UNIVERSITY OF CALIFORNIA

Los Angeles

A Time Series Analysis of Major U.S. Construction Industry Material Commodity Prices

A thesis submitted in partial satisfaction

of the requirements for the degree

Master of Applied Statistics

by

Brayden Conroy Nicholl

2022

ABSTRACT OF THE THESIS

A Time Series Analysis of Major U.S. Construction Industry Material Commodity Prices

by

Brayden Conroy Nicholl
Master of Applied Statistics
University of California, Los Angeles, 2022
Professor Frederic R. Paik Schoenberg, Chair

Employing over 10 million people, with an annual expenditure averaging over $1.2 trillion, the U.S. construction industry is massive in both scale and reach. Every day, thousands of construction builds are underway across the country, each with a unique set of costs contributing to its overall budget. For many, the cost of commodities, or materials used to build a project, is a huge variable that changes constantly and can negatively impact a project both financially and temporally. Being able to better understand and predict the way the prices these commodities change over time could lead to huge savings of time and money for anyone at the planning stages of a construction project. Statistical time series modeling is one effective way to increase this understanding and prediction abilities. This thesis models and forecasts the price indices of three construction commodities integral to the industry: lumber, iron/steel, and concrete. For each commodity, the month-to-month percentage changes of nearly 100 years of index values are modeled using Autoregressive Moving Average (ARMA) models, which are selected based on multiple model selection criteria. The models are then used to forecast an 8-month period directly following the end of the data series, which corresponds to the first 8 months of 2022. In doing so, it is

found that the lumber model using ARMA methods makes forecasts indicative of what has happened in reality in the first half of 2022, where as the iron/steel and concrete models have much more difficulty in doing so.

The thesis of Brayden Conroy Nicholl is approved.

Yingnian Wu

Oscar Hernan Madrid Padilla

Frederic R. Paik Schoenberg, Committee Chair

University of California, Los Angeles

2022

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

Introduction

Contractors, project managers, and consultants across the construction industry all rely heavily on the price of raw materials when making decisions regarding future project builds. How the prices of materials escalate or de-escalate over time is a huge variable that needs to be accounted for, regardless of a project's time duration. Increased accuracy in the forecasting of these material prices could lead to changes in decision-making that would see temporal savings, as well as financial savings in the millions, if not billions, of dollars. For a financier of a project, such savings would come in the form of reduced contingency, the amount of money that a budget often includes in order to prevent cost overruns related to unforeseen, unexpected or underestimated events [9]. Contingency considerations are typically made at the concept and planning stages of the project, and this can lead to cost overruns if the forecasts needed are either inaccurate or not performed at all. Also, because most mid-project cost reports make the unrealistic assumption that future unstarted work will be performed at budget, cost overruns are underestimated until late in the life of a project when there's little that can be done to control them [18]. As a potential solution to this lack of forecasting ability, a time series analysis of the past prices of major material commodities, could lead to findings that enable project managers and consultants to give greater assurance to financiers that project cost estimates are both valid, and sustainable long term.

Time series analysis is a type of statistical research that deals with data observed at different points in time, and benefits from sampling the observations in a chronologically adjacent manner. Such a sampling method restricts the applicability of the many conventional statistical methods traditionally dependent on the assumption that these adjacent observations are independent and identically distributed [17]. For this research specifically, a time series modeling method known as Autoregressive Moving Average (ARMA) models are used to model U.S. government compiled price indices for the material commodities lumber, iron/steel, and concrete. These three commodities were specifically chosen due to their roles as the main materials used as the infrastructure of nearly every building nationally.

ARMA modeling was introduced by Box and Jenkins in their 1970 paper *Time Series Analysis, Forecasting and Control* and has since become of the most popular ways for working with time series [11]. Using a linear combination of both a set number of preceding series value and a set number of preceding errors, ARMA models have the ability to forecast a univariate time series, given the series is stationary, meaning it has constant variance over time.

While predicting the price of material commodities has long been of interest to those in the construction industry, any statistical time series analysis of the field is somewhat unprecedented. For many firms, current forecasting involves either the surveying of industry experts on general expected trends, or the implementation of a simple moving average. Individually, such results contain information valuable to those in decision-making positions for project builds. But because they are very simplistic in the forecasts they produce, they are often relegated to the role of thought provocation or sense checking for many people that, ideally, should be making much more educated, mathematically-backed decisions. Thus, modeling methods like ARMA have the ability to provide this backing, and potentially change the way construction project contingency and time escalation are handled.

While the three commodities in question are commonly used in tandem during the building of a structure, and therefore are not truly independent in how a change in pricing of one affects another, an assumption of independence is made for this analysis, on the basis that such dependencies are unique to specific areas, and any variation across regions would not translate to the national level, which the series are based on.

CHAPTER 2

Data

2.1 Background

The data used for analysis was sourced from the U.S. Bureau of Labor Statistics' Producer Price Indexes [4]. Although the site contains data series of hundreds of unique construction commodities, the three chosen for analysis, lumber, iron/steel, and concrete, are considered some of, if not the, most widespread and vital commodities in use across the industry. These series describe the price of a given commodity over a month, based on the entire U.S. market as a whole. Due to this, regional fluctuations in a market can influence an index as a whole, but very rarely, and often on a relatively minor scale. Thus, this analysis should not necessarily be assumed as an accurate representation of the price of a given commodity in a specific location.

All of the commodities share the same format, which is that of monthly releases dating back to January of 1926 [7]. As some of the oldest economic time series compiled by the U.S. government, these indices are one of the only ways to analyze and understand the changes in the state of the national construction industry that have occurred over the course of more than 100 years. The series are numerical in nature and rounded to the tenths decimal place, and have a base value of 100 set to 1982 [5].

Often released around the middle of the following month, the data indicates the aggregate price of a commodity over the course of a month. The two to three week lag on reporting from the U.S. BLS is due to the surveying and compilation required before publication of a national-level index can occur.

