well life determined using variable prices to value returned using deterministic prices.

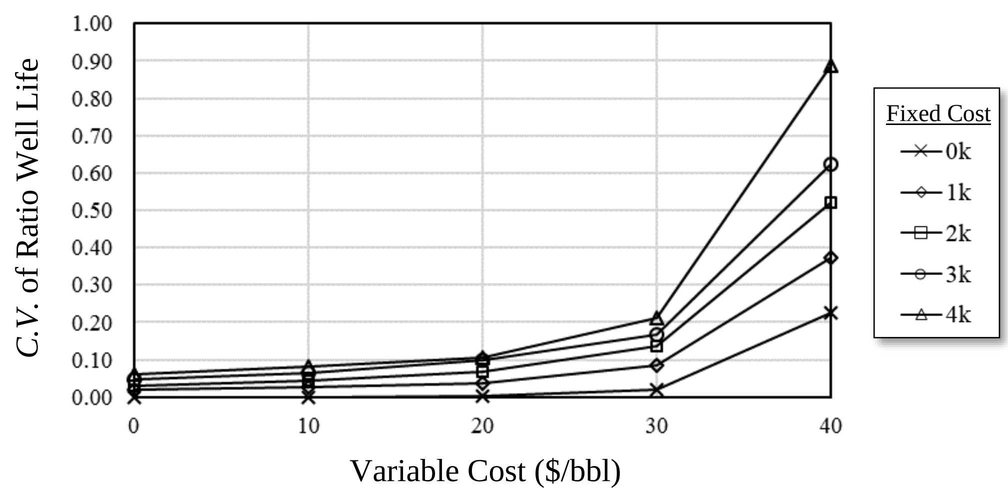

Fig. 18—Coefficient of variation of the ratio of well life determined using variable prices to value returned using deterministic prices.

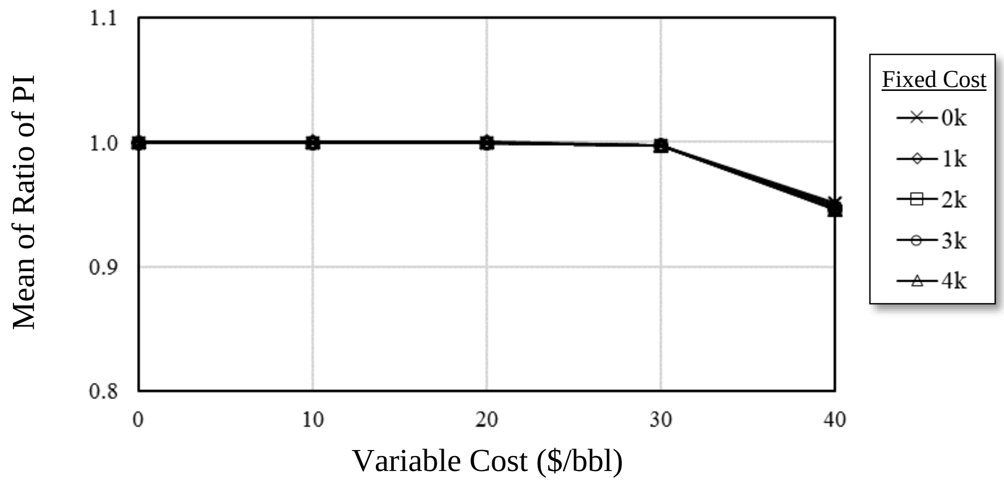

Fig. 19—Average of the ratio of profitability index (*PI*) determined using variable prices to value returned using deterministic prices.

