Forecasting Drilling Bits Demand:

A New Horizon for an Energy Global Technology Company's S&OP

by

Thiago Pinheiro Faury
Bachelor of Science in Industrial Engineering, University of São Paulo / École Centrale de Nantes, 2009

SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE IN SUPPLY CHAIN MANAGEMENT
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

Signature of Author: ____
Department of Supply Chain Management
May 10, 2024

Certified by: ____
Dr. Ilya Jackson
Postdoctoral Associate
Capstone Advisor

Accepted by: ____
Prof. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

Forecasting Drilling Bits Demand: A New Horizon for an Energy Global Technology Company's S&OP

by

Thiago Pinheiro Faury

Submitted to the Program in Supply Chain Management
on May 10, 2024 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

In response to the challenges of forecasting demand for drilling bits, this capstone project aimed to enhance the accuracy and efficiency of demand prediction for a global technology company in the energy sector. The goal was to replace outdated manual forecasting methods with automated causal and time-series models, optimizing the company's Sales and Operations Planning (S&OP) processes in preparation for larger automation within the company's Integrated Business Planning software. Utilizing historical data on rig counts and market shares, these models predicted future bit runs and estimated associated revenues with significantly improved accuracy, reducing global Mean Absolute Percentage Error (MAPE) from 6–8% to 2–3%. The analysis revealed that the performance of these models varied considerably among different geographic units ("geounits"), highlighting the necessity for customized forecasting strategies. Importantly, the best-performing models correlated with geounits' business models, whether rental or bulk sales. Simpler time-series models frequently outperformed more complex causal models, suggesting that complexity does not always yield better forecasting results. This project streamlined the forecasting process and laid a strong foundation for ongoing improvement and adaptation to emerging market conditions. The integration of these models into the company's S&OP practices represents a significant step forward in leveraging technological advancements to maintain a competitive edge in the energy sector.

Capstone Advisor: Dr. Ilya Jackson
Title: Postdoctoral Associate

ACKNOWLEDGMENTS

First, I would like to thank God for giving me this lifetime opportunity to study at MIT, which materialized a long-awaited dream.

A dream is not achieved by a single person, but it reflects the joint effort of a series of people involved to make it happen. My family is a cornerstone in this journey, so I would like to profoundly thank my wife Jackelline and my kids Pedro and Helena who have been my pillar, my joy, and my strength in the ups and downs of the last few years, starting at the online studies for MicroMasters and culminating in the in-person semester at MIT. Your unwavering support and sacrifices made this journey possible, and your encouragement sustained me throughout, thank you!

I am deeply grateful for my parents Sidney and Sandra, and my brother Caio, who have always incentivized my willingness to learn and encouraged me to pursue new challenges. This dream is also yours!

I am equally thankful to my parents-in-law Edson and Laíse, who supported my wife, kids, and me when we moved to Cambridge and during the MIT SCM Study Trek. Thanks for making our lives run smoothly!

I would like to express my deepest gratitude to Rym Khelil, Mohamed Esmat, Dr. Ilya Jackson, and writing coach Toby Gooley, whose guidance, insights, and support were indispensable in delivering this capstone project. Their contributions were truly above and beyond. Additionally, I am thankful for the valuable input and assistance from Gabriel Ossa, Massoud Ahmadi, Steve Dunton, and Vanice Hoffmeister, who were all directly involved and played crucial roles in the project's success.

My big thanks as well to my extended family (siblings-in-law, nieces, nanny) and friends who helped my wife, kids, and me, especially during the two years of MicroMasters and during the semester at MIT campus.

I would also like to thank my company for all the support in selecting and sponsoring me during the entire MicroMasters and SCM Blended programs.

And finally, I am thankful to the MIT SCM Family for the friendships that will always hold a special place in my heart. I am so grateful to be surrounded by such a talented group of people.

TABLE OF CONTENTS

1. Introduction

This capstone project focuses on demand forecasting for oil and gas (O&G) drilling bits used by our sponsor, a global energy technology company with 98,000 employees operating in 120 countries. The company's core business is to supply products and services for production systems, reservoir performance, and well construction.

One of the sponsor's keys to success is demand forecast accuracy. Why does that matter? The sponsor relies on accurate forecasts to optimize drilling schedules, ensure the availability of the right drilling bits at the right time, and maintain efficient supply chain management (preventing equipment shortages or excess inventory). The challenge is to do it well and better than its competitors.

1.1 Motivation

Over the last five years, the sponsor has been professionalizing its planning and supply chain function, structuring its organization and work processes in a demand-to-delivery framework. To support this strategy the company has been investing in digital solutions to facilitate the sales and operations planning process (S&OP). Although the digital backbone is there for reporting demand forecasts, in some business lines the company still relies on people to manually calculate forecasts without any common methodology.

This situation creates an opportunity for certain people in some countries to over-forecast activity and revenue to secure extra resources. For the company, this practice translates into additional assets, inventory, transportation, personnel costs, and less cash flow. Conversely, if resources are misallocated due to over-forecasting in one country, and another country that genuinely needs these resources is unable to complete a project, the company might face penalties and market share loss. This situation can lead to dissatisfied clients who may choose competitors for future projects.

Because headquarters (HQ) is aware of this behavior, it is obliged to adjust any misalignments of activity and revenue, which creates a vicious circle of forecast overrides/underrides. This is where forecasting automation comes into play. Enhanced automation reduces the need for manual oversight, fostering confidence that the projected demand aligns closely with the actual market demand.

Looking from another perspective, the sponsor is gradually migrating to an Enterprise Resource Planning software (ERP), which will allow the company to go to the next level of supply chain automation and integration with its suppliers and customers. One of the key enablers for this strategy is an ERP solution called IBP (Integrated Business Planning). For the sponsor to run IBP it is necessary to have a forecasting model designed in the background so that IBP can automatically calculate the upcoming demand and from that derive production schedules. IBP is expected to be rolled out for the Bits business line by mid-2024, hence this reinforces the urgency and importance of having a forecast model supporting this implementation.

1.2 Problem Statement and Research Question

In oil and gas drilling operations, a drilling bit is the first piece of equipment that goes downhole and therefore is a key component for the success of drilling operations. Every time a bit goes downhole this is called a "bit run." The bit runs reflect the level of activity of the sponsor and are utilized as a demand signal to calculate the number of bits that need to be manufactured.

Forecasting the demand for these bits is quite complex, as they are produced in tens of different diameters, from 3.5 to 36 inches, and with thousands of active designs. Each year hundreds of new designs are introduced. This cascades into tens of thousands of order lines per year and hundreds of thousands of cutting elements (Sponsor Company, personal communication, 2023a).

Some of the bits are manufactured on a build-to-order (BTO) system (similar to make-to-order), while others are engineered-to-order (ETO). BTO bits follow a standard bill of materials and do not require any customization. ETO bits are customized for specific applications and can have additional features. The main types of bits are roller-cone and PDC (polycrystalline diamond compact). Fabrication lead time depends on bit size, type, and bill of materials and can range from two to 16 weeks (Sponsor Company, personal communication, 2023b).

As of today, the sponsor company handles demand forecasting for drilling bits in a manual S&OP process requiring the input of people from 30 geographic business units, called "geounits." A geounit is a set of one or more nearby countries with a similar type of oil and gas operations. The accuracy of this forecast is good enough for the upcoming three months, but this period is not enough to manufacture some drilling bits and to plan long-term capacity and production schedules with the bit manufacturers.

Despite all the complexities described above, while looking at the global historical data every quarter, it is possible to find some correlations between bit runs and certain parameters that determine the level of drilling activity (e.g., rig count, market share, oil price), as presented in Section 3.3. Oil and gas operations are typically cyclical, with yearly seasons and multi-year trends that tend to occur every 7–8 years based on oil price. These level, trend, and seasonality factors can be captured in a time-series forecasting model to predict future demand.

In brief, this capstone addresses the following question:

- How can the sponsor forecast drilling bits activity and associated revenue based on a set of lagging and leading indicators?

1.3 Project Objectives and Expected Outcomes

The main objective of this project is to design a model that calculates the forecast for the upcoming 4–6 quarters of bit runs and associated revenue.

Sponsor stakeholders have clearly outlined the need to have a set of forecasting methods trained and tested for accuracy, across multiple geounits, to find a standard solution that would fit most of the geounits. This standardized solution would subsequently be used to input data into the algorithm for automated forecasting calculations in IBP.

The principal deliverable of this project is the forecasting model for the upcoming 4–6 quarters of bit runs and associated revenue. Expected to serve as a decision-support tool, this model will be utilized by global (HQ), regional (Operations Control Center – OCC), and local (geounit) sales, operations, and supply chain managers to enhance activity and revenue calibration. This, in turn, will lead to better resource planning, improved customer service levels, reduced inventory, and enhanced cash flows.

Building on this pivotal advancement, Chapter 1 highlights the importance of accurate demand forecasting for drilling bits, essential for streamlining operations and strategic planning. This initial perspective serves as a basis for the detailed exploration in the following chapters, culminating in Chapter 4 which scrutinizes the forecasted outcomes and connects these findings to their practical applications within the industry. This approach bridges theoretical methods with practical insights, paving the way for the discussions on implications, limitations, and recommendations in later chapters.

2. State of the Practice

This chapter explores the oil and gas supply chain, breaking it down into upstream, midstream, and downstream sectors, to explain their roles and connections. It addresses how investment influences the industry, challenges common perceptions about its future, particularly regarding renewable energy, and underscores the critical role of demand forecasting. It also provides a detailed look at forecasting, discussing various models and factors that affect predictability in the industry, such as market trends, geopolitical events, and technological advancements. By examining these elements, the chapter provides essential knowledge to understand the detailed forecasting methodology presented in Chapter 3.

2.1 Oil and Gas Supply Chain

The oil and gas supply chain is traditionally divided into three main branches: upstream, midstream, and downstream (as shown in Figure 1). Some companies work within specific portions of a branch, whereas others are vertically integrated and cover most products and services along the chain. *Upstream*, also known as Exploration and Production (E&P), involves locating reserves and extracting oil and gas. *Midstream* handles the processing, storage, and transport of these resources, often across great distances to refineries and markets, using pipelines, trucks, tankers, and rail cars. *Downstream* deals with refining raw materials into products like fuel, plastics, detergents, and synthetic fibers for various consumer applications. Some firms specialize in certain sectors, while vertically integrated companies manage multiple stages of this complex supply chain. (Lisitsa et al., 2019; Hussain et al., 2006).