2.2 Data Reformatting

The aforementioned producer price indices were obtained in the format of a raw numerical indices, but due to the nature of the time series analysis required, the desired data was not the index values for the respective months, but actually the percentage change from the previous month. To make this change, each commodity index was transformed using the equation

$$change_i = \frac{index_i - index_{i-1}}{index_{i-1}} * 100$$

where i was a given month's chronological place in the series. Thus, each point the new series was the percentage change from the previous month. For example, the new first entry in the series, denoted February 1926, is the percentage change from January 1926 to February 1926.

This reformatting of the data set allowed for the ability to control for inflation, since the original data was not adjusted for inflation, and therefore the scale of the changes in the original indices were not normalized.

CHAPTER 3

Exploratory Data Analysis

3.1 Visual Time Series Assessment

As can be seen in Figure 3.1, the various indices grow in value over time in a generally exponential fashion. This exponential growth was one reason behind the switch to monthly change percentages for the analysis data.

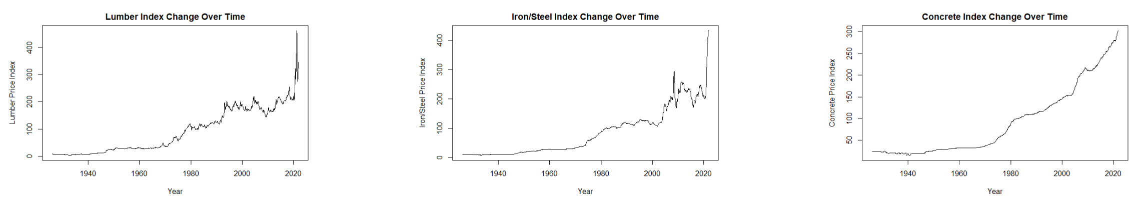

(a) Lumber (b) Iron/Steel (c) Concrete

Figure 3.1: Commodity Indices' Values Over Time

Initially, the data used was going to be monthly index value change, as opposed to percentage change. However, the amount of inflation seen over the course of approximately 95 years caused the scale of the indices to change vastly across the duration of the data set. Visualized in Figure 3.2, the problem with this change in scale can be seen when comparing the raw changes and percentage changes of the concrete index.

When looking solely at the raw monthly changes, it seems as if the price of concrete has had its most volatile swings in recent years, with similarly large swings seen in the late 2000s and late 1930s. But when normalizing for the scale of the data, the most recent price changes seem relatively tame, and the highest volatility by far occurs in the late 1930s, with

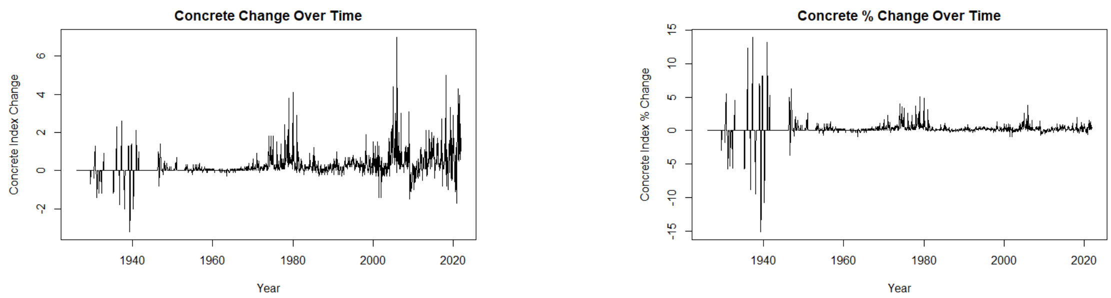

(a) Monthly Raw Change (b) Monthly Percentage Change

Figure 3.2: Concrete Index's Changes Over Time

the next largest swings coming in the late 1970s. This normalization allows for comparison across eras, and negates any possible time-dependent bias in the data.

3.2 Seasonality

Because construction, as an industry, is inherently dependent on the weather, is was a necessity to check for possible cyclicity or seasonality in the data. For seasonality, The percentage changes for the commodities were plotted on polar coordinates, shown in Figure 3.3.

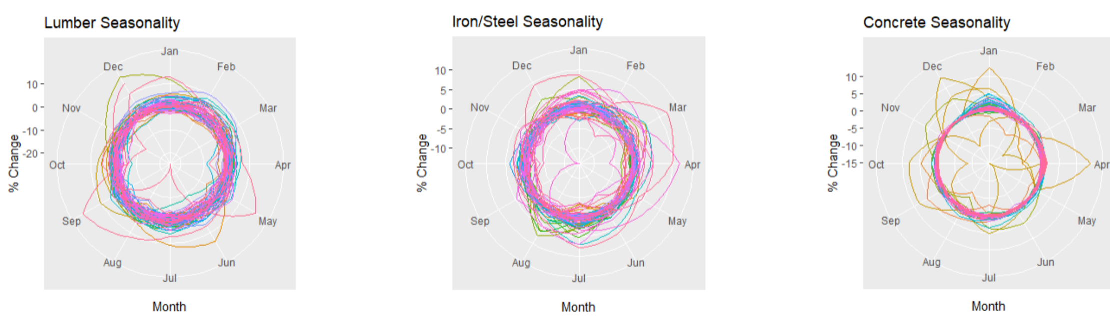

(a) Lumber (b) Iron/Steel (c) Concrete

Figure 3.3: Commodity Indices' Values Over Time

None of the three plots indicated any sort of seasonal trend in the data, since the monthly percentage changes hover around the same mean values, regardless of month. It is interesting to note, however, that the plots seem to indicate that concrete has a tighter spread for its monthly changes, albeit with more common occurrences of abnormal, outlying values when compared to either lumber or iron/steel.

3.3 AR Spectrum

To check for non-annual cyclicity that may not appear in a seasonal plot, an individual spectral analysis was performed for each of the three commodities.

The process began with fitting autoregressive models of varying orders to the monthly percentage change data, and calculating the spectral density of the optimal model, which was chosen based on AIC (both autoregressive (a.k.a. AR) models and AIC are explained in the following ARMA Modeling and Model Selection chapters, respectively). The obtained spectral density was then plotted using a periodogram.