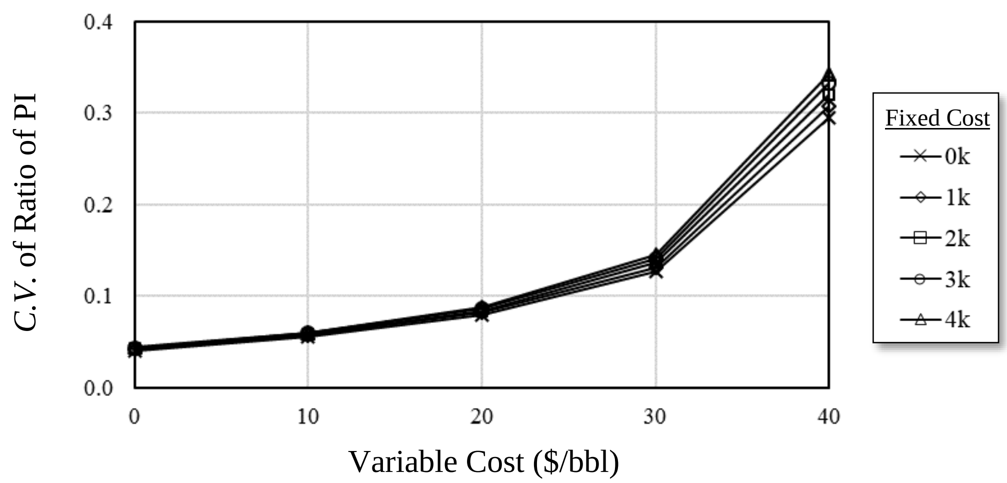

Fig. 20—Coefficient of variation of the ratio of profitability index (*PI*) determined using variable prices to value returned using deterministic prices.

CHAPTER VI

CONCLUSION

- The standard density function of incremental returns can be modeled by the proposed double Laplace distribution, independent of the defined time step, using only the coefficient of variation as input.
- Day-of-the-week effects cause returns with a time step of one day to form an improper set. Monday-to-Tuesday returns are not equivalent to, and cannot be grouped with, Friday-to-Monday returns.
- Day-of-the-week effects can be eliminated by redefining the time-step of returns to one week.
- The primary driver of mis-assessment of mean and variance of project metrics under a deterministic-price scenario is a failure to account for earlier well-shut-in events due to low-price environments. This conclusion may be impacted by an imposed maximum project life, which is a common practice in project evaluation. The impact of not imposing a maximum project life was not assessed.

REFERENCES

17 C.F.R. § 210.4 (2010).

Akilu, N.Y., McVay, D.A., Lee, J. 2006. Inverted Hockey Stick Method Quantifies Price uncertainty in Petroleum Investment Evaluation. *Petroleum Science and Technology* **24**: 167-185. https://doi.org/10.1081/LFT-200044500.

Arps, J.J. 1945. Analysis of Decline Curves. *Transactions of the AIME* **160** (01): 228-247. https://doi.org/10.2118/945228-G.

Bathaee, Y. 2018. The Artifical Intellegence Black Box and the Failure of Intent and Causation. *Harvard Journal of Law & Technology* **31** (2): 889-938.

Behr, A., Pötter, U. 2007. Alternatives to the normal model of stock returns: Gaussian mixture, generalised logF and generalised hyperbolic models. *Annals of Finance* **5**: 49-68. https://doi.org/10.1007/s10436-007-0089-8.

Bollerslev, T. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* **31**: 246-327. https://doi.org/10.1016/0304-4076(86)90063-1.

Buhrmester, V., Münch, D., Arens, M. 2019. Analysis of Explainers of Black Box Deep Neural Networks for Computer Vision: A Survey. *ArXiv* **abs/1911.12116**.

Charles, A. 2010. The day-of-the week effects on the volatility: The role of the asymmetry. *European Journal of Operational Research* **202** (1): 143-152.

Council, National Research. 1992. *The National Energy Modeling System*. Washington, DC: The National Academies Press.

EIA. 2014. Annual Energy Outlook 2014, United States Department of Energy, (April 2014).

EIA. 2016. Trends in U.S. Oil and Natural Gas Upsream Costs, United States Department of Energy, Southwest, Washington, DC, (March 2016).

EIA. Cushing, OK WTI Spot Price FOB (Dollars per Barrel), https://www.eia.gov/dnav/pet/hist/RWTCD.htm (downloaded 1/1/2020).

Engineers, Society of Petroleum. SPE Petroleum Resources Management System Guide for Non-Technical Users https://www.spe.org/industry/docs/PRMS-Guide-for-Non-Technical-Users-2007.pdf.

Fattouh, B., Killian, L., Mahadeva, L. 2012. *The Role of Speculation in Oil Markets: What Have We Learned So Far?* . Oxford, UK (Reprint).

Hou, A., Suardi, S. 2011. A Nonparametric GARCH Model of Crude Oil Price Return Volatility. *Energy Economics* **34** (2): 618-626. https://doi.org/10.1016/j.eneco.2011.08.004.

Kang, S.H., Kang, S.-M., Yoon, S.-M. 2009. Forecasting volatility of crude oil markets. *Energy Economics* **31**: 119-125. https://doi.org/10.1016/j.eneco.2008.09.006.