Figure 1

Oil and Gas Branches

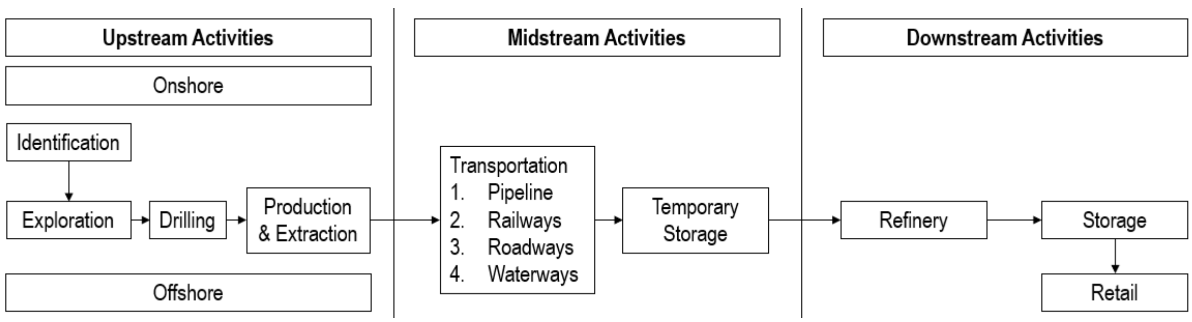

From "Supply-chain management in the oil industry," by S. Lisitsa, A. Levina, and A. Lepekhin, 2019. *E3S Web of Conferences*, *110*, Article 02061 (https://doi.org/10.1051/e3sconf/201911002061).

Understanding the intricate structure of the oil and gas supply chain is crucial for comprehending its operational dynamics and financial implications. Each segment of the supply chain, from upstream to downstream, involves specific activities and contributes to the overall financial health of companies within the industry. The interconnectedness of these segments means that decisions made in one area, such as exploration or transportation, have ripple effects throughout the entire chain (see Figure 2).

Zhu et al. (2020) expand on the oil and gas supply chain's structure to demonstrate the occurrence of the bullwhip effect, a scenario where demand fluctuations are magnified further up the supply chain. They highlight the capital-intensive nature of the industry, noting that exploration and production (E&P) companies, as well as integrated oil and gas firms, allocate significant capital expenditure (CAPEX) budgets annually for exploration, drilling, and equipment. This underscores the importance of accurate demand forecasting at the initial tiers of the supply chain. Accurate forecasts are crucial because they influence investment decisions and operational strategies, which can have substantial financial consequences across all tiers of the supply chain, given the magnitude of the investments involved.

Figure 2

Oil and Gas Tiers – Extended Version

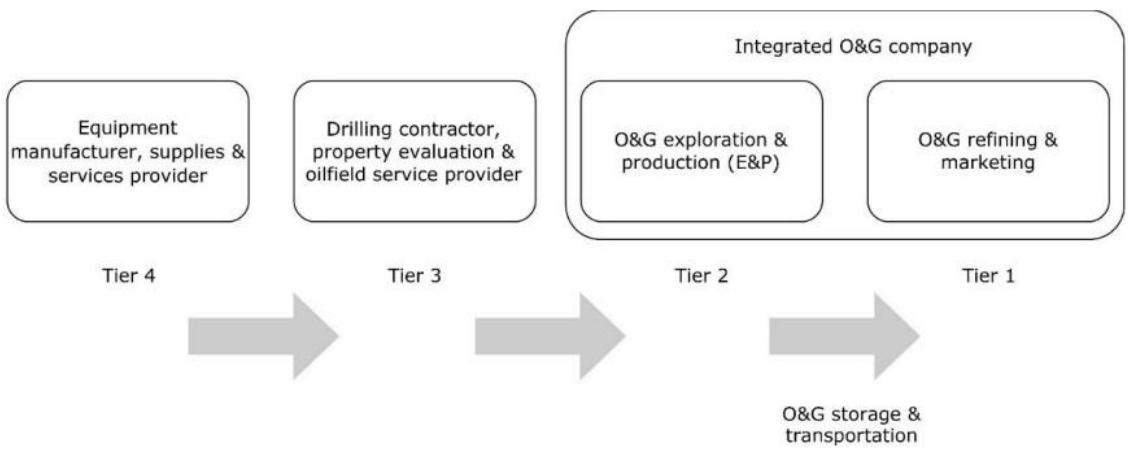

Note. As the demand moves up the tiers, small changes amplify across the supply chain, causing larger variations upstream (this is known as the bullwhip effect). From "Bullwhip Effect in the Oil and Gas Supply Chain: A Multiple-Case Study," by T. Zhu, J. Balakrishnan, and G.J.C. da Silveira, 2020. *International Journal of Production Economics*, *224*, 107548 (https://doi.org/10.1016/j.ijpe.2019.107548).

The inherent volatility of oil prices further compounds this need for precision in forecasting. As commodities, oil and gas prices are susceptible to many influencing factors, including global supply and demand dynamics, economic indicators, geopolitical events, and decisions from the Organization of the Petroleum Exporting Countries (OPEC), as shown in Figure 3. Such price fluctuations introduce additional layers of complexity and risk to the supply chain, magnifying the impact of the bullwhip effect and emphasizing the interconnectedness of accurate forecasting and market price dynamics in guiding strategic and operational decisions within the industry.

Figure 3

How oil prices have reacted to political and economic events

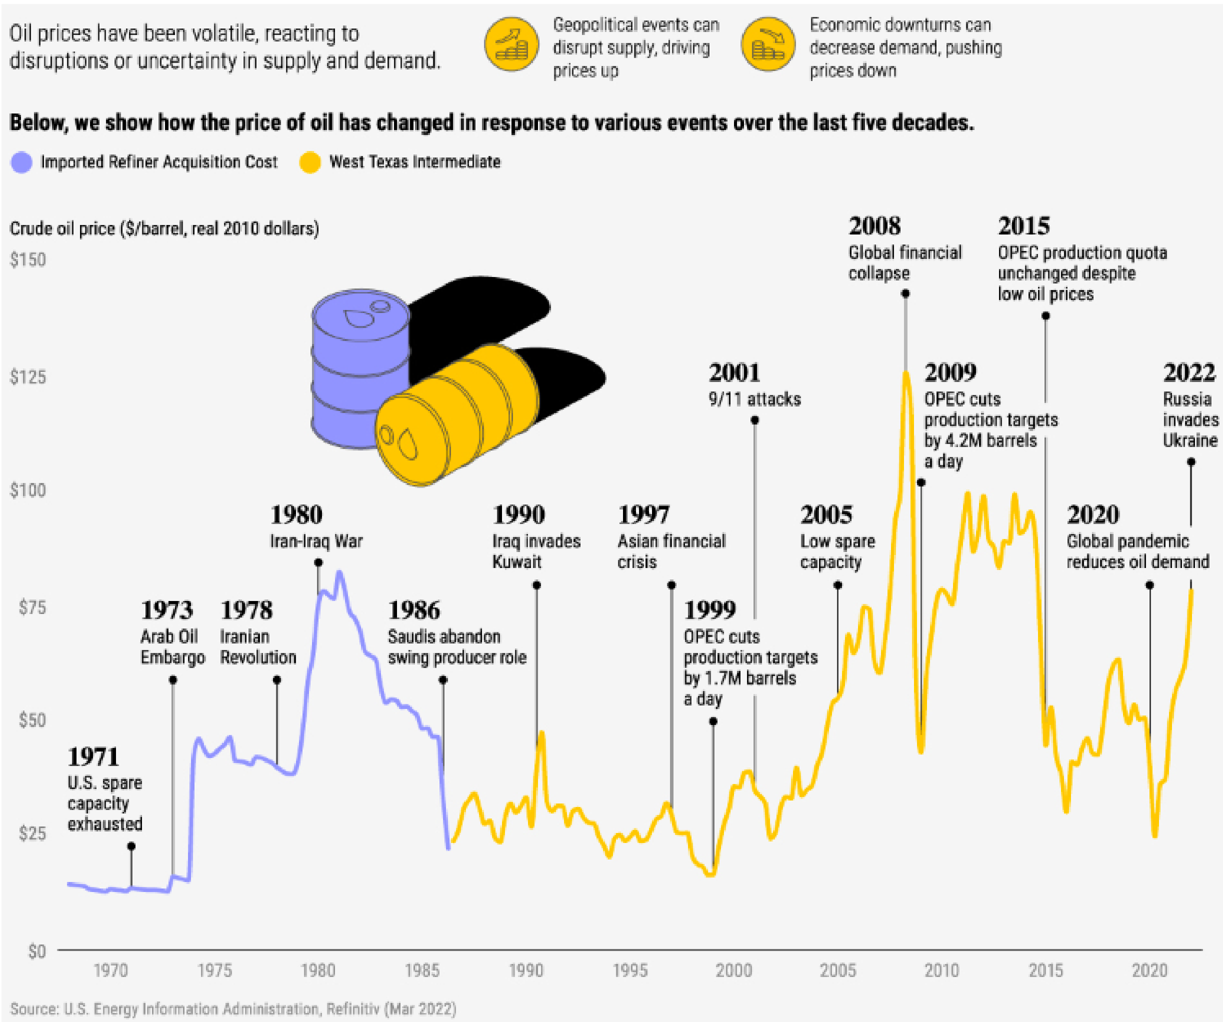

From "How oil prices have reacted to political and economic events," by J. Ross, 2022. *Advisor Channel – Visual Capitalist*. Poster presentation (https://advisor.visualcapitalist.com/historical-oil-prices/).

Another myth is that drilling activities would vanish as renewables take over global energy production. Although in recent years there has been a considerable focus on shifting from fossil fuels to renewable energy sources, as of 2022 oil and gas were responsible for 54% of the global energy supply, and as of 2050 they are predicted to account for 37% (DNV, 2023), hence its contribution will remain relevant over the next 30 years (see Figure 4). As part of decarbonization efforts, oil and gas companies are investing in technologies such as geothermal energy and carbon capture, utilization, and sequestration (CCUS), which have the advantage of utilizing the expertise and equipment already in place for drilling oil and gas wells (Sponsor Company, personal communication, 2022). Therefore, a nuanced and data-driven approach to demand forecasting for drilling activities is essential, as these operations are expected to remain significant in the energy landscape despite the shift toward renewable sources.

Figure 4

World primary energy by source – EJ/year (DNV, 2023)

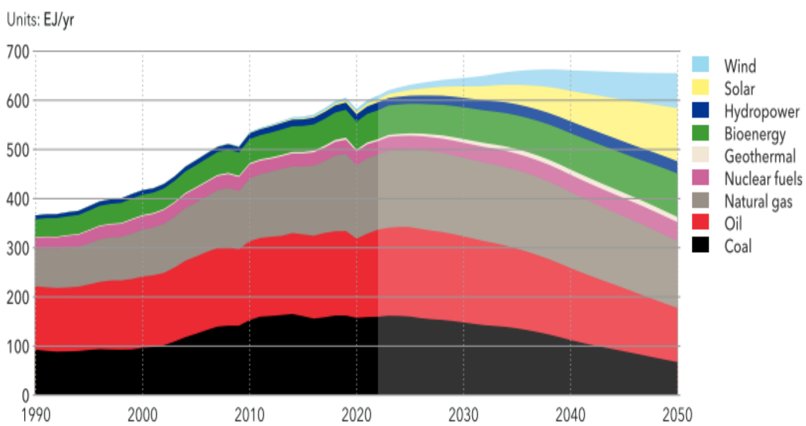

Source	2022	2030	2040	2050
Wind	7	17	42	71
Solar	7	23	62	109
Hydropower	16	19	24	25
Bioenergy	59	65	74	88
Geothermal	4	7	11	11
Nuclear	27	31	35	38
Natural gas	153	160	155	137
Oil	177	175	146	110
Coal	160	148	112	67
Total	**610**	**645**	**661**	**656**

From "Energy Transition Outlook 2023," by DNV, 2023. (https://www.dnv.com/energy-transition-outlook/download.html).

