DETERMINISTIC AND STOCHASTIC ANALYSES TO QUANTIFY THE RELIABILITY OF UNCERTAINTY ESTIMATES IN PRODUCTION DECLINE MODELING OF SHALE GAS RESERVOIRS

By Brent L. Johanson

A thesis submitted to the Faculty and the Board of Trustees of the Colorado School of Mines in partial fulfillment of the requirements for the degree of Master of Science (Petroleum Engineering).

Golden, Colorado

Date: ________________

Signed: ________________________

Brent L. Johanson

Signed: ________________________

Dr. Yu-Shu Wu
Advisor

Golden, Colorado

Date: ________________

Signed: ________________________

Dr. William Fleckenstein
Professor and Head
Department of Petroleum Engineering

ABSTRACT

Decline curve analysis seeks to predict the future performance of oil and gas wells by fitting a mathematical function to historical production data and extrapolating its trend into the future. A recurring issue associated with decline curve analysis centers on the reliability of these predictions. To evaluate the reliability of decline curve predictions, stochastic analyses, such as the bootstrap method, can be used to generate a range of cumulative production or estimated ultimate recovery outcomes. By comparing whether the actual cumulative production of a series of wells falls between certain probabilistic estimates generated by the bootstrap method, the reliability of the bootstrap method and the mathematical models used to fit a curve to the historical production data can be evaluated. In addition, the reliability of predications based on certain numbers of months of production data can be examined. Although prior research implemented the bootstrap method with the Arps model to predict future production or evaluate and improve the reliability of stochastic estimates in conventional wells, this research is the first to also use the Duong model as well, use various ranges of months of production data to determine decline model parameter values, and use shale gas production data to evaluate the reliability of both deterministic and stochastic estimates based on decline curve analyses.

Based on an examination of historical production data from horizontal and vertical gas wells in the Barnett Shale, this research finds that when 12 or more months of production data are available, the bootstrap method can be used to reasonably predict the next five years of production for both vertical and horizontal wells. Increasing the number of months of available generally improves both the reliability and precision of

predictions. The range of values by which deterministic and stochastic predictions vary from the actual cumulative production of groups of wells deceases by about 5% to 10%, and the median values of such distributions of predictions become about 10% to 15% closer to the actual production values, as an additional 6 to 12 months of data become available for modeling decline curve parameter values. The research methodology undertaken in this research can be easily applied to production data from other formations to assist engineers in making predictions about the future performance of wells in other geologic and geographic contexts.

TABLE OF CONTENTS

ABSTRACT ii

LIST OF FIGURES viii

LIST OF TABLES x

ACKNOWLEDGEMENTS xi

CHAPTER 1 INTRODUCTION 1

CHAPTER 2 BACKGROUND INFORMATION 5

2.1 What Methods Can be Used to Calculate Reserves Estimates? 5

2.2 Deterministic and Stochastic Approaches 11

2.3 Manual Fit vs. Computer Regression 12

2.4 Choice of Mathematical Model 13

2.5 How Effectively do the Arps and Duong Models Predict Future Production? 20

2.6 Which Data Will Be Used: All, or Just a Subset? 21

2.7 The Bootstrap Method 22

CHAPTER 3 METHODOLOGY 29

3.1 Selection of Field Production Data Sets 29

3.2 The Four Principal Decisions 32

3.2.1 Deterministic vs. Stochastic Approach 32

3.2.2 Manual Fit vs. Computer Regression 33

3.2.3 Choice of Mathematical Model 33

3.2.4 Selection of Various Ranges of Months of Production Data 33

3.3	Creation of Synthetic Data Sets from Field Data	36
3.4	Sampling of Individual Production Data Points vs. Blocks of Residuals	36
3.5	Sorting of Data to Permit Better Curve Fitting	39
3.6	Determination of Number of Synthetic Data Cases	41
3.7	Summary of Methodology	44
CHAPTER 4	RESULTS AND DISCUSSION	45
4.1	Verification Procedures	45
4.2	Effect of Data Availability Upon the Reliability of Decline Curve Predictions	46
4.3	Effect of Eliminating Early Time Data Upon the Reliability of Decline Curve Predictions	61
4.4	Dimensionless Bracket Widths	62
4.5	Comparison of This Research with Cheng's Modified Bootstrap Method	71
CHAPTER 5	CONCLUSIONS	75
5.1	Conclusions Concerning the Number of Months of Production Data Used to Model Decline Curve Properties	76
5.2	Conclusions Concerning Deterministic Predictions	76
5.3	Conclusions Concerning Stochastic Predictions	77
5.4	Comparison of the Performance of the Base Case and the P50 Case	77
5.5	Conclusions Concerning the Precision of Stochastic Estimates	78
5.6	Application of the Approach in This Research to Other Shale Formations	78
CHAPTER 6	FUTURE RESEARCH	79

REFERENCES CITED 81

APPENDIX A HORIZONTAL WELL DATA USED IN THIS RESEARCH 83

APPENDIX B VERTICAL WELL DATA USED IN THIS RESEARCH 87

LIST OF FIGURES

Figure 2.1: Graphical Representation of Duong Method to Find q_1 18

Figure 2.2: Probability Case Determinations 23

Figure 3.1: Barnett Vertical Well Restimulations 30

Figure 3.2: Examples of Production Rate Data, Before and After Cleanup....34

Figure 3.3: Example Residuals from Data Sets of Horizontal Wells 38

Figure 3.4: Example Residuals from Data Sets of Vertical Wells 38

Figure 3.5: Example Synthetic Data Sets 40

Figure 3.6: Comparison of Synthetic Data Sets and Original Data Set....42

Figure 3.7: Cumulative Distribution Frequencies with Various Numbers of Synthetic Data Sets 43

Figure 3.8: Overlapping Nature of Cumulative Frequency Distributions with 500 Synthetic Data Sets 43

Figure 4.1: Example Decline Curve, Arps Model 47

Figure 4.2: Example Cumulative Frequency Curve, Arps Model 47

Figure 4.3: Example Decline Curve, Duong Model 48

Figure 4.4: Example Cumulative Frequency Curve, Duong Model 43

Figure 4.5: Histograms from Horizontal Well Set, 72 Months of Production Data....50

Figure 4.6: Histograms from Vertical Well Set, 72 Months of Production Data....50

Figure 4.7: Example Cumulative Production Curves, Horizontal Well, Arps Model, Various Numbers of Months of Production Data 53

Figure 4.8: Example Cumulative Production Curves, Horizontal Well, Duong Model, Various Numbers of Months of Production Data 54

Figure 4.9: Histograms from Horizontal Well Set, Arps Model, Various Numbers of Months of Production Data 55

Figure 4.10: Histograms from Horizontal Well Set, Duong Model, Various Numbers of Months of Production Data 56

Figure 4.11: Histograms from Vertical Well Set, Arps Model, Various Numbers of Months of Production Data 57

Figure 4.12: Histograms from Vertical Well Set, Duong Model, Various Numbers of Months of Production Data 58

Figure 4.13: Histograms from Horizontal Well Set, Arps Model, Various Numbers of Initial Months of Production Data Eliminated 63

Figure 4.14: Histograms from Horizontal Well Set, Duong Model, Various Numbers of Initial Months of Production Data Eliminated 64

Figure 4.15: Histograms from Vertical Well Set, Arps Model, Various Numbers of Initial Months of Production Data Eliminated 65

Figure 4.16: Histograms from Vertical Well Set, Duong Model, Various Numbers of Initial Months of Production Data Eliminated 66

Figure 4.17: Log Production Rate vs. Log Time Plot, Horizontal Wells 67

Figure 4.18: Log Production Rate vs. Log Time Plot, Vertical Wells 67

Figure 4.19: Dimensionless Bracket Width, Horizontal Wells, Arps Model 69

Figure 4.20: Dimensionless Bracket Width, Horizontal Wells, Duong Model 69

Figure 4.21: Dimensionless Bracket Width, Vertical Wells, Arps Model 70

Figure 4.22: Dimensionless Bracket Width, Vertical Wells, Duong Model 70

LIST OF TABLES

Table 3.1: Number of Wells for Which Arps Model Converged Upon a Solution ...35

Table 4.1: Histogram Centers 60

Table 4.2: Ranges of Dimensionless Bracket Widths 72

ACKNOWLEDGMENTS

I would like to thank Dr. Yu-Shu Wu of the Energy Modeling Group in the Petroleum Engineering Department at the Colorado School of Mines for the generous guidance and support during my time as a graduate student at the School of Mines.
I would also like to thank Tom McCoy at Cimarex Energy for providing the production data used in this research, and Dr. Todd Hoffman and Dr. Jennifer Miskimins for their valuable feedback on the methodology and results presented in this thesis. In addition, I would like to thank all of my committee members for stimulating classes they taught at the Colorado School of Mines.

CHAPTER 1

INTRODUCTION

In several petroleum engineering contexts, great value is placed on adequately predicting what will happen with future oil and gas production rates. Decline curve analysis (DCA) is one method commonly used to predict these production rates. The basic premise behind DCA is that past production history can be used to predict future production performance. Since production rates tend to follow a falling trajectory over time, a mathematical equation can be used to generate a curve that follows the overall decline trend. The parameters within that mathematical expression are tuned to fit the curve to the actual production data. By then extrapolating that trend into the future, we can estimate not only how much the well will produce, but also the timeframe over which this production will be realized.

Although it is common to apply DCA to an individual well, the method is actually much more versatile. Decline curve analysis has been applied to a group of analog wells with similar geologic properties or within close geographic proximity to one another. It has also been used to model production decline in an entire field or formation. Knowledge about the production decline behavior of any of these producing entities - individual well, group of analog wells, field, or formation – can be gleaned from DCA methods.

Decline curve analysis serves a number important purposes within the petroleum industry. Examples include reserves estimates for SEC reporting requirements, where estimated ultimate recovery (EUR) values are modeled to guide a corporation's managers, current shareholders, and prospective investors. Credible and consistent reserves estimates allow managers to allocate capital for competing alternative uses. Good reserves estimates also enable investors to better evaluate the financial position of various operators, which serves to bolster investor confidence in the capital markets. In addition, decline curve analysis assists the economic appraisal of producing entities. Since the discounted present value of the entity's future income stream ultimately determines its worth, a better understanding of the production that will flow out of oil and gas wells over their lifetime will lead to more accurate determinations of value. Decline curve analysis is also used in exploration activities, where the lifetime production potential of a proposed well, or target, is evaluated. Finally, DCA finds application in decisions about whether various completion or maintenance options are economically justified.

In all of these cases, a more advanced understanding of decline curve methods and models will enable more reliable estimates about the future output of producing entities. These better estimates will, in turn, assist decision makers in designing more informed courses of action.

Engineers and other analysts so commonly use DCA to peek into the future because it is easy to apply in practice. In terms of data, DCA requires only production rate data (often considered on a monthly basis) as a function of time. In addition, DCA tends to not be unreasonably time consuming. This aspect of the method is critical when elaborate data collection is cost prohibitive, as in the case of wells with relatively low production rates, or when many data sets must be re-evaluated on a periodic basis, such as reserves estimates for SEC reporting.

Despite all the potential benefits that DCA offers, its application may also bear some pitfalls. Practitioners today use DCA extensively to model production decline in shale formations. But DCA, and its classic mathematical equation, the Arps model, were developed empirically to describe conventional reservoirs (Arps 1945). The physical attributes of producing shale formations, and the flow within them, are fundamentally different than those found in conventional reservoirs. As will be discussed in the next section of this paper, most of the methods for estimating reserves in conventional reservoirs don't work particularly well when applied to unconventional reservoirs, particularly ultra-tight shale reservoirs. The assumption that the old paradigm can be validly applied to shale reservoirs is therefore suspect, and merits further investigation.

This research seeks to answer three main questions in the context of decline curve analysis of production data from the Barnett Shale formation:

1. How effectively can two decline curve models, the classic Arps hyperbolic model and the more recently proposed Duong power law model (Duong 2011), predict future production decline trends in groups of vertical and horizontal Barnett Shale wells?
2. How reliably can the bootstrap method, which creates synthetic production data sets by randomly sampling the original production data set for the purpose of modeling alternative outcomes (Efron et al. 1998), predict the actual cumulative production that will occur?
3. How many months of production data are needed to generate reliable cumulative production estimates?

In examining these themes, this research seeks to develop a better understanding of how well the Arps and Duong models work, in practice, when applied to actual production data from wells in the Barnett Shale. It also attempts to shed light upon issues related to the reliability of probabilistic reserves estimates generated in stochastic models. As will be seen, an important element of such a process is its ability to incorporate data from many well sets, including wells with production decline data points that deviate significantly from their overall trend line.

CHAPTER 2

BACKGROUND INFORMATION

A number of methods exist to estimate reserves in conventional formations. A brief review of some of these methods follows, including an examination of the challenges associated with applying them to shale formations. The decline curve method is then examined, with a review of its history, some of the various models proposed to fit curves to production data, and other variations in its implementation.

2.1 What Methods Can Be Used to Calculate Reserves Estimates?

Reserves estimates in shale resource plays can be difficult to calculate because of the features that characterize shale reservoirs. Shale formations tend to occur on a basin-wide scale, but this does not mean that all areas within a shale formation possess the attributes needed to produce at economic rates. The poorly defined physical boundaries of sweet spots, or favorable target locations, make even an initial determination of reservoir volume problematic. In addition, heterogeneity of rock properties such as directional permeability and pore storage capacity complicate matters further. Considerations related to the in-place hydrocarbon fluids, such as the amount of organic kerogen and its thermal maturity, the degree of gas adsorption, and the ratio of free to adsorbed gas, throw additional ambiguity into the mix. Finally, geomechanical concerns (stress anisotropy) influence how well fluids might be produced. The rock's in-situ stress state drives its ability to form complex fracture networks, which in turn may greatly impact hydrocarbon production. All of these factors bear upon the volume and rate at which hydrocarbons can be produced from a shale formation.

The challenge to petroleum engineers is that information about most of these influences is difficult and expensive to obtain in practice, especially in the context of a particular producing location. Various methods of estimating reserves have been developed for conventional reservoirs, where they have proven reliable and effective. Problems exist, however, when attempting to apply these techniques to shale formations. Lee et al. (2010) reviewed a number of these reserves estimation methods and evaluated their applicability to shale reservoirs.

The volumetric method calculates the total volume occupied by oil/gas and estimates a recovery factor. Lee et al. observed that this technique has been successful in conventional reservoir systems, but that in less well-developed plays it is difficult to determine a producing zone's areal extent and net pay thickness, due to the geologic complexity over the basin- wide. In unconventional reservoirs, Lee et al. noted that low permeability may greatly affect the recovery factor, since gas molecules relatively close to a conductive fracture may not be able to migrate to the fracture at an economically reasonable rate.

When thinking about how or whether to apply the volumetric method to shale gas formations, it is not clear that the concept of a recovery factor even makes much sense. In conventional systems, the magnitude of the recovery factor is driven primarily by considerations of relative permeability, fluid viscosity, and wetting phase. All of these factors that propel or impair hydrocarbon flow in a conventional system are fluid properties. But in a shale formation, the greatest hindrance to gas flow is low permeability, a rock property. The main distinction between the two cases is that a recovery factor driven more by fluid properties may play out much faster than one driven by low permeability. Whereas the time frame of production drawdown in conventional oil fields (to the point where oil no longer flows unless artificial lift is applied) is usually on the order of months or years, the corresponding time for tight gas and shale gas wells is much longer. Because the recovery factor in a shale formation is therefore driven by different forces and operates on an expanded time scale, crudely borrowing and applying the recovery factor concept from conventional contexts may be more misleading than instructive.

Another method to estimate gas reserves in conventional formations is a material balance technique that compares reservoir pressure (normalized by the prevailing z-factor) with cumulative production. Relating p/z with cumulative production enables detection of a linear trend (ideally) as the reservoir pressure drops with continued gas production (Craft et al. 1991, Lee et al. 2010). This technique relies upon the assumption of a constant drainage volume, such that the observed pressure drop relays useful information about the total amount of gas initially present in the reservoir. It also assumes boundary dominated flow, where pressure equilibrium has stabilized throughout the reservoir. Lee et al. note that this material balance technique cannot be legitimately applied to shale reservoirs, since pressure equilibrium within the reservoir volume takes so long to establish, and drainage volumes may even increase with time as the pressure wave slowly extends through the tight rock matrix.

A third reserves estimation technique is to compare performance of similar (analog) wells to the well under evaluation, the target well. This analogy method requires that the analog well or reservoir must be at more advanced stage of depletion than the target, have undergone the same completion techniques as the target, and possesses recovery mechanisms no more or less favorable than those of the target (Lee et al. 2010). To apply this method to shale reservoirs presents several challenges. Analogs can be difficult to identify due to the relatively short development history of shale plays. In addition, a "moving target" dynamic, in the form of technology advancements in the ways that wells are drilled and completed, means that it is difficult in practice to draw analogies between previously developed and current shale wells.

Next, reservoir simulation can assist with developing reserves estimates. This technique uses equations which model the physical processes that occur in a reservoir. Parameter values of various reservoir properties are calibrated until the production decline trend predicted by the simulator matches the historical production data. The simulation then predicts the reservoir and wellbore pressure drops, saturation changes, and production rate over some period into the future. Reservoir simulation can provide insights into reservoir dynamics, and enables modeling of real systems through the history matching of actual production data. Because it can be time consuming and therefore expensive to undertake, reservoir simulation is a good method for estimating reserves in big hydrocarbon deposits (e.g. offshore leases), or where several operators are willing to work together to share the cost of the simulation and coordinate their production activities to take advantage of the insights that it provides. Reservoir simulation presents the disadvantage of non-uniqueness of parameter values: a number of different sets of parameter values may equally well describe the observed production history, yet may yield different predictions of future reservoir behavior. Another problem is that for small hydrocarbon deposits, full-scale simulation is too costly in relation to the estimated value of the well's future production. Finally, when numerous operators control small, fragmented acreage blocks in close proximity to one another, they may be unwilling to cooperate and share the expense of reservoir simulation.

Most of the methods developed to estimate reserves in conventional formations therefore possess significant shortcomings when applied to shale reservoirs. With regard to another method, decline curve analysis, it is not difficult to apply a trend line to historical production data. The real question though, is how reliably that trend, when extrapolated into the future, can predict future production and therefore generate an acceptable estimate of reserves. Even though DCA is thus easily applied in a technical sense, it is not clear at the outset whether or not it works very well generally or as applied to specific production decline scenarios.

The last several years have seen great debate within the petroleum engineering community with regard to the propriety of applying traditional DCA techniques to model shale formation production trends. Led notably by Art Berman (2012), critics of this practice have claimed that many operators significantly overestimate their reserves by engaging in untenable modeling practices. Berman studied production decline in hundreds of Haynesville Shale wells and concluded that the reserves estimates of a number of the major operators in that region were highly questionable.

Berman proposes a fluid flow/production model in which incremental production over time originates from ever-smaller fractures, until it ultimately becomes dominated by production from the matrix. Production begins with induced hydraulic fractures and large natural fractures, then flows primarily from microfractures, and finally originates from a slow matrix “bleed”. Berman argues that we should not include in an Arps-based DCA the early time production, as it represents an unstable transient flow regime, dominated by flow in the large fractures. The effect of including this early time production is to introduce curvature into the decline curve that causes the curve's tail to pitch upward (or to pitch downward less slowly). This curvature extends the production tail unreasonably, implying greater reserves than will actually be realized. In addition, Berman argues that formation type curves, generated by aggregating production rates from multiple wells, introduce a strong survivorship bias into the analysis. A survivorship bias exists when the better (more prolific) wells survive a longer period of time and bias the ever-smaller group of wells that define the type curve.

Berman contends that the actual commercial life of Haynesville wells is 8-10 years, compared with the 40-65 years projected by some operators. Moreover, the discounted present value of the production that will occur so far into the future is negligible compared to that which will occur in the first 20 years. Thus, Berman argues, reserves estimates generated by several major operators are misleading and overstated, by as much as 450%.

This controversy about reserves projections based on DCA illustrates that more study of the method is needed. A number of approaches in how to apply DCA have been developed, and yet more may appear with future attempts to refine (or reinvent) the technique. Currently, there are four main axes of consideration when applying decline curve analysis to production decline data:

1. whether to use a deterministic or stochastic approach,
2. when modeling a best fit decline curve to the historical production data, whether to manually fit the line or to use computer regression,
3. which mathematical model to use, and

4. which production data will be used: all of it, or only some subset.

Each of these elements is next considered in turn.

2.2 Deterministic and Stochastic Approaches

Several methods have historically been used to implement a DCA approach, including deterministic and stochastic ones. The deterministic approach creates three EUR scenarios, each of which uses a single value for each parameter in the mathematical equation that fits a curve to the production decline trend. These scenarios constitute the very probable, most likely, and less probable estimates of how production will fall off over time. The implicit assumption is that the actual future decline trend will be between the high and low estimates, and somewhat close to the most likely estimate, most of the time (SPE-PRMS 2011). An even simpler version of the deterministic approach would be to use nonlinear regression techniques to determine the modeling equation's parameter values, and use those values to extrapolate a single EUR value. In this case, the resulting EUR value would, by default, constitute the most likely estimate.

The deficiency of the deterministic approach is that it does not incorporate or quantify uncertainty as part of its production decline evaluation. The very probable, most likely, and less probable scenarios it creates cannot be said to constitute the P90, P50, and P10 cases, respectively. In other words, a deterministic approach does not predict, in a statistical sense, that 10% of the time the actual EUR will exceed the P10 case, that 90% of the time the actual EUR will exceed the P90 case, and that 80% of the time the actual EUR will fall between the high and low estimates.

A stochastic analysis approach, on the other hand, does incorporate this type of uncertainty evaluation. Also known as a Monte Carlo analysis, the stochastic approach employs a probability simulator that conducts multiple iterations where a value for each parameter is randomly selected, based on pre-assigned parameter value ranges and distributions. The simulator then compiles statistics about EUR outcomes, including the frequencies of various outcomes and the relative influence of the various parameters. In practice, it can be difficult to reliably execute a stochastic analysis because the ranges and distributions of the various parameter values must be known or estimated, and this information is hard to determine when only a small set of analog wells is available or when the cost of collecting this data is high.

2.3 Manual Fit vs. Computer Regression

Decline curves can be fit to production data either manually or by automated computer regression algorithms. Jochen et al. (1996) reprint a quote attributed to Steven Holditch that suggests the inherent weaknesses of the manual-fit method:

> "The estimator…will then pick up his pencil and straight-edge, squint through one eye, stick his tongue out the corner of his mouth, and rely on his experience to make a reasonable pick of the decline."

Although technology has advanced to enable an analyst to drag portions of the computer generated decline curve with a mouse, the process is still highly subjective. Different analysts are likely to pick different curve paths and trajectories, so the technique of manually fitting decline curves suffers from the potential problem of non-reproducibility.

Nonlinear regression techniques executed by computer algorithms solve the problem of subjectivity. One commonly used nonlinear regression technique, called the sum of least squares, considers a series of possible curves that might be fit to the data. For a given mathematical model that describes a nonlinear function, the difference between the squares of the respective data points and the corresponding points on the proposed curve are summed and this sum is compared with the sums similarly generated with respect to other proposed curves. The curve with the minimum sum of squares is chosen as the one that best fits the original data points. The benefit of nonlinear regression techniques is their repeatability: the same solution will occur across multiple executions of the computer algorithm as applied to the same data set. The disadvantage of computer-driven regression is that the computerized algorithm may not be able to converge upon a reasonable solution for some data sets.

2.4 Choice of Mathematical Model

In addition to multiple analysis methods, a variety of decline curve models have been proposed. Each of these models uses distinct mathematical equations to describe production decline trends. The classic (and most prominent) formulation is the Arps model which describes decline as a hyperbolic function that depends on initial decline rate, initial production rate, and a factor that describes the curvature of the decline trend, known as the Arps b-factor. This model was first described in the 1940's and has found widespread acceptance for conventional reservoir applications. It consists of three forms, categorized on the basis of the value of the b-factor. The most general of these, when the value of b is greater than zero but less than unity, is called the hyperbolic form and is as follows:

$$q = q_i \frac{1}{(1+bD_i t)^{\frac{1}{b}}} \tag{2.1}$$

The exponential form of the equation arises when the value of b equals zero, and is expressed as:

$$q = q_i \frac{1}{e^{D_i t}} \tag{2.2}$$

Finally, the harmonic form of the Arps equation, when the value of b equals unity, takes the form:

$$q = q_i \frac{1}{1+D_i t} \tag{2.3}$$

Although the Arps model was developed in a purely empirical manner, it was later shown that the exponential form of the equation could be developed analytically for slightly compressible systems, provided that the conditions of constant bottomhole pressure, production of the well at or near the reservoir capacity, and a constant drainage area (such that boundary dominated flow has developed) are assumed (Fetkovich et al. 1996; Ahmad 2006). The requirement of boundary-dominated flow poses a major challenge to successfully applying the Arps hyperbolic decline function to shale reservoirs. Whereas conventional systems will likely change from a transient flow regime to a boundary-dominated one within a relatively short period of time (on the order of days or weeks), their shale counterparts will often require years or even decades to make this transition. Since the first several years of production data collected from a shale reservoir are likely based on a transient flow regime, these data cannot be used to accurately determine the Arps b-factor. When such a practice has been attempted, the value of the Arps b-factor has often been found to exceed unity. This is a physically impossible result because it suggests that cumulative production can increase without limit (Lee & Sidle 2010), a situation that likely arises due to the use of transient flow data when only boundary-dominated data can be legitimately used (Fetkovich et al. 1987). Over time, it is usually observed that the b-factor decreases to sub-unity values as the flow regime transitions from transient to boundary-dominated. At this point the Arps method becomes more reliable in predicting future production trends.