Kolman, E., Margaliot, M. 2007. Knowledge Extraction from Neural Networks using the All-Permutations Fuzzy Rule Base. *IEEE Trans Neural Netw*. **18** (3): 925-931. 10.1109/TNN.2007.891686.

Küchler, U., Neumann, K. 1999. Stock Returns and Hyperbolic Distributions. *Mathematical and Computer Modelling* **29**: 1-15. https://doi.org/10.1016/S0895-7177(99)00088-6.

Lee, C.Y., Huh, S.Y. 2017. Forecasting Long-Term Crude Oil Prices Using a Bayesian Model with Informative Priors. *Sustainability* **9** (190). https://doi.org/10.3390/su9020190.

Li, J., Thompson, H. 2010. A Note on the Oil Price Trend and GARCH Shocks. *International Association for Energy Economics* **31** (3): 159-165. 10.5547/ISSN0195-6574-EJ-Vol31-No3-8.

Lopez, J.A. 2001. Evaluating the Predictive Accuracy of Volatility Models. *Journal of Forecasting* **20**: 87-109.

Maberly, E.D. 1995. Eureka! Eureka! Discovery of the Monday Effect Belongs to the Ancient Scribes. *Financial Analysts Journal* **51** (5): 10-11. https://doi.org/10.2469/faj.v51.n5.1931.

Mandelbrot, B. 1960. The Pareto-Levy Law and the Distribution of Income. *Intern. Econ. Rev.* **1** (2): 79-106. 10.2307/2525289.

McDonald, R.L. 2006. *Market Derivatives*, 2nd edition. Boston, Massachusetts: Pearson Education Inc.

McMichael, C. 1999. The Fallacy of the Hockey Stick Projection. Presented at the SPE Annual Technical Conference and Exhibition, Houston, Texas, 3-6 October. SPE-56454-MS.

Mian, M.A. 2011. *Project Economics and Decision Analysis. Volume 2: Probabilistic Models*, 2nd edition. Tulsa, Oklahoma: PennWell.

Nadarajah, S. 2007. The linear combination, product and ratio of Laplace random variables. *Statistics* **41** (6): 535-545.

Olsen, G.T., McVay, D.A., Lee, W.J. 2005. Inverted Hockey Stick Method Reveals Greater Uncertainty in Petroleum Investment Evaluations. Presented at the SPE Western Regional Meeting, Irving, California, 30 March - 1 April. SPE-93999-MS. https://doi.org/10.2118/93999-MS.

Patton, A.J. 2011. Volatility Forecast Using Imperfect Volatility Proxies. *Journal of Econometrics* **160** (1): 246-256. https://doi.org/10.1016/j.jeconom.2010.03.034.

Sadorsky, P. 2006. Modeling and Forecasting Petroleum Futures Volatility. *Energy Economics* **28**: 467-488. https://doi.org/10.1016/j.eneco.2006.04.005.

Taleb, N. 2010. *The Black Swan: The Impact of the Highly Improbable*, 2nd edition. United States: Random House Trade Paperbacks.

Wachtmeister, H., Lund, L., Aleklett, K. et al. 2017. Production Decline Curves of Tight Oil Wells in Eagle Ford Shale. *Natural Resources Research* **26** (3): 365-377. https://doi.org/10.1007/s11053-016-9323-2.

Wei, Y., Wang, Y., Huang, D. 2010. Forecasting Crude Oil Market Volatility: Further Evidence Using GARCH-class Models. *Energy Economics* **32**: 1477-1484. https://doi.org/10.1016/j.eneco.2010.07.009.

Yanagisawa, A. 2009. *Usefulness of the Forward Curve in Forecasting Oil Prices*. Tokyo, Japan: Energy Working Papers (Reprint).

Yu, H.-C., Chiou, I., Jordan-Wagner, J. 2008. Does the weekday effect of the yen/dollar spot rates exist in Tokyo, London, and New York? An analysis of panel probability distribution. *Applied Economics* **40** (20): 2631-2643. http://dx.doi.org/10.1080/00036840600970294.

APPENDIX A

DENSITY CURVE TABLES

A density curve table was generated using observed *c.v.* values for 1, 7, 14, and 70 days of holding. Values in **Table 6** are presented in increments of 1%. Latin hypercube sampling was chosen over Monte Carlo, due to increased stability when sampling from exponential-type functions. The Mersenne Twister was used as a random number generator, and a unique seed was randomly generated for each Laplace distribution. One million samples were taken from each Laplace distribution to ensure stability in the tails of the distributions.