The insights given by literature presented in this chapter were discussed with the company sponsor and helped to frame the project's methodology as follows:

- For the time horizon of this project (next 4 to 6 quarters) assume that the effect of the energy transition will be negligible as oil and gas demand is expected to continue to grow for the next five years before it starts declining and shifts to other renewable sources.
- As oil prices are volatile, the project will start with simpler forecasting models that do not account for oil prices. As the models mature, oil prices will be introduced.

2.2 Demand Forecasting

"Forecasting has fascinated people for thousands of years, sometimes being considered a sign of divine inspiration, and sometimes being seen as a criminal activity." (Hyndman & Athanasopoulos, 2021, Chapter 1, para. 1).

Today, business forecasting does not rely on supernatural powers but can be done through various formal models. These models range from those purely based on human-theoretical knowledge to others that are purely based on artificial intelligence, as presented in Figure 5 (Wilson, 2019).

Figure 5

Forecasting Models

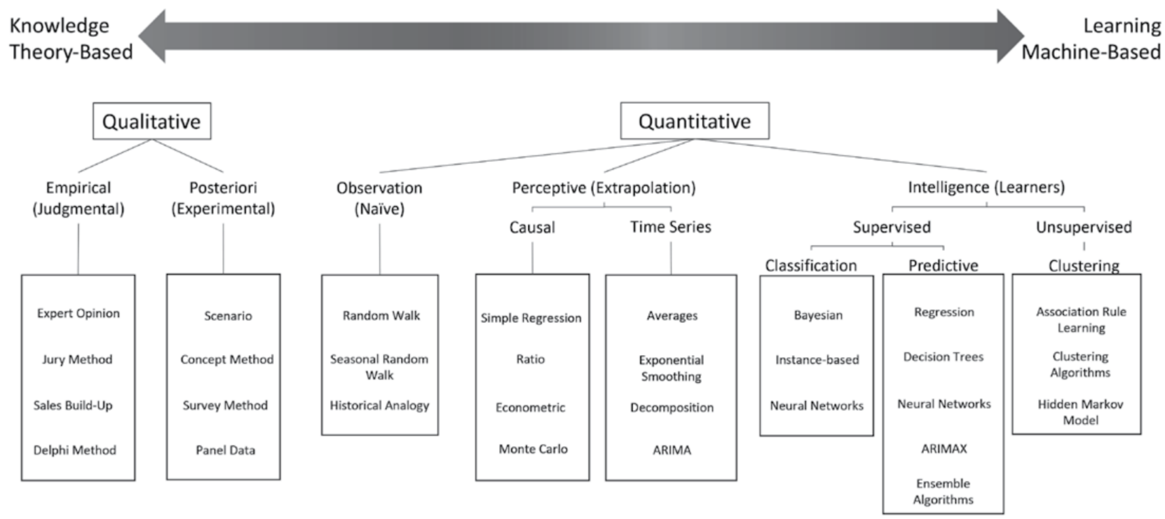

From "The Predictive Analytics Toolbox — Understanding the Tools and Algorithms," by E. Wilson, 2019. *Journal of Business Forecasting, 38*(2), 9–15. (https://www.proquest.com/docview/2273136825/fulltextPDF/E080588C4176400EPQ/1?accountid=12492).

Despite the multitude of forecasting models, there is no one-size-fits-all method, hence the selection of the appropriate model(s) is crucial. In fact:

> The predictability of an event or a quantity depends on several factors including (1) how well we understand the factors that contribute to it, (2) how much data is available, (3) how similar the future is to the past, and (4) whether the forecasts can affect the thing we are trying to forecast (Hyndman & Athanasopoulos, 2021, Section 1.1, para. 2).

Applying these checks to the bits demand forecast resulted in the following insights:

1) *Understanding the contributing factors:* Stakeholders know that bit runs demand is driven by rig count and market share. Some exogenous factors (like oil price) also contribute, but the correlation is not clear.

2) *Data availability:* There are 3–4 years of historical data, which is enough to run some causal and time-series models (e.g., regression, exponential smoothing), but not others (e.g., ARIMA, machine learning).
3) *Similarity between past and future:* The oil and gas industry operates in cyclical patterns every 7–8 years, sometimes disrupted by geopolitical events, so overall the future can be forecasted based on the past.
4) *Are forecasts self-fulfilling?* No, oil price is not directly dictated by the volume of bit runs, but by global supply-demand levels and geopolitical decisions, hence bit runs have negligible effect on the oil price.

These insights, combined with the literature recommendations (Hyndman & Athanasopoulos, 2021; Vandeput, 2021; Wilson, 2019) helped to narrow down the models utilized in this capstone, as summarized in Table 1.

Table 1

Models and applicability to this capstone

Approach	Model Class	Model Type	Model Sub-Type (examples)	Utilized on this project? Why?
Qualitative	Empirical		Expert, Jury, Delphi	No, there is enough quantitative data; qualitative would only be needed in the absence of quantitative data
	Posteriori		Survey, Panel	
Quantitative	Observation		Random Walk (Naïve)	Partially, only for benchmarking accuracy
	Perceptive (Extrapolation)	Causal	Regression, LASSO, Ridge, ElasticNet	Yes, able to account for multiple contributing factors
			Dynamic (Econometric)	No, this is not a closed-loop case (self-fulfilling)
		Time-Series	Averages, Decomposition	No, these models over smooth rapid data rise and falls
			Exponential Smoothing (ETS), Theta, FourTheta, Croston	Yes, able to handle level, trend, and seasonality. Croston specifically applies for intermittent demand.
			ARIMA	No, it needs at least 30 quarters of historical data
	Intelligence (Learners)	Supervised	Neural Networks, Random Forest, Gradient Boosting	No, it requires much more data
		Unsupervised	Clustering	No, it does not apply to this type of forecast

3. Methodology

As an outcome of the literature review, this capstone will focus on generating forecasts via quantitative causal and time-series models. Hyndman & Athanasopoulos (2021) propose a step-by-step methodology for producing these forecasts, which is shown in Figure 6 and further detailed below.

Figure 6

Methodology for producing forecasts

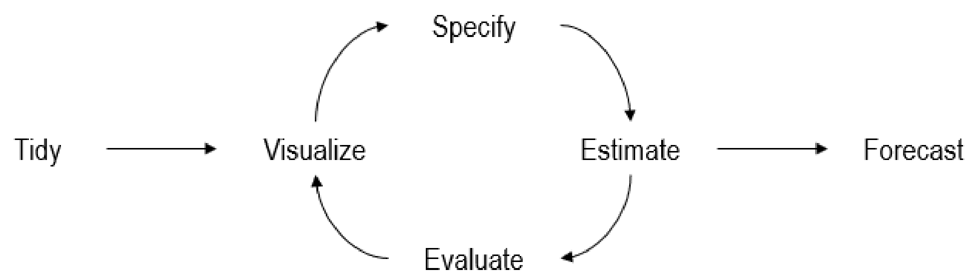

From "Forecasting: principles and practice," by R. J. Hyndman and G. Athanasopoulos, 2021. 3rd edition, OTexts. (https://otexts.com/fpp3/).

The roadmap for executing this methodology is the following:

1) **Data preparation (Tidy):** Here the data is gathered and checked for completeness. It is also the moment when the data is treated so that the merging of different data sets is feasible.
2) **Plot the data (Visualize):** The focus at this stage is to have an initial view of how the data is behaving and based on this behavior define a model to fit the data into.
3) **Define a model (Specify):** In this step, each model is configured with its set of parameters allowing the forecaster to fine-tune the models to produce better forecasts.
4) **Train the model (Estimate):** After setting up the model, it is run with a subset of the data (training set).
5) **Check model performance (Evaluate):** As soon as the model is fitted, its performance can be checked against other models via accuracy measures (e.g., MAE, MAPE, RMSE). Steps 2 to 5 are iterated until the model is accurate enough as agreed with project stakeholders.
6) **Produce forecasts (Forecast):** Once the model is properly set up and sufficiently accurate, forecasts can be generated with this model.

Before going through each step of the methodology in detail, it is worth having a high-level view of the whole process. Understanding the overall methodology before examining the specifics provides clear insights into how each phase of the process informs and shapes subsequent steps. This high-level perspective ensures that detailed analyses are strategically aligned with our broader goals, leading to more effective decision-making throughout the project.

3.1 High-level Overview

The process began by consolidating various pieces of information, which had been pre-processed in Excel, into a unified dataset within Python. The primary objective was to create a comprehensive table that contained both lagging and leading indicators, which could include independent variables such as "rig count" and "market share", as well as dependent variables. These dependent variables were derived by calculating ratios of other variables, like "bit runs per rig count" and "revenue per bit run".

This dataset was then divided into two segments: one for training and the other for testing. This division was crucial for the development of forecasting models. Models based on time series were applied solely to ratio data, while causal models were used to examine the relationships between both lagging and leading indicators and the ratios, thereby generating forecasts.

The forecasting process involved making predictions about future ratios, which were subsequently adjusted by incorporating leading indicators (for example, "future rig count") to compute the sought outcomes (such as "future bit runs").

The logical sequence and relationships among the variables collected and computed are succinctly depicted in Figure 7.

Figure 7

Logical relationship among variables collected and computed

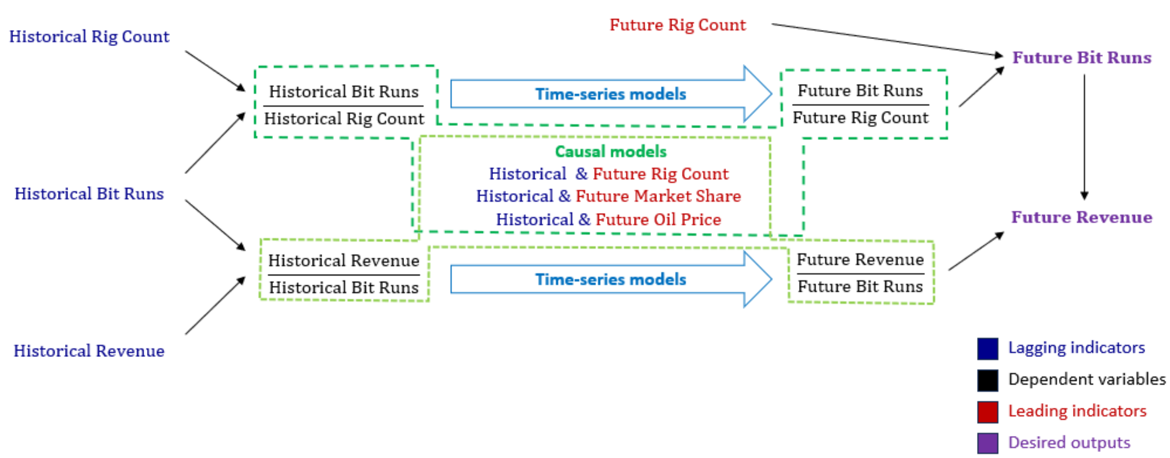

The subsequent sections of this chapter will detail how each step of the methodology was applied in this capstone.