The real utility of decline curve analysis, though, is to enable a prediction of future production based on whatever rate data is available, and often this is limited to rate data from relatively early times. When a lengthy time period must be endured in order for a model to obtain adequate predictive power, the usefulness of the tool has been diminished. The Arps method may therefore inadequately predict future production rates in low permeability reservoirs that require longer time periods to exhibit boundary-dominated flow, as it may tend to misestimate the future production, estimated ultimate recovery, and reserves in place when only transient period production data are available.

To correct an overly optimistic prediction of future events that occurs when the Arps b-factor exceeds unity, an engineer applying a hyperbolic function to production decline should recognize that the value of the b-factor must decrease at some future, unknown point in production. He or she must then declare *when* that point in time will be, and *how* the value of the b-factor will change. In practice this often consists of predicting a transition from hyperbolic to exponential decline once the instantaneous decline rate crosses below some specified threshold, and also includes specifying the value of that exponential decline rate from the point of transition to the end of production. This practice amounts to estimating when and how the b-factor will change, a process based on neither scientific nor empirical understanding, but instead upon intuition or mere guesswork.

In addition, it is relevant to consider that the empirical origins of the Arps model are production data from conventional reservoir systems. There is simply no reason to assume that this framework can find valid extension to shale reservoirs, where the physical processes associated with production and pressure depletion are markedly different. For example, shale reservoirs are notable for their distinctly low permeability and the presence of a high amount of adsorbed gas, conditions not observed in conventional contexts.

In an attempt to overcome the deficiencies of the Arps decline curve model, several other models have been developed over the past few years in response to the greater interest in and importance of shale reservoirs. Among these are the Duong power law model, an empirical formulation that considers the dominance of fracture flow over matrix flow in shale reservoir systems that have been subjected to extensive hydraulic fracturing (Duong 2011). Since production data of shale reservoirs are commonly observed to exhibit fracture-dominated flow regimes (linear and bilinear flows) rather than matrix-dominated ones (pseudo-radial and boundary-dominated flows), the contribution of matrix flow to overall production is thought to be small compared with that of fracture flow. The concept of drainage area used in conventional systems therefore cannot be properly applied to shale ones to estimate future production and EUR. Instead, Duong postulated that fracture densities within producing shale formations will increase over time due to changes in local stresses that reactivate fractures and faults. This process of fracture creation occurs due to changes in effective stress that result from the pressure depletion that accompanies fluid production, and has the effect of opening new reservoir volumes to production. The reactivation of fractures and faults constitutes the creation of higher permeability channels that can contribute to production. Duong's analysis of data from several formations indicated that plotting gas production rate (normalized by cumulative gas production) vs. time on a log-log chart reveals a straight line that takes the form

$$\frac{q}{G_p} = at^{-m} \tag{2.4}$$

where q is the production rate, G_p is the cumulative gas production, a is the y-intercept and m is the slope of the line. The slope is always negative, and m takes a value greater than unity for production data from shale formations (Duong 2011). Using equation (2.4) along with production data and nonlinear regression techniques, values for parameters a and m can be calculated. Next, an initial production rate, q_1, can be calculated according to the relationship

$$q_1 = \frac{q}{t(a,m)} \tag{2.5}$$

where

$$t(a,m) = t^{-m} e^{\left(\frac{a}{1-m}\right)(t^{1-m}-1)} \tag{2.6}$$

With values for a, m, and q1 thus determined, the decline curve can be plotted as a function of time, as

$$q(t) = q_1\, t(a,m) \tag{2.7}$$

When the line generated by equation (2.7) does not pass though the origin due to wellbore operating conditions, an adjustment factor, q_∞, can be added:

$$q(t) = q_1\, t(a,m) + q_\infty \tag{2.8}$$

Figure 2.1 shows a graphical interpretation of this concept.

From equation (2.8), an expression for cumulative production as a function of time can be derived as

$$G_P(t) = \frac{q_1}{a} e^{\left(\frac{a}{1-m}\right)(t^{1-m}-1)} \tag{2.9}$$

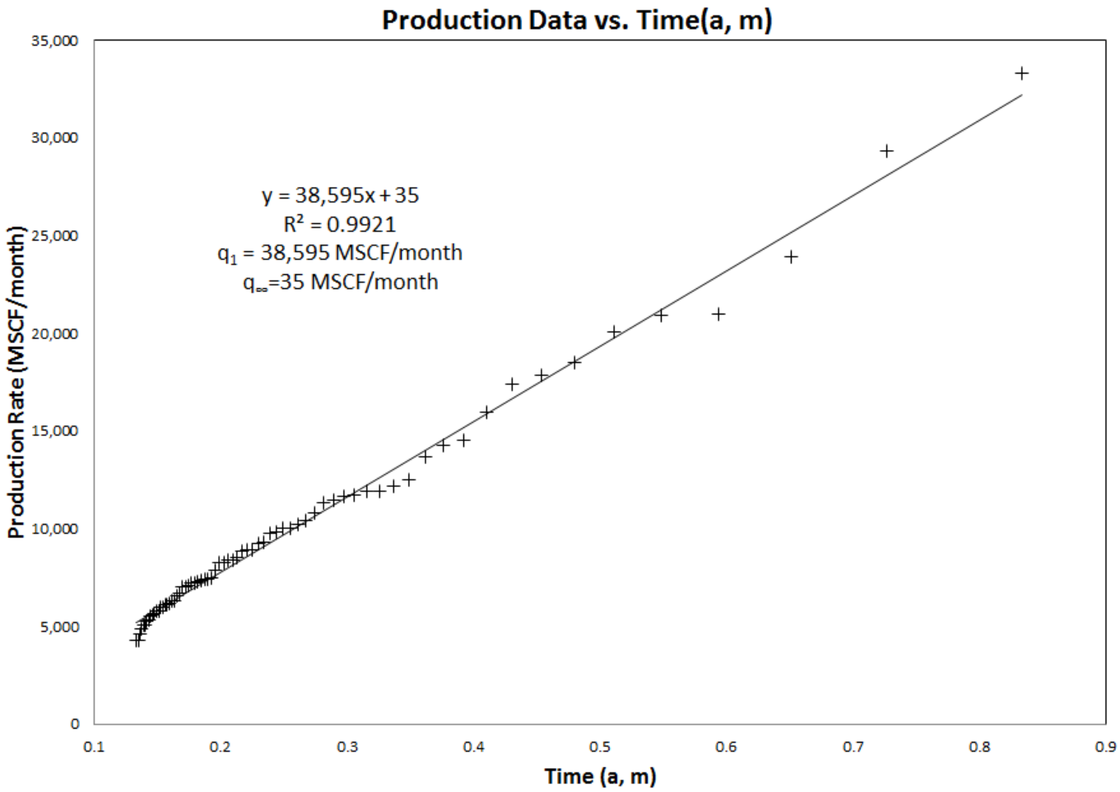

Figure 2.1: Using the Duong approach to model production decline, an initial production rate, q_1, is calculated as the slope of the line when production rate is plotted against time expressed as a function of parameter values *a* and *m*.

Duong found that these empirical relationships held good predictive power when applied to several field examples, an assertion later confirmed independently (Lee 2012, as cited in Mishra 2012), thereby bolstering confidence in the method's efficacy when applied to producing shale formations contexts.

In addition to the Arps hyperbolic and Duong power law models, other models have been proposed to describe production decline in shale reservoirs. These include Ilk's power law model (Ilk 2008), Valko's Stretched Exponential Production Decline (SEPD) model (Valko 2010), Weibull's Growth Model (Weibull 1951), as well as the Generalized Likelihood Uncertainty Estimation (also known as GLUE), which combines the predictions of several of these models (Mishra 2012). When put to the task of predicting future production rates, the problem of non-uniqueness arises. Different models predict varying reserves quantities, but it is not clear whether one model is better than the others in a particular reservoir application (Mishra 2012). To identify which of the models may be more contextually appropriate than the others, some of the creators of these models have suggested flow regimes or other physical phenomena associated with the operation of their models. For example, Duong cites successively opening natural fractures, and Ilk postulates the presence of a linear flow regime. If it could somehow be known that these phenomena were occurring, one model might be identified as preferable over the others. A fair criticism of this thought process is that these models are entirely empirical, and lack a "direct link with reservoir engineering theory, other than via analogy" (Okouma et al. 2012).

Without a sound theoretical basis to justify a model choice, some researchers have proposed the use of cross plots to detect or infer information about flow regimes. The trends displayed in various charts might inform which model is the most appropriate in a particular context (Okouma 2012). While this concept seems appealing because it attempts to link production decline with a reservoir flow regime, in practice the process remains highly subjective and muddled. A close reading of the proffered analysis techniques fails to identify dispositive methodologies and workflows that can guide practitioners to independently replicable conclusions. The theoretical process of how to select the most appropriate production decline model therefore remains elusive. Perhaps it is better to examine how the decline curve models can be applied in practice to achieve consistent, repeatable predictions.

2.5 How Effectively do the Arps and Duong Models Predict Future Production?

The Arps and Duong models both do an acceptable job of fitting a curve to existing production decline; the execution of this task is neither difficult nor controversial. The more important issue is how reliably these models can predict what will happen with future production rates.

It is intuitive that engineers should apply the model that best conforms with the physical reservoir conditions present in a particular shale reservoir that are most responsible for driving production rates. It is not readily apparent, however, which model will best fit a specific case, especially since a sound theoretical basis does not exist to link various in situ conditions with a particular model, and detailed information about field conditions is often lacking anyway. This is doubly true for new plays where lengthy production history is unavailable. An added complication is that in shale reservoirs, stimulation programs (especially hydraulic fracturing) ideally create complex fracture networks. The presence of low matrix permeability will mean that most flow occurs within rock volumes that are close to fractures, ensuring complex flow realities that are difficult to characterize as a particular flow regime or pattern. In addition, the low matrix permeability makes it hard to conclude that boundary-dominated flow has been attained over a relatively short period of time. It is hard to say whether boundary dominated flow is achievable at all, in a practicable timeframe, in the context of very tight formations. All that may exist for a number of years is something resembling a transient flow regime. Therefore, a sound theoretical basis does not exist to justify the application of the Arps method to model production decline in shale reservoirs.

As previously mentioned, Lee and Sidle (2010) showed that a b-factor in excess of unity will cause cumulative production to increase in an unbounded manner. A practical problem commonly arises when an attempt is made to fit a hyperbolic curve to the production data from a producing shale reservoir. It is often found that the curve that best fits the production decline trend has a b-factor value exceeding unity, and that a curve with a b-factor value less than or equal to unity poorly models the decline. Applying the Arps model to shale reservoir production data therefore suffers from a practical problem as well.

2.6 Which Data Will Be Used: All, or Just a Subset?

The Arps empirical model was developed for use with data from times with boundary dominated flow. Other models may similarly benefit from exclusion of transient flow data, since it may tend to lead to unreliable predictions. In addition, data from early times may include significant influences from sources other than those indicating reservoir deliverability, such as flowback of fracturing fluid, liquid loading within the wellbore, and choking of valves at the surface. The dilemma confronting a practitioner is to how make use of as much of the data as possible without including data corrupted by these other influences. Therefore, if not much data is available to project future production decline, one must carefully consider a decision to discard some of it. The use of charts showing production rate vs. time, plotted on a log-log coordinate axes, can help to identify the emergence of flow regimes and lend confidence to decisions about which production data to include or discard.

2.7 The Bootstrap Method

As previously discussed, a Monte Carlo simulation is a type of stochastic method that generally requires knowledge about parameter value ranges and distributions before

it can be executed. The conventional bootstrap method is a special type of Monte Carlo simulation. This method creates a series of sampled data sets from an original data set and draws certain statistical inferences from that series of sampled sets. Unlike a conventional Monte Carlo simulation, the bootstrap method does not require knowledge or estimates about reservoir or decline curve properties, their statistical distributions, or reservoir flow regimes. These properties are not needed because the bootstrap method relies only upon rearrangements of the original data set to generate outcomes.

The bootstrap method operates in the following manner. The original data set is sampled with replacement (i.e., when a randomly selected data point from the original data set is selected for incorporation into the sampled data set, it is still available for selection from the original data set in future sampling iterations). The sampled data sets, which contain the same number of data points as the original data set, may therefore contain duplicate values from the original data set, and may omit some of the values in the original data set (Efron et al. 1998). When a sufficient number of bootstrap realizations are executed, the distribution qualities of the sampled data set should resemble those of the original data set, and the resulting series of synthetic data sets is reproducible.

The bootstrap method is applied to a production data set for an individual well by undertaking nonlinear regression analysis on the various synthetic data sets, using a specified mathematical model to describe the best fit curve. Based on the model's resulting parameter values for each of the synthetic data sets, a cumulative production value is calculated by integrating the mathematical function over a specified range of months. The set of cumulative production values for the various synthetic data sets forms a distribution of cumulative production values, which serve as proxies for reserves estimates for that well. From this distribution of cumulative production values, the P10, P50, and P90 cases can be determined. These probability cases are defined by their relationship to other values within the distribution. The P10 case, for example, is the cumulative production value for which it is the case that 10% of all values in the distribution exceed this value. Figure 2.2 graphically illustrates the concept of how the P10, P50, and P90 probability cases are determined from the distribution of cumulative production values created by the bootstrap method.

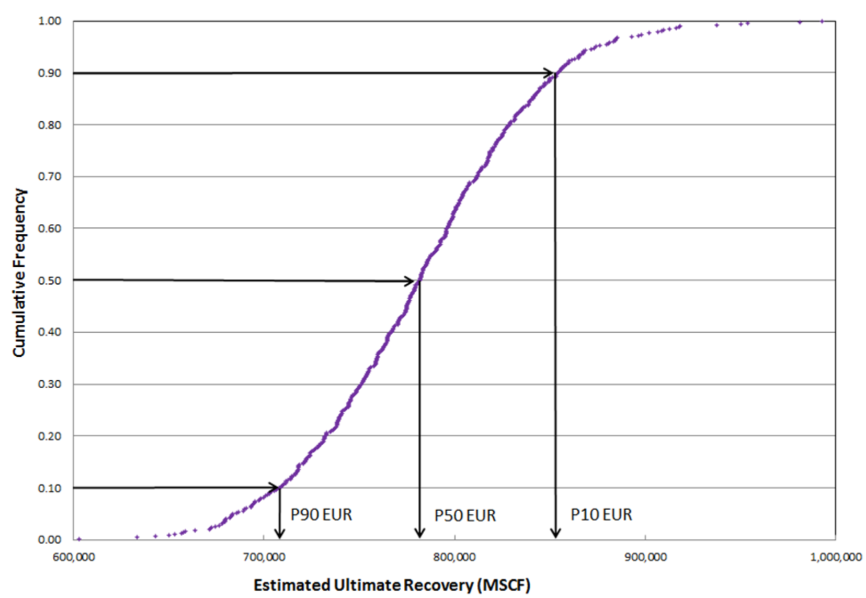

Figure 2.2: The P10, P50, and P90 probability cases are determined according to their relative positions within the cumulative frequency distribution of EUR or cumulative production values generated by the bootstrap method. For example, the P90 case is that EUR value for which it can be said that 90% of the EUR values in the distribution exceed that value.

The bootstrap method is able to generate a confidence interval, defined as the percentage (or fraction) of all outcomes between specified array positions in the distribution that results from the synthetic data sets. For example, when P10 and P90 cases are determined, there is an 80% (or 0.80) confidence interval. In the context of well production data subjected to bootstrap sampling, this indicates that for 80% of the wells in a group we would expect the actual cumulative production to fall between the P10 and P90 estimates. The coverage rate is defined as the observed incidence of the actual cumulative production falling between the P10 and P90 estimates. The reliability of the bootstrap method as applied to the data under consideration can be evaluated by comparing the coverage rate with the confidence interval.

Jochen at al. (1996) applied the bootstrap method to develop stochastic estimates for the future production of an oilfield, using the Arps hyperbolic model and aggregated production data. Jochen addressed some theoretical considerations underlying the bootstrap method, including the method's assumption that data in the original sample are independent and identically distributed. Jochen interpreted this requirement to imply that fluctuations in data are due to measurement error rather than changes in operating conditions. They concluded that since changes in operating conditions are more likely to account for observed noise or fluctuations in production data, the bootstrap method as applied to production decline analysis suffers from an inherent weakness. But to conceive of data point fluctuations as either errors in measurement or changes in operating conditions may be to cast the issue too narrowly. A broader interpretation of these fluctuations might be that they simply represent the reservoir's deviation from the smooth production output that is represented by a decline curve regressed to that production data. This fluctuation could stem from management decisions and logistical considerations at the surface, the effects of liquid loading, or natural reservoir modulations we don't yet fully understand. We should also keep in mind that it is the physical conditions and processes of the reservoir that cause production rates to occur, and not the production rate in one month that causes the production rate in the next month. Viewed in this light, the presence of apparent noise in production data does not call into question the validity of the bootstrap method.

Despite these theoretical concerns, Jochen et al. applied the bootstrap method to probabilistic EUR estimates for an entire field. They created a probability distribution of reserves estimates, from which P10, P90, and P50 values were drawn. Jochen et al. note the practical difficulties of validating their bootstrap-derived production forecasts with actual production outcomes observed in the fields they considered. In practice, field operating conditions change with technology advancements and additional development (e.g. infill drilling, well workovers, etc.), altering the fields' production trajectories and ultimate recoveries. They concluded that it was therefore hard to evaluate how well the bootstrap estimates of EUR actually would predict cumulative production over time, and they were not able to verify their predictions with field data.

More recently, Cheng et al. (2010) applied a modified bootstrap method, and pioneered the technique of comparing group's coverage rate with its theoretically predicted confidence interval. Cheng examined the decline trends of production data from a group of 100 wells, using the Arps hyperbolic equation to model production decline. For each well data set, the research involved fitting a curve, using nonlinear regression techniques, to the production decline data over a given period of time (for example, 70 months). This time demarcation can be thought of as a cutoff point, such that known production data beyond it was not used to model the decline curve parameter values. But rather than randomly sampling the production data points to construct the synthetic data sets, Cheng sampled residual values. A residual value is the distance between an actual production data point and the corresponding data point on the best fit line regressed to the set of production data points.

A set of residual data values was thus determined for each well data set. In addition, Cheng's modified bootstrapping method sampled not individual residuals, but blocks of them. The residual data set was therefore divided into blocks that each contained a number of data points that bore a statistically significant correlation to one another, as determined with an autocorrelation function. These blocks were then randomly sampled and concatenated in the order they were sampled, to form a sampled residuals data set. The sampled residuals data set was then added to the original production data set to form a synthetic data set, and a best-fit curve was regressed to that synthetic data set.

This process of creating synthetic data sets was repeated multiple times for each well. At each instance, the EUR predicted by the synthetic data set was determined, and the set of EURs predicted by the set of synthetic data sets was evaluated to define the predicted P10, P50, and P90 cases for that well over the total period of time for which historical production data was available (for example, 200 months). These cases were then compared with the actual cumulative production over the full production data set for that well, and a question was asked: does the well's actual cumulative production fall between the EURs predicted by the P10 and P90 cases? This process was repeated for each well in the data set. The percentage of wells where the actual cumulative production fell within the P10 and P90 estimates was the coverage rate for that set of wells.

The P10 and P90 cases define the confidence interval for a given well. If the P10 and P90 values that are generated for each well are reliable (that is, if the uncertainty estimate is reliable), we would expect that for 80% of the wells the EUR resulting from the decline curve fit to the original production data (the base case) will fall between the P10 and P90 values for that well. Cheng found that when all of the production data (including early time data) are equally weighted to generate best-fitting curves using the Arps model, the confidence interval was not reliable. In other words, the coverage rate did not approximate the confidence interval.

Cheng found that the confidence interval could be greatly improved, however, by preferentially weighting the data that occurred in closer proximity to the cutoff date, to the point where the coverage rate does approximate the confidence interval. When the coverage rate and confidence interval converge, it is possible to conclude that the confidence interval is reliable. Cheng et al. thus devised a method to evaluate and improve the reliability of the bootstrap method in generating EUR reserves distributions, including P10, P50, and P90 probability cases.

My research extends the prior work done by Jochen and Cheng by applying the bootstrap method using not only the Arps model, but the Duong model as well. In addition, I examine the effect of using different numbers of months of historical production data to model the synthetic data sets, and compare the prediction results from horizontal and vertical well data sets. These research inquiries are in pursuit of a larger goal to develop a practical method that engineers can implement in order to make more reliable cumulative production predictions.

CHAPTER 3

METHODOLOGY

The research undertaken in this project consisted of obtaining production data from wells in the Barnett Shale and conducting various decline curve analysis upon it. This chapter details the steps taken to generate the results that will be described in the next chapter.

3.1 Selection of Field Production Data Sets

The production data (production rate vs. time) used in this research consisted of a group of horizontal wells and a group of vertical wells in the Barnett Shale. The vertical wells were drilled between 2000 and 2006 and are located in Wise, Denton, Parker, Tarrant, and Johnson Counties in east Texas. The horizontal wells were drilled between 2005 and 2006, and are located in Wise and Denton counties. All of these wells are still active and currently producing. All of the data used in this research is publically available production rate data.

At least 72 months of historical production data were available for each of the wells in these sets. Although it had been hoped that a longer time span would be available, many vertical wells in the Barnett appear to undergo restimulation or shut-in approximately 8-10 years into their lifetimes, as seen in Figure 3.1. An ideal scenario would be to have many wells available for analysis, with all wells possessing a long production history, but these considerations can conflict with one another. In addition, it was considered advantageous for processing purposes to use a consistent number of months of production data for all wells in the research. A balance was therefore struck between the timespan over which data was available for a given well, and the number of wells that possessed some minimum number of months of production data. Since there were numerous wells that contained at least 72 months of production history, this timespan was selected for analysis.

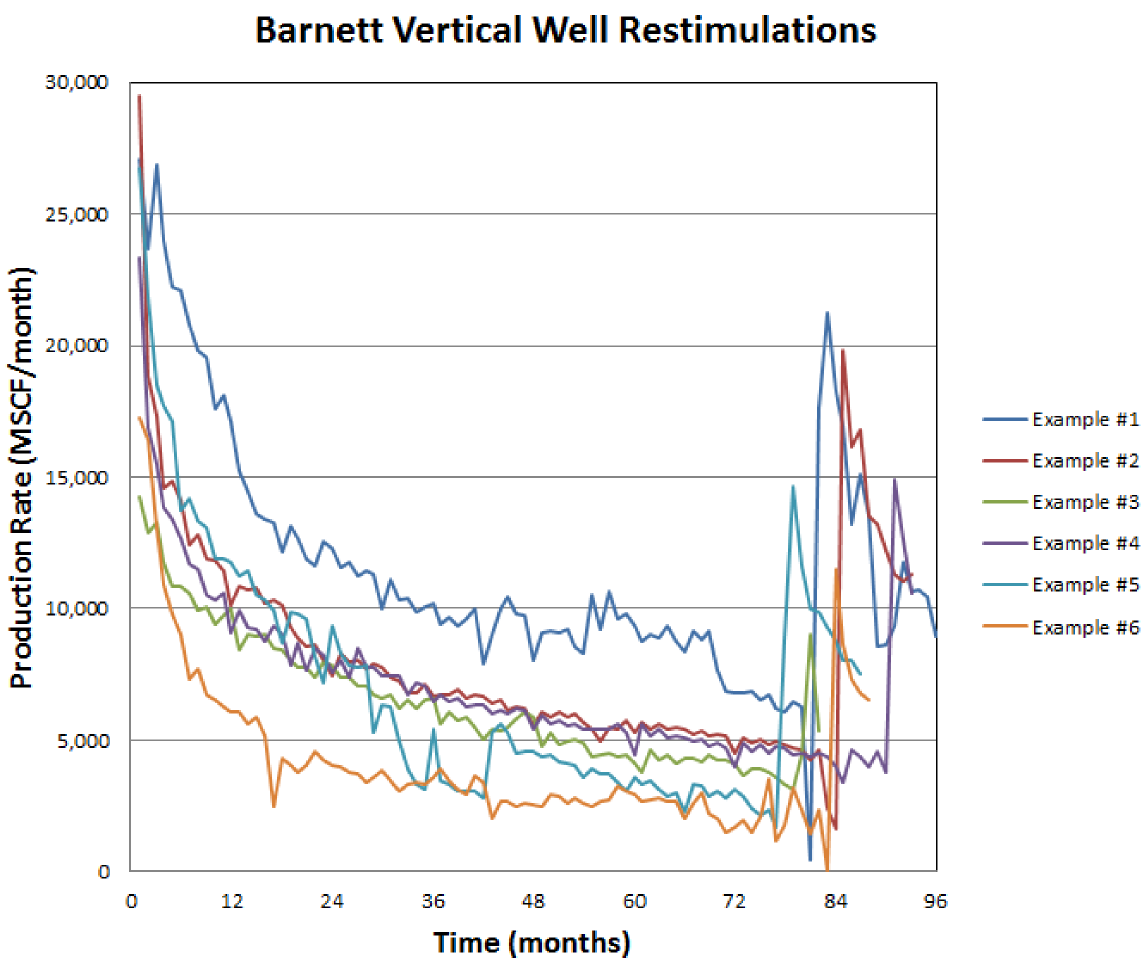

Figure 3.1: Vertical wells in the Barnett Shale are commonly observed to undergo restimulation approximately six and a half to seven years into their producing lifetimes. This restimulation has the effect of dramatically altering the reservoir flow characteristics, and renders the decline curve trend inapplicable in predicting the well's future production. Because a large number of wells contained at least 72 months of historical data, this timeframe was used to generate cumulative production values against which decline curve projections were compared.

The concept of decline curve analysis relies upon the assumption that production rate will decline according to some identifiable trend over time. The trend need not be smooth and explicit, but it must be present in at least a general form. Wells that did not exhibit this trend were eliminated from consideration in this research. By visually inspecting a set of 396 vertical wells in the Barnett Shale, 169 wells were found to possess the requisite overall decline trend. A similar analysis was performed upon a set of 218 horizontal wells, where 116 of them made the cut. Therefore, the well data used in this research analysis came from a total of 285 wells. The names and lease numbers of these wells are listed in Appendix A (horizontal wells) and Appendix B (vertical wells).