In the periodogram, the x-axis is frequency over a given time period, in this case 1 year (e.g. a frequency value of 2 means a period occuring twice per year, having a length of 6 months). The y-axis is the spectral amplitude of the model, and the line explains how the spectral density changes as the frequency of the periods change. A peak, or local max, in the line indicates a possible cyclicity in the data at that frequency [14].

3.3.1 Lumber

Starting with lumber, an AR model of order 16 was chosen as optimal. Figure 3.4 displays the periodogram of the $AR(16)$ model for lumber.

The lumber periodogram shows there may be cycles with frequencies near 1 or 4. A frequency of 1 means a yearly cycle with similar changes in the index occurring around the

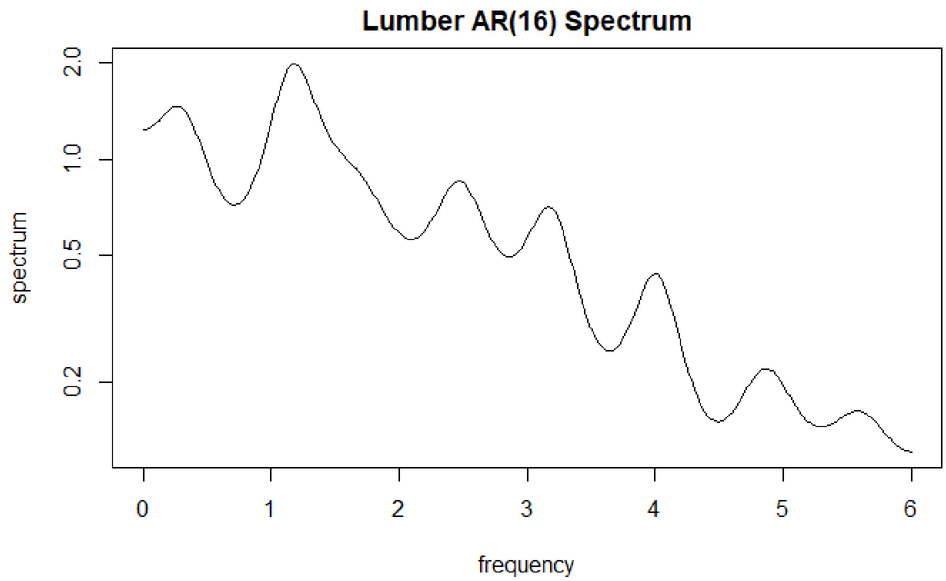

Figure 3.4: Lumber AR Spectrum Periodogram

same time every year. Although possible, this scenario is unlikely, given the lack of yearly seasonality seen in Figure 3.3. A frequency of 4 means a cycle occurring every 3 months, which would not necessarily show up in a plot like Figure 3.3. A solution to combat this potential cyclicity is discussed in the Monthly Data Averaging section later in this chapter.

3.3.2 Iron/Steel

Similarly to the process for lumber, the iron/steel index had its monthly percentage changes modeled and its optimal model's periodogram plotted. For iron/steel, an $AR(29)$ model was chosen by AIC to be the best. The periodogram for Iron/Steel can be seen in Figure 3.5.

Unlike lumber, the curve of the iron/steel periodogram has no peaks that stand out among the rest. Fortunately, this indicates that there is most likely no true cycle in the index, and that no further corrective action is needed in this case.

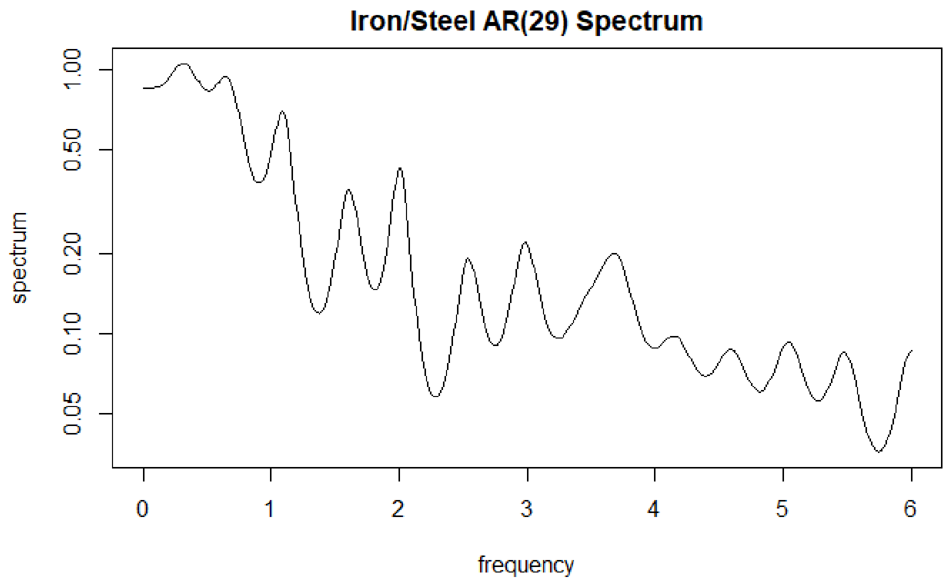

Figure 3.5: Iron/Steel AR Spectrum Periodogram

3.3.3 Concrete

The same was also done for the concrete data, and this time, an $AR(16)$ model was selected. Figure 3.6 is the periodogram for the monthly percentage changes for concrete.

The spectral density of the concrete data is extremely variable, as can be seen in the plot. Although there are a high number of local maxima occurring at various frequency levels, the most notable with respect to their adjacent local minima are at frequencies of around 4 and 5. These frequencies correspond to periods of approximately 3 months and 2.5 months, respectively. Similarly to lumber, action will thus need to be taken to deal with any cycles that may exist in the data.