3.2 Data Preparation (Tidy)

This section outlines data cleaning and preparation steps like consolidating datasets, addressing missing values, and ratio calculations to ensure data quality for subsequent analysis and modeling.

Initially, the sponsor provided Excel and PowerBI sources detailing rig count, bit runs, revenue, and market share data, mostly available quarterly from Q1-2019, except for market share data from Q1-2020. The focus was on drilling rig data relevant to forecasting drilling bit demand, excluding workover rigs used for well maintenance.

Challenges included transforming PowerBI pivot tables into a format suitable for Python analysis and reconciling geographic reorganization in the data. Market share data discrepancies across sources were harmonized by prioritizing more accurate data from 2021 while retaining 2020 data for a broader dataset.

Post-cleanup, the data was combined into a single Python dataset. Table 2 details the relevant columns of this dataset. Initial considerations included oil price as an influential variable; however, due to its volatility, it was excluded from this model but is considered for future model enhancements.

Table 2

Relevant columns in a single dataset after Data Preparation (Tidy)

Variable	Model Use	Description
GU-YYYY-QQ	Primary Key	Unique identifier of a pair of geounit and quarter/year, utilized as the primary key for merging datasets.
Rig Count	Independent (Feature)	Sum of drilling rigs actively operating in a geounit in a specific quarter/year.
GEOUNIT	Filtering / Primary Key	Primary key for building dictionaries of datasets (see section 3.4) and for filtering data of a specific geounit.
YYYY-QQ	Independent (Feature)	Time unit for the time series. Although it is originally stored as a text, it is treated to generate timestamps and forecast features (see section 3.4).
Bit Runs	Final desired output	One of the final desired outputs calculated as $\frac{BitRuns}{RigCount} \times Rig\ Count$. It is the sum of bit runs for a geounit in a specific quarter/year.
Revenue	Final desired output	One of the final desired outputs calculated as $\frac{Rev}{BitRun} \times Bit\ Runs$. It is the sum of bits-related revenue for a geounit in a specific quarter/year.
BitRuns/RigCount	Dependent	One of the variables aimed to be predicted. It is the ratio of bit runs divided by rig count for a geounit in a specific quarter/year.
Rev/Bit Run	Dependent	One of the variables aimed to be predicted. It is the ratio of revenue divided by bit runs for a geounit in a specific quarter/year.
All Cpy's Rev	Independent	Sum of bits-related revenue for Sponsor company and all its competitors operating in a geounit in a specific quarter/year.
Market Share	Independent (Feature)	Percentage of revenue captured by the Sponsor Company in a geounit in a specific quarter/year. It is calculated as $\frac{Revenue}{All\ Cpy's\ Rev}$.
OilPrice	Independent	Average Brent price of a barrel of oil in a specific quarter/year.

3.3 Plot the Data (Visualize)

Before building forecasting models in Python, the data was compiled in an Excel spreadsheet, facilitating easier interpretation and communication with stakeholders. The first analyses were done on top of global data, without going into the details of each geounit. One of the first interesting findings was to discover a statistically significant correlation between rig count/market share and bit runs (see Figure 8). This helped to choose *causal models* as one of the categories of models to forecast bit runs and revenue.

Figure 8

Statistically significant correlation between rig count/market share and bit runs (global data)

Regression Statistics	
Multiple R	0.922234227
R Square	85.05%
Adjusted R Squar	80.07%
Standard Error	439.5414505
Observations	9

	Coefficients	*Standard Error*	*t Stat*	*P-value*
Intercept	-13064.95131	3458.189805	-3.777974	0.9%
Rig Count	1.504521606	0.491133149	3.063368	2.2%
Market Share	51241.23683	16615.5714	3.083929	2.2%

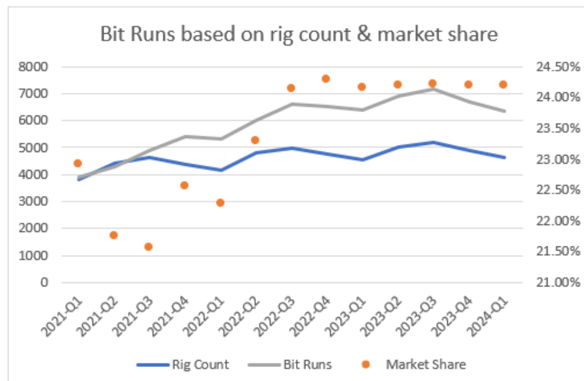

Running exponential smoothing via Excel Forecast Sheet on global data revealed an upward trend and repeating wavy pattern, indicating seasonality. This pattern was especially pronounced in rig count and bit runs, and less so in market share (see Figure 9). These pieces of evidence suggested the usage of *time-series models* in the subsequent steps of this capstone.

Figure 9

Upward trend and seasonal pattern for rig count, bit runs, and market share (global data)

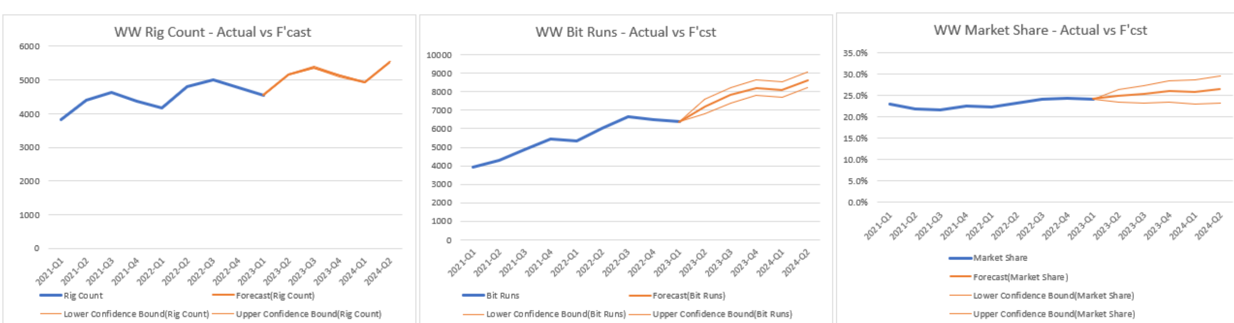

3.4 Define a Model (Specify)

For running predictive models, especially with time-series data, one of the first steps was to preprocess the data. This stage prepared the data by filling in missing values and normalizing numeric values. A pipeline helped automate these steps, ensuring consistency when the model was applied to new data.

Data was split into training and testing sets, with the former used to train the model and the latter to assess its accuracy. Training/testing and accuracy metrics are further explained in Sections 3.5 and 3.6.

The models built were the most applicable to the data available as presented in Table 1. The overall mechanism of how each model treated the data is shown in Table 3.

Table 3

Models utilized in this capstone and their overall functionality

Model Type	Model Sub-Type	Overall Functionality
Time-Series	Exponential Smoothing	It gives more weight to recent observations.
	Theta	It deconstructs the series into trend and residual elements.
	FourTheta	It is an advancement of Theta, catering to complex patterns.
	Croston	It is ideal for intermittent demand, predicting demand level and timing.
Causal	Linear Regression	It finds the best linear fit using all variables.
	Ridge	It reduces overfitting by penalizing large coefficients.
	LASSO	It selects variables, shrinking some coefficients to zero.
	ElasticNet	It combines LASSO's sparsity with Ridge's stability.
Time-Series and Causal	Hybrid	It combines a top-performing time-series model with a causal one.
Random Walk	Naïve	It assumes the future is identical to the past. Utilized for benchmarking.

3.5 Train the Model (Estimate)

Before executing the model, the dataset was divided into two subsets: a training set containing the initial portion of the historical data (80% of the rows) and a test set containing the final portion (20%). At this stage, the historical data was from 2020 to 2022 and the testing data was from 2023. The purpose of this split was to have the models trained against historical data (2020-2022) to generate forecasts for 2023 that could be compared to actual data from 2023 (coming from the test set). As agreed with the Sponsor, this perspective was referred to as a "backward-looking" approach, since it looks at the past to validate if the models are accurate.

For each model, the following steps were executed on an iterative basis (one code run per geounit):

- Capping outliers to ceiling and floor values within a normal distribution (minimum cap of 5% probability and maximum cap of 95%).
- Fitting the model to the data i.e. adjusting the model to best represent the patterns or relationships present in the data.
- Forecasting the values of the data points within the same period of the test set, by utilizing the previously fitted model. In case the forecast yielded negative values these would be replaced by 0.
- Evaluating the model performance, by calculating the corresponding error metrics and comparing the forecasted values against the test values. Further details in Section 3.6.
- Plotting the results, to visualize if the model ran as expected and, if not, take corrective actions accordingly.

Figure 10 shows an example of a plot generated after running the above steps via Exponential Smoothing for one specific geounit and one specific dependent variable (bit runs per rig count).