Within individual well data sets, minimal effort was made to clean up the data in preparation of computer regression analysis, as the guiding principle in this regard was to alter the original production data as little as possible. Production values of zero, located between months of relatively high production, were eliminated because these values were thought to reflect mechanical or logistical difficulties rather than constitute credible indications of reservoir deliverability. In addition, months where production spiked dramatically then resumed a clear trend were cleaned up to smooth the overall trend. Examples of production data sets before and after such a cleanup are shown in Figure 3.2. It may have actually been unnecessary to eliminate these outliners, since their limited numbers have a negligible effect upon the decline curve fit to the entire data set. With regard to their possible inclusion within synthetic data sets, the act of sorting the synthetic data sets similarly negates their effect. In future applications of the techniques presented in this research, the process of smoothing the historical production data sets by eliminating outliers could be eliminated from the workflow, since the presence or removal of this data points are not thought to have any significant impact upon simulation results.

Finally, initial production months where the production rate had not yet reached a maximum rate were eliminated, as these data points were thought to be influenced by flowback of fracturing fluid or other factors and thus not indicative of the reservoir's early ability to produce hydrocarbons.

3.2 The Four Principal Decisions

As previously discussed in Chapter 2, an analyst using DCA methods to generate estimates about future production must consider four principal choices:

1. Whether to generate deterministic or stochastic estimates,
2. Whether to manually fit the decline curves to the production data, or to use computer regression techniques,
3. Which mathematical model to use, and
4. Which production data will be used to model decline curve parameter values, all of it or just some subset.

The choices made with respect to the decline curve analysis conducted in this research are next considered in turn.

3.2.1 Deterministic vs. Stochastic Approach

This research sought to understand how effectively decline curve analysis can generate cumulative production predictions in groups of wells. Since both deterministic and stochastic methods are commonly used within the petroleum industry, the performance of both approaches were evaluated as part of this research. Deterministic estimates were generated by fitting decline curves to the original production data, and stochastic estimates were generated by fitting decline curves to multiple synesthetic data cases created with the bootstrap method.

3.2.2 Manual Fit vs. Computer Regression

A large number of well data sets (285) were used as the production data to which decline curve analysis was applied. For each of these data sets, 500 synthetic data sets were created and evaluated, for each of 2 mathematical models, and for each of 11 simulations that varied the number of months of production data used to determine the values of the decline curve parameters. As a result, over 3.1 million synthetic data sets were created and evaluated. Due to this large number of synthetic data sets, computer regression was used to fit the decline curves, calculate cumulative production values, and compare the results with the historical production data. These processes were implemented with MatLab code created specifically for this project.

3.2.3 Choice of Mathematical Model

The classic Arps hyperbolic model was used in this research. This model's widespread use in the industry, along with its uncertain functional behavior as applied to shale reservoirs, made it a compelling choice for further investigation. In addition, the Duong model was also used due to its empirical development with shale gas data and the industry interest it has generated since its introduction in 2011.

3.2.4 Selection of Various Ranges of Months of Production Data

This research seeks to better understand how effectively the Arps and Duong models extrapolate the future performance of horizontal and vertical gas wells. One question it seeks to explore is whether the amount of available data (in terms of number of months) drives the predictive performance of these models. Accordingly, this research uses several ranges of months of production data to model anticipated future production.

It then compares the models' predictions about what would occur over a 72 month period with what actually did occur.

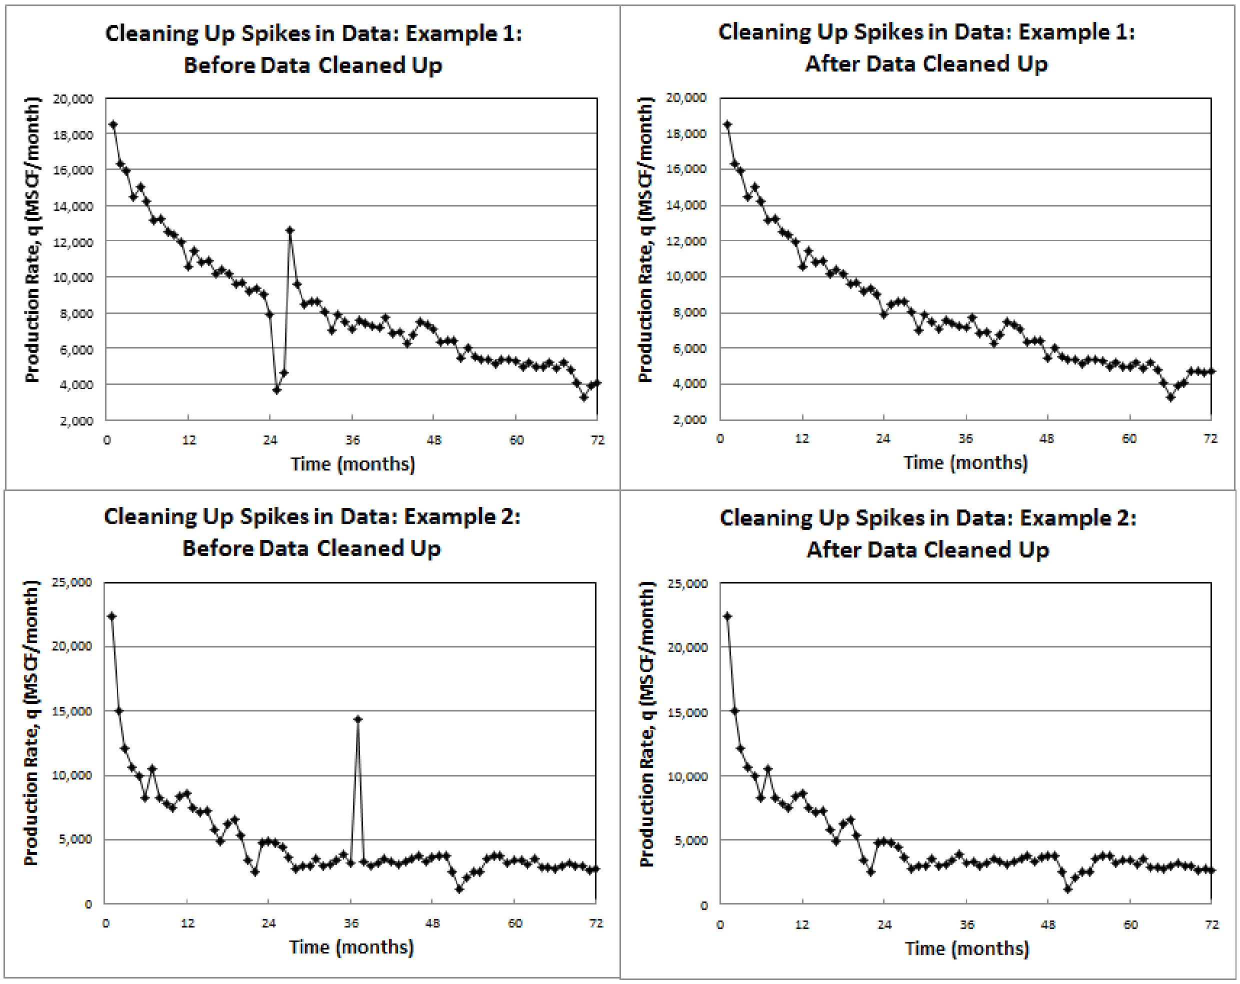

Figure 3.2: Production data sets used in this research were cleaned up to eliminate large spikes in the data not thought to be representative of reservoir deliverability.

The ranges of months selected for analysis were months 1-6, 1-12, 1-18, 1-24, 1-36, 1-48, 1-60, and 1-72. As an example, for the month range 1-6, the nonlinear regression analysis used only data from months 1-6 to find parameter values for the deterministic approach (the base case, consisting of the decline curve fit to the original production data) and the stochastic approach (the P10, P50, and P90 cases). Using these model parameters, EUR values for each of the cases were found as the cumulative production over a 72 month period. These EUR estimates were then compared with the actual cumulative production observed in the field data.

When running the various simulations based on production from different ranges of months, it was found that ranges with fewer months exhibited fewer wells for which the computer regression process was able to converge upon a reasonable solution. Table 3.1 shows the numbers of wells for which reasonable decline curve parameters could be obtained, based upon the number of months of production data used to model those parameters. Arps b-factor values were considered to be reasonable if between zero and three, inclusive, and initial decline rate values were considered to be reasonable if between zero and unity, inclusive. The Duong model tended to more consistently converge upon reasonable solutions, but this research applied Duong analysis to only

Table 3.1: When only six months of production data were used to model Arps decline curve parameters, the computer regression algorithm was able to converge upon reasonable solutions for only about one third of the original production data sets. For example, only 56 of the 116 wells in the vertical data set, and only 35 of the 116 wells in the horizontal set, would submit to sample set would s1ubmit to computer regression analysis. These ratios improved dramatically when twelve months of data or more were used to model the decline curve parameters.

Months of Production Data Used	Number of Wells for Which Arps Model Converged Upon a Solution	
	Vertical Well Set	Horizontal Well Set
1 to 6	56	35
1 to 12	140	101
1 to 18	155	113
1 to 24	157	115
1 to 36	166	116
1 to 48	168	116
1 to 60	169	116
1 to 72	169	116

those wells for which reasonable Arps solutions could be obtained, in order to enable the generation of comparisons between common sets of wells.

Additional simulations were conducted to determine EUR values modeled on production data from months 6-36, 12-36, and 18-36. The results from these simulations were compared with those modeled on production data months 1-36 to determine whether any benefit resulted from disregarding early production data.

3.3 Creation of Synthetic Data Sets from Field Data

As previously discussed, synthetic data cases, also known as bootstrap realizations, were created in order to estimate the P10, P50, and P90 values (also called probability cases, or P-cases). A couple of considerations exist with regard to the creation of these synthetic data sets, including whether to sample the actual production data points or their residuals, and how many synthetic data cases to generate for each well.

3.4 Sampling of Individual Production Data Points vs. Blocks of Residuals

The two prior studies that used the bootstrap method to generate EUR distributions and probability cases differed in their sampling approach, with Jochen et al. (1996) using individual production data points and Cheng et al. (2010) using blocks of residuals. It was therefore necessary to determine, for purposes of this research, which of these approaches to adopt. A process was first devised for examining residuals values over the course of the 72 month production data time span by normalizing the residuals sets from individual wells. The purpose of this normalization procedure was to enable detection of recurrent patterns or trends that might be present in multiple wells. The method for normalizing residuals recognizes that residuals which exist above the base case's regression line take on positive values, whereas those below the regressed line are negative. A common pattern observed in production data, when the magnitude of residuals is high, are for the data points oscillate above and below the regressed base case line. Therefore, the absolute value of each of the residuals was taken to highlight patterns of deviation from the regressed base case line. These magnitude values were then divided by the initial production rate for that well to create a normalized residuals data set.

Figure 3.3 shows the result of normalizing the residuals data sets of six of the horizontal wells used in this research, and Figure 3.4 shows the equivalent result for six of the vertical wells. The horizontal wells tend to display residuals of stronger magnitudes at early and late times within the 72 month window under consideration. The vertical wells do not show these early and late time patterns, but do exhibit distinct periods within each well data set where the residuals values alternate between higher and lower values. If Cheng's method of sampling residuals and then concatenating them were applied here to the horizontal wells, the effect would be to introduce information about production variability into the data set. For example, residuals points or blocks from early or late times could make their way into the middle sections of the synthetic data sets. And it would be equally possible that production variability information could be deleted, if residuals points or blocks from the middle sections of the residuals data set became concatenated within the early or late portions of the synthetic data sets. A similar phenomenon is possible with regard to the vertical well data sets. This significant possibility of introducing or deleting information about production rate variability as a

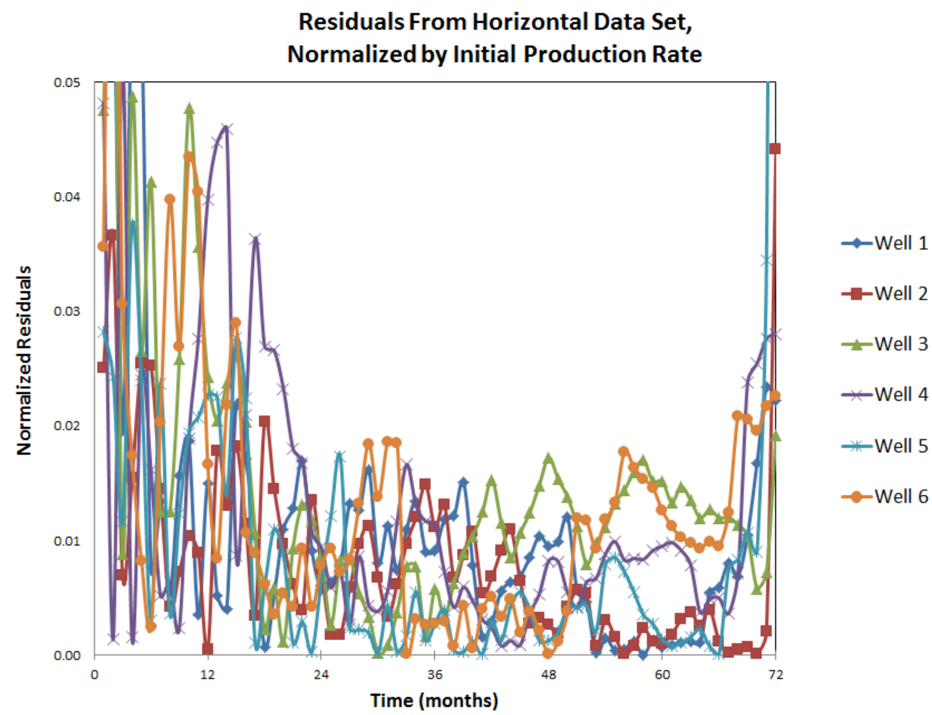

Figure 3.3: Horizontal wells tend to show increased residuals magnitudes in early and late times during their first 72 months of production. The presence of this phenomenon suggested against sampling in residuals to create synthetic data sets.

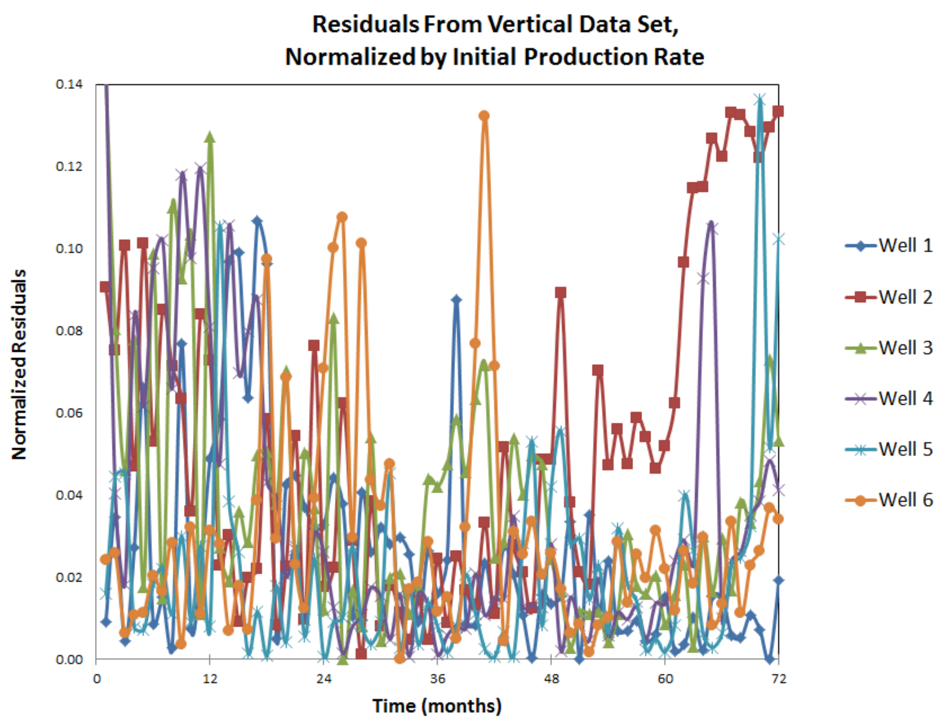

Figure 3.4: Vertical wells tend to show alternating periods of high and low residuals magnitudes during various periods in their first 72 months of production.

function of time strongly suggests that sampling in residuals (either as individual points or in blocks) should be avoided. Instead, it was decided to sample individual data points since this process better preserves the fidelity that the synthetic data sets will have to the original production data set.

Sampling individual production data points is not without complications, however. With shale formations, steeply declining production rates mean that successive data points (even over a relatively short time frame, on the order of several months) can vary significantly. The danger in sampling production data points is that if the randomly chosen data points are concatenated in the order they are sampled, the resulting synthetic data set can show a flat or even upwardly sloping trend. Measures must therefore be adopted to prevent such an occurrence, which is not representative of the decline trend commonly seen in production data sets.

3.5 Sorting of Data to Permit Better Curve Fitting

Examples of six synthetic data cases are shown in Figure 3.5, which plots both the unsorted and sorted (in descending order) versions of these data sets. From visual inspection, it is readily apparent that the unsorted versions of these data sets do not exhibit decline trends representative of those in field production data sets, but the sorted versions do display such trends. The issue has a practical dimension as well. The unsorted data sets, with their erratic trajectories, result in regressed parameter values that are negative or complex (containing an imaginary component, i, equal to the square root of negative one) when subjected to nonlinear regression analysis. To avoid this outcome, the synthetic data sets are sorted before analysis.

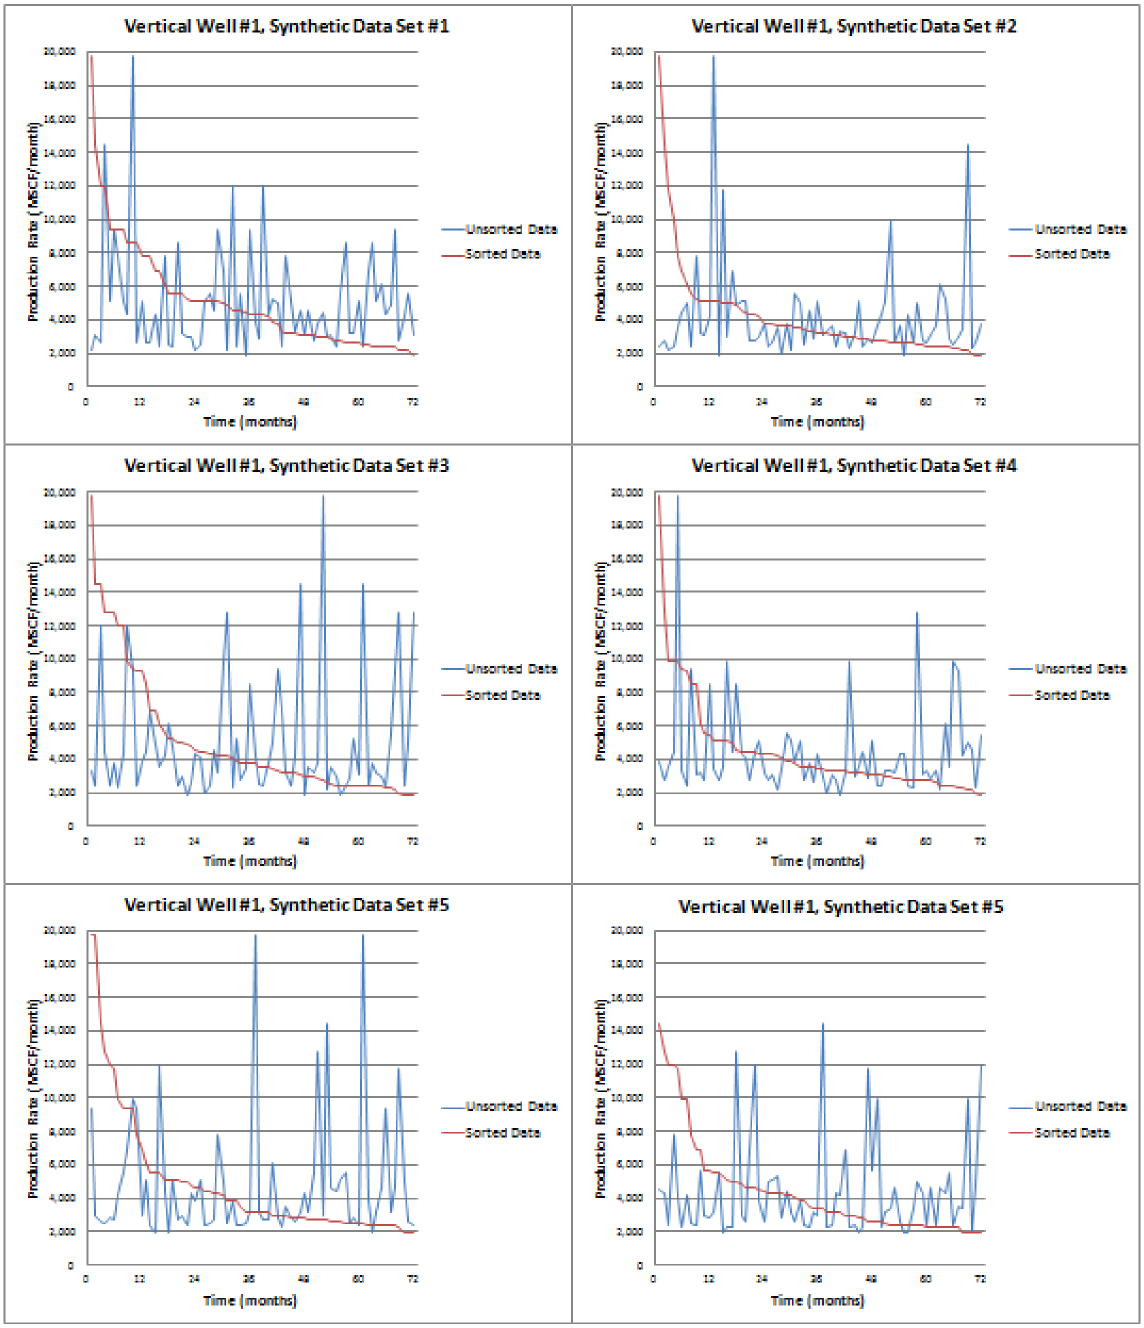

Figure 3.5: Example synthetic data sets, as sampled, and sorted in descending order of magnitude. Synthetic data sets were created by randomly sampling, with replacement, the original production data sets.

Multiple (sorted) synthetic data sets are plotted along with the original (sorted) production data set in Figure 3.6 to show the qualitative relationship between the original and synthetic sets. The synthetic data sets tend to roughly approximate the original, with slight production departures above and below the path taken by the original data set. The synthetic data sets are therefore representative of the overall decline trend seen in the original data, and can be thought of as credible, alternative scenarios of the actual field case.

3.6 Determination of Number of Synthetic Data Cases

Consideration must be given to how many bootstrap realizations are necessary and sufficient. The guiding objective is to balance considerations of repeatability and computing overhead. Generating too few bootstrap realizations will cause the resulting EUR distribution to vary enough that sufficient confidence cannot be placed in its repeatability. Too many bootstrap realizations, on the other hand, will demand excessive computing time, which increases roughly linearly with the number of bootstrap realizations. An optimum balance is achieved when the cumulative frequency function of EUR values approaches a smooth, S-shape, indicating a normally distributed and repeatable distribution.

Figure 3.7 shows example cumulative frequency functions associated with various numbers of bootstrap realizations. It can be seen that generating only 25 bootstrap realizations causes the cumulative frequency function to fail to attain a smooth shape. Increasing the number of bootstrap realizations to 500 results in the desired S-shaped curve. Figure 3.8 shows the EUR distributions for a single well obtained from multiple simulation runs. The overlapping nature of these distribution profiles indicates repeatability over successive simulations. From this analysis it was concluded that 500 bootstrap realizations would be necessary and sufficient to generate repeatable P-cases for each of the wells in the sample sets.