3.4 Monthly Data Averaging

With Figures 3.4 and 3.6 exposing potential cycles in their respective indices, cyclicity must be removed from the data before any modeling can take place. To achieve this, the mean

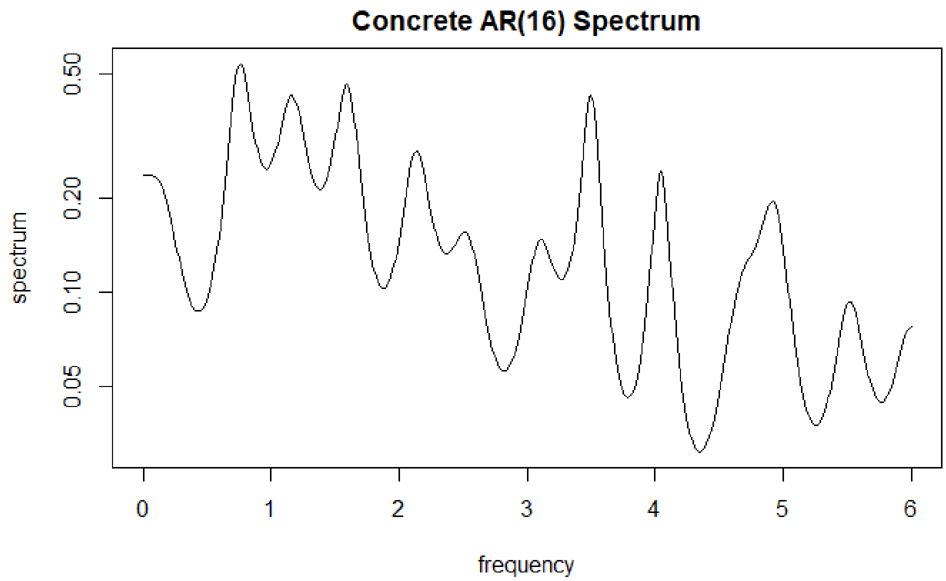

Figure 3.6: Concrete AR Spectrum Periodogram

percentage change was found for each of the 12 months of the year, across the entire length of the series. Each observation then had its corresponding monthly mean subtracted to normalize the data across months. In doing so, a cycle of any duration would then be removed, and each data point now was the percentage change above or below the monthly mean for that specific month of the year. After this subtraction was performed, the time series were ready for modeling using ARMA methods.

Note that because no cyclicity problems were found within the iron/steel data, this monthly averaging was only performed on the lumber and concrete series.

CHAPTER 4

ARMA Modeling

In order to model and, in following chapters, forecast the time series values of the Producer Price Index commodities, $ARMA(p, q)$, short for Autoregressive Moving Average models were used. An $ARMA(p, q)$ is a combination of $AR(p)$ and $MA(q)$ models, where p and q are the orders of their respective models [10].

An AR model, also known as an Autoregressive model, is a model that calculates the regression of past time series and uses the regression to estimate a present or future value in a univariate time series [16]. For a given time t in the series, the $t-1, t-2, \ldots, t-p$ values are all used to find the estimate for t, in combination with their corresponding coefficients β_i. The equation of an $AR(p)$ model can be written as

$$y_t = \beta_1 y_{t-1} + \beta_2 y_{t-2} + \cdots + \beta_p y_{t-p}$$

where y_i is the value of the time series at time i, and β_i is the coefficient corresponding to that series value.

An MA, or Moving Average, model calculates the residual errors of past entries in the time series, and uses these errors to estimate a present or future value in a univariate time series [16]. The $t-1, t-2, \ldots, t-q$ points in the time series have their impact on the estimate of the series at time t decided by their coefficients α_i. The equation of an $MA(q)$ model can be written as

$$y_t = \alpha_1 \epsilon_{t-1} + \alpha_2 \epsilon_{t-2} + \cdots + \alpha_q \epsilon_{t-q}$$

where ϵ_i is the error of the value of the time series at time i, and α_i is the coefficient corresponding to that series value's error.

ARMA models make use of both the previous lags from AR modeling and the errors from MA modeling to forecast the next value of a univariate time series. An $ARMA(p, q)$ model is written as

$$y_t = \beta_1 y_{t-1} + \beta_2 y_{t-2} + \cdots + \beta_p y_{t-p} + \alpha_1 \epsilon_{t-1} + \alpha_2 \epsilon_{t-2} + \cdots + \alpha_q \epsilon_{t-q}$$

where $y_i, \beta_i, \epsilon_i, \alpha_i$ all have the same the same usage as in the AR and MA models.

A time series, by its very nature, contains observations that cannot be independent of each other, due to the chronological nature of their ordering. Therefore, to make up for the lack of independence in the data, any time series must be stationary in order to be ARMA modeled. Stationarity is the assumption that the mean and variance are constant throughout the data [17].

The original dataset did not meet these assumptions, however. To assure that the mean remained constant across the series, the data was transformed from index values to month-to-month changes. Then, to deal with the variance, the raw monthly changes were changed to percentage value changes as described in section 2.2. After these transformations were performed, each of the three series were ready to be modeled.

4.1 Model Selection Methods

4.1.1 ACF and PACF

To determine the order of an ARMA model there are multiple methods that can be used. One is the combination of the ACF, also known as the Auto Correlation Function, and the PACF, the Partial Auto Correlation Function.

The ACF is a function that take all of the preceding observations in a series, and calculates the Pearson Correlation between the time to be estimated t, and each of the intervals, in this case time lags, from 0 to t [16]. ACF, however, does not factor in the effects that the past observations have on the observation at time t, and therefore there are no weighted coefficients in the function. ACF can be used to determine the optimal number of terms in an MA model, or the number of MA terms in an ARMA model.

The PACF finds the partial correlation between the time t and a given time $t-p$. Unlike ACF, PACF only looks at the correlation between two times, and not the intervals between them. This allows for the ability to observe the relationship of a set interval, without worrying about the effect that intermediate intervals could have on the relationship. PACF is used to determine the optimal number of terms in an AR model, or the number of AR terms in an ARMA model.

4.1.2 AIC

AIC, which stands for Akaike's Information Criterion, is an estimator of prediction error named for its creator, Hirotugu Akaike [2]. It is a model selection criterion that, when calculated for multiple models, can be used to compare the models head-to-head to determine which is best. The formula for AIC is

$$AIC = 2k - 2ln(\hat{L})$$

where k is the number of estimated parameters in the model, and $\hat{L}$ is the log-likelihood, a measure of model fit [8]. Model fit improves as log-likelihood increases, so therefore the model with the lowest AIC value is deemed as being the best.