	Days Held, τ			
	1	7	14	70
Percentile	**Coefficient of Variation, *c.v.***			
p	0.02379	0.053024	0.111362	0.157801
0.01	-2.583	-2.553	-2.536	-2.466
0.02	-2.179	-2.158	-2.147	-2.107
0.03	-1.935	-1.924	-1.921	-1.888
0.04	-1.767	-1.754	-1.754	-1.727
0.05	-1.631	-1.622	-1.621	-1.601
0.06	-1.520	-1.514	-1.511	-1.495
0.07	-1.423	-1.420	-1.418	-1.404
0.08	-1.337	-1.336	-1.335	-1.326
0.09	-1.264	-1.262	-1.260	-1.254
0.10	-1.197	-1.197	-1.195	-1.189
0.11	-1.136	-1.135	-1.133	-1.129
0.12	-1.080	-1.079	-1.078	-1.074
0.13	-1.028	-1.027	-1.026	-1.024
0.14	-0.978	-0.979	-0.978	-0.977
0.15	-0.931	-0.933	-0.933	-0.933
0.16	-0.889	-0.889	-0.891	-0.891
0.17	-0.847	-0.848	-0.850	-0.851
0.18	-0.808	-0.809	-0.811	-0.813
0.19	-0.770	-0.772	-0.774	-0.777
0.20	-0.735	-0.736	-0.739	-0.742
0.21	-0.700	-0.703	-0.705	-0.709
0.22	-0.667	-0.670	-0.673	-0.676
0.23	-0.635	-0.639	-0.641	-0.645
0.24	-0.605	-0.608	-0.610	-0.615
0.25	-0.576	-0.579	-0.580	-0.586
0.26	-0.547	-0.550	-0.552	-0.558
0.27	-0.519	-0.522	-0.524	-0.530
0.28	-0.492	-0.495	-0.497	-0.504
0.29	-0.465	-0.469	-0.470	-0.478
0.30	-0.439	-0.443	-0.445	-0.452

Table 6—Z-score curves describing the possible gain and loss of asset value. Coefficient of variation values correspond to those observed for WTI spot price returns at 1, 7, 14 and 70 days of holding.

	Days Held, τ			
	1	7	14	70
Percentile	**Coefficient of Variation, *c.v.***			
p	0.02379	0.053024	0.111362	0.157801
0.31	-0.414	-0.418	-0.419	-0.427
0.32	-0.389	-0.393	-0.395	-0.403
0.33	-0.365	-0.369	-0.371	-0.379
0.34	-0.341	-0.345	-0.347	-0.356
0.35	-0.317	-0.322	-0.323	-0.333
0.36	-0.295	-0.299	-0.301	-0.310
0.37	-0.272	-0.277	-0.279	-0.288
0.38	-0.250	-0.255	-0.257	-0.266
0.39	-0.229	-0.233	-0.235	-0.244
0.40	-0.207	-0.211	-0.213	-0.223
0.41	-0.186	-0.190	-0.193	-0.202
0.42	-0.165	-0.169	-0.171	-0.181
0.43	-0.144	-0.149	-0.151	-0.160
0.44	-0.123	-0.128	-0.130	-0.140
0.45	-0.103	-0.108	-0.110	-0.120
0.46	-0.083	-0.087	-0.090	-0.099
0.47	-0.063	-0.067	-0.069	-0.079
0.48	-0.042	-0.047	-0.049	-0.059
0.49	-0.022	-0.027	-0.029	-0.039
0.50	-0.003	-0.006	-0.009	-0.019
0.51	0.017	0.014	0.011	0.001
0.52	0.037	0.034	0.031	0.021
0.53	0.057	0.054	0.051	0.041
0.54	0.078	0.074	0.071	0.061
0.55	0.098	0.094	0.092	0.081
0.56	0.118	0.115	0.112	0.101
0.57	0.139	0.135	0.133	0.122
0.58	0.160	0.156	0.154	0.143
0.59	0.181	0.176	0.175	0.164
0.60	0.202	0.198	0.195	0.185

Table 6 Continued—Z-score curves describing the possible gain and loss of asset value. Coefficient of variation values correspond to those observed for WTI spot price returns at 1, 7, 14 and 70 days of holding. Continued.