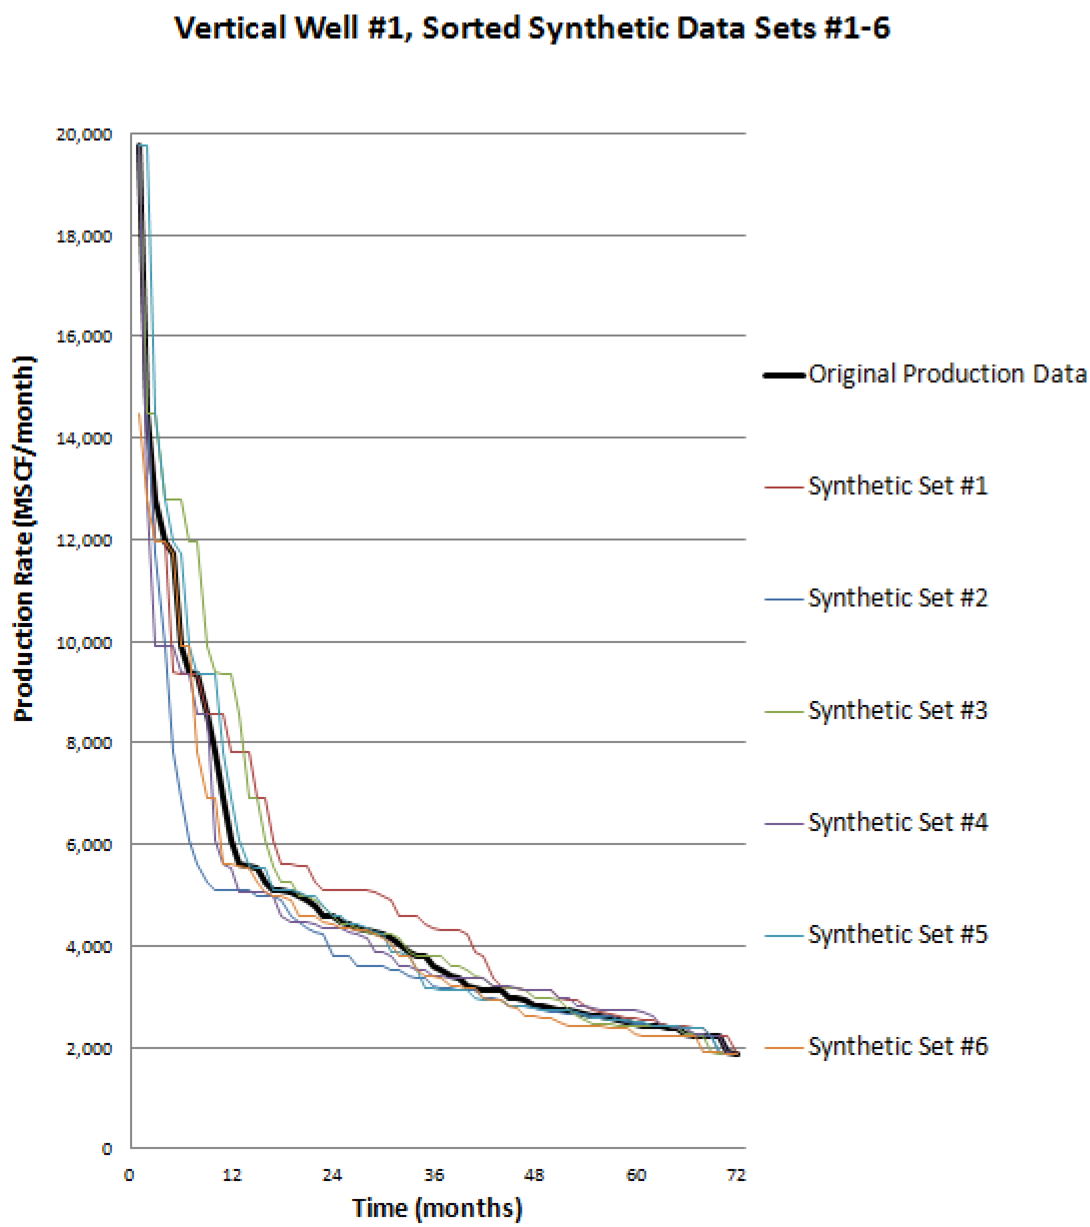

Figure 3.6: Synthetic data sets are alternative versions of the original production data set. The synthetic data sets comprise scenarios with higher or lower cumulative productions over the 72 month time frame considered in this research.

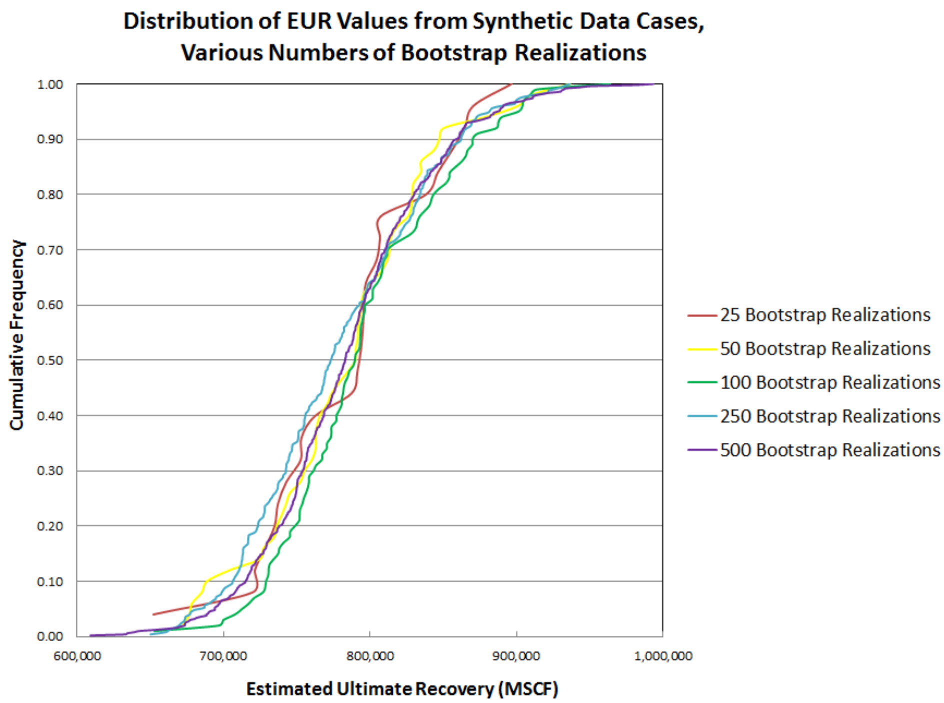

Figure 3.7: Increasing the number of synthetic data sets, or bootstrap realizations, generated for each production data set causes the corresponding EUR distribution to approach a normal distribution, as seen in the S-shaped curve that results from higher numbers of synthetic data sets within the distribution.

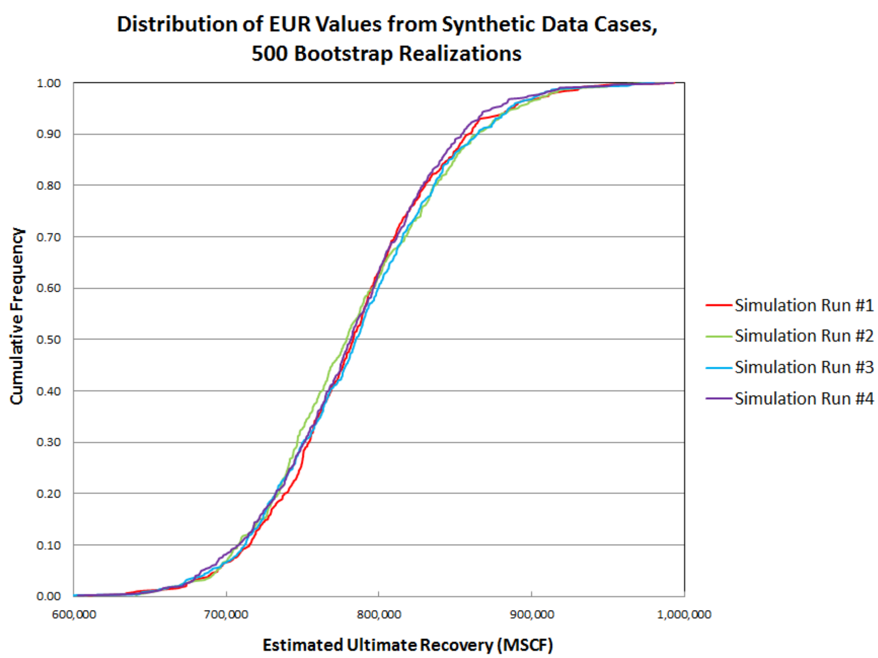

Figure 3.8: When 500 bootstrap realizations are created for each well, the resulting EUR values approach a repeatable distribution across simulation runs, as seen in the overlapping nature of the various cumulative frequency functions.

3.7 Summary of Methodology

The methodology used in this research to generate decline curve estimates consisted of the following workflow:

- For each of the 285 wells:
    - For each model: Arps and Duong
        - Using various ranges of months of production data:
            - The Deterministic Approach: Evaluate the base case EUR
                - regress a best-fit curve to the original production data
                - determine the model parameter values
                - extrapolate that curve to 72 months to determine the cumulative production
                - compare that extrapolated cumulative production with the actual historical production
            - The Stochastic Approach: Evaluate the P10, P50, and P90 probability cases
                - create 500 synthetic data cases for each well,
                - regress a best fit curve to each synthetic data case,
                - determine decline curve parameter values,
                - extrapolate that curve to 72 months to determine cumulative production
                - compare that extrapolated cumulative production with the actual historical production

CHAPTER 4

RESULTS AND DISCUSSION

Once the EUR values of the base case and the three P-cases (P10, P50, and P90) were obtained (over various ranges of months to model the decline curve parameters, for each of the two mathematical models, as applied to each of the wells in the two data sets), they were compared to the actual cumulative production seen in the field data. A series of histograms was created to display the results in a manner that permitted detection of trends.

4.1 Verification Procedures

Several steps were taken to verify that the computer code written to undertake this research was working properly. For several example wells in each simulation, the original production data, the base case decline curve regressed to that data, and the decline curves for the P-cases were plotted against time. Examples for the Arps and Duong models are shown in Figures 4.1 and 4.3, respectively. Charts were also created to represent the cumulative production over time associated with these decline curves, and examples are shown in Figures 4.2 and 4.4. These charts show the results of using production data from months 1-72 to model the decline curve parameters over the entire 72 month span. As expected, the base case regressed curves in Figure 4.1 (for the Arps model) and Figure 4.2 (for the Duong model) nicely track the actual production data. With respect to the total cumulative production at month 72, the P10 EUR case lands above this value, the P50 case lands right on top of it, and the P90 case lands below it. These results conform to expectations.

Careful analysis of the decline curve for the Duong model (Figure 4.3) reveals that at some points in time during production, the projected P10 decline curve path falls below the other paths, especially during early times. This phenomenon occurs due to the normal functioning of the model, and does not represent a programming error. The P-cases are defined according to the EUR values (cumulative production at 72 months) that they predict. These values are determined by integrating the appropriate mathematical equation over the range of time represented in the analysis (months 1 to 72), Visually, this cumulative production for each of the cases is equal to the area under its decline curve. Therefore, even though the P10 case may at times proceed along a path beneath those of the other curves, its cumulative production over the course of the analysis period will be greater. This result is confirmed by inspection of Figure 4.4, the cumulative production chart that corresponds to the decline curves in Figure 4.3.

4.2 Effect of Data Availability Upon the Reliability of Decline Curve Predictions

Charts that show decline curves and cumulative production over time represent the actual and modeled behavior of *individual* wells. Since there were numerous wells in the horizontal and vertical well sets, a series of histograms were created to convey information about the behavior of the *group* of wells. The histogram bins describe normalized percent differences between the actual cumulative production and the cumulative productions predicted by the base case and P-cases for each well. These normalized percent differences for the base case and P-cases of each well are calculated

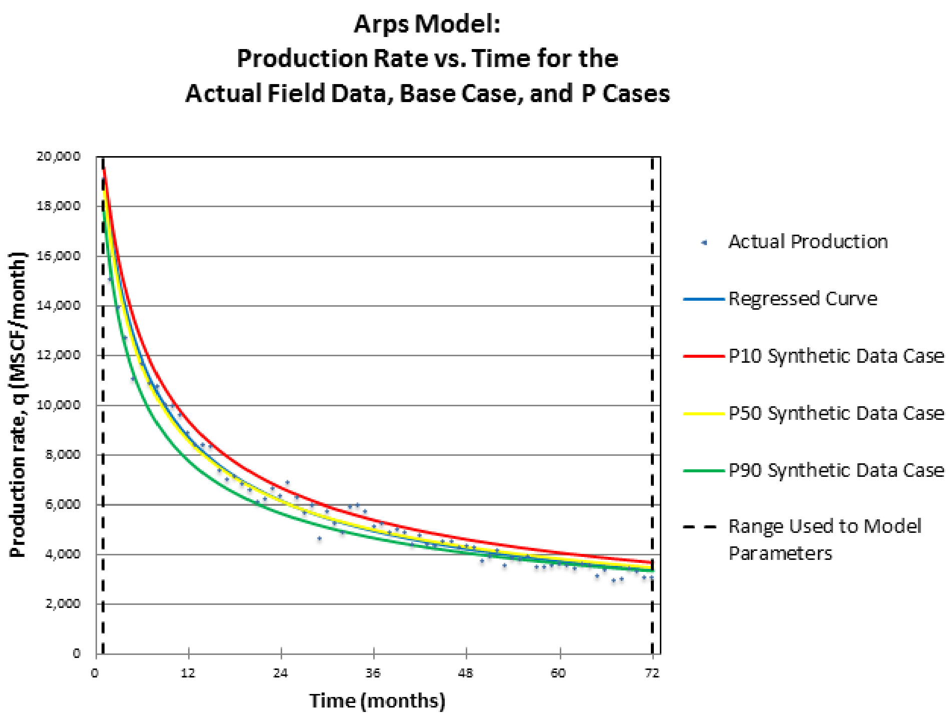

Figure 4.1: Example decline curves, using the Arps hyperbolic model,
fit to production data over a range of 72 months, and selected probability cases.

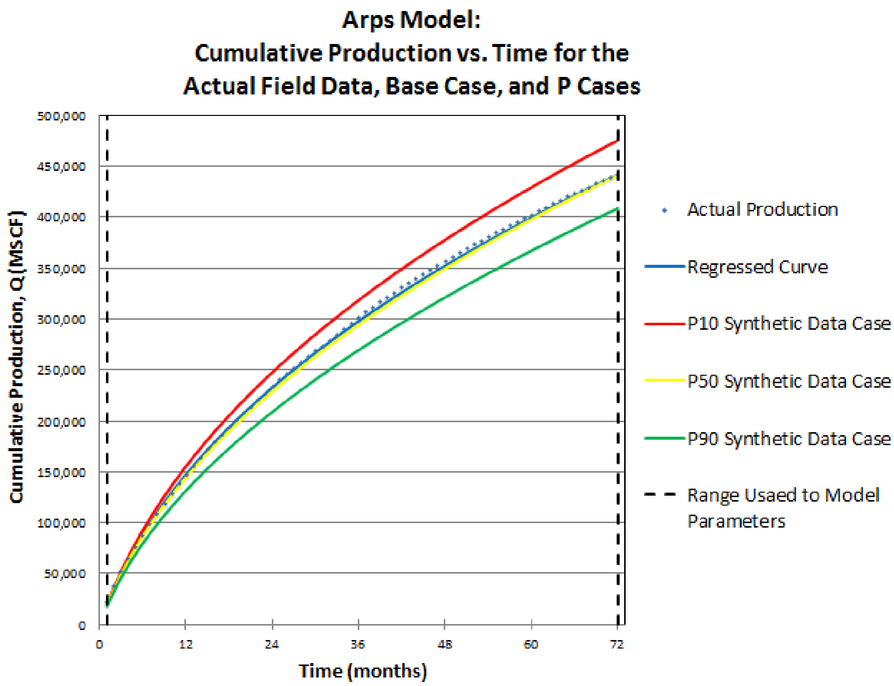

Figure 4.2: Cumulative Production curves corresponding
to the base case and probability cases presented in Figure 4.1.

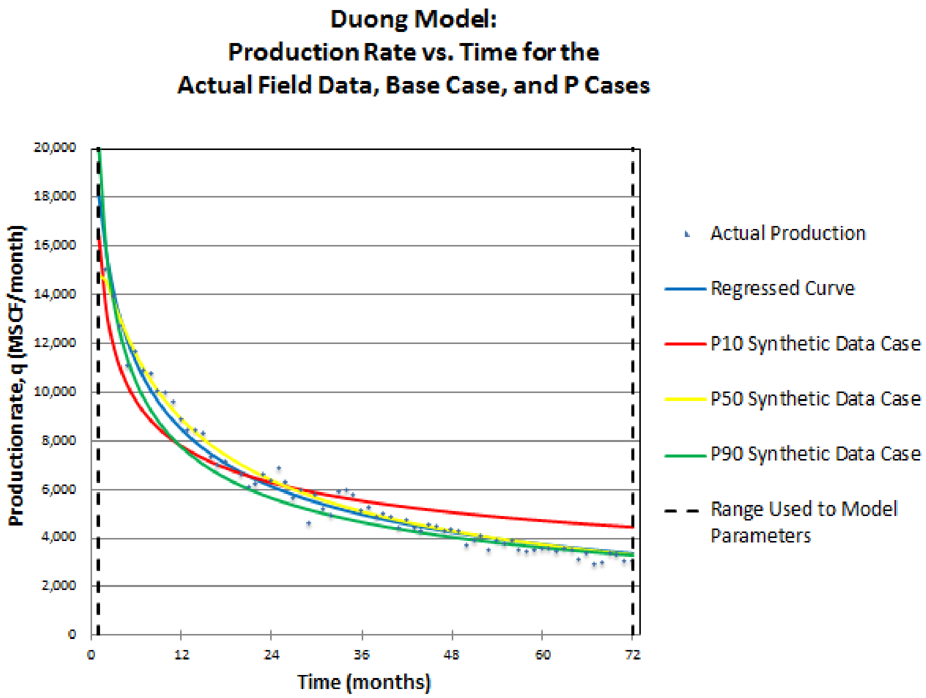

Figure 4.3: Example decline curve, using the Duong power law model, fit to production data over a range of 72 months, and selected probability cases.

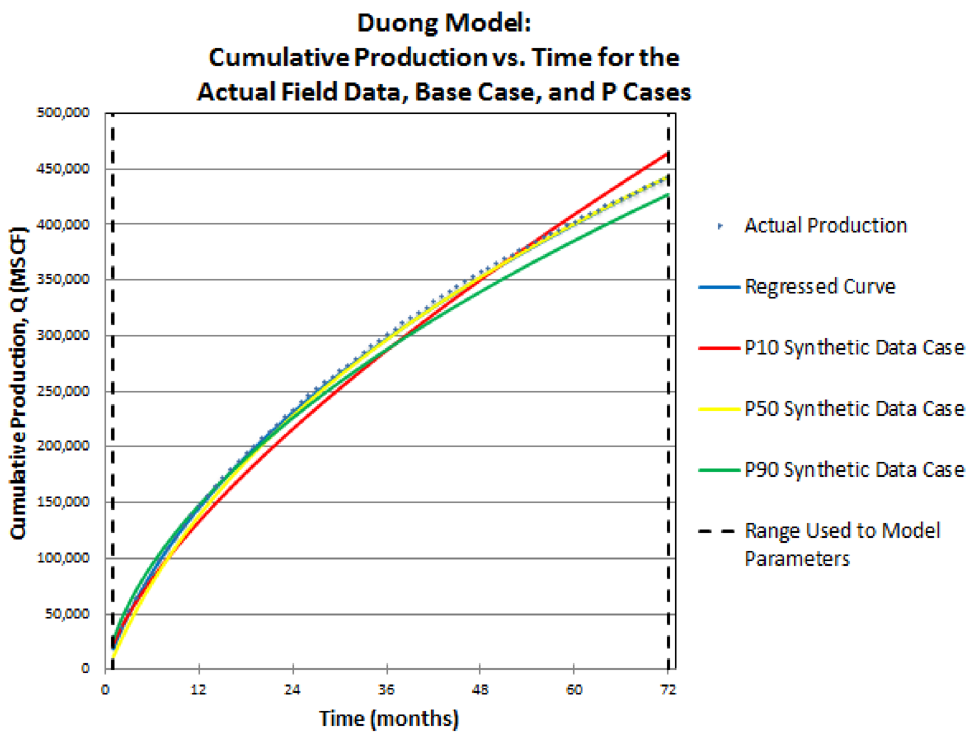

Figure 4.4: Cumulative Production curves corresponding to the base case and probability cases presented in Figure 4.3.

by the following formula:

$$percent\ difference = \frac{x - actual\ cumulative\ production}{actual\ cumulative\ production} \qquad (4.1)$$

where

$$x = \{\text{base case EUR, P10 EUR, P50 EUR, P90 EUR}\}.$$

Figures 4.5 and 4.6 show the histograms generated by applying the Arps and Duong models to the horizontal and vertical data sets, respectively. Since all 72 months of production data were used to generate these base and P-cases, we would anticipate that the variations between the predictions and the actual cumulative production would be small. We observe this result in these histograms. For both of the models, and in each of the well sets, the base case and P50 case show very little variation from the actual cumulative production. The P10 case tends to show a small positive variation, indicating that the P10 values tend to be slightly larger than the actual cumulative production values for the various wells. The P90 values tend to show a small negative variation, reflecting their slightly smaller values compared with the actual cumulative productions. It is not problematic that some of the P10 or P90 values fall outside a 10% variation from the actual cumulative production. It should be recalled that the P-cases are defined by their relationships to one another, not by their relationship to the actual cumulative production. Instead, the normalized percent variation of the base and P-cases is an indication of how well the Arps or Duong models predict the actual cumulative production in the context of application as synthetic data cases. When we see low percent variation between the

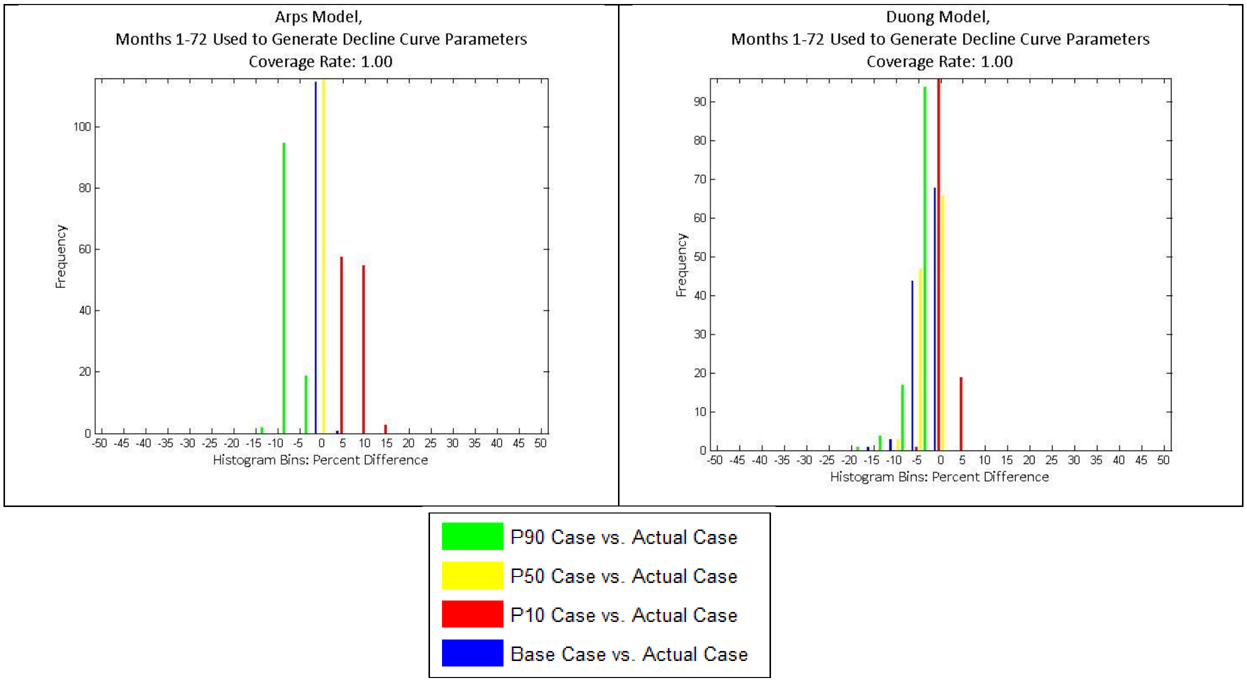

Figure 4.5: Horizontal Wells: Histograms (from simulations that used all 72 months of horizontal well production data to model decline curve parameters) provide evidence that computer code is running properly and models are behaving as predicted.

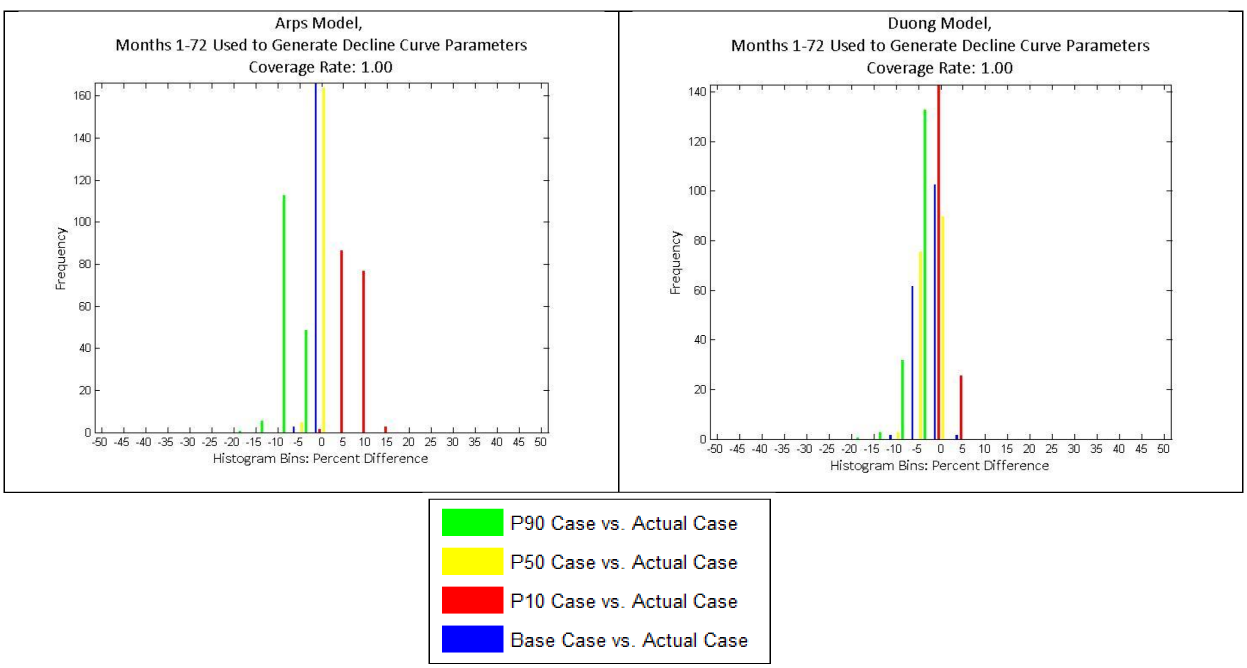

Figure 4.6: Vertical Wells: Histograms (from simulations that used all 72 months of vertical well production data to model decline curve parameters) provide evidence that computer code is running properly and models are behaving as predicted.

various projected cases and the actual case, we can have confidence that the process of using synthetic data sets to generate P-cases is valid. In fact, if we did not observe some variation, we would question the functional validity of the bootstrap method, as we expect it to generate some variability.