4.1.3 RMSE

Another model selection method is RMSE, or root mean square error. Put simply, RMSE is the amount of error that the average prediction made by the model has when compared to its actual point in the data. It can be found by squaring the absolute value of each residual error, squaring each absolute value, summing over the squares, dividing by the number of observations to find the mean square, and taking its square root [6]. Arithmetically, this can be expressed as

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \hat{x}_i)^2}{N}}$$

where x_i is the actual observation, $\hat{x}_i$ is the observation estimate, and N is the number of observations. Because RMSE is an evaluation of model error, it is ideally minimal as possible, and the model with the lowest RMSE would be considered best.

4.2 Lumber

For the lumber series, the data used was the monthly averaged percentage changes time series, for the reasons described in section 3.4.

Firstly, the ACF and PACF values for the data were plotted, as can be seen in Figure 4.1. The ACF plot has large dropoffs in ACF values after both the first and second lags (Note that the "Lag" on the x-axis is listed in years, so each column corresponds to a number of months). This means that the recommended number of MA terms is either 1 or 2. The PACF plot has the largest drop off after the first lag, so 1 AR term is recommended. Therefore, the plots in Figure 4.1 indicate either $ARMA(1,1)$ or $ARMA(1,2)$ would be best.

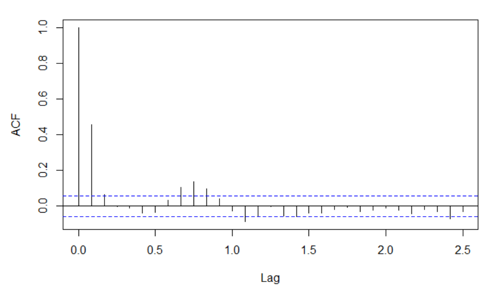

(a) ACF

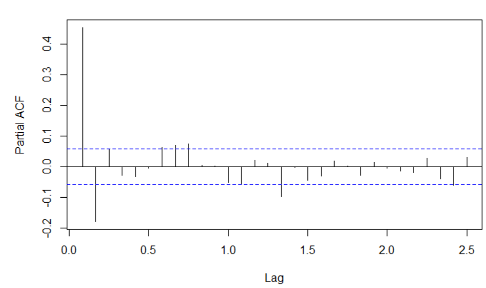

(b) PACF

Figure 4.1: Lumber ACF and PACF Plots

To evaluate models based on AIC and RMSE, ARMA models of every possible combination of AR and MA terms from 0 to 4 was fit, and both their AICs and RMSEs were recorded. These results are shown in Table 4.1, sorted by AIC.

	AR_term	MA_term	AIC	RMSE
23	4.00	2.00	5186.64	2.28
25	4.00	4.00	5187.72	2.28
24	4.00	3.00	5187.86	2.28
20	3.00	4.00	5188.14	2.28
3	0.00	2.00	5199.25	2.30
7	1.00	1.00	5199.42	2.30
8	1.00	2.00	5201.24	2.30
4	0.00	3.00	5201.24	2.30
12	2.00	1.00	5201.27	2.30
16	3.00	0.00	5202.08	2.30
21	4.00	0.00	5202.32	2.30
5	0.00	4.00	5203.23	2.30
17	3.00	1.00	5203.24	2.30
9	1.00	3.00	5203.24	2.30
13	2.00	2.00	5203.28	2.30
22	4.00	1.00	5203.29	2.30
11	2.00	0.00	5205.12	2.31
18	3.00	2.00	5205.18	2.30
10	1.00	4.00	5205.24	2.30
14	2.00	3.00	5205.24	2.30
2	0.00	1.00	5206.31	2.31
19	3.00	3.00	5207.24	2.30
15	2.00	4.00	5207.24	2.30
6	1.00	0.00	5242.90	2.35
1	0.00	0.00	5515.68	2.65

Table 4.1: Lumber ARMA Results

The model chosen by AIC was an $ARMA(4,2)$ model, and the model chosen by RMSE was an $ARMA(4,4)$, albeit marginally over $ARMA(4,2)$. Interestingly, the $ARMA(1,1)$

and $ARMA(1,2)$ models chosen by ACF and PACF were only the sixth and seventh best models, respectively, of the 25 created, based on AIC. Given these results, the $ARMA(4,2)$ model was selected as the best, and used to make the lumber forecasts seen in the Model Forecasting chapter.

Interestingly, the worst model by far was that of $ARMA(0,0)$. This model with neither AR nor MA terms is interpreted as modeling strictly white noise in the data with no trend whatsoever. The fact that this model in particular performed so poorly is a clear indicator that there is indeed a time-dependent trend in the series.

The $ARMA(4,2)$ model created from the lumber data was

$$y_t = 1.55y_{t-1} - 1.56y_{t-2} + 0.72y_{t-3} - 0.25y_{t-4} - 1.00\epsilon_{t-1} + 0.80\epsilon_{t-2} + 0.01$$

with the 0.01 being the intercept value. The most interesting term was the first AR term, which had a coefficient value of 1.55. The fact that this coefficient was both positive and significantly greater than 0 could be interpreted as meaning that there is true continuity in the series. This means that based on the model, the next percentage change observation in the lumber series is more likely to be in the same direction (positive or negative) as the previous observation than it is to be in the opposite direction. This aligns with what was seen in Figure 3.4, because a steadily falling AR spectrum also indicates continuity in a series. Such a finding could, in theory, be used to more accurately predict the change in the price of lumber in upcoming months.

When examining the model residuals, the errors had a generally normal distribution, centered around 0, which can be seen in the lower right of Figure 4.2. The top plot of the same figure showed no obvious trend over time, with the only abnormality being the scale of the residuals near the end of the series, which was not surprising given the huge swings caused by the COVID-19 pandemic, seen in Figure 3.1a. The lower right plot, the ACF of the residuals, was small in scale for all lags up to 36 months. All of these findings were very encouraging, because they did not raise any concerns about the fit of the model to the data.