Figure 10

Training and testing an Exponential Smoothing model for one specific geounit for Bit Runs per Rig Count

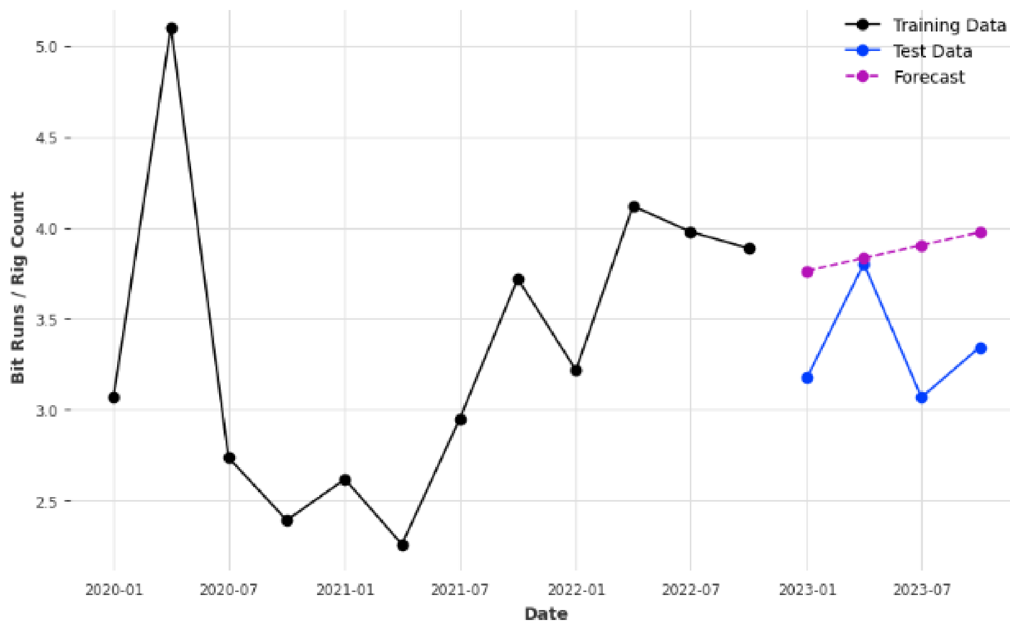

3.6 Check Model Performance (Evaluate)

To assess how accurate each model was, initially the following metrics were calculated:

- Mean Absolute Error (MAE)
- Mean Absolute Percentage Error (MAPE)
- Root Mean Squared Error (RMSE)

The metrics were recorded in a dictionary of data frames ("tables"), each data frame associated with one geounit. After running all the models, each geounit had a data frame with its accuracy metrics. Another way to visualize the same information was to plot the errors in scatter plots, with a red horizontal line drawn against the Naïve model (see example in Figure 11). Naïve being the simplest model, it serves as a baseline for accuracy. Any model with an error lower than Naïve is a good performer, any model with an error higher than Naïve is a bad performer.

After seeing the multitude of graphs generated by the scatter plots, the project team has decided to streamline the analyses by having one single error metric: MAPE. This metric was selected mainly because it is scale-independent, which allows for comparison across different datasets or models, and because it is more intuitive, especially for communicating with non-technical stakeholders, which are the main decision owners at the Sponsor's S&OP-related projects.

For a full comparison between accuracy metrics and why MAPE was selected for this project, please refer to Appendix A.

Figure 11

Error metrics for one specific geounit for Bit Runs per Rig Count

	Model	MAE	MAPE	RMSE
0	ETS	0.521949	0.163784	0.601228
1	Theta	0.344632	0.105537	0.359303
2	FourTheta	0.334257	0.102025	0.347209
3	LinReg	0.584323	0.168770	0.747936
4	LASSO	0.434979	0.135809	0.485580
5	Ridge	0.541505	0.157324	0.646777
6	Elastic Net	0.484375	0.151783	0.556081
7	Hybrid	0.891544	0.255703	1.136463
8	Croston	0.225485	0.065236	0.282504
9	Naive	0.539832	0.169023	0.608387

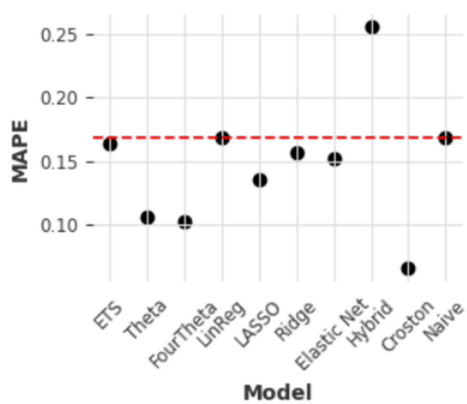

3.7 Produce Forecasts (Forecast)

At this stage, the models were utilized against the full actual data (from 2020 to 2023) to generate the forecasts for the upcoming 5 quarters (from Q1-2024 to Q1-2025). The number of 5 quarters was chosen to ensure that at least one year ahead was being forecasted, but also to align with the quality of the leading indicators being reported by the geounits. For most geounits the visibility of which rigs will operate and when is good enough only for the next 5 quarters, beyond this period the rig count sharply decreases, making the forecasts considerably inaccurate. As agreed with the Sponsor, this perspective was referred to as a "forward-looking approach", since the focus is to prepare for future activity.

The results obtained were plotted for each geounit against each one of the two desired outputs: Bit Runs and Revenue (see example in Figure 12). Further details of overall results and discussions are in Chapter 4.

Figure 12

Actual and forecasted Bit Runs (by model) for one specific geounit

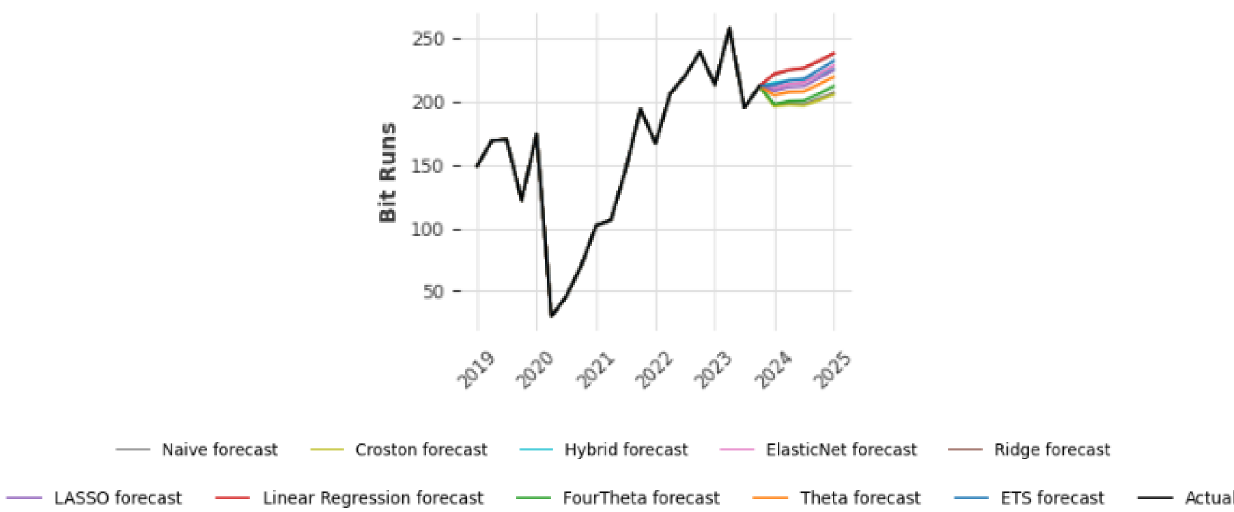

4. Results and Discussion

This chapter focuses on the results of forecasting models to predict the dependent variables, specifically the ratios of Bit Runs per Rig Count and Revenue per Bit Run, alongside the desired outcomes, i.e., Bit Runs and Revenue as introduced in Section 3.1, Figure 7. The results are detailed below, with a particular focus on the quantitative data produced through the application of the methodology outlined in Chapter 3. A more comprehensive exploration of the implications, recurring trends, and potential strategies is presented towards the end of the chapter.

4.1 Results

As explained in Section 3.3, at a global level there is a correlation between rig count, market share, and bit runs. The geounits however have different operating environments so the challenge here was to find which forecasting model(s) would better represent such relationships at a local level. To start, the methodology described in Chapter 3 was applied to train and test models to predict Bit Runs per Rig Count.

The original models applied were time-series (ETS, Theta, FourTheta) and causal (Linear Regression, LASSO, Ridge, ElasticNet). A dashboard with 30 MAPE scatter plots (one per geounit) was presented to the Sponsor and MIT advisor. This dashboard showed some geounits performing better with time-series models and some with causal models. The MIT advisor then recommended creating a hybrid model with the best-performing time-series model and the best-performing causal model. This combination could potentially improve accuracy, but there was no guarantee.

On the other hand, some geounits exhibited very high MAPE values (greater than 100%). After consultation with the Sponsor, it was hypothesized that this inaccuracy stemmed from the business model employed by these geounits. They operate with infrequent bulk sales of drilling bits, creating peaks in Bit Runs per Rig Count that deviate from typical operations, where bits are usually rented based on the corresponding drilling activity. While reviewing this demand pattern with the MIT advisor, the recommendation was to run a Croston model, which is very suitable for intermittent demand, like this case.

Another model that was missing from the original analysis was Naïve, which serves as a benchmark for performance for any other forecasting model. While redesigning the scatter plots, a red line was drawn across the graph at the MAPE level obtained by the Naïve model, making it easier to compare which models were performing well or poorly.

Looking at the new dashboard revealed some interesting patterns in different geounits, which are depicted in Figure 13:

1. For 9 geounits out of 30: Croston outperformed any other forecasting model – Figure 13(a)
2. For 7 geounits out of 30: most forecasting models outperformed Naïve – Figure 13(b)
3. For 3 geounits out of 30: no forecasting model outperformed Naïve – Figure 13(c)
4. For 11 geounits out of 30: some models outperformed Naïve, while others did not – Figure 13(d)

Figure 13

MAPE for Bit Runs per Rig Count (by model)

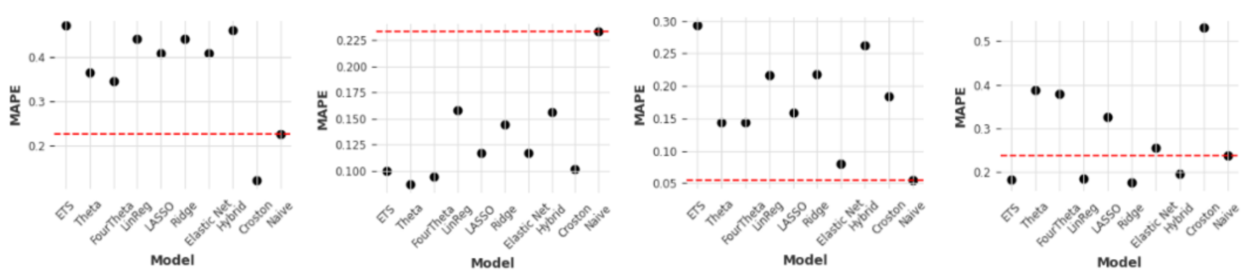

(a) Croston beats all

9/30 geounits

(b) All beat Naïve

7/30 geounits

(c) Naïve beats all

3/30 geounits

(d) Mixed performance

11/30 geounits

A similar dashboard was created for Revenue per Bit Run. Observing MAPE values across geounits, two geounits stood out with MAPE values around 10^{16} whereas all other geounits had a MAPE between 0 and 3. This happened because the bit runs reached a level of zero in some specific quarters. This yielded an infinite number while calculating the ratio of revenue per bit run. To correct this, the models were updated to replace infinite ratios with the average of the other quarters.

While plotting MAPE scatter plots for Revenue per Bit Run one of the main changes was the increase in the number of geounits in which Naïve outperforms all the other models, as shown in Figure 14(c), compared to Figure 13(c). This could be interpreted as the models being underfitted and perhaps the addition of some exogenous variables could improve the forecast accuracy of Revenue per Bit Run.