The coverage rate for this validation simulation was 1.00 for both models in both the horizontal and vertical data sets. This result indicates that the actual cumulative production for each of the wells fell between the P10 and P90 EUR values generated from its synthetic data sets, and is further confirmation that the computer code that implements the research design was functioning properly.

With this confirmation that the research method and the code implementing it were operating as expected, various simulations were run using the two models (Arps and Duong) as applied to different ranges of months to model the decline curve parameters of the wells in the horizontal and vertical well data sets. Figures 4.7 and 4.8 show the cumulative production vs. time charts for an example horizontal well using the Arps and Duong models for simulations using months 1-6, 1-12, 1-18, 1-24, 1-36, and 1-48. All charts in these figures display simulation results based on a common example well.

The trend apparent in each of these figures is that increasing the number of months used to model the decline curve parameters causes the difference between the P10 and P90 values to decrease. In other words, the P10 to P90 bracket becomes more narrow. Yet this chart represents only the behavior of a single well. Figures 4.9 and 4.10 show the aggregated results of all the horizontal wells in various simulations using the Arps and Duong models. When only months 1-6 are used to model decline curve properties, normal distributions of percent difference variations do not result for the various cases. As successively more production data is considered, however, normal distributions form and become centered more closely to the center of the x-axis at zero percent variation. Similar results occur in the application of the Arps and Duong models to the vertical well data set, as seen in Figures 4.11 and 4.12. The convergence of the histograms toward zero percent variation with the actual cumulative production shows that when more months of data are used to model future performance, the range of the probabilistic outcomes decreases.

When only the initial 6 months of data are available to project the future, the bootstrapping method does not tend to give estimates of the various probability cases, as the P10 to P90 brackets are very wide compared with simulations in which 12 or more months of production data were available. But what about simply regressing a curve to the original production data from the first 6 months? This outcome is revealed in the base case variation (shown in blue) in Figures 4.9, 4.10, 4.11, and 4.12. The Duong model at least exhibits a normally distributed histogram for percent variation between the base case and actual case, but neither model shows histogram distributions that are nicely consolidated near zero percent variation.

These results indicate that using only 6 months of field data, with either a simple decline curve regression or a probabilistic approach using the bootstrap method, one cannot reliably predict a well's actual cumulative production over a 72 month time horizon. By the 12 month mark, however, better consolidation of histograms into bell-shaped patterns and migration of these histograms toward zero percent variation has occurred. The Duong model's base case is especially noteworthy, centered on 0% to -5%

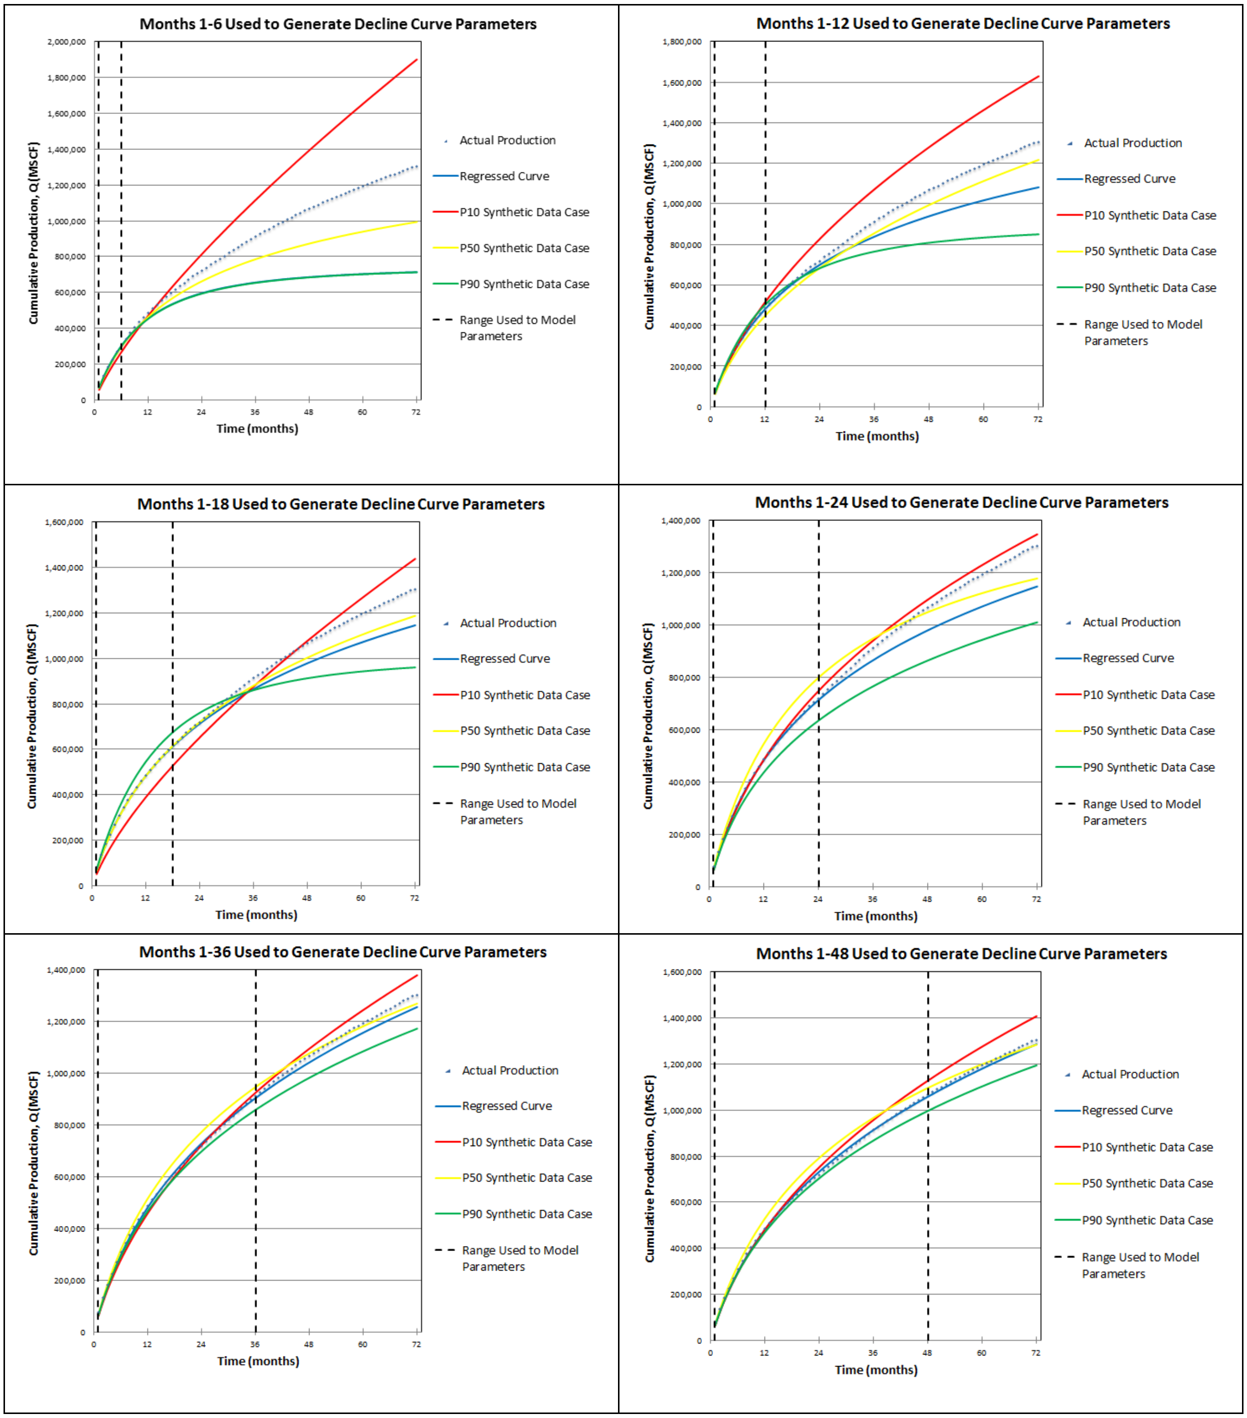

Figure 4.7: Cumulative production curves from an example horizontal well, using the Arps model, show a narrowing of the width between the P10 and P90 values as successively increasing numbers of months are used to model production decline parameters.

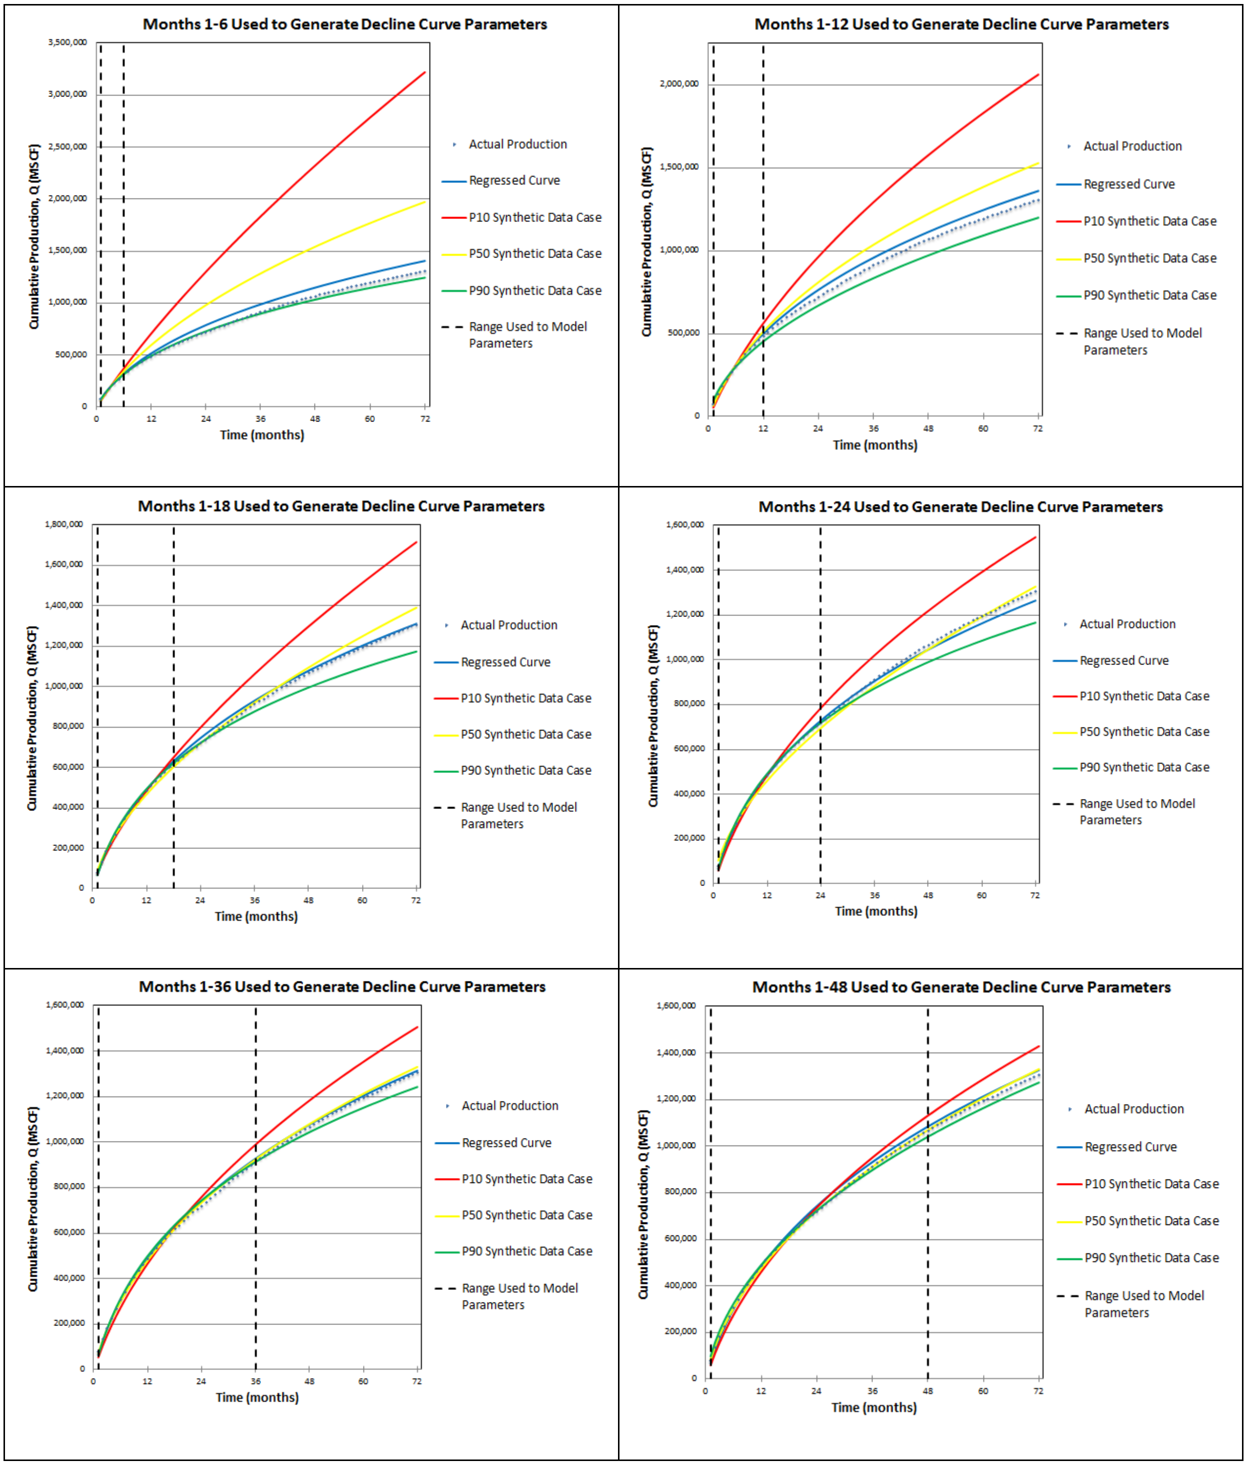

Figure 4.8: Cumulative production curves from an example horizontal well, using the Duong model, show a narrowing of the width between the P10 and P90 values as successively increasing numbers of months are used to model production decline parameters.

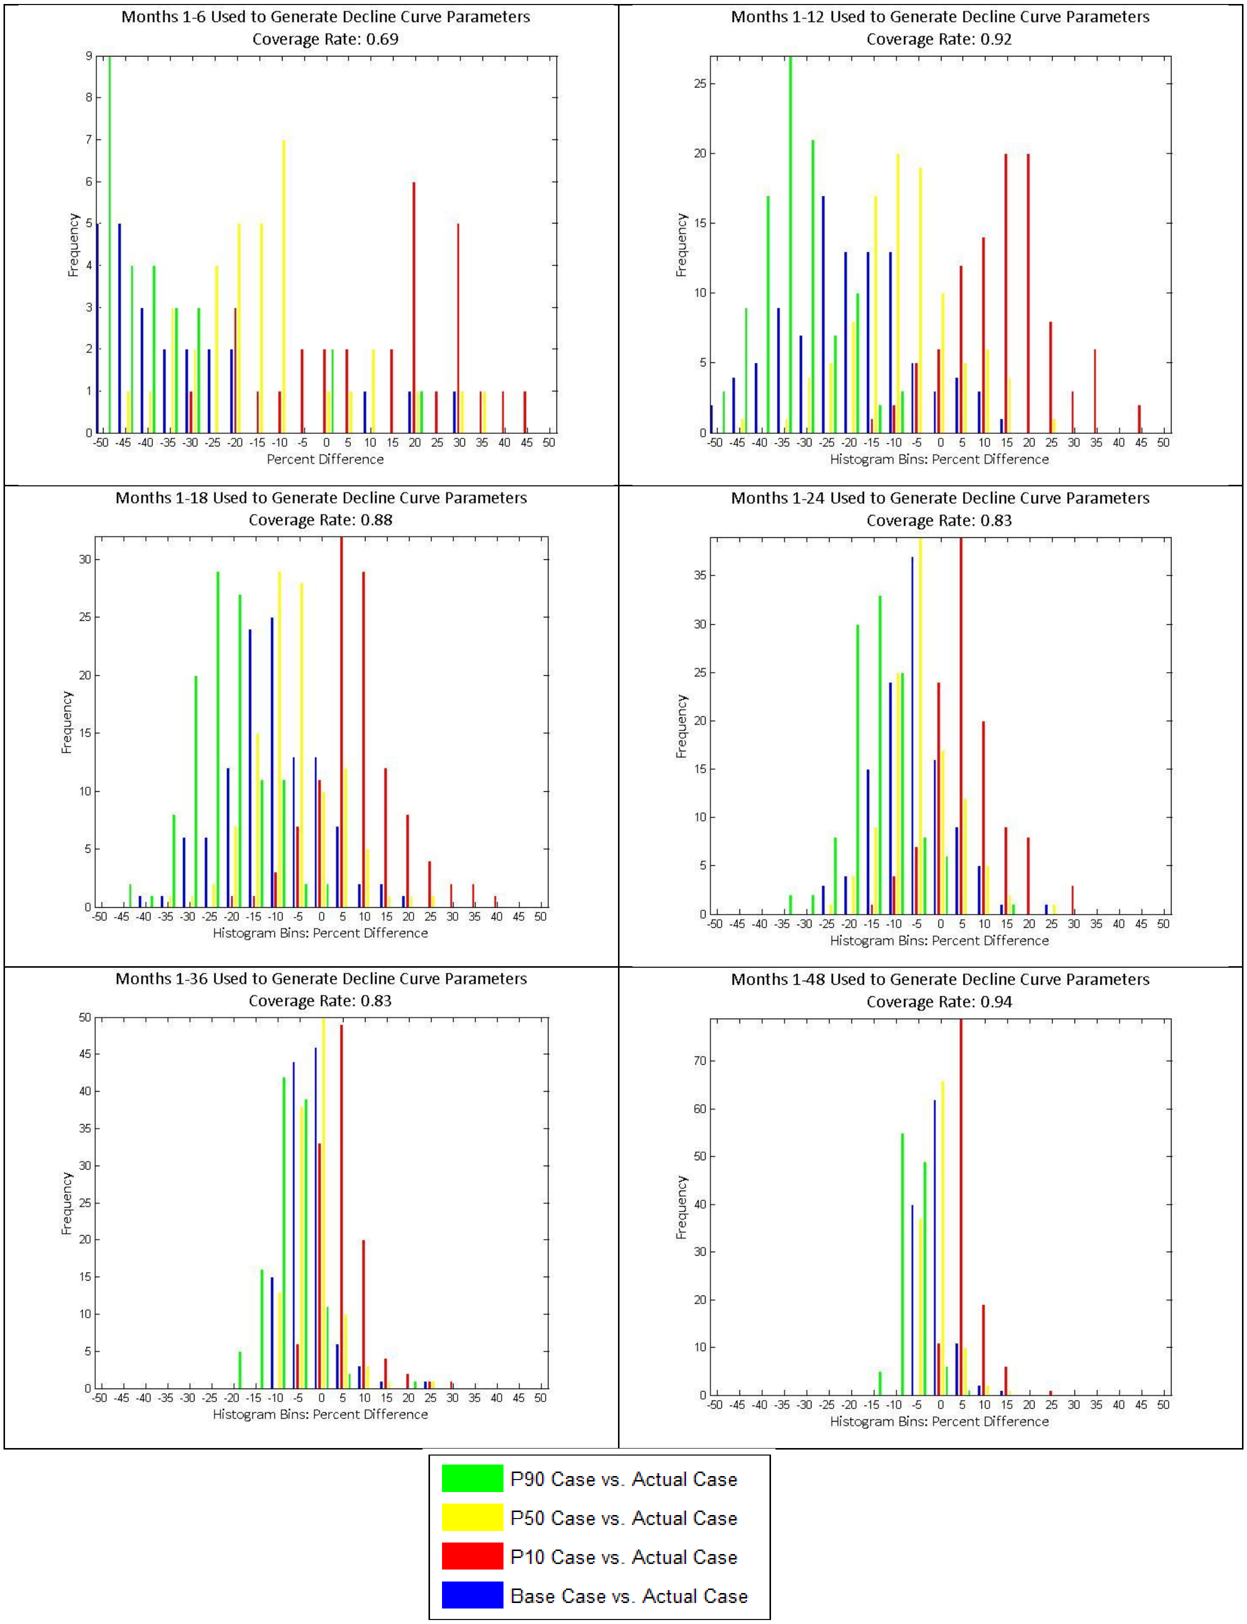

Figure 4.9: Horizontal Well Set, Arps Model. Histograms showing coverage rates and frequency distributions of percent variation between the actual case and various other cases.

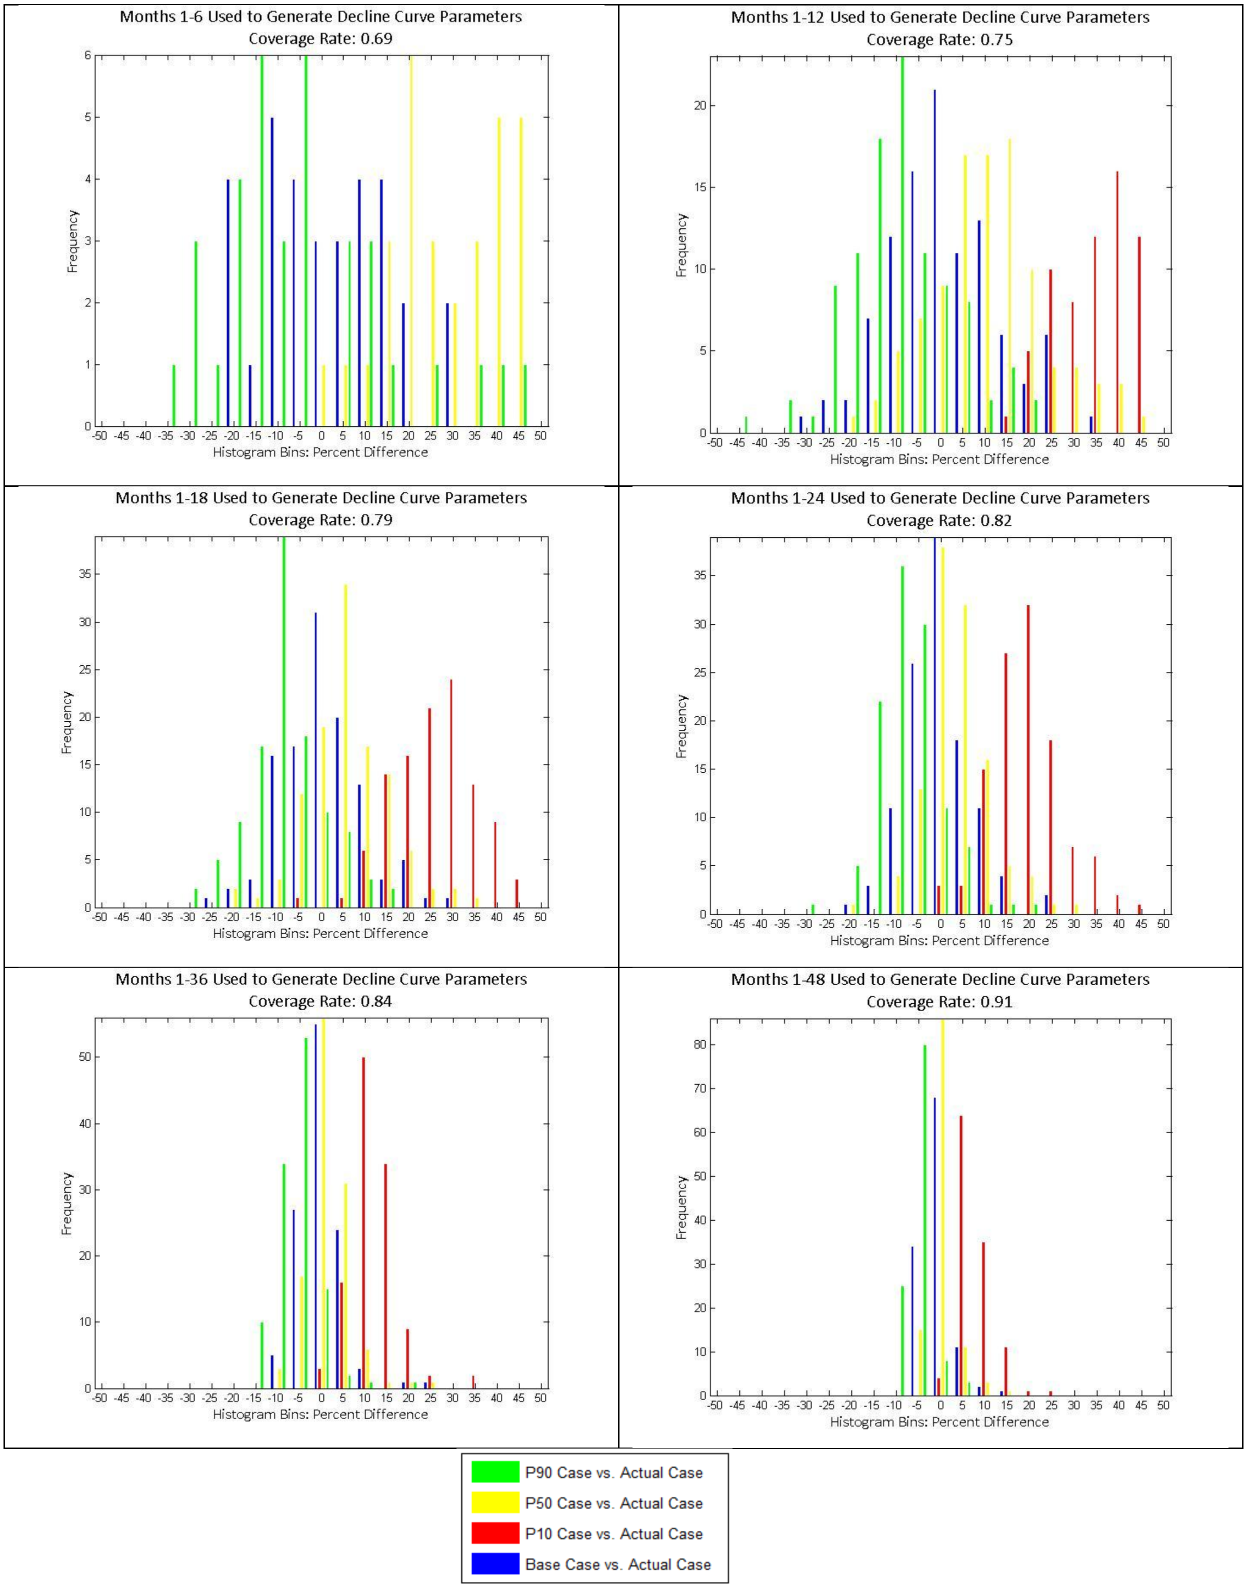

Figure 4.10: Horizontal Well Set, Duong Model. Histograms showing coverage rates and frequency distributions of percent variation between the actual case and various other cases.