	Days Held, τ			
	1	7	14	70
Percentile	**Coefficient of Variation, *c.v.***			
p	0.02379	0.053024	0.111362	0.157801
0.61	0.224	0.219	0.217	0.207
0.62	0.245	0.241	0.239	0.228
0.63	0.267	0.262	0.261	0.251
0.64	0.290	0.285	0.284	0.273
0.65	0.313	0.308	0.306	0.296
0.66	0.336	0.331	0.330	0.320
0.67	0.360	0.354	0.354	0.343
0.68	0.384	0.379	0.378	0.367
0.69	0.409	0.404	0.403	0.393
0.70	0.434	0.430	0.428	0.418
0.71	0.460	0.456	0.454	0.444
0.72	0.487	0.483	0.481	0.471
0.73	0.514	0.510	0.508	0.499
0.74	0.542	0.538	0.537	0.527
0.75	0.570	0.567	0.566	0.557
0.76	0.601	0.597	0.596	0.587
0.77	0.632	0.628	0.627	0.619
0.78	0.664	0.661	0.658	0.651
0.79	0.697	0.694	0.692	0.684
0.80	0.732	0.729	0.727	0.720
0.81	0.767	0.766	0.763	0.756
0.82	0.804	0.804	0.801	0.795
0.83	0.843	0.845	0.842	0.835
0.84	0.885	0.887	0.884	0.878
0.85	0.929	0.931	0.929	0.924
0.86	0.976	0.978	0.976	0.972
0.87	1.026	1.026	1.026	1.024
0.88	1.079	1.080	1.079	1.079
0.89	1.137	1.138	1.138	1.138
0.90	1.199	1.201	1.201	1.203

Table 6 Continued—Z-score curves describing the possible gain and loss of asset value. Coefficient of variation values correspond to those observed for WTI spot price returns at 1, 7, 14 and 70 days of holding. Continued.

	Days Held, τ			
	1	7	14	70
Percentile	**Coefficient of Variation, *c.v.***			
p	0.02379	0.053024	0.111362	0.157801
0.91	1.268	1.272	1.271	1.273
0.92	1.344	1.347	1.348	1.353
0.93	1.430	1.433	1.435	1.444
0.94	1.529	1.532	1.534	1.546
0.95	1.641	1.646	1.651	1.668
0.96	1.779	1.786	1.794	1.812
0.97	1.954	1.966	1.977	2.002
0.98	2.201	2.217	2.228	2.264
0.99	2.611	2.642	2.654	2.721

Table 6 Continued—Z-score curves describing the possible gain and loss of asset value. Coefficient of variation values correspond to those observed for WTI spot price returns at 1, 7, 14 and 70 days of holding. Continued.

APPENDIX B

DENSITY MATCHING

Fig. 21—Probability density of R_1 generated using observed mean and variance.

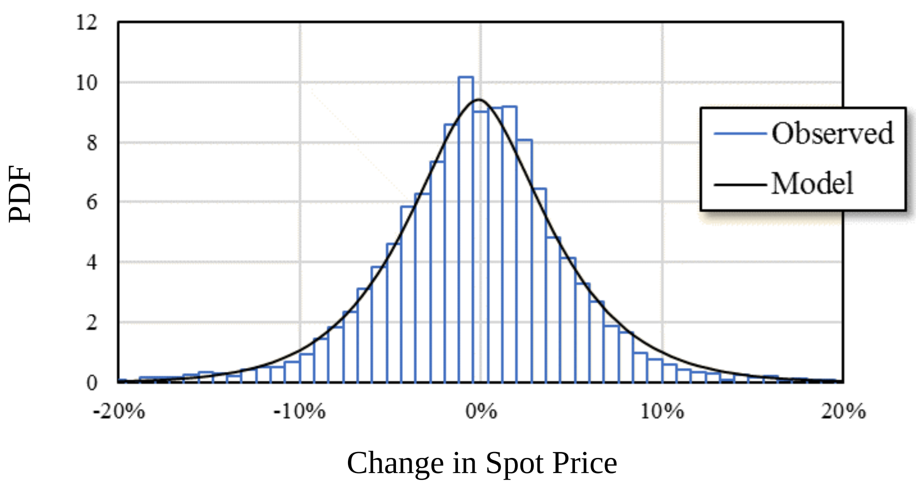

Fig. 22—Probability density of R_7 generated using observed mean and variance.