Figure 14

MAPE for Revenue per Bit Run (by model)

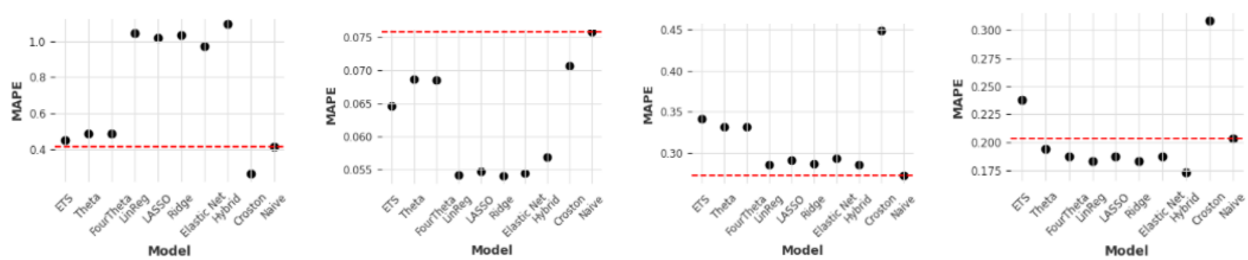

(a) Croston beats all

9/30 geounits

(b) All beat Naïve

4/30 geounits

(c) Naïve beats all

11/30 geounits

(d) Mixed performance

6/30 geounits

To enhance the visualization of model performance globally, Excel pivot tables with conditional formatting were utilized. Two pivot tables were developed: one for Bit Runs per Rig Count and another for Revenue per Bit Run, detailed in Appendix B. These tables were organized with rows in descending order based on the average MAPE for each geounit and columns in ascending order from the model with the lowest to the highest average MAPE. Conditional formatting was used to highlight data patterns more effectively.

Below are the main findings from these pivot tables:

- Bit Runs per Count:
  - Globally, only three time-series models outperform Naïve: Croston, Theta, and FourTheta.
  - Few geounits have very high MAPE across most models, with Croston outperforming all.
  - Theta and FourTheta have very similar global performance.
  - Theta and Croston seem to complement each other, alternating the geounits in which each one performs the best.

- Revenue per Bit Run:
  - Globally, Naïve outperforms all other models, followed by Croston, Theta, and FourTheta
  - The dynamics of Revenue per Bit Run do not seem to be effectively captured by either time-series or causal models. This issue could be improved with additional exogenous factors (e.g., oil price, consumer price index) or feature engineering (e.g., lagging).

4.2 Discussion

From the above findings, one question stood out: Why was Croston superior to the models for around 30% of the geounits? The Sponsor had the answer: This was most likely linked to the business model of these geounits, which had intermittent bulk sales of drilling bits. These sales do not happen very often (e.g., once a year, every two years) but every time they occur this results in a spike of bit runs and revenue. Since Croston is designed for intermittent demand, it was the best model to capture this pattern.

Still, the majority of the geounits operate under a rental business model, with the bit runs and revenue closely related to the activity drivers (e.g., rig count, market share). For these geounits, Theta had a better global performance.

As presented in Section 1.1, this project aims to integrate a forecasting model into IBP, which typically prefers a single model for simplicity. However, given the distinct needs of different geounits, the Sponsor has agreed to consider using more than one model if justified. Looking at the findings described in previous paragraphs, the Sponsor agreed to have some flexibility on the number of models, but it will be at the Sponsor's discretion to decide how many models will provide input to IBP. This capstone recommends the usage of Theta, Croston, and/or Naïve, as detailed in Table 4.

Table 4

Key business models for drilling bits and corresponding best-performing forecasting models

				Best Performing Model	
Business Model	**% of geounits**	**Pros**	**Cons**	**For Bit Runs**	**For Revenue**
Rental Model	70%	• Steady revenue, easier demand forecasting. • Better customer relationship opportunities.	• Higher operational complexity for bit maintenance and logistics. • Revenue influenced by long-term market trends and external factors.	Theta	Naïve
Bulk Sales Model	30%	• Simple sales, larger revenue per transaction. • Immediate market, less customer engagement post-sale.	• Hard to predict demand, leading to inventory issues. • Sensitive to market fluctuations and external factors like oil prices.	Croston	Naïve

To further assist the Sponsor in scrutinizing which parameters to consider for utilizing one model or another, a decision tree was done with the support of the machine learning software Orange. The software was inputted a small table with 30 rows (one per geounit) and columns shown in Table 5. Rig Count, Bit Runs, and Revenue had the latest quarterly actual number. All the other numeric values applied over the training period (i.e., from 2020 to 2022). The 10 models were grouped into four model types: (1) *Causal*: Linear Regression, LASSO, Ridge, ElasticNet, Hybrid; (2) *Intermittent*: Croston[1]; (3) *Random Walk*: Naïve; and, (4) *Time-Series*: Exponential Smoothing (ETS), Theta, FourTheta.

Table 5

Relevant columns of table inputted in machine learning software Orange

Variable	Model Use	Description
Rig Count	Numeric (Feature)	Sum of drilling rigs actively operating in a geounit in a specific quarter/year.
GEOUNIT	Text (Metadata)	Three-letter code to designate the corresponding geounit.
Bit Runs_Actual	Numeric (Feature)	Sum of bit runs for a geounit in a specific quarter/year.
Revenue_Actual	Numeric (Feature)	Sum of bits-related revenue for a geounit in a specific quarter/year.
BR/RC_Average	Numeric (Feature)	Average of quarterly ratios of Bit Run/Rig Count over the training period.
BR/RC_StdDev	Numeric (Feature)	Standard deviation of quarterly ratios of Bit Run/Rig Count over the training period.
BR/RC_CoeffVar	Numeric (Feature)	Coefficient of variation of Bit Run/Rig Count. Calculated as $\frac{BR/RC\_Average}{BR/RC\_StdDev}$. Reflects how volatile the Bit Run per Rig Count is over the training period.
Rev/BR_Average	Numeric (Feature)	Average of quarterly ratios of Revenue/Bit Run over the training period.
Rev/BR_StdDev	Numeric (Feature)	Standard deviation of quarterly ratios of Revenue/Bit Run over the training period.
Rev/BR_CoeffVar	Numeric (Feature)	Coefficient of variation of Revenue/Bit Run. Calculated as $\frac{Rev/BR\_Average}{Rev/BR\_StdDev}$. Reflects how volatile Revenue per Bit Run is over the training period.
MktShare_Average	Numeric (Feature)	Average of quarterly market shares over the training period.
MktShare_StdDev	Numeric (Feature)	Standard deviation of quarterly market shares over the training period.
MktShare_CoeffVar	Numeric (Feature)	Coefficient of variation of Market Share. Calculated as $\frac{MktShare\_Average}{MktShare\_StdDev}$. Reflects how volatile Market Share is over the training period.
BR/RC_ModelType_ lowestMAPE	Categorical (Target)	Type of model that has the lowest MAPE for the corresponding geounit over the testing period: Causal, Intermittent, Random Walk, or Time-Series
BR/RC_CountOfFcst ModTypesInTop3	Numeric (Feature)	Count of different model types within the top 3 performing forecasting models for Bit Run/Rig Count for a specific geounit over the testing period. For example, if the 3 top performing were 2 causal + 1 time-series, the count of types will be 2.

The machine learning software managed to fit the data in a decision tree with parameters that helped to segregate the best-performing models for 60% of the geounits (18 out of 30), as depicted in Figure 15. Some variables showed good results segregating models' performance, for instance:

[1] Although Croston is a time-series, given its outperformance it is being tracked separately here as Intermittent.

- If Top3 model types for a geounit are very diverse (*BR/RC_CountOfFcst ModTypesInTop3* = 3), then the Intermittent model is very likely to be the best.
- Volatility helps to differentiate whether time-series or causal models best fit a geounit's Bit Runs per Rig Count (e.g., *BR/RC_CoeffVar* and *MktShare_StdDev* in Figure 15).

Figure 15

A decision tree with parameters to segregate the best-performing models by geounit for Bit Runs/Rig Count

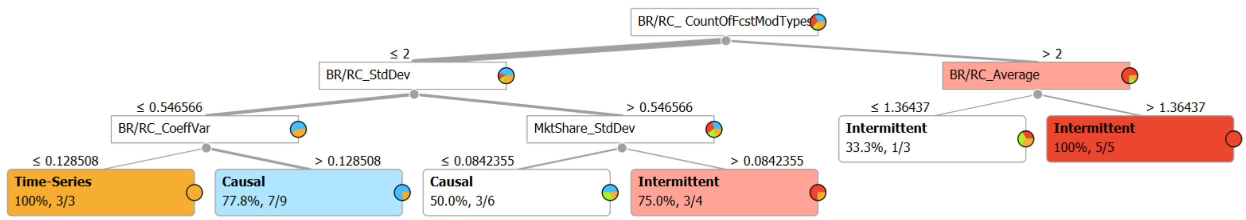

Another approach to delimiting parameters was done via scatter plots in Orange. The scatter plots shown in Figures 16(a) and 16(b) were some of the most insightful ones:

- Most of the geounits that best fit against the Intermittent model seem to have a high market share (greater than 40%) and intermediate level of bit runs per rig count (around 2 to 4). These are the geounits that operate under a bulk sales model.
- Geounits with a very low coefficient of variation (less than 0.2) or very high (greater than 0.8) seem to have a better fit against time-series models, which quickly adjust for fluctuations.

Figure 16

Identifying parameters to cluster the best-performing models by geounit for Bit Runs/Rig Count

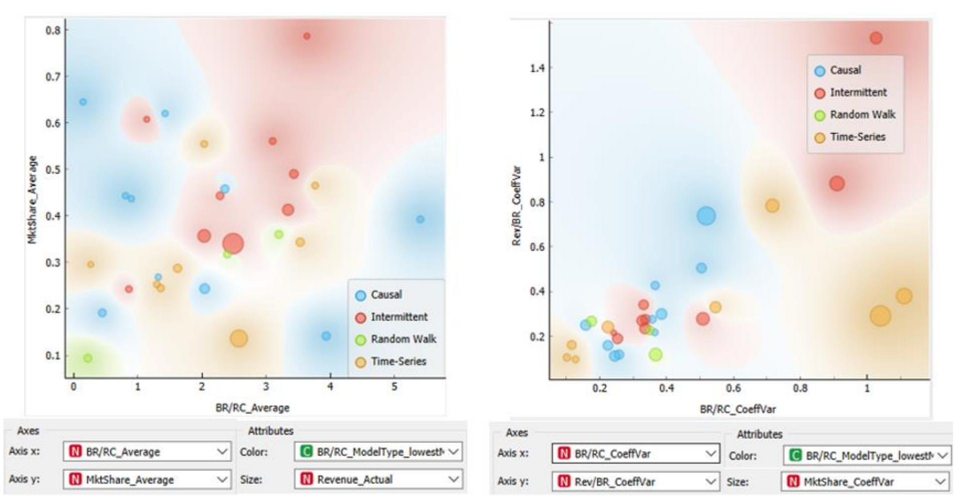

(a) BitRuns/RigCount_avg vs MarketShare_avg (b) BitRuns/RigCount_CV vs RevPerBR_CV

After presenting these results/analyses to the Sponsor, the stakeholders requested visual comparisons of top-performing models against actual and manually forecasted values by geounits as of the quarter preceding the testing period (i.e., Q4-2022, since the testing period was the four quarters of 2023). This aimed to validate whether the models outperformed manual forecasts.