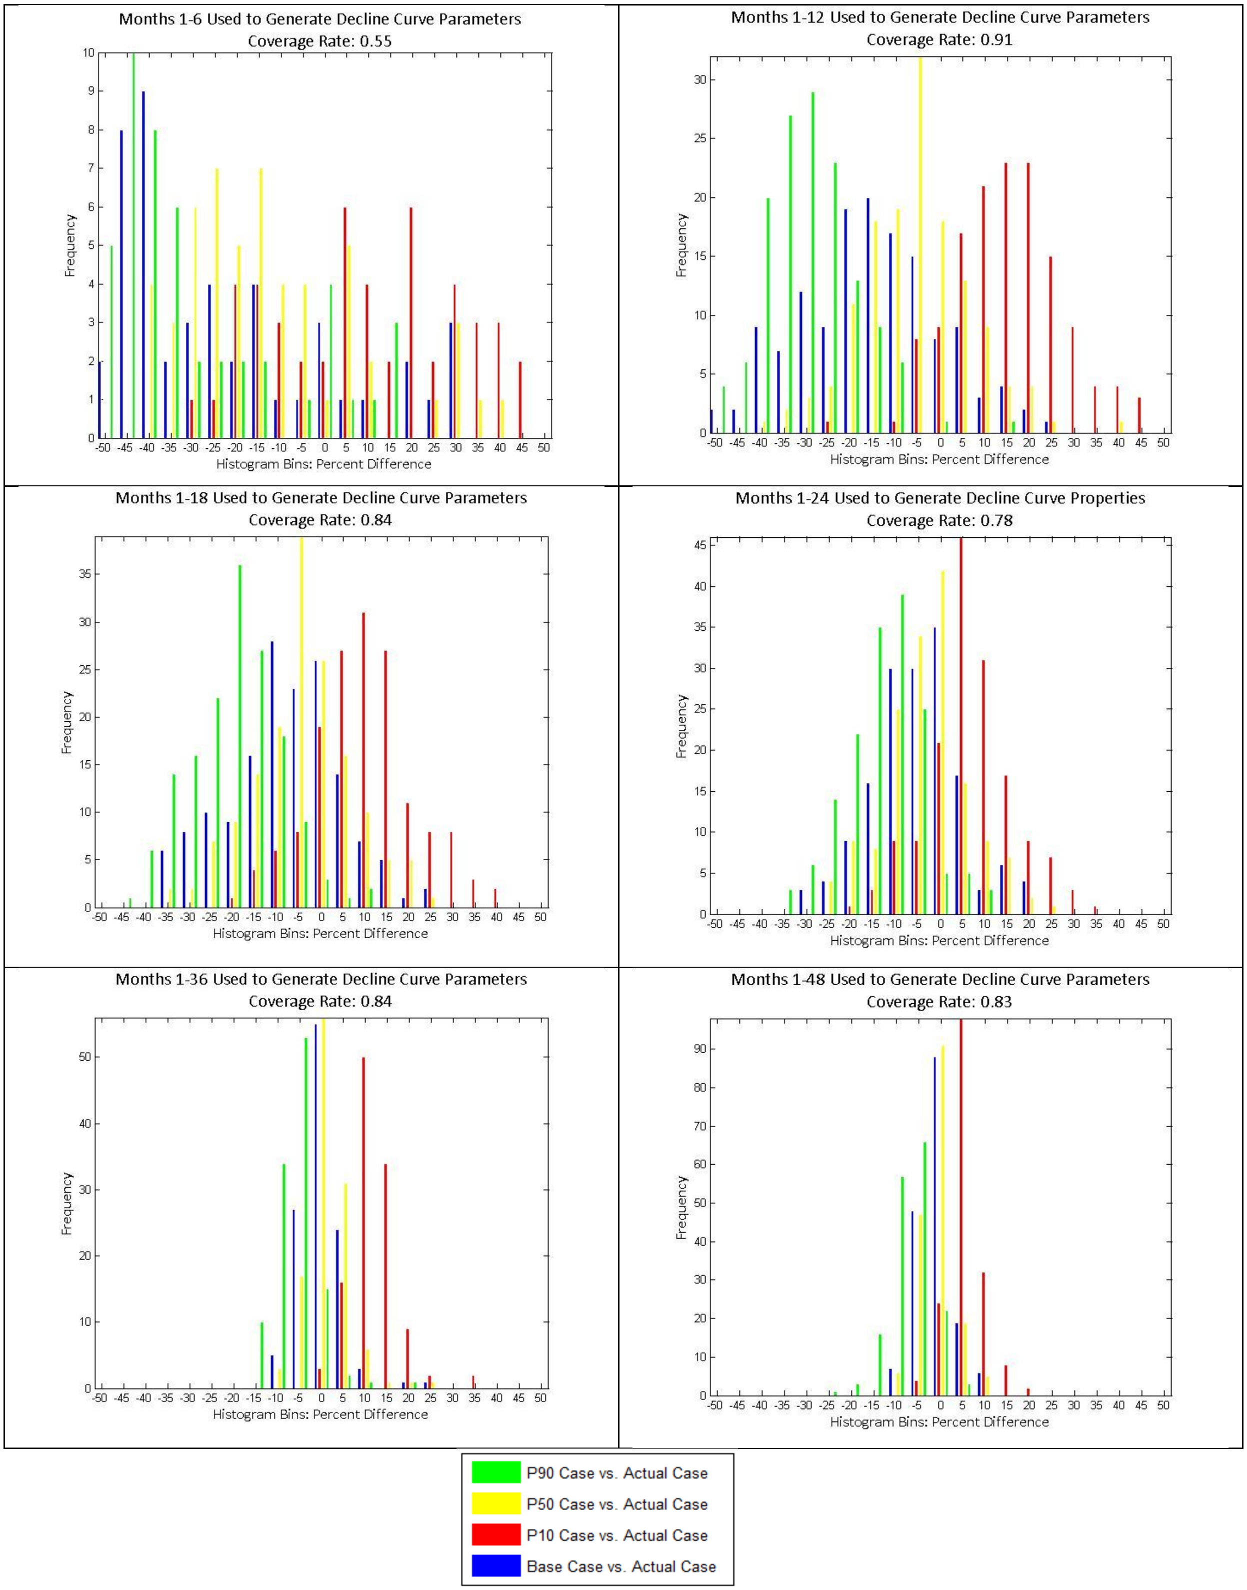

Figure 4.11: Vertical Well Set, Arps Model. Histograms showing coverage rates and frequency distributions of percent variation between the actual case and various other cases.

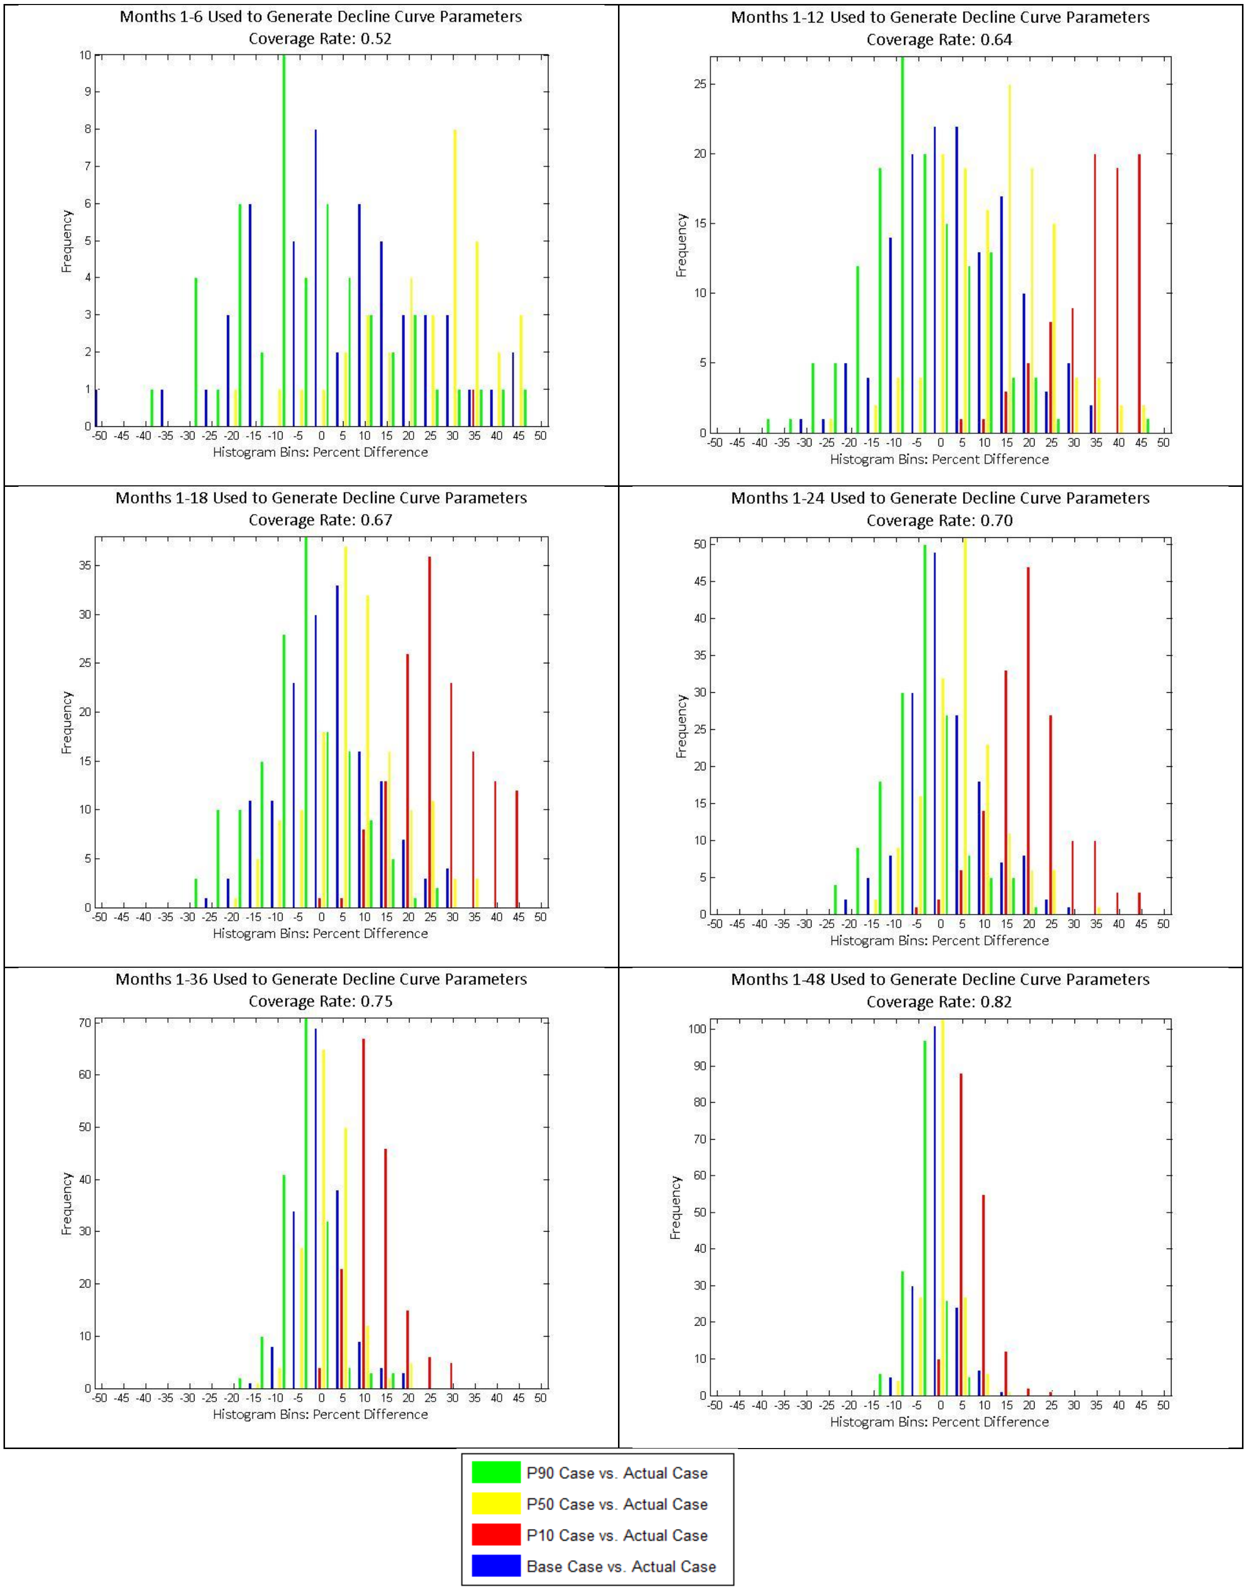

Figure 4.12: Vertical Well Set, Duong Model. Histograms showing coverage rates and frequency distributions of percent variation between the actual case and various other cases.

of the actual case for both well sets. This outcome is better than the Arps model's base case, which is around -15% to -30% of the actual case. The P50 case from the Arps model may also be a good indicator of future performance, as wells within this group tend to show a percent variation from the actual cumulative production centered around 0% to -10%. As the number of months of available field data increases, consistent trends are seen in the horizontal and vertical data sets with both the Arps and Duong models. The base case and P-case predictions continue to improve, as seen in the narrowing and centering of their histogram distributions. In addition, with a minimum of 12 months of production data, the coverage rates tend to approach or exceed the theoretically predicted 0.80 point. These results indicate that more production data lead to predictions about future production that are more accurate (higher coverage rates) and more precise (narrower brackets).

Table 4.1 summarizes trends seen in Figures 4.9, 4.10, 4.11, and 4.12 concerning the positions at which the histograms in those charts are centered. Generally, it can be seen from looking at the charts and Table 4.1 that for both Arps and Duong models, in both horizontal and vertical wells, the P50 case does a better job of predicting production than does the base case, by about 5% to 10%. Cumulative production estimates made using the deterministic approach, the Duong model, and at least 12 to 18 months of production data will result in base case histograms that center on zero percent variation with the actual historical production data. Under these circumstances, the Arps model will tend to underestimate production by about 15-20% for both horizontal and vertical wells. When deterministic estimates are generated with either model with 36 months of production data or more, the percent difference histograms are centered on zero. With

Table 4.1: Various trends can be observed in this summary of the percent difference bins upon which the histograms in Figures 4.9, 4.10, 4.11, and 4.12 are centered, according to classifications categorized by the horizontal and vertical well sets, using either the Arps or Duong model, when various ranges of months of production data are used to model decline curve properties.

			Histogram Centers (Percentage Difference Bin)			
Well Data Set	Model	Month Range Used to Model Decline Parameters	Base Case	P10 Case	P50 Case	P90 Case
Horizontal Wells	Arps	1 to 6	-	-	-	-
		1 to 12	-20	20	-10	-30
		1 to 18	-15	5	-10	-25
		1 to 24	-5	5	-5	-15
		1 to 36	0	5	0	-10
		1 to 48	0	0	0	-10
	Duong	1 to 6	-	-	-	-
		1 to 12	0	40	15	-10
		1 to 18	0	30	5	-10
		1 to 24	0	20	0	-10
		1 to 36	0	10	0	0
		1 to 48	0	5	0	0
Vertical Wells	Arps	1 to 6	-	-	-	-
		1 to 12	-15	20	0	-25
		1 to 18	-10	10	0	-20
		1 to 24	0	5	0	-5
		1 to 36	0	10	0	0
		1 to 48	0	0	0	-5
	Duong	1 to 6	0	-	-	-10
		1 to 12	0	20	15	-5
		1 to 18	0	10	10	-5
		1 to 24	0	5	5	0
		1 to 36	0	10	0	-5
		1 to 48	0	0	0	0

this threshold of data, both models provide comparable and reasonable coverage rates for both the horizontal and vertical well sets.

With regard to using the stochastic approach to generate P10 estimates, the Arps model tends to overestimate the actual case by about 15-20% less than does the Duong model. This effect is more pronounced when less production data is used to model the decline curve parameters. Concerning P50 estimates generated for horizontal wells, the Arps model tends to underestimate the actual case by about 5-10% and Duong tends to overestimate it about 5-15%. The P50 results generated for vertical wells reveal that the Arps model tends to provide highly reliable estimates (approximately 0% difference), whereas the Duong model overestimates the actual production data by approximately 10% when 24 months or less data is used, but provides highly precise estimates when 36 months data or more are used to model decline curve parameters. With regard to P90 estimates, the Duong model underestimates the actual case by about 5% to 20% less than the Arps model does, generally making the Duong model a better option to analyze production data from both horizontal and vertical wells across all ranges of months.

4.3 Effect of Eliminating Early Time Data Upon the Reliability of Decline Curve Predictions

The presence of transient flow data may cause anomalies in the determination of the base case or the P-cases. An additional set of simulations was therefore executed to determine whether elimination of early data might improve coverage rates and histogram EUR distributions. Figures 4.13, 4.14, 4.15 and 4.16 show the results of eliminating the first 6, 12, or 18 months of data in the horizontal and vertical data sets, and applying the Arps and Duong models. These simulations are compared with the corresponding cases where such data was not eliminated. In all of these simulations, the first 36 months of data were presumed to be available. The results from this inquiry show decreased coverage rates and expanding histogram distributions when a successively increasing number of months of early data are eliminated. Both of these trends are disadvantageous to making good predictions. The most probable explanation for this outcome is that eliminating early time data does not remove transient data from the data sets. Transient data was likely already removed when the initial 3 or 4 months of data were eliminated, as part of the process of determining the initial production rate, when the data sets were cleaned up prior to their analysis. On the other hand, eliminating additional months of data enables the models to utilize less information upon which to base their predictions.

For these reasons, it is better not to eliminate additional months of early data. Further support for this proposition is seen in Figures 4.18 and 4.19, which plot production rate versus time on log-log coordinate axes. The generally flat lines with half slopes indicate that linear flow has developed.

The results in Figures 4.9, 4.10, 4.11, and 4.12 can be used by an engineer to develop predictions, depending on the available data. For example, if 12 months of data are available for a group of vertical wells in the Barnett Shale, an engineer can discern from an examination of Figure 4.12 that the curved created by the Duong model from regression to original production data will generally yield reliable predictions about the next five years of production. Empirical observation thus drives this pragmatic method, which can be easily replicated on various sets of analog wells in order to generate predictions about groups of wells. Although this method could be used to predict the behavior of individual wells also, it should be recognized that the histogram distributions represent a range of outcomes observed in the group of analog wells, and that the behavior of an individual well is less predictable than are the tendencies of the group.

4.4 Dimensionless Bracket Widths

In order to more easily compare the widths of the P10 to P90 brackets shown in the various percent variation histograms, a dimensionless width parameter was developed as follows:

$$\text{dimensionless width} = \frac{\text{P10 EUR} - \text{P90 EUR}}{\text{actual cumulative production}} \tag{4.2}$$

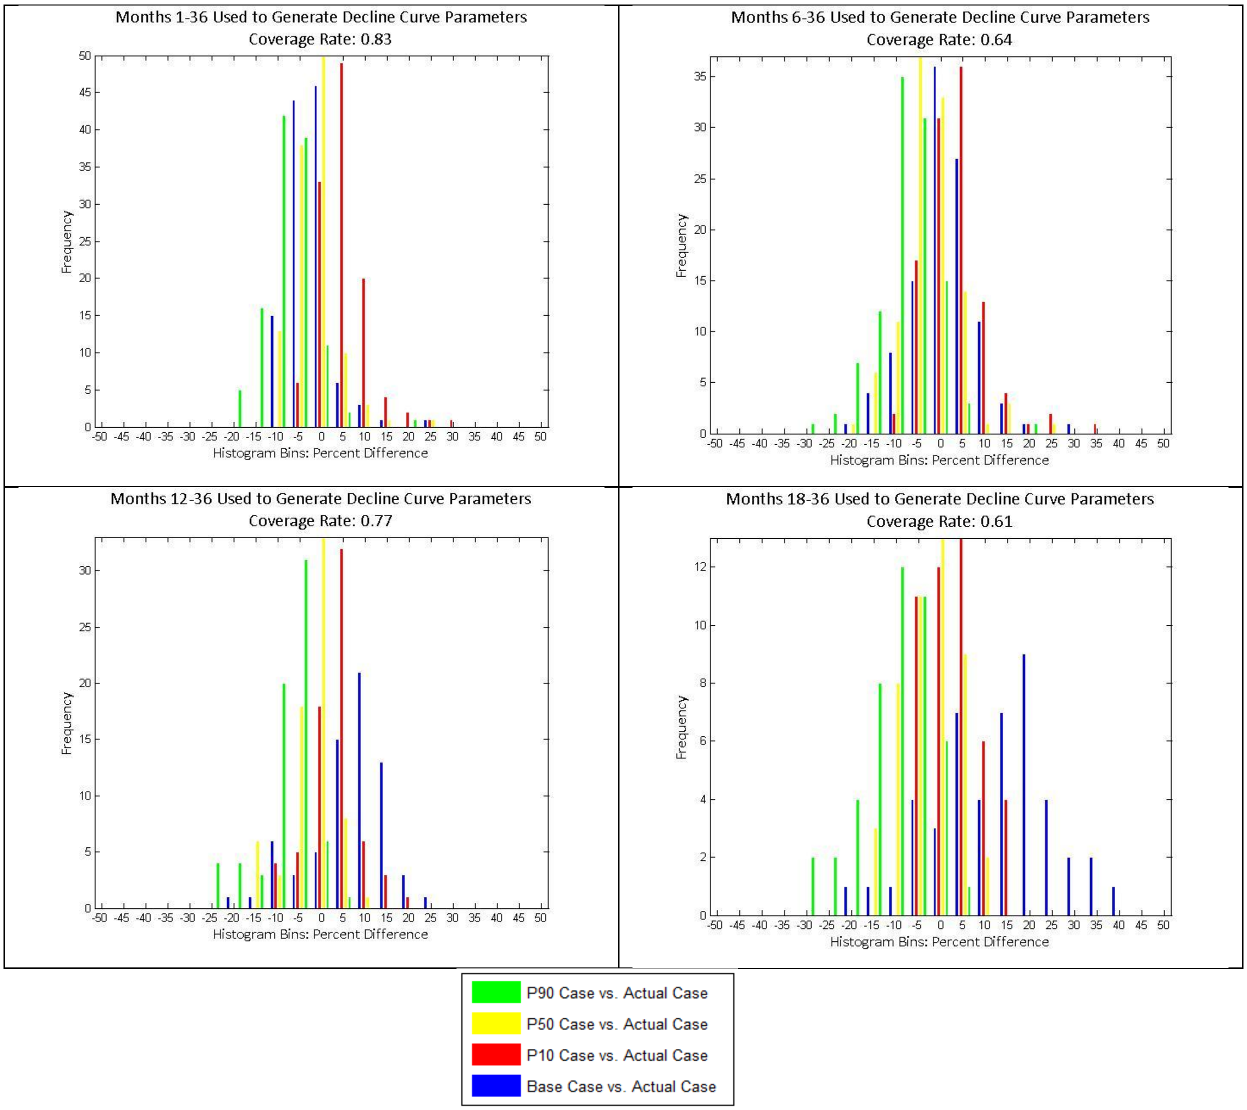

Figure 4.13: Horizontal Well Set, Arps Model. Elimination of data from various numbers of early months causes histogram distributions to become less centered and more dispersed.