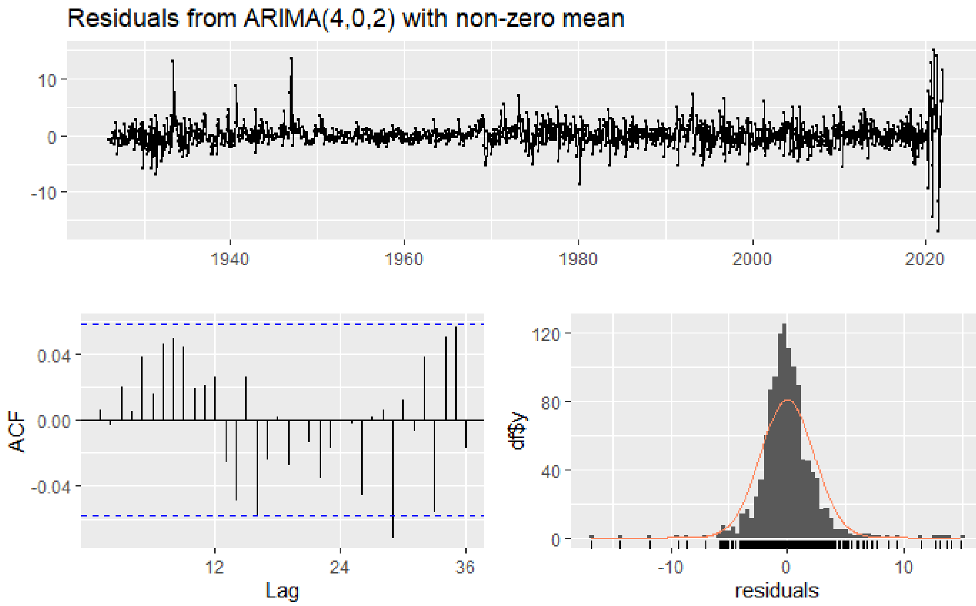

Figure 4.2: Lumber ARMA Model Residuals

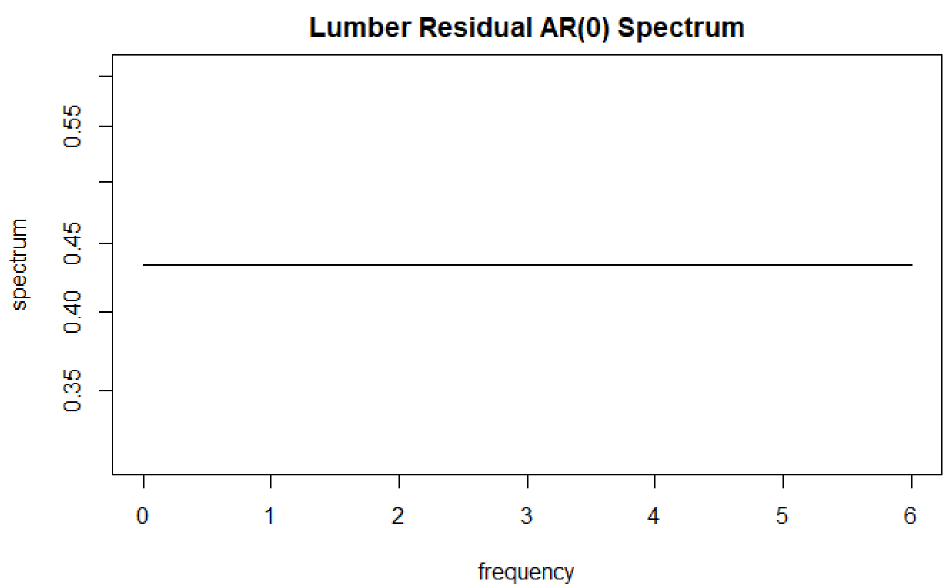

Figure 4.3: Lumber Model Residual AR Spectrum

Figure 4.3 is the plotted periodogram of the model residuals' AR spectrum. At first glance, it may seem somewhat uninformative. However, the fact that the optimal AR order chosen for the residuals is 0 can be interpreted as meaning that they have a distribution resembling pure white noise. This was an indicator that the model fits the data well, as an optimal model should have residuals with such a distribution. Thus, the choice of $ARMA(4, 2)$ model by AIC is backed by the residual data seen above.

The model residuals also have an RMSE of 2.28. This can be interpreted as meaning that the model is an average of 2.28% off when predicting the month-to-month percentage change of the lumber index. Given that lumber is historically the most volatile of the three commodities analyzed, being off only 2.28% on average when making predictions is quite impressive.

4.3 Iron/Steel

For the iron/steel series, the data used was the percentage change series. This series was not averaged by the monthly means however, since it was deemed unnecessary in section 3.4.

Similarly to the lumber data, the ACF and PACF plots were created, and can be seen in Figure 4.4. For the ACF plot, there was a steady dropoff from the first lag to around the fifth or sixth. Because of this, the amount of MA terms needed for modeling was unclear. PACF showed a massive decrease after the first lag, so 1 AR term was the clear favorite in this case. While leaving room for interpretation for the MA terms, the plots of Figure 4.4 indicate the optimal model as being somewhere in the range of $ARMA(1, 1)$ and $ARMA(1, 6)$.

Model selection via AIC and RMSE was performed using the same method as lumber, with the fitting of all possible term number combinations from 0 to 4. The recorded AICs and RMSEs are listed in Table 4.2, sorted by AIC.