For confidentiality purposes, the graphs and numbers on global bit runs and revenue will not be displayed in this report. However, the MAPEs will be shown, as depicted in Figure 17.

The main outcomes were the following:

- **Globally**: Across the year, the top three models (Croston, Theta, Naïve) consistently outperformed geounit forecasts, reducing yearly MAPE from 6–8% to 2–3%. However, MAPE values indicated that manual forecasts were occasionally more accurate within the near-term horizon of 1-2 quarters.
- **Locally**: The top models generally outperformed manual forecasts in 30% of the geounits, matched them in 50%, and underperformed in 20%. For revenue, these models matched the manual forecasts in 40% of the cases and underperformed in 30%

Overall, automated forecasts proved more advantageous globally than locally, particularly for predicting Bit Runs compared to Revenue.

Figure 17

Global and local performance of Top3 automated forecasting models against geounits manual forecasting

(a) Globally, Top 3 beat GU

Bit Runs

Geounit (All)	AbsPercError_BitRuns_Croston	AbsPercError_BitRuns_Theta	AbsPercError_BitRuns_Naive	AbsPercError_BitRuns_GU-Fcst
2023-Q1	0.3%	2.8%	2.0%	4.9%
2023-Q2	1.9%	2.6%	0.6%	0.1%
2023-Q3	1.7%	1.0%	0.6%	7.2%
2023-Q4	5.8%	8.0%	8.5%	12.7%
MAPE	2.4%	3.6%	2.9%	6.2%

Revenue

Geounit (All)	AbsPercError_Revenue_Croston	AbsPercError_Revenue_Theta	AbsPercError_Revenue_Naive	AbsPercError_Revenue_GU-Fcst
2023-Q1	10.6%	3.1%	1.6%	0.6%
2023-Q2	10.3%	1.8%	2.3%	5.0%
2023-Q3	10.5%	2.5%	1.4%	11.4%
2023-Q4	7.7%	1.7%	5.0%	14.8%
MAPE	9.8%	2.3%	2.6%	8.0%

(b) GU f’cst beats all Top3

6/30 geounits

Geounit GU01	AbsPercError_BitRuns_Croston	AbsPercError_BitRuns_Theta	AbsPercError_BitRuns_Naive	AbsPercError_BitRuns_GU-Fcst
2023-Q1	21.3%	6.5%	16.8%	5.6%
2023-Q2	10.7%	9.3%	6.7%	1.6%
2023-Q3	20.4%	1.9%	16.0%	4.2%
2023-Q4	17.6%	3.6%	13.3%	4.0%
MAPE	17.5%	5.3%	13.2%	3.9%

9/30 geounits

Geounit GU04	AbsPercError_Revenue_Croston	AbsPercError_Revenue_Theta	AbsPercError_Revenue_Naive	AbsPercError_Revenue_GU-Fcst
2023-Q1	10.0%	0.8%	4.3%	28.7%
2023-Q2	14.6%	19.3%	21.8%	7.1%
2023-Q3	30.3%	31.9%	25.9%	30.7%
2023-Q4	97.5%	80.0%	110.0%	61.6%
MAPE	38.1%	33.0%	40.5%	32.0%

(c) GU beats some Top3

15/30 geounits

Geounit GU02	AbsPercError_BitRuns_Croston	AbsPercError_BitRuns_Theta	AbsPercError_BitRuns_Naive	AbsPercError_BitRuns_GU-Fcst
2023-Q1	16.8%	18.4%	3.1%	6.6%
2023-Q2	16.1%	3.1%	3.6%	0.0%
2023-Q3	11.2%	13.8%	7.8%	3.9%
2023-Q4	29.3%	22.3%	7.3%	34.6%
MAPE	18.3%	14.4%	5.4%	11.3%

12/30 geounits

Geounit GU05	AbsPercError_Revenue_Croston	AbsPercError_Revenue_Theta	AbsPercError_Revenue_Naive	AbsPercError_Revenue_GU-Fcst
2023-Q1	17.4%	24.2%	26.6%	1.8%
2023-Q2	30.5%	5.3%	6.4%	17.9%
2023-Q3	47.8%	29.0%	20.0%	35.1%
2023-Q4	48.1%	36.1%	20.5%	36.5%
MAPE	36.0%	23.7%	18.4%	22.8%

(d) All Top3 beat GU

9/30 geounits

Geounit GU03	AbsPercError_BitRuns_Croston	AbsPercError_BitRuns_Theta	AbsPercError_BitRuns_Naive	AbsPercError_BitRuns_GU-Fcst
2023-Q1	6.5%	7.4%	15.1%	43.7%
2023-Q2	68.5%	57.7%	82.1%	68.9%
2023-Q3	38.7%	47.1%	33.7%	88.0%
2023-Q4	40.9%	11.4%	52.3%	72.9%
MAPE	38.7%	30.9%	45.8%	68.4%

9/30 geounits

Geounit GU06	AbsPercError_Revenue_Croston	AbsPercError_Revenue_Theta	AbsPercError_Revenue_Naive	AbsPercError_Revenue_GU-Fcst
2023-Q1	11.7%	0.4%	9.5%	11.4%
2023-Q2	10.7%	10.7%	12.4%	1.9%
2023-Q3	13.4%	3.1%	11.2%	32.9%
2023-Q4	18.1%	4.8%	15.9%	45.7%
MAPE	13.5%	4.7%	12.3%	23.0%

5. Conclusion

Throughout this project, forecasting models were successfully developed and tested, not only predicting future bit runs but also estimating the associated revenue with greater accuracy than the current manual forecasting done by 30 geounits. This accomplishment facilitates the smooth incorporation of our findings into the company's Integrated Business Planning system, thereby marking a notable improvement in the Sales and Operations Planning (S&OP) practices of the sponsor company.

A key insight from this study was that the best-performing forecasting models correlated with the corresponding geographic unit's business model for drilling bits (i.e. rental and bulk sales). However, it is noted that the performance of these models varied across different geographic units, highlighting the need for tailored strategies. Interestingly, simpler time-series models often outperformed more complex causal models in terms of accuracy, suggesting that complexity does not always yield better forecasting results. Overall, automated forecasts proved more advantageous globally than locally, particularly for predicting Bit Runs compared to Revenue.

Moving forward, the project recommends several next steps to enhance the forecasting models. It is advised to incorporate additional external factors such as oil price and consumer price index into the models to improve their robustness. Furthermore, the implementation of feature engineering techniques, such as the use of lagging indicators, is proposed to increase predictive accuracy by more accurately reflecting historical trends and patterns.

The strategic implications of improved forecasting accuracy are significant. Enhanced accuracy is expected to lead to superior inventory management, reduced operational costs, optimized resource allocation, and better overall operational efficiency and customer satisfaction. Moreover, having accurate demand forecasts allows the company to quickly adapt to market fluctuations, thus maintaining a competitive edge.

This project establishes a robust groundwork for advancing demand forecasting within the energy industry. It provides not only a practical tool for immediate implementation but also a framework for continuous improvement. The future steps include scaling these solutions across various business lines, further refining the models, and broadening the scope to encompass predictive analytics for other vital components of the supply chain. Incorporating this proactive approach into the company's S&OP practices marks a major advancement in utilizing technological innovations to sustain a competitive advantage in the energy sector.

References

Andrés, D. (2023). Error Metrics for Time Series Forecasting. *Machine Learning Pills*. https://mlpills.dev/time-series/error-metrics-for-time-series-forecasting/

Chima, C.M. (2007). Supply-Chain Management Issues In The Oil And Gas Industry. *Journal of Business & Economics Research*, *5*(6), 27–36. https://doi.org/10.19030/jber.v5i6.2552

DNV (2023). *Energy Transition Outlook 2023*. https://www.dnv.com/energy-transition-outlook/download.html

Hussain, R., Assavapokee, T., & Khumawala, B. (2006). Supply Chain Management in the Petroleum Industry: Challenges and Opportunities. *International Journal of Global Logistics & Supply Chain Management*, *1*(2), 90–97. https://www.researchgate.net/publication/261912419_Supply_Chain_Management_in_the_Petroleum_Industry_Challenges_and_Opportunities

Hyndman, R.J., & Athanasopoulos, G. (2021). *Forecasting: principles and practice* (3rd ed.). OTexts. https://otexts.com/fpp3/

Hyndman, R.J., & Koehler, A.B. (2006). Another Look at Measures of Forecast Accuracy. *International Journal of Forecasting*, *22*(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001

Lisitsa, S., Levina, A., & Lepekhin, A. (2019). Supply-chain management in the oil industry. *E3S Web of Conferences*, *110*, Article 02061. https://doi.org/10.1051/e3sconf/201911002061

Rink, K. (2021). Time Series Forecast Error Metrics You Should Know. *Towards Data Science*. https://towardsdatascience.com/time-series-forecast-error-metrics-you-should-know-cc88b8c67f27

Ross, J. (2022). How oil prices have reacted to political and economic events. [Poster presentation]. *Advisor Channel – Visual Capitalist*. https://advisor.visualcapitalist.com/historical-oil-prices/

Sponsor Company (2022). 2022 Sustainability Report. Corporate Report.

Sponsor Company (2023a). Bits fleet management [PowerPoint slides].

Sponsor Company (2023b). Bits S&OP Monthly Forecast [Data set].

Vandeput, N. (2021). *Data Science for Supply Chain Forecasting* (2nd ed.). De Gruyter.

Wilson, E. (2019). The Predictive Analytics Tool Box — Understanding the Tools and Algorithms. *Journal of Business Forecasting*, *38*(2), 9–15. https://www.proquest.com/docview/2273136825/fulltextPDF/E080588C4176400EPQ/1?accountid=12492

Zhu, T., Balakrishnan, J., & da Silveira, G.J.C. (2020). Bullwhip Effect in the Oil and Gas Supply Chain: A Multiple-Case Study. *International Journal of Production Economics*, *224*, 107548. https://doi.org/10.1016/j.ijpe.2019.107548

Appendix A

Table A-1 outlines the pros and cons of major error metrics used in business and academia. After discussions with the sponsor, MAPE was chosen as the project's metric because it facilitates comparison across geounits and models and is intuitive, making it easy to communicate to stakeholders.