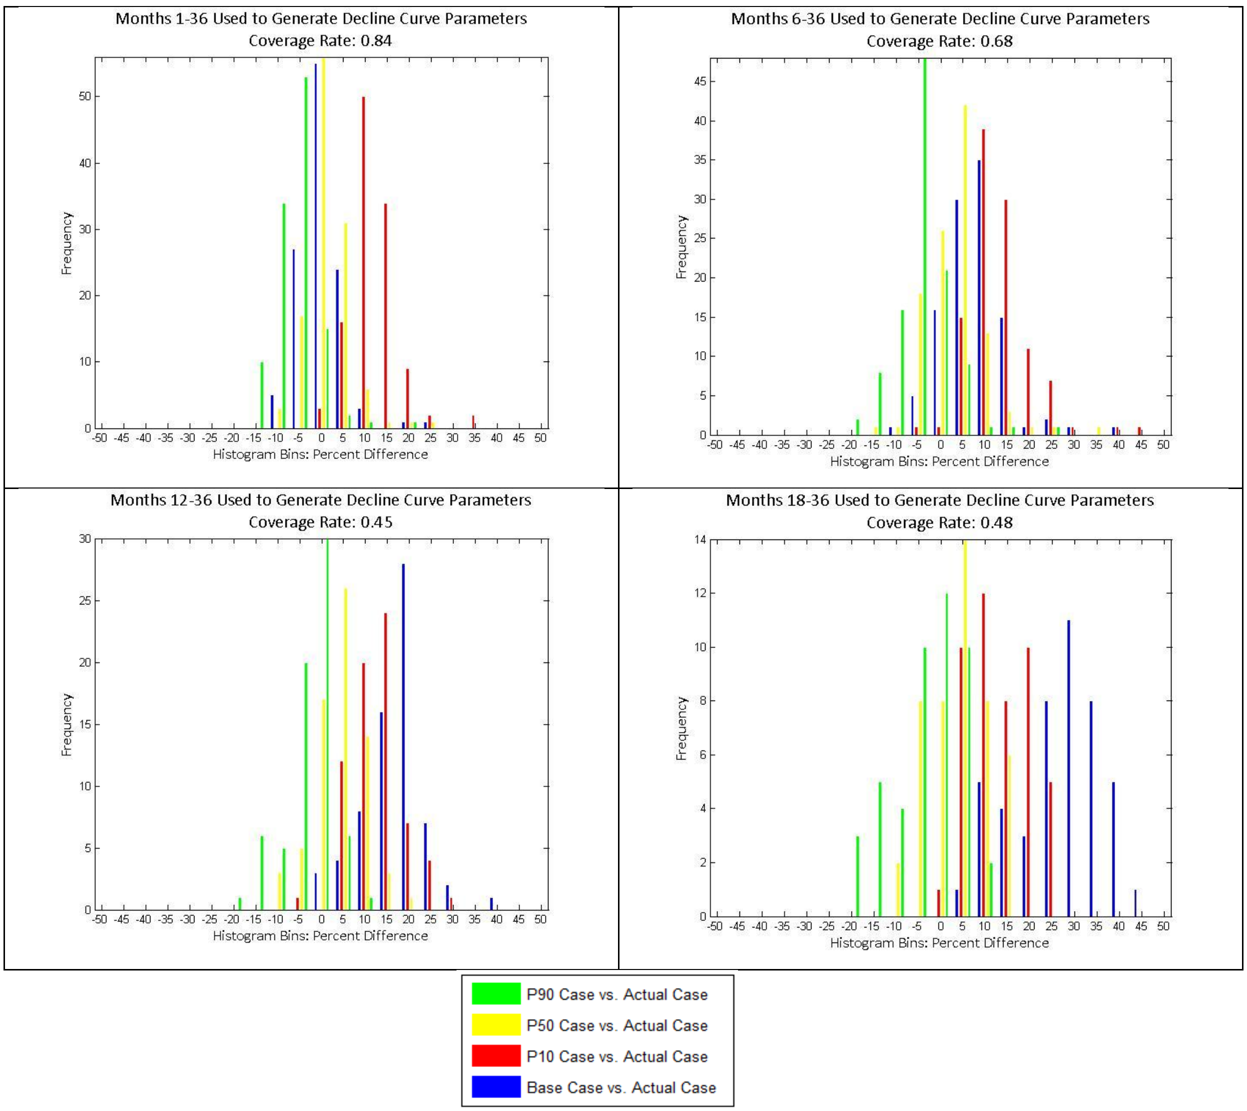

Figure 4.14: Horizontal Well Set, Duong Model. Elimination of data from various numbers of early months causes histogram distributions to become less centered and more dispersed.

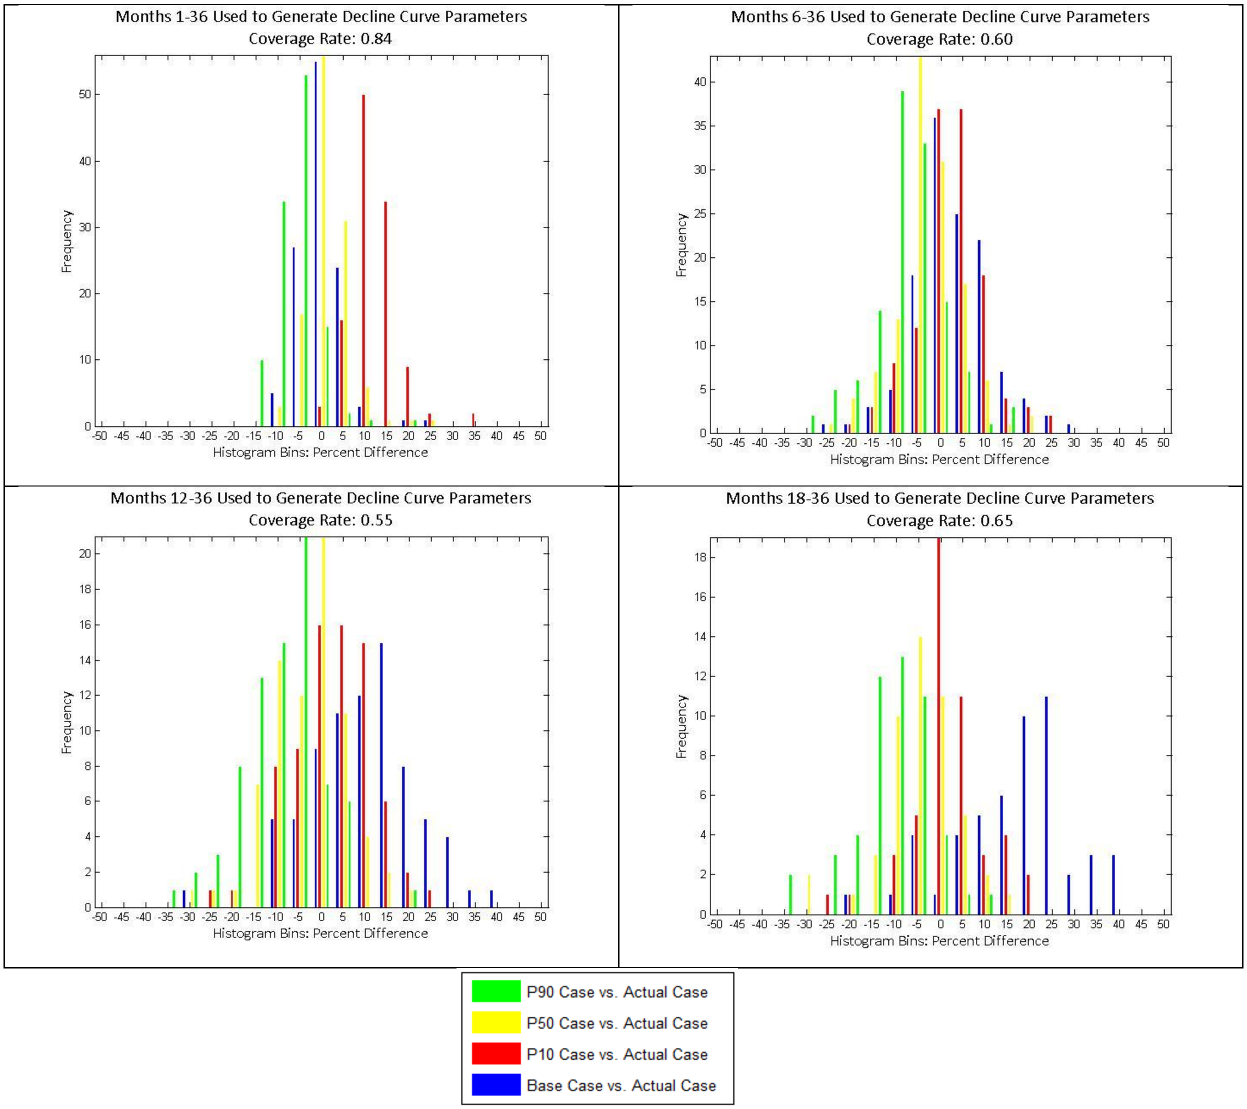

Figure 4.15: Vertical Well Set, Arps Model. Elimination of data from various numbers of early months causes histogram distributions to become less centered and more dispersed.

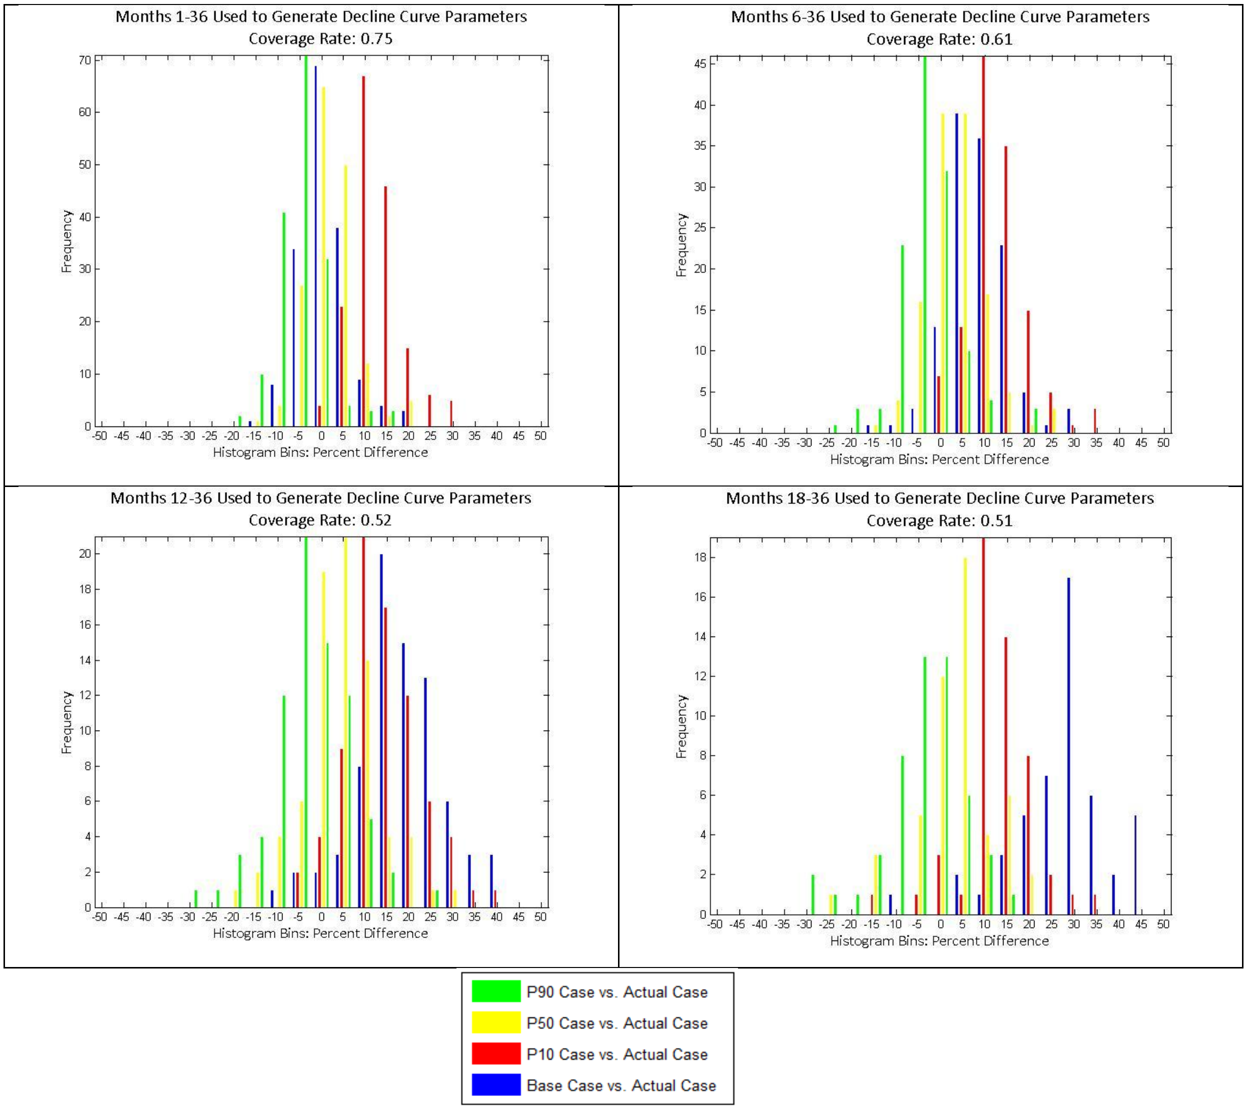

Figure 4.16: Horizontal Well Set, Duong Model. Elimination of data from various numbers of early months causes histogram distributions to become less centered and more dispersed.

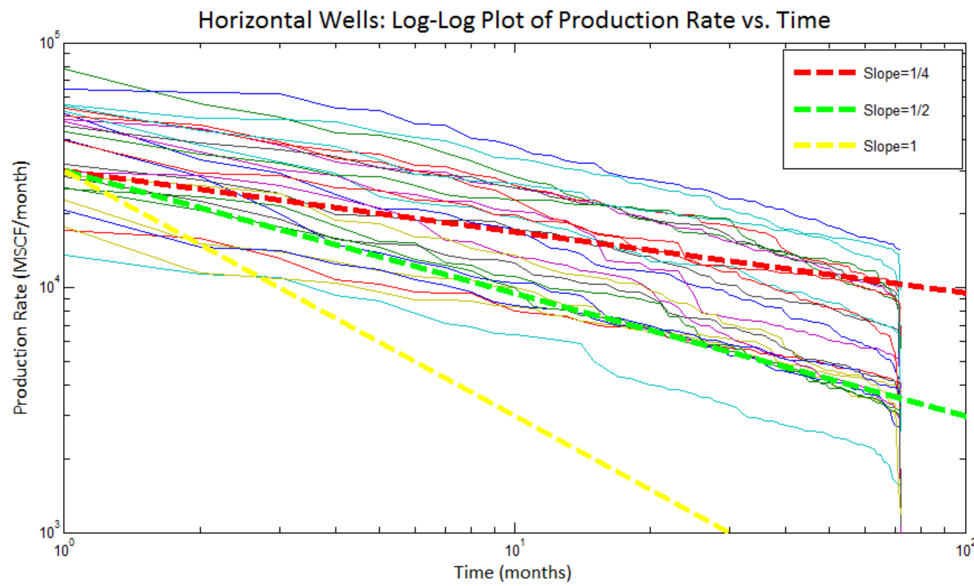

Figure 4.17: The half-slope lines in a log-log plot of production rate vs. time indicate a stabilized, linear flow regime for the horizontal wells analyzed in this study. This stability conforms with Duong's suggestion that only data from stabilized flow periods should be used to determine parameter values.

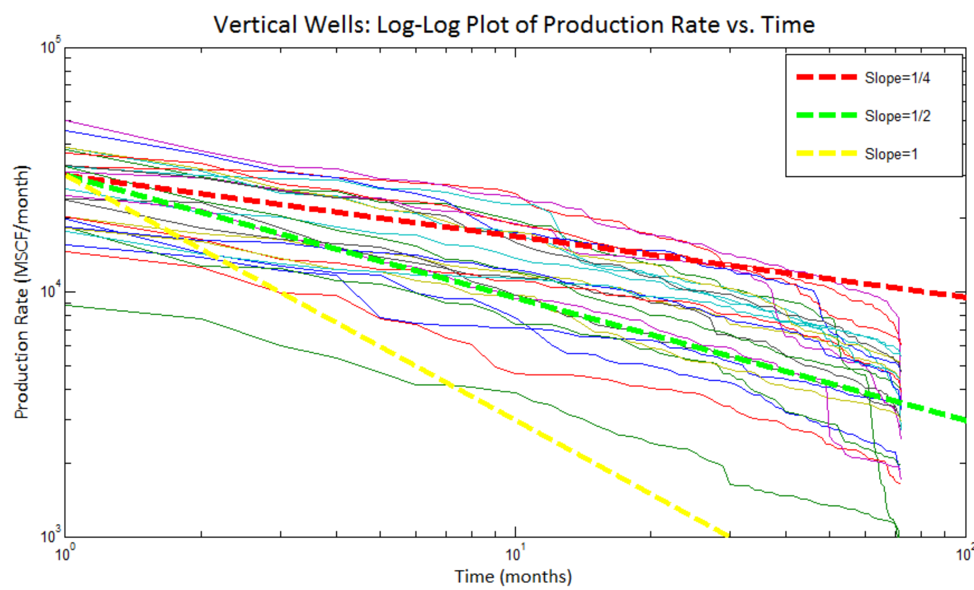

Figure 4.18: The half-slope lines in a log-log plot of production rate vs. time indicate a stabilized, linear flow regime for the vertical wells analyzed in this study.

The results of plotting the dimensionless parameters for each well in the horizontal and vertical well sets are shown in Figure 4.19 (horizontal wells, Arps method) Figure 4.20 (horizontal wells, Duong method), Figure 4.21 (vertical wells, Arps method) and Figure 4.22 (vertical well, Duong method). These charts reveal that for both horizontal and vertical wells, the Arps model is able to generate slightly narrower dimensionless widths than is the Duong model, up to the point where 24 months of production data are used to model the decline curve parameters. Past this time frame, the Arps model loses its ability to generate ever-narrower brackets, but the Duong model continues to generate narrower brackets when additional months of data are used to model the decline curve parameters. This phenomenon is seen in the mottled, overlapping marker symbols associated with the various month ranges in the Arps model, and the sorted, ever-decreasing marker symbols in the Duong model. The Duong model appears to provide more consistent results than does the Arps model when more production data is available. The coverage rates from the various simulations generally approach the confidence interval when approximately 18 months of data are used to model decline curve parameters. If both models thus tend to provide reliable predictions after approximately 18 months, the Arps model may be a better choice when 24 months of production data (or less) are available, and Duong may be a better choice when more than this threshold number of months of production data is available. A general rule of thumb may be that the Arps model should be used when less data are available, and the Duong model should be used to later fine tune cumulative production estimates. The ranges over which the normalized bracket widths tend to appear are summarized in Table 4.2.

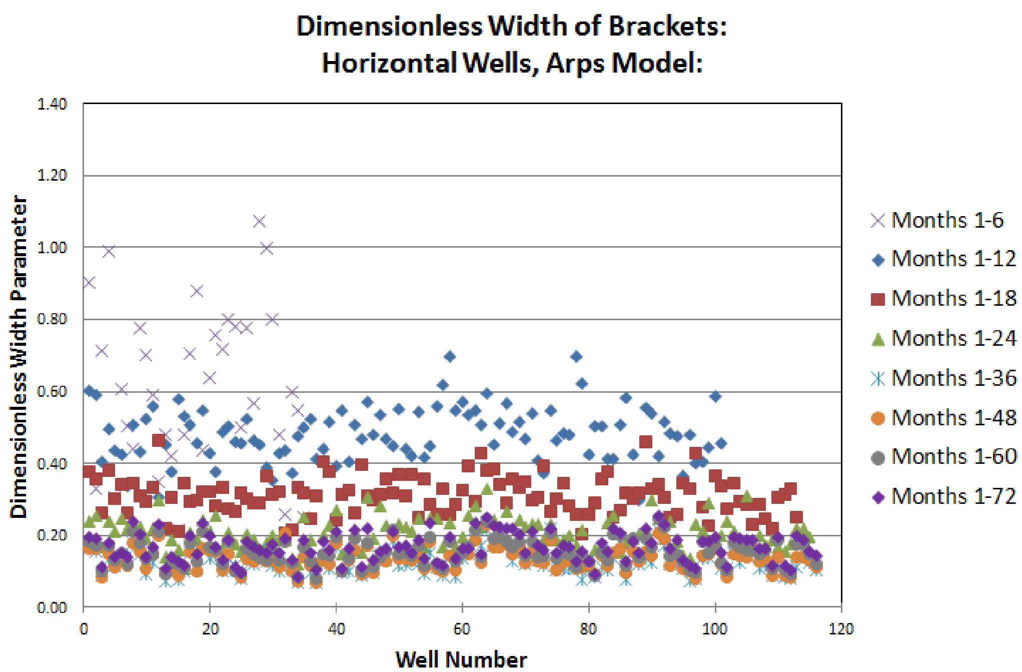

Figure 4.19: Horizontal Wells, Arps model. The dimensionless width of brackets decreases as the number of months of production data used to model decline curve properties increases until the transition between months 1-36 and 1-48.

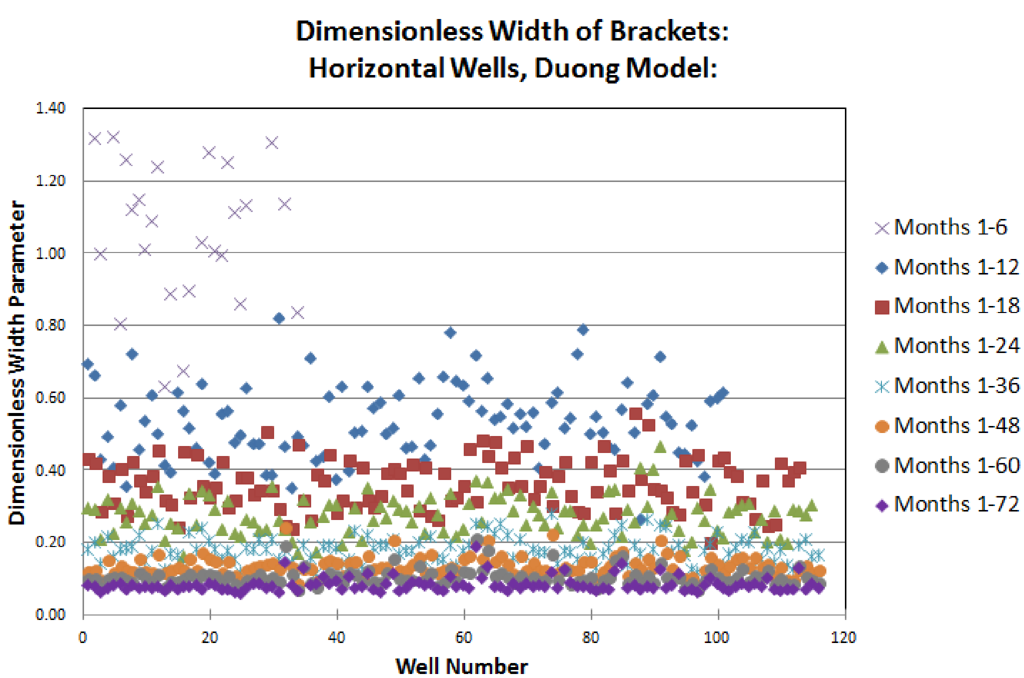

Figure 4.20: Horizontal Wells, Duong model. The dimensionless width of brackets continues to decrease through all successive increases in number of months used to model production decline parameters.

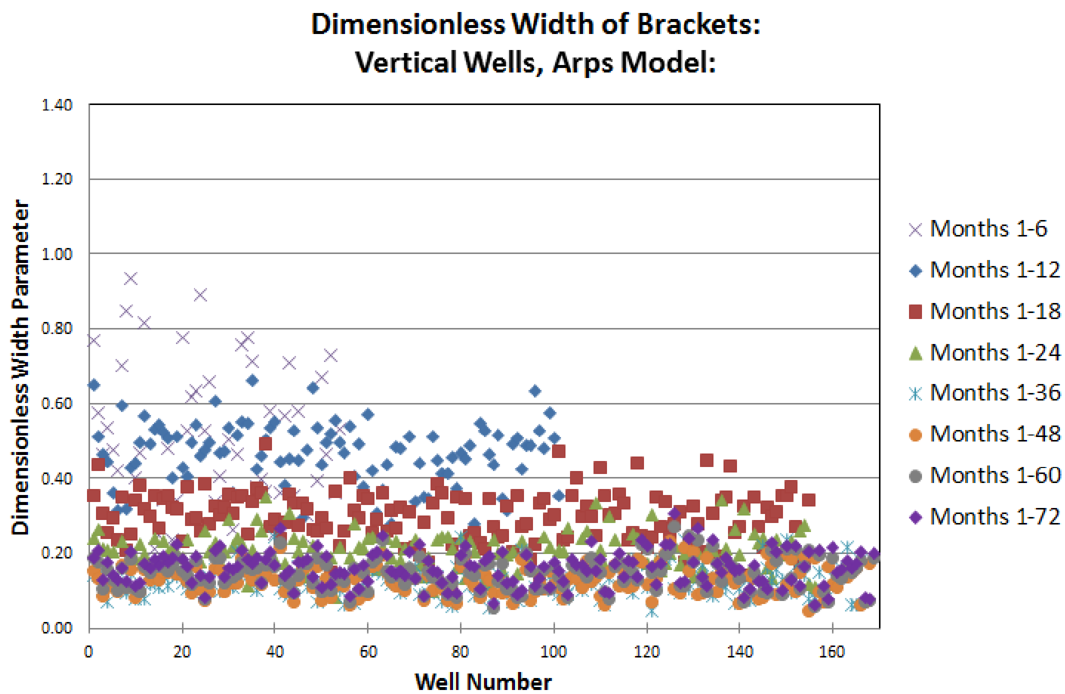

Figure 4.21: Vertical Wells, Arps model. The dimensionless width of brackets decreases as the number of months of production data used to model decline curve properties increases until the transition between months 1-36 and 1-48.

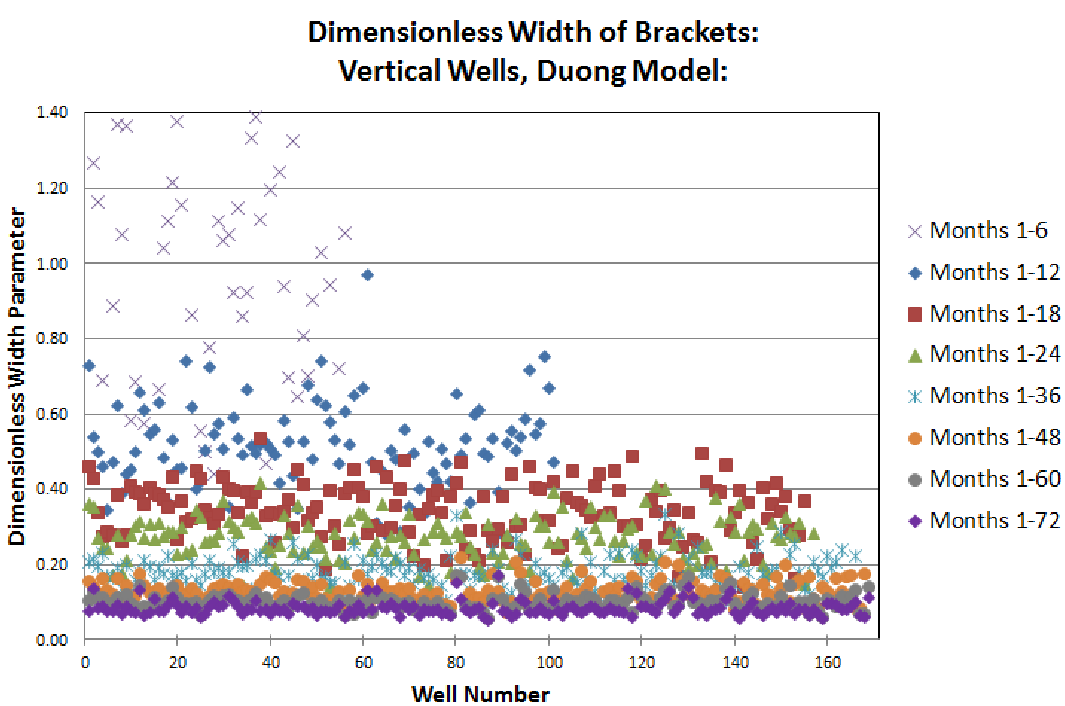

Figure 4.22: Horizontal Wells, Duong model. The dimensionless width of brackets continues to decrease through all successive increases in number of months used to model production decline parameters.