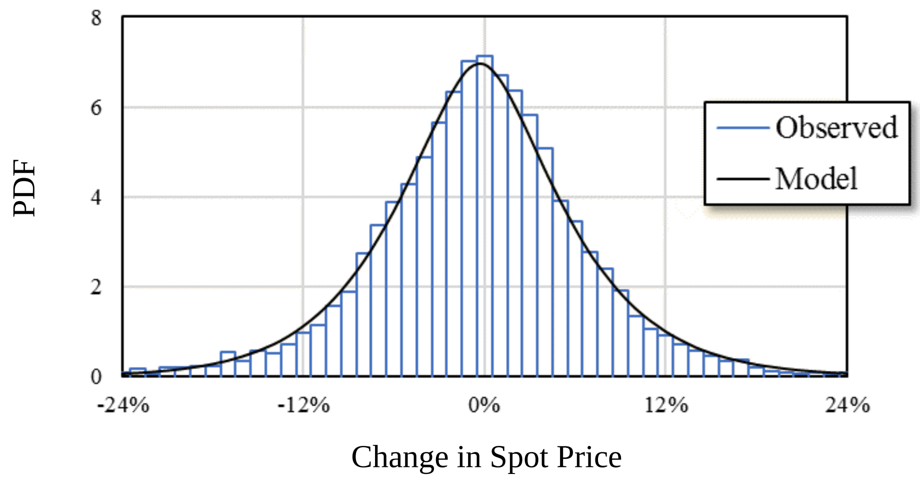

Fig. 23—Probability density of R_{14} generated using observed mean and variance.

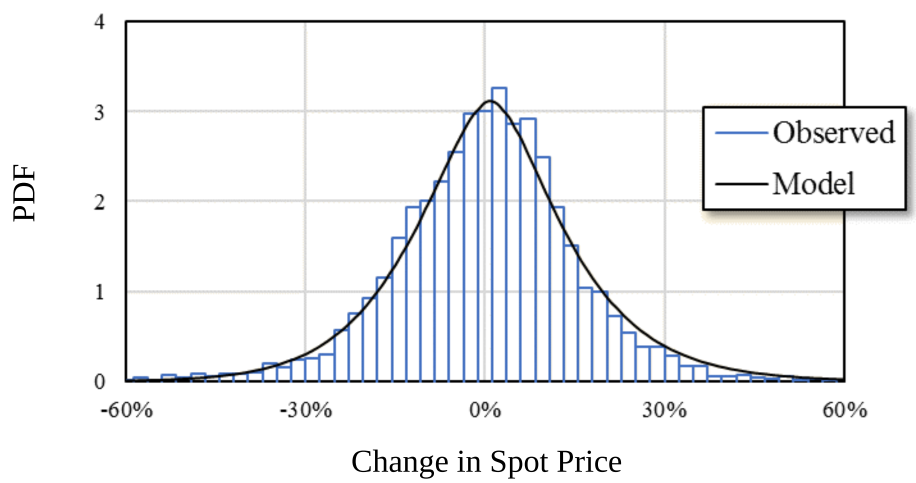

Fig. 24—Probability density of R_{70} generated using observed mean and variance.

APPENDIX C

RELATIONS TABLE

A relational table consisting of subscripts used for variables is contained in this appendix. A price being predicted from specifies a column, the price being predicted to specifies the row, and the entry for that coordinate in the matrix is the subscripts for $r_{\tau,n}$. The set R_τ is a set of elements collected along the diagonal. Only when $\tau = 7k$ for integer k is a diagonal filled. Above the identity diagonal there are no prices to compare to (they would occur in negative τ). If there were only 15 spot prices available for analysis, there would be 18 values in R_1 (including 6 empty values), but only 17 values in R_2. Increasing holding period by one day causes a loss of a data point. Although for a large number of spot prices the data sets are similar, they diverge *slightly* in mean and variance.

$\tau \backslash n$	0	1	2	3	4		7	8	9	10	11		14	15	16	17	18
0	x																
1	1,1	x							$r_{\tau,n}$								
2	2,1	1,2	x														
3	3,1	2,2	1,3	x		↑											
4	4,1	3,2	2,3	1,4	x												
		←				**Weekend**	→										
7	7,1	6,2	5,3	4,4	3,5	↓	x										
8	8,1	7,2	6,3	5,4	4,5		1,8	x									
9	9,1	8,2	7,3	6,4	5,5		2,8	1,9	x								
10	10,1	9,2	8,3	7,4	6,5		3,8	2,9	1,10	x							
11	11,1	10,2	9,3	8,4	7,5		4,8	3,9	2,10	1,11	x						
14	14,1	13,2	12,3	11,4	10,5		7,8	6,9	5,10	4,11	3,12		x				
15	15,1	14,2	13,3	12,4	11,5		8,8	7,9	6,10	5,11	4,12		1,15	x			
16	16,1	15,2	14,3	13,4	12,5		9,8	8,9	7,10	6,11	5,12		2,15	1,16	x		
17	17,1	16,2	15,3	14,4	13,5		10,8	9,9	8,10	7,11	6,12		3,15	2,16	1,17	x	
18	18,1	17,2	16,3	15,4	14,5		11,8	10,9	9,10	8,11	7,12		4,15	3,16	2,17	1,18	x

Fig. 25—Demonstrating the relationship between values of $r_{\tau,n}$. For example, the first 14-day return ($r_{14,1}$) divided by the first one-day return ($r_{7,1}$) is equal to the eighth 7-day return ($r_{7,8}$).