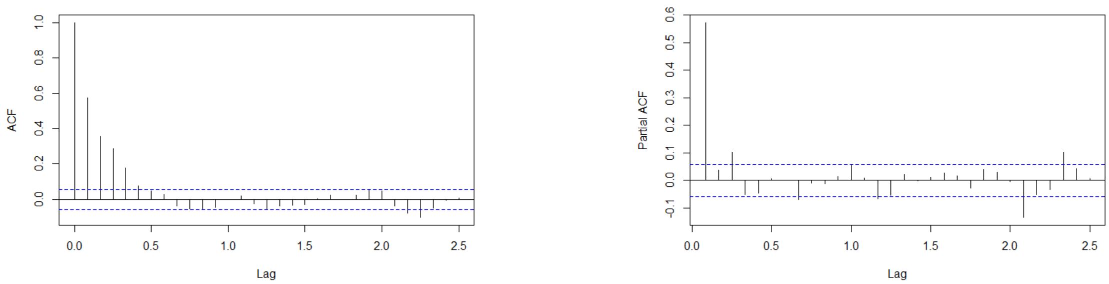

(a) ACF (b) PACF

Figure 4.4: Iron/Steel ACF and PACF Plots

	AR_term	MA_term	AIC	RMSE
20	3.00	4.00	4034.08	1.38
19	3.00	3.00	4034.87	1.38
23	4.00	2.00	4035.15	1.38
24	4.00	3.00	4035.77	1.38
18	3.00	2.00	4035.95	1.39
22	4.00	1.00	4036.81	1.39
21	4.00	0.00	4037.51	1.39
9	1.00	3.00	4037.72	1.39
25	4.00	4.00	4038.02	1.38
16	3.00	0.00	4038.55	1.39
10	1.00	4.00	4038.99	1.39
17	3.00	1.00	4039.01	1.39
14	2.00	3.00	4039.27	1.39
15	2.00	4.00	4040.77	1.39
5	0.00	4.00	4041.19	1.39
13	2.00	2.00	4041.50	1.39
8	1.00	2.00	4043.48	1.39
7	1.00	1.00	4048.20	1.40
6	1.00	0.00	4048.64	1.40
11	2.00	0.00	4049.02	1.40
12	2.00	1.00	4052.64	1.40
4	0.00	3.00	4064.00	1.41
3	0.00	2.00	4103.58	1.43
2	0.00	1.00	4150.09	1.46
1	0.00	0.00	4505.30	1.71

Table 4.2: Iron/Steel ARMA Results

Both AIC and RMSE selected $ARMA(3,4)$ as best for modeling the iron/steel data. This turns out to be quite contradictory to the findings of the ACF and PACF plots, since while $ARMA(1,3)$ performs reasonably well (ranked eighth out of 25 by AIC), the rest of the plots' recommended models did not fare well, consistently ranking in the poorer half of the models tested when judged by the other selection methods. Due to the fact that two of the three methods agreed so closely, and that the ACF did not have a definitive answer on visual inspection, the $ARMA(3,4)$ model was chosen to be used for the forecasting performed later.

Again, the the $ARMA(0,0)$ model was far and away the worst model in terms of both AIC and RMSE. This lends validity to the claim that there is a trend in the data, and that the movement of the iron/steel index is not purely white noise.

The $ARMA(3,4)$ model created from the iron/steel data was

$$y_t = -0.02y_{t-1} - 0.48y_{t-2} + 0.71y_{t-3} + 0.58\epsilon_{t-1} + 0.80\epsilon_{t-2} - 0.20\epsilon_{t-3} - 0.08\epsilon_{t-4} + 1.92$$

with 1.92 being the intercept value. In this case, the first AR term does not tell much of a story. It was slightly negative, but it is so close to zero that conclusions drawn from it would not be of much value. Based on the coefficient value of -0.02, the index is not more likely to continue in the same direction it has prior, or move in the opposite direction. What makes this finding interesting is that such an explanation is reminiscent of a pure white noise distribution, but the $ARMA(0,0)$ model that models this scenario performs terribly when applied to the iron/steel series. Also, the iron/steel AR spectrum depicted in Figure 3.5 shows an overall decline as the frequency increases, suggesting that there is some continuity to be found in the data. A possible explanation could be that while the iron/steel time series does not have the point-to-point continuity that would come with a large first AR term, the overall positive trend seen in Figure 3.1b makes a white noise distribution, such as that of $ARMA(0,0)$, a very poor fit.

For iron/steel, the residual plots in Figure 4.5 are positive in what they say about the chosen model. The top residual plot does not reveal any unknown trends over time in the