Table A-1

Comparison of Error Metrics

Metric	Pros	Cons
Mean Absolute Error (MAE)	• Easy to understand and interpret. • Gives an average of the absolute errors, providing a clear picture of forecast accuracy. • Less sensitive to outliers compared to RMSE.	• Does not indicate the direction of the errors (over or under-forecasting). • Treats all errors equally, which may not be ideal in all contexts.
Mean Squared Error (MSE)	• Punishes larger errors more severely than smaller ones, which can be useful when large errors are particularly undesirable. • Commonly used, making it easier to compare results across different studies or models.	• The scale of the errors can be difficult to interpret due to squaring. • More sensitive to outliers compared to MAE.
Root Mean Squared Error (RMSE)	• Provides errors in the same units as the forecasted values, making interpretation more straightforward than MSE. • Heavily penalizes larger errors, which might be beneficial in some contexts.	• Like MSE, it is sensitive to outliers. • Can be more challenging to interpret compared to MAE since it is not a direct average of the errors.
Mean Absolute Percentage Error (MAPE)	• Expresses error as a percentage, which can be more intuitive, especially for communicating with non-technical stakeholders. • Scale-independent, which allows for comparison across different datasets or models.	• Can be misleading if there are zero or near-zero actual values in the data. • Asymmetric, since overestimates and underestimates are not treated equally.
Symmetric Mean Absolute Percentage Error (sMAPE)	• Addresses some of the asymmetry issues in MAPE, providing a more balanced view of overestimates and underestimates. • Still provides error in percentage terms, facilitating interpretability.	• Can still be problematic when actual values are close to zero. • Allows error to be negative. • Not easily interpretable, which might complicate comparisons with existing literature or results.

Note: Table prepared by capstone author based on “Time Series Forecast Error Metrics You Should Know,” by K. Rink, 2021. *Towards Data Science*. (https://towardsdatascience.com/time-series-forecast-error-metrics-you-should-know-cc88b8c67f27); “Error Metrics for Time Series Forecasting,” by D. Andrés, 2023. *Machine Learning Pills*. (https://mlpills.dev/time-series/error-metrics-for-time-series-forecasting/); and, “Another Look at Measures of Forecast Accuracy”, by R.J. Hyndman and A.B. Koehler, 2006. *International Journal of Forecasting*, *22*(4), 679–688 (https://doi.org/10.1016/j.ijforecast.2006.03.001).

Appendix B

Figure B-1

MAPE Pivot Table for Bit Runs per Rig Count [2]

Croston	Theta	FourTheta	Naive	Elastic Net	LASSO	ETS	Ridge	LinReg	Hybrid	Grand Total
30.6%	18.8%	113.4%	27.9%	167.4%	162.1%	173.9%	180.8%	181.2%	245.9%	130.2%
48.9%	133.2%	112.9%	253.9%	126.0%	121.3%	146.8%	120.1%	119.8%	115.9%	129.9%
54.5%	100.0%	85.4%	100.0%	157.4%	152.3%	100.0%	169.4%	170.1%	170.4%	125.9%
39.7%	91.4%	89.4%	102.8%	78.2%	72.2%	154.2%	86.9%	86.6%	90.7%	89.2%
38.7%	30.9%	30.9%	45.8%	122.4%	119.8%	90.8%	131.8%	132.6%	54.2%	79.8%
40.4%	31.0%	30.9%	47.5%	67.1%	58.2%	38.2%	82.8%	83.8%	105.8%	58.6%
47.0%	92.1%	92.1%	44.0%	43.6%	47.2%	58.1%	36.6%	36.3%	49.7%	54.7%
25.5%	55.5%	54.7%	62.5%	57.2%	54.6%	47.8%	59.7%	59.7%	60.3%	53.7%
73.0%	27.0%	30.4%	28.1%	30.9%	45.2%	36.7%	77.6%	79.0%	86.6%	51.5%
25.8%	76.4%	62.2%	24.9%	21.6%	21.7%	27.6%	62.5%	65.9%	70.5%	45.9%
45.9%	46.1%	45.0%	17.9%	40.0%	39.9%	47.6%	30.3%	29.9%	29.9%	37.3%
11.9%	36.5%	34.5%	22.5%	40.9%	40.9%	47.1%	44.2%	44.2%	46.1%	36.9%
36.9%	28.2%	31.9%	25.8%	22.4%	33.6%	65.1%	28.7%	26.3%	45.5%	34.4%
15.0%	34.7%	25.1%	23.7%	31.0%	31.6%	34.1%	31.3%	31.4%	30.2%	28.8%
53.0%	38.8%	37.8%	23.9%	25.6%	32.8%	18.2%	17.6%	18.5%	19.6%	28.6%
14.8%	17.2%	16.4%	65.5%	14.5%	12.7%	14.2%	36.6%	37.2%	48.5%	27.8%
21.5%	25.6%	15.1%	30.0%	15.0%	15.7%	30.2%	23.2%	25.2%	22.7%	22.4%
8.1%	13.2%	13.4%	10.8%	26.7%	26.8%	18.0%	26.2%	26.0%	20.2%	18.9%
24.7%	29.5%	29.5%	28.3%	18.6%	18.6%	26.9%	3.1%	3.5%	4.1%	18.7%
20.3%	16.8%	17.9%	20.1%	19.9%	16.8%	19.9%	18.4%	18.5%	17.8%	18.6%
18.3%	14.4%	14.4%	5.4%	8.0%	15.9%	29.3%	21.7%	21.6%	26.3%	17.5%
44.5%	17.8%	25.3%	4.0%	15.0%	17.5%	15.6%	11.6%	11.5%	12.0%	17.5%
20.3%	5.6%	8.1%	13.6%	15.6%	17.5%	14.3%	17.5%	19.6%	33.9%	16.6%
6.5%	10.6%	10.2%	16.9%	15.2%	13.6%	16.4%	15.7%	16.9%	25.6%	14.7%
5.6%	8.7%	8.8%	19.3%	4.0%	5.0%	5.3%	26.6%	27.6%	26.5%	13.7%
10.2%	8.7%	9.4%	23.3%	11.7%	11.7%	10.0%	14.4%	15.8%	15.6%	13.1%
5.2%	2.9%	2.9%	4.3%	5.5%	2.8%	2.6%	24.9%	26.1%	21.1%	9.8%
31.6%	6.9%	6.7%	6.9%	7.1%	6.9%	7.2%	6.5%	6.5%	6.5%	9.3%
17.5%	5.3%	7.9%	13.2%	3.2%	3.9%	5.4%	6.7%	6.9%	11.6%	8.2%
28.8%	35.3%	36.6%	38.4%	41.8%	42.0%	44.9%	48.7%	49.2%	52.2%	41.8%

Figure B-2

MAPE Pivot Table for Revenue per Bit Run [2]

Naive	Croston	FourTheta	Theta	Elastic Net	ETS	LASSO	Ridge	LinReg	Hybrid	Grand Total
77.2%	65.9%	94.9%	94.9%	147.6%	128.5%	163.6%	161.8%	166.1%	166.4%	126.7%
60.7%	187.2%	68.9%	66.2%	105.2%	97.2%	141.9%	135.8%	144.1%	132.7%	114.0%
68.1%	29.2%	112.4%	113.1%	60.5%	124.2%	54.2%	54.6%	53.4%	60.6%	73.0%
41.4%	26.1%	48.4%	48.3%	97.3%	44.9%	102.0%	103.2%	104.3%	109.4%	72.5%
129.6%	32.2%	47.9%	52.9%	51.1%	60.9%	50.5%	50.5%	50.4%	50.0%	57.6%
36.2%	24.6%	43.1%	43.2%	49.0%	50.2%	49.2%	48.7%	48.7%	51.4%	44.4%
37.4%	35.9%	40.1%	40.0%	44.8%	34.5%	47.5%	47.2%	47.8%	55.5%	43.1%
30.1%	33.9%	43.8%	52.5%	42.6%	44.5%	43.0%	41.3%	41.0%	39.9%	41.3%
24.9%	25.3%	30.8%	30.5%	46.9%	49.8%	46.6%	46.6%	46.5%	35.5%	38.4%
22.0%	18.5%	32.5%	32.1%	33.0%	20.1%	41.5%	41.3%	43.4%	46.3%	33.1%
27.2%	44.9%	33.2%	33.2%	29.4%	34.1%	29.1%	28.6%	28.5%	28.5%	31.7%
12.9%	83.6%	13.6%	10.6%	12.8%	35.0%	22.1%	28.1%	32.6%	43.5%	29.5%
25.0%	25.7%	26.6%	26.6%	31.3%	34.3%	31.5%	30.8%	30.8%	30.4%	29.3%
34.4%	14.0%	23.6%	23.7%	32.5%	25.7%	33.2%	33.6%	33.8%	30.5%	28.5%
59.3%	16.7%	18.9%	19.3%	20.4%	20.6%	22.2%	24.5%	25.2%	29.7%	25.7%
31.9%	10.0%	37.6%	37.6%	16.8%	41.5%	13.1%	16.3%	15.9%	24.4%	24.5%
19.0%	35.1%	27.6%	23.5%	19.0%	24.6%	20.9%	22.0%	22.9%	22.7%	23.7%
10.1%	7.5%	11.2%	11.2%	37.0%	10.7%	37.4%	37.3%	37.4%	20.2%	22.0%
20.4%	30.8%	18.8%	19.4%	18.7%	23.8%	18.8%	18.3%	18.3%	17.3%	20.5%
22.3%	17.4%	16.1%	16.1%	18.8%	19.3%	18.7%	17.2%	17.0%	27.7%	19.1%
17.8%	28.7%	17.8%	16.2%	18.3%	18.0%	18.3%	18.0%	17.9%	15.5%	18.6%
11.9%	12.7%	12.8%	12.5%	17.3%	12.8%	19.4%	23.5%	25.0%	25.3%	17.3%
9.6%	7.3%	17.2%	17.2%	16.1%	22.8%	15.3%	14.9%	14.6%	22.5%	15.8%
14.1%	26.5%	24.9%	22.7%	10.5%	16.1%	9.9%	9.8%	9.6%	8.6%	15.3%
12.6%	28.4%	14.3%	14.3%	13.3%	13.3%	12.7%	12.4%	12.2%	12.3%	14.6%
9.5%	10.3%	16.0%	16.2%	14.2%	15.1%	14.6%	13.3%	13.3%	14.8%	13.7%
13.6%	16.2%	14.0%	14.0%	6.4%	9.0%	6.5%	7.8%	8.0%	10.6%	10.6%
2.8%	13.5%	2.8%	2.8%	13.7%	7.8%	12.0%	19.0%	19.6%	6.9%	10.1%
7.6%	8.4%	7.7%	7.7%	6.5%	7.4%	6.9%	6.2%	6.4%	6.2%	7.1%
7.6%	7.1%	6.9%	6.9%	5.4%	6.5%	5.5%	5.4%	5.4%	5.7%	6.2%
29.9%	30.8%	30.8%	30.8%	34.6%	35.1%	36.9%	37.3%	38.0%	38.4%	34.3%

[2] Note: For confidentiality purposes, the three-letter codes of the geounits are not being displayed on this report. Each row on each pivot table corresponds to a specific geounit.