APPENDIX D

SUMMARY OF FINANCIAL METRICS

		Reserves, MMbbl							
FC	**VC**	**Variable Price**				**Static Price**			
$M	$/bbl	Mean	P90	P50	P10	Mean	P90	P50	P10
0	0	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
0	10	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
0	20	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
0	30	5.3	3.9	5.2	6.8	5.4	4.0	5.3	6.8
0	40	4.6	1.8	4.9	6.6	5.4	4.0	5.3	6.8
1	0	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
1	10	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
1	20	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
1	30	5.3	3.9	5.2	6.8	5.4	4.0	5.3	6.8
1	40	4.6	1.8	4.8	6.6	5.4	4.0	5.3	6.8
2	0	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
2	10	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
2	20	5.3	4.0	5.3	6.8	5.4	4.0	5.3	6.8
2	30	5.2	3.8	5.2	6.8	5.4	4.0	5.3	6.8
2	40	4.5	1.7	4.8	6.5	5.4	4.0	5.3	6.8
3	0	5.4	4.0	5.3	6.8	5.4	4.0	5.3	6.8
3	10	5.3	4.0	5.3	6.8	5.4	4.0	5.3	6.8
3	20	5.3	4.0	5.2	6.8	5.4	4.0	5.3	6.8
3	30	5.2	3.8	5.2	6.8	5.4	4.0	5.3	6.8
3	40	4.4	1.6	4.7	6.5	5.3	4.0	5.3	6.8
4	0	5.3	4.0	5.3	6.8	5.4	4.0	5.3	6.8
4	10	5.3	4.0	5.2	6.8	5.4	4.0	5.3	6.8
4	20	5.3	3.9	5.2	6.8	5.4	4.0	5.3	6.8
4	30	5.2	3.7	5.1	6.7	5.3	4.0	5.3	6.8
4	40	4.4	1.6	4.6	6.4	5.3	4.0	5.2	6.8

Table 7—Mean, P90, P50 and P10 of reserves determined using real oil-price movements (variable price) as well as with deterministic pricing (deterministic price). Results presented for various fixed (FC) and variable (VC) operating costs.

PROJECT NPV, $MM									
FC	**VC**	**Variable Price**				**Static Price**			
$M	$/bbl	Mean	P90	P50	P10	Mean	P90	P50	P10
0	0	76	5	67	159	76	25	73	132
0	10	44	-23	34	122	44	-1	41	91
0	20	11	-51	2	85	11	-26	9	51
0	30	-21	-80	-30	48	-21	-51	-23	11
0	40	-57	-124	-63	11	-53	-76	-55	-29
1	0	74	2	64	157	74	22	70	129
1	10	41	-25	32	119	41	-3	38	89
1	20	9	-54	0	82	9	-28	6	49
1	30	-24	-83	-33	45	-24	-53	-25	8
1	40	-60	-126	-66	9	-56	-78	-57	-32
2	0	71	0	62	154	71	20	68	127
2	10	39	-28	30	117	39	-5	36	87
2	20	6	-56	-3	80	7	-30	4	46
2	30	-26	-85	-35	43	-26	-55	-28	6
2	40	-62	-126	-69	6	-58	-81	-60	-34
3	0	69	-2	60	152	69	18	65	125
3	10	37	-30	27	115	37	-7	34	84
3	20	4	-58	-5	77	4	-33	2	44
3	30	-28	-87	-37	40	-28	-58	-30	4
3	40	-64	-127	-71	4	-60	-83	-62	-36
4	0	67	-4	58	149	67	16	63	122
4	10	34	-32	25	112	34	-10	31	82
4	20	2	-60	-7	75	2	-35	-1	42
4	30	-31	-89	-39	38	-30	-60	-32	1
4	40	-66	-128	-74	2	-63	-85	-64	-38

Table 8—Mean, P90, P50 and P10 of project NPV (10% discount) determined using real oil-price movements (variable price) as well as with deterministic pricing (deterministic price). Results presented for various fixed (FC) and variable (VC) operating costs.