4.5 Comparison of This Research with Cheng's Modified Bootstrap Method

As previously discussed, Cheng et al. (2010) implemented a modified bootstrap method to model production decline using the Arps function, and verified the reliability of probability predictions by comparing them with actual production data. Cheng thus developed an insightful framework for evaluating the effectiveness of the Arps model and the ability of the modified bootstrap method to generate probability cases.

Cheng first examined the operation of the conventional bootstrap method as applied to their data sets, and found that the coverage rate of their well set actually decreased as more and more of the historical production data was used to model the future production decline. This highly counterintuitive result was observed even when transient flow data was excluded from the stochastic analysis. One would normally expect that using more data to model the future trend would lead to better rather than worse predictions about the probability cases. My research did not replicate this result, as I observed coverage rates generally holding steady or improving with expansion of the months of data used to model the decline curve parameters. This discrepancy may be explained in the character of the data sets used in the two research inquiries. Cheng disclosed that they used 100 well data sets, but provide no further information about these wells. We do not know whether they were vertical or horizontal, whether they were drilled into the same formation or in close proximity to one another, or when they were developed. In fact, we do not even know whether these wells were in conventional or unconventional reservoirs. From the figures presented in their article, Cheng appears to have used nicely behaved wells with small fluctuations in production rate from month to month, and therefore low residuals values. Their example wells also seem to span very long production time periods, on the order of 200 months. It thus appears that Cheng

Table 4.2: Summary of the ranges of values exhibited by the dimensionless bracket widths of horizontal and vertical wells, as seen in Figures 4.19, 4.20, 4.21, and 4.22.

Well Data Set	Model	Month Range Used to Model Decline Parameters	Dimensionless Bracket Width Range	
			Low	High
Horizontal Wells	Arps	1 to 6	0.30	1.00
		1 to 12	0.40	0.65
		1 to 18	0.20	0.45
		1 to 24	0.15	0.25
		1 to 36	0.10	0.20
		1 to 48	0.10	0.20
		1 to 60	0.10	0.20
		1 to 72	0.10	0.25
	Duong	1 to 6	0.60	1.40
		1 to 12	0.35	0.80
		1 to 18	0.25	0.50
		1 to 24	0.20	0.40
		1 to 36	0.15	0.25
		1 to 48	0.15	0.20
		1 to 60	0.10	0.14
		1 to 72	0.08	0.10
Vertical Wells	Arps	1 to 6	0.20	0.90
		1 to 12	0.30	0.65
		1 to 18	0.20	0.45
		1 to 24	0.18	0.35
		1 to 36	0.10	0.20
		1 to 48	0.05	0.20
		1 to 60	0.10	0.20
		1 to 72	0.10	0.25
	Duong	1 to 6	0.40	1.40
		1 to 12	0.35	0.75
		1 to 18	0.20	0.50
		1 to 24	0.20	0.40
		1 to 36	0.15	0.30
		1 to 48	0.15	0.20
		1 to 60	0.10	0.12
		1 to 72	0.08	0.10

may have used wells distributed far and wide, in order to gather a sufficiently large collection of wells that met certain qualitative characteristics.

Cheng's examples of the conventional bootstrap method are based on using 70 months of historical production data in order to achieve a 68% coverage rate (when excluding transient data) or 69% coverage rate (when including transient data). The process of using a conventional bootstrap method to sample data points from nicely behaved wells with a long span of available data would likely result in probability predictions with narrow brackets, since this method does not allow for much variation in the synthetic data sets generated in this manner. Narrow brackets, in turn, would explain why the coverage rates suffered so dramatically. Cheng then (presumably) used these same data sets to evaluate the modified bootstrap described in their paper.

To model decline curve parameters in their modified bootstrap method, Cheng recommend using monthly production data from periods that represent 20%, 30%, and 50% of the available production data. But Cheng et al. do not state whether this method has been validated in cases other than when long periods of production data are available. For example, it is not clear whether the process of disproportionately weighting data will lead to reliable predictions when only 12 or 18 months of production data are available.

The weakness in relying upon production data sets with low residuals values over very long time spans is that the character of these data sets is not representative of most data sets from shale formations. We should therefore call into question general principles gleaned from such a process, as they may not find valid application to more typical data sets that exhibit natural production fluctuations from month to month. In addition, since most shale wells will probably need restimulation within their first decade of production in order to produce at economic rates, Cheng's modified bootstrap method may not work well when applied to producing shale formations.

In contrast to the data used in Cheng's research, the well data sets used in this research project showed higher production fluctuations from month to month, and were thus representative of the production decline commonly observed. In addition, the data sets spanned only 72 months, and are thus representative of typical shale wells. The generated probability brackets tended to be wider and the predictions were more reliable than in Cheng's investigation of the conventional bootstrap method. While an ideal scenario would be to have the best of both worlds – narrow brackets and high reliability – these features are tradeoffs where the question becomes one of balance. With the current state of this research into using the bootstrap method to generate stochastic estimates, this balancing objective is met by recognizing that there are minimum requirements of the input data to achieve narrow brackets and good coverage rates. This research reveals that at least 12 to 18 months of production data are generally needed in order to adequately predict a well's production behavior during the remainder of its first 6 years. Although analysts might prefer to have better predictions even sooner, this minimum number of months likely represents a reasonable timeframe in many contexts.

CHAPTER 5

CONCLUSIONS

This research project used computer regression techniques to examine the performance of the Arps and Duong decline curve models, in conjunction with deterministic and stochastic analysis techniques, as applied to 285 Barnett Shale gas wells. The bootstrapping method randomly samples a production data set to create synthetic data sets, to which decline curves can be fit to generate cumulative production outcomes. From the resulting distribution of cumulative production values, probability cases can be determined. The bootstrapping method offers an advantage over a traditional Monte Carlo simulation: parameter ranges and their distributions need not be known in advance.

By comparing historical production data from sets of horizontal and vertical wells with decline curve predictions about the future production of those wells, the reliability of predictions generated by the Arps and Duong models were evaluated when various ranges of months of production data were used to model decline curve parameters.

Prior research has applied the bootstrap method to model production decline. Jochen et al. (1996) used the bootstrap technique to generate stochastic predictions about field performance, but did not evaluate the resulting probabilistic predictions by comparing them with actual production data. Cheng et al. (2010) used the bootstrap method to generate probabilistic predictions and evaluated the results with field data, but focused primarily upon improving probabilistic predications for conventional wells with substantial production histories. Both the Jochen and Cheng studies used only the Arps model in their approaches.

The research described in this thesis is the first to use both deterministic and stochastic methods to compare the performance of the both the Arps and Duong models as applied to shale gas production data. Deterministic predictions were created by fitting decline curves to historical production data, whereas stochastic estimates were generated with the bootstrapping technique. Various ranges of months of production data were used to model decline curve parameter values so that the reliability of cumulative production predictions generated from those month ranges could be gauged. Based on these analyses, this research draws the following conclusions:

5.1 Conclusions Concerning the Number of Months of Production Data Used to Model Decline Curve Properties

When only 6 months of production data are available, neither deterministic nor stochastic methods can yield reliable predictions about a well's future production.

When 12 or more months of production data are available, both deterministic analysis, and stochastic analyses using the bootstrap method, can reasonably predict a well's cumulative production over the first 72 months of its producing lifetime.

Increasing the number of months of production data that are used to model decline curve parameters will generally improve a model's predictive performance. Successive additions of 6 to 12 months of the production data used to model decline curve properties typically cause the range of percent differences to decrease approximately 10% to 15%, and the histogram centers to shift toward zero percent variation by approximately 5% to 10%.

5.2 Conclusions Concerning Deterministic Predictions

When the Duong model is used with at least 12 to 18 months of production data, the histograms that show the normalized percent by which the base case varies from the actual case will center upon zero percent difference, whereas the Arps model will tend to underestimate production by about 15%-20%. If 36 months of data are available to model decline curve properties, both models generate histograms centered on zero percent difference, and both models provide comparable and reasonable coverage rates. These trends hold true for sets of both horizontal and vertical wells.

5.3 Conclusions Concerning Stochastic Predictions

P10 estimates: The Arps model tends to overestimate the actual case by about 15-20% less than does the Duong model. This tendency of the Duong model to overestimate the actual case is exacerbated when less production data is used to model the decline curve parameters, such that 6 to 12 months of data will not enable reliable estimates.

P50 estimates: For horizontal wells, the Arps model tends to underestimate the actual case by about 5-10% and Duong tends to overestimate it about 5-15%. For vertical wells, that the Arps model provides highly reliable estimates (approximately 0% difference), whereas the Duong model overestimates the actual production data by approximately 10% when 24 months or less data is used, but provides highly precise estimates when 36 months data or more are available.

P90 estimates: When any range of months of production data are available, the Duong model underestimates the actual case by about 5% to 20% less than does the Arps model, making the Duong model a superior choice for P90 estimates.

5.4 Comparison of the Performance of the Base Case and the P50 Case

In both horizontal and vertical wells, the P50 case generally yields better predictions about cumulative production than does the base case, by about 5% to 10% across all ranges of months of production data.

5.5 Conclusions Concerning the Precision of Stochastic Estimates

Examining the normalized distance between the P10 and P90 estimates for each well within the horizontal and vertical well sets, according to the number of months of production data used to generate decline curve properties, reveals certain trends. The Arps model generates more precise predictions (narrower brackets) when 36 months of production data or less are available, whereas the Duong continues to generate ever-narrowing brackets when a successively increasing number of months of production data are used to model decline curve properties. Each additional 12 month increment reduces the bracket width by about half of its prior value, indicating substantial improvements in precision. Therefore, fewer than 36 months of production data are available, the Arps method may be a better model choice to predict the next three years of production, but when more than 36 months of production data are available, the Duong model may serve to better fine tune cumulative production predictions.

5.6 Application of the Approach in This Research to Other Shale Formations

More generally, this research methodology can be extended to groups of wells in other producing formations to characterize the performance of different decline curve models as applied to varying ranges of months of production data. From the results of such an analysis, knowledge can be gained about which model may perform better and how many months of production data are necessary to achieve good predictions about future cumulative production in a specific geologic and geographic context.

CHAPTER 6: FUTURE RESEARCH

The decline curve analysis method can provide valuable insights into the future production potential of oil and gas wells, yet there is still much to learn about this method and its practical application. A number of additional research questions merit further inquiry. The research described in this thesis used data from wells in the Barnett Shale. Although it is believed that similar trends will appear data from other producing formations, the extent to which the base and probability cases differ from the historical producing history may vary. Additional research using production data from other shale formations can be helpful in quantifying these differences, and may provide additional insights into the causes of these variations.

In addition to extending the research conducted for this thesis, a number of other questions about decline curve analysis remain to be explored. For example, many operators rely upon average or typical decline curves for a particular formation, which are based on aggregated production data, to predict the performance of a new well when only its initial production rate is known. Although this practice is common because may offer insights into the performance of a well based on the past performance of analog wells, it relies upon the assumption that a well's initial production and other decline curve parameters are independent of one another. Further research should be conducted to confirm the validity of this assumption.

Another topic of interest stems from the fundamental assumption within decline curve analysis that production rates will decline steadily and predictably with time. Good theoretical bases justify this expectation, yet such an outcome is not consistently observed in practice. In the research conducted for this thesis, for example, the group of vertical wells selected for study contained only 169 wells, out of the 396 wells originally in the set, that exhibited a visually identifiable decline trend. If fewer than half of the wells in a set will readily submit to decline curve analysis techniques, some method for analyzing and predicting the behavior of the other wells would be very valuable.

A final thought regarding future research derives from the practice of restimulating wells to extend their producing lifetimes. In the Barnett vertical wells used in this research, it was not uncommon to observe the effects of a restimulation, in the form of dramatically increased production, at about 72 to 84 months into a well's lifetime. The significant changes to reservoir flow dynamics that result from restimulation efforts render the prior decline curve trajectory incapable of directly predicting a well's future performance . Yet it is not well understood whether the properties of the succeeding decline trend will match or depart from those of the prior one. To better understand the second (or third, or fourth) decline trend is fertile topic for investigation, as this knowledge could assist in developing better predictions about the reserves potential contained within a reservoir.

REFERENCES CITED

Ahmed, T., 2006. Reservoir Engineering Handbook (Third Edition). Elsevier, Amsterdam.

Arps, J.J., 1945. Analysis of Decline Curves. Trans. AIME (1945) 160, 228-247.

Berman, A. and Pittinger, L., 2012. Shale Play Assessment Methods: A Discussion. Presentation made in Tulsa, Oklahoma, March 9.

Cheng, Y., Wang., Y., McVay, D., and Lee, J., 2010. Practical Application of a Probabilistic Approach to Estimate Reserves Using Production Decline Data. SPE Journal of Economics and Management, April.

Craft, B., Hawkins, M., and Terry, R., 1991. Applied Petroleum Reservoir Engineering (2nd Edition), Prentice Hall.

Duong, A., 2011. Rate-Decline Analysis for Fracture-Dominated Shale Reservoirs. SPE Journal of Reservoir Evaluation and Engineering, June 2011.

Efron, B. and Tibshirani, R., 1998. An Introduction to the Bootstrap. Monographs on Statistics and Probability 57. Chapman & Hall/CRC. Boca Raton, FL.

Fetkovich, M.J., Vienot, M., Bradley, M. Kiesow, U., 1987. Decline Curve Analysis Using Type Curves: Case Histories. SPE Joural of Formation Evaluation. Volume 2, Number 4. December 1987. SPE 13169-PA.

Fetkovich, M.J., Fetkovich, E., and Fetkovich, M.D., 1996. Useful Concepts for Decline-Curve Forecasting, Reserve Estimation, and Analysis. SPE Journal of Reservoir Engineering, February.

Jochen, V., Holditch, S., and Spivey, J., 1996. Probabilistic Reserves Estimation Using Decline Curve Analysis with the Bootstrap Method. SPE 36633.

Ilk, D., Perego, A., Rushing, J., and Blasingame, T., 2008. Exponential Versus Hyperbolic Decline in Tight Gas Sands – Understanding the Origin and Implications for Reserves Estimates Using Arps' Decline Curves. SPE 116731.

Lee, J. and Sidle, R., 2010. Gas-Reserves Estimation in Resource Plays. SPE Journal of Economics and Management. October.

Mishra, S., 2012. A New Approach to Reserves Estimation in Shale Gas Reservoirs Using Multiple Decline Curve Analysis Models. SPE 161092.

Okouma, V., Hosseinpour, N., and Blasingame, T., 2012. Practical Considerations for Decline Curve Analysis in Unconventional Reservoirs – Application of Recently Developed Time-Rate Relations. SPE 162910.

Society of Petroleum Engineers, 2011. Guidelines for Application of the Petroleum Resources Management System (SPE-PRMS). Richardson, Texas.

Valko, P., 2010. A Better Way to Forecast Production From Unconventional Gas Wells. SPE 134231.

Weibull, W., 1951. A Statistical Distribution Function of Wide Applicability. Journal of Applied Mechanics. Vol. 18, pp. 293-297.

APPENDIX A

Production Data Used in This Research: Barnett Horizontal Wells

Lease Name	Well Number
AGF RANCH ET AL B	1H
NOBLES UNIT	2H
HAM J A	5
BEARD	1
COLE TRUSTS 576 A	14H
COLE TRUSTS 576 A	13H
CALLAGHAN JOHN C	1
DUNN HOBBY D	1
DUNN HOBBY D	2
SEWELL RANCH A	6
BROWN AGNES	4
SULLIVAN VENA	1
EVERS	9H
FRANK J J GAS UNIT B	3
CARAWAY MONROE A	1
SHOOP A	10
GIBBS HOWARD	5H
FRANK J J GAS UNIT B	4
BEECHWOOD	D1H
CATES PARTNER UNIT	2
PARKEY UNIT	3H
JONES	6H
CHARALAMBOPOULOS	1
DCCO 1-SULLIVAN P G A	15H
SLAY C W	11
EWING G	7H
FORTENBERRY	2H
DCC0 1-HOSEK SHIRLEY M	12H
DCCO 1- DCTC	17H
DCCO 1-PETERSON RUBY	10
WOOD JAMES L GU	11H
DCCO 1-MORRIS TED	21H
WCCO 4-LOGAN H H	11
DCCO 1-SULLIVAN P G A	16
REED-HATHORN UNIT	11H

Lease Name	Well Number
HARPER CAROLYN	1H
BURNS ANNA BETH	20H
BILTMORE	4H
CASTO-CHIEF	1H
SHRYACK	1H
JONES	5H
GRAHAM HEIRS	18H
ROANOKE RANCH	1
BURNS ANNA BETH	19H
FARMER-RUSSELL	1H
CHADWICK RESOURCES GU B	1
SEALS G T	18H
CARDWELL-SATER	1
TAYLOR R B	6
AGF RANCH ET AL C	1
YARBROUGH B ET AL	1
LAU	11
TINDLE C UNIT	2H
MCKAMY-LITTLE UNIT	4
EAGLE FARMS UNIT 11	1
PUGH GAS UNIT	3
HOLLAND C GAS UNIT A	1
WHITEHALL	1H
WILKES ROY M	4
L B M BRYAN GAS UNIT	3
BROWN AGNES	5
ALLIANCE	E2H
DUNN CLARA LYLES	14
CHADWICK RESOURCES LTD	5
OBENCHAIN	B3H
DILG M J ESTATE	1
HARDEMAN SAVOIE	1H
NOBLES CHARITABLE TRUST	1
RLM	2H
LAU	12

Lease Name	Well Number
PENLEY W E	1H
KEESE UNIT	1H
BILTMORE	6H
READ	1
FOX ZINA GU 2	6
ALLIANCE	E1H
WAGGONER CRYSTELLE GU	34
VISTA MANUFACTURING	9H
HARRIS UNIT	4H
KNAPP	10
KNAPP	11
DENTON CJW	1
MCRAE L DUESSEN	3H
BROWN	1H
SMITH H A	2H
SMITH MILDRED	1
DOWNE	4H
COLE TRUSTS 576 A	10
AARON-STAR	1H
STONE S C	1
DUNN HOBBY ESTATE	7
LAWRENCE UNIT	1H
GIBBS CROCKETT	1H
FARMER BETTY	1H
NORTH-SCHLUTER UNIT	1H
COLE M T TRUST TWO A	3
NORTH JERRY	5H
LIPSCOMB K P GU	14
FURST RANCH	C2H
GIBBS HOWARD	3H
STANCIL-ROBERTS	3
DUNN A	3H
MITCHELL	1
TALLEY ONETA	4H
OBENCHAIN	D1H

Lease Name	Well Number
HICKORY PARK UNIT	1H
LIPSCOMB K P GU	16
MACKEY W L GU 3 (SA)	12
DCCO 2-SULLIVAN P G	21
CHADWICK RESOURCES GU C	1
SANDERS FRED	11
LIPSCOMB K P GU	15
TROPHY CLUB 8	8205SL
GREEN	1H
PORTER-RILEY	1H
KANN MCDONALD	1

APPENDIX B

Production Data Used in This Research: Barnett Vertical Wells

Lease Name	Well Number
G A D PARTNERS	2
BONDS RANCH	B4
VINSON FEE D	2
COLE TRUST FOUR	4
COLE TRUST FOUR	5
FARLOW J F GAS UNIT A	3
BOAZ W O	14
BLAKLEY EST GU B	7
LAU	7
BLAKLEY VILLA ESTATE A	3
SCHWARZ DIETER GAS UNIT A	15
WISE JAMES E	6
MASSEY GEORGIA GAS UNIT	4
MASSEY GEORGIA GAS UNIT	6
DCCO 2-JONES LEWIS	10
NICHOLSON	8
G A D PARTNERS	5
PIRKLE LORENE	12
NALER MAE C	11
RIGGS J H GAS UNIT	8
TRINITY INDUSTRIES	1
WISE JAMES E	8
GAGE REX	1
BLAKLEY EST GU B	6
WCCO 4-MORRIS ADA	15
HARRINGTON	2
VANN CATTLE YARDS	A1
JOHNSON LOTTIE BARTON	20
VINSON FEE G	8
VINSON FEE G	14
WCCO 4-MILLER WM GU 2	11
BONDS RANCH	E3
FOSSIL HILL	9
HARRINGTON	3
PAYTON J O	4

Lease Name	Well Number
HOYL DALE	4
BLAKLEY EST GU E	6
GRAHAM-SHOOP	5
RLM	1H
WCCO 1-SIMS THOMAS P A	3
BUREL UNIT	3D
COPELAND UNIT D2	D2
CURRIE FLORENCE	9
GRANT HEIRS POOLED UNIT	8
BRYAN D	8
FOSTER WANDA	2
WCCO 4-WM MILLER GU 2	8
SANDFIELD L P	16
MCRAE	20
LOGAN-MILLER	6
COLE TRUSTS 576	2
DCCO 2-HARDEMAN C J	9
SANDFIELD L P	14
SANDFIELD L P	15
CURRIE FLORENCE	8
MCCURDY FLOYD GAS UNIT B	4
CRABTREE	5
YOUNG STELLA GU	12
TXI GAS UNIT	5
SANDFIELD L P	13
BONDS RANCH	F2
KRUM SOUTH	1
G A D PARTNERS	1
CHAPEL HILL	C2
CHAPEL HILL	C3
DCCO 2-WILLIAMS RAY	13
HARMONSON MORRIS	4
PEEK-CRABTREE GAS UNIT	5
MCKELVEY CLARK W FLP	5
GARRETT RUFUS GAS UNIT A	8

Lease Name	Well Number
WYATT	1
CRABTREE	6
SCHWARZ DIETER	9
GARRETT RUFUS GAS UNIT A	7
SHIFFLETT B UNIT	5D
BONDS RANCH	4018
SHOOP GLENN P	4
BLAKLEY EST GU E	11
ASKEY W A B	10
WELDON YOUNG UNIT	7D
TEDROW WILLIAM GAS UNIT	9
SCHOOLFIELD UNIT	10D
CADDELL UNIT	9D
KELLEY JIM BILL GU	6
C W B	22D
BONDS RANCH	504
COX	7
INDIAN CREEK	R5
INDIAN CREEK	I6T
GRANT HEIRS POOLED UNIT	5
WILKES ROY M GAS UNIT A	3
GRANT HEIRS POOLED UNIT	3
PEEK-CRABTREE GAS UNIT	7
SEWELL RANCH INC	4
COPELAND	B3
SHOOP A	11
COX	6
TADLOCK MARGARET	2
CADDELL UNIT	8D
CLOUGH ADELE GU A	9
CRABTREE	8
MCCURDY A C	3D
PATTERSON M L	8
BONDS RANCH	1302
MILLER-CUFFMAN UNIT	13D

Lease Name	Well Number
BURKHALTER EWELL	12D
PARKEY B	9D
COPELAND	A3
BANKS A UNIT	10D
TEDROW WILLIAM GAS UNIT	10
BLAKLEY EST GU E	9
BOAZ W O	22D
LANGLEY A UNIT	6D
LANGLEY UNIT	11D
LANGLEY A UNIT	7D
C W A	28D
BANKS A UNIT	11D
BURKHALTER EWELL	13D
BEARD	4A
STRAIGHT WAYNE	1
BONDS RANCH	508
WCCO 4-MILLER WM GU 2	12
CHISM ROYCE	7
JEWEL B	1
H35	1
BONDS RANCH A	107
WILKES ROY M GAS UNIT A	4
VINSON FEE G	19
INDIAN CREEK	K4
WCCO 4-FITCH GEORGIE	9
BONDS RANCH A	115
MOORE GORDON F	12
WCCO 4-FITCH GEORGIE	8
BONDS RANCH	507
SEALS G T	15
GRANT HEIRS POOLED UNIT	6
GRIFFIN S H ESTATE	15
PIRKLE LORENE	13
GRAHAM-SHOOP	9
MALONE SAMUEL	9

Lease Name	Well Number
FINLAYSON O T A	7
LOGAN-MILLER	8
WCCO 4-MILLER WILLIAM GU B-1	12
COPELAND	C1
DCCO 1-MORRIS TED	12
COPELAND	B2
MCCURDY A C	4
JARVIS-COHEN	4
E P R I	4D
HBB UNIT	1D
VINSON FEE H	11
COPELAND	B1
COPELAND	A2
DUNN HOBBY ESTATE	9
YOUNG STELLA GU A	17
SWAFFORD DON	4
COLE TRUST FOUR	7
BONDS RANCH A	108D
BLAKLEY EST GU E	10
LOGAN-MILLER	9
WHITE AMY	2
MAEYERS DAN B GAS UNIT A	7
WALLACE FRANK	9
CHAPEL HILL	D3
GRIFFIN S H ESTATE	14
SULLIVAN-MUSE	1D
WISE JAMES E	9
KNAPP	9
KNAPP	8
MCALISTER O H	10