PROJECT PI									
FC	**VC**	**Variable Price**				**Static Price**			
$M	$/bbl	Mean	P90	P50	P10	Mean	P90	P50	P10
0	0	1.5	1.0	1.4	2.0	1.5	1.2	1.5	1.9
0	10	1.3	0.8	1.2	1.8	1.3	1.0	1.3	1.6
0	20	1.1	0.7	1.0	1.6	1.1	0.8	1.1	1.3
0	30	0.9	0.5	0.8	1.3	0.9	0.7	0.8	1.1
0	40	0.6	0.1	0.6	1.1	0.6	0.5	0.6	0.8
1	0	1.5	1.0	1.4	2.0	1.5	1.1	1.5	1.8
1	10	1.3	0.8	1.2	1.8	1.3	1.0	1.3	1.6
1	20	1.1	0.6	1.0	1.5	1.1	0.8	1.0	1.3
1	30	0.8	0.5	0.8	1.3	0.8	0.6	0.8	1.1
1	40	0.6	0.1	0.6	1.1	0.6	0.5	0.6	0.8
2	0	1.5	1.0	1.4	2.0	1.5	1.1	1.4	1.8
2	10	1.3	0.8	1.2	1.8	1.3	1.0	1.2	1.6
2	20	1.0	0.6	1.0	1.5	1.0	0.8	1.0	1.3
2	30	0.8	0.4	0.8	1.3	0.8	0.6	0.8	1.0
2	40	0.6	0.1	0.5	1.0	0.6	0.5	0.6	0.8
3	0	1.5	1.0	1.4	2.0	1.5	1.1	1.4	1.8
3	10	1.2	0.8	1.2	1.8	1.2	1.0	1.2	1.6
3	20	1.0	0.6	1.0	1.5	1.0	0.8	1.0	1.3
3	30	0.8	0.4	0.8	1.3	0.8	0.6	0.8	1.0
3	40	0.6	0.1	0.5	1.0	0.6	0.5	0.6	0.8
4	0	1.4	1.0	1.4	2.0	1.4	1.1	1.4	1.8
4	10	1.2	0.8	1.2	1.7	1.2	0.9	1.2	1.5
4	20	1.0	0.6	1.0	1.5	1.0	0.8	1.0	1.3
4	30	0.8	0.4	0.7	1.3	0.8	0.6	0.8	1.0
4	40	0.6	0.1	0.5	1.0	0.6	0.4	0.6	0.7

Table 9—Mean, P90, P50 and P10 of project PI determined using real oil-price movements (variable price) as well as with deterministic pricing (deterministic price). Results presented for various fixed (FC) and variable (VC) operating costs.

PROJECT WELL LIFE									
FC	**VC**	**Variable Price**				**Static Price**			
$M	$/bbl	Mean	P90	P50	P10	Mean	P90	P50	P10
0	0	454	480	480	480	454	480	480	480
0	10	454	480	480	480	454	480	480	480
0	20	453	480	480	480	454	480	480	480
0	30	443	480	480	480	454	480	480	480
0	40	358	12	480	480	454	480	480	480
1	0	391	48	480	480	401	48	480	480
1	10	388	44	480	480	400	45	480	480
1	20	382	39	480	480	398	41	480	480
1	30	362	30	480	480	397	35	480	480
1	40	278	9	480	480	394	28	480	480
2	0	367	35	480	480	383	36	480	480
2	10	362	33	480	480	382	33	480	480
2	20	353	28	480	480	379	30	480	480
2	30	329	22	480	480	376	25	480	480
2	40	246	8	162	480	372	19	480	480
3	0	348	30	480	480	368	30	480	480
3	10	342	27	480	480	366	28	480	480
3	20	331	24	480	480	363	25	480	480
3	30	304	18	480	480	359	21	480	480
3	40	221	7	83	480	354	14	480	480
4	0	332	26	480	480	355	26	480	480
4	10	324	24	480	480	352	24	480	480
4	20	311	21	480	480	349	21	480	480
4	30	283	16	403	480	344	18	480	480
4	40	203	6	58	480	337	12	480	480

Table 10—Mean, P90, P50 and P10 of project well life determined using real oil-price movements (variable price) as well as with deterministic pricing (deterministic price). Results presented for various fixed (FC) and variable (VC) operating costs.