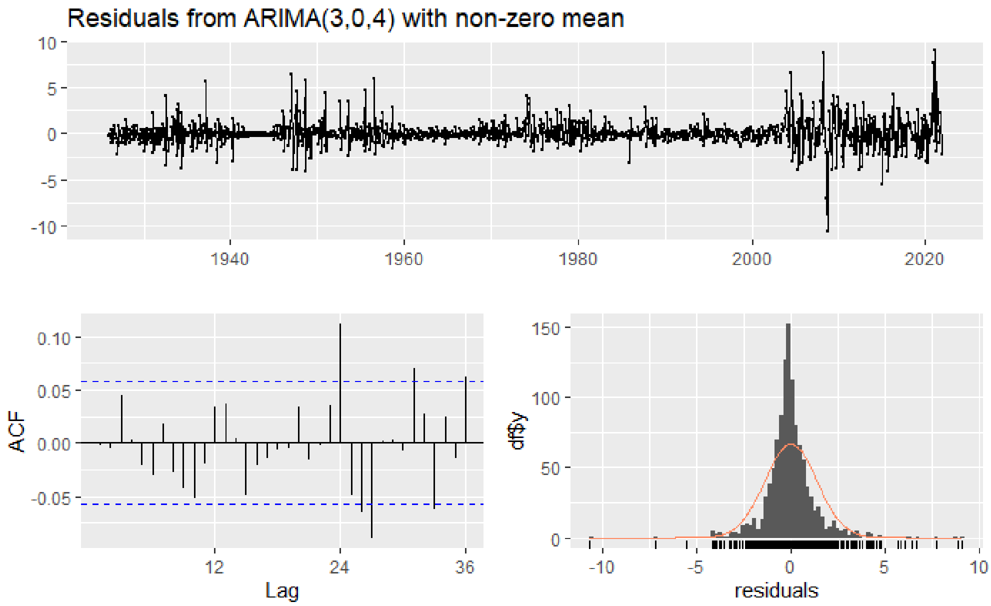

Figure 4.5: Iron/Steel ARMA Model Residuals

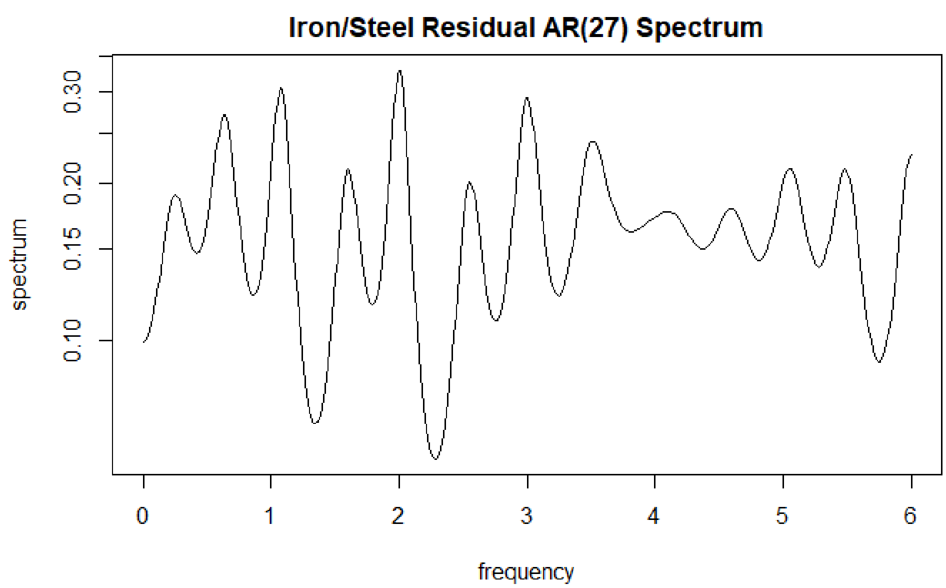

Figure 4.6: Iron/Steel Model Residual AR Spectrum

data, although the residuals appear to often be larger than normal in the 2010s, which could indicate that the model did not perform as well during this period as it had done previously.

The residuals appear to have a bell-shaped distribution as desired, albeit with a smaller standard deviation than a true normal distribution would. The ACF plot has generally small magnitudes, with no obvious trends over adjacent lags, and thus raises no concerns in particular.

The AR spectrum periodogram for the residuals, shown in Figure 4.6, did not choose 0 as the optimal order for iron/steel, instead opting for an $AR(27)$ spectrum. While not as ideal a choice as with lumber, the figure contains no peaks standing out from the rest, meaning that there are no apparent cycles in the residuals, a positive finding. Also, across the up and down fluctuations, the spectrum values remain generally centered in the 0.15-0.20 range, which indicates a lack of continuity from one residual to the next, like what occurs in an ideal white noise distribution.

Over the entire series, the $ARMA(3, 4)$ model chosen had a residual RMSE of only 1.38. Being only 1.38% off on average when making predictions is very promising, especially given the amount of volatility the iron/steel industry has seen in recent years. Such a low RMSE is an indicator that the model is indeed a good fit for the time series data.

4.4 Concrete

The monthly averaged percentage changes were used to model the concrete series, due to reasons similar to that of the lumber series.

The ACF and PACF plots were also made for concrete, seen in Figure 4.7. For ACF, there are large dropoffs after both the first and second lags, indicating a potential choice of either 1 or 2 MA terms. The PACF plot, however, is not nearly as helpful. While it does reduce in some areas, the Partial ACF value also has large increases in multiple areas, and has no overarching trend as the number of lags increases. Thus, there is no good estimate for the number of AR terms based on the PACF plot.

AIC and RMSE were once again tested for each possible term combination numbering

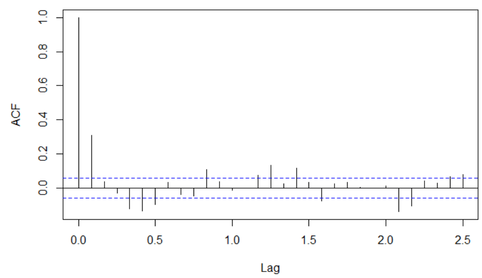

(a) ACF

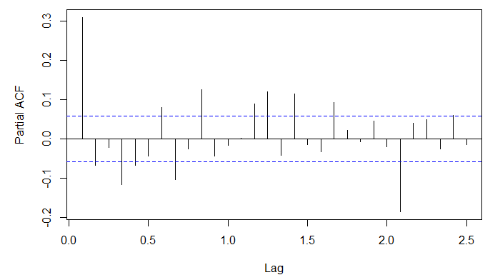

(b) PACF

Figure 4.7: Concrete ACF and PACF Plots

from 0 to 4, with the results listed in Table 4.3, sorted by AIC.

	AR_term	MA_term	AIC	RMSE
14	2.00	3.00	3803.80	1.25
22	4.00	1.00	3807.70	1.26
19	3.00	3.00	3808.19	1.25
15	2.00	4.00	3809.05	1.26
20	3.00	4.00	3810.19	1.25
24	4.00	3.00	3810.19	1.25
23	4.00	2.00	3810.56	1.26
21	4.00	0.00	3811.32	1.26
10	1.00	4.00	3811.93	1.26
25	4.00	4.00	3812.19	1.25
18	3.00	2.00	3815.11	1.26
17	3.00	1.00	3815.92	1.26
5	0.00	4.00	3821.31	1.26
2	0.00	1.00	3823.26	1.27
9	1.00	3.00	3823.59	1.27
11	2.00	0.00	3823.62	1.27
12	2.00	1.00	3823.67	1.27
7	1.00	1.00	3823.83	1.27
3	0.00	2.00	3824.02	1.27
4	0.00	3.00	3824.50	1.27
13	2.00	2.00	3824.72	1.27
16	3.00	0.00	3825.08	1.27
8	1.00	2.00	3825.82	1.27
6	1.00	0.00	3826.73	1.27
1	0.00	0.00	3941.39	1.34

Table 4.3: Concrete ARMA Results

$ARMA(2,3)$ was chosen by both methods for the best modeling of the concrete series. The runner-up was chosen to be $ARMA(4,1)$, which could potentially validate the choice made earlier via ACF, but MA term counts of either 1 or 2 are not common near the top of the table, so this second place selection could also be an outlier in that regard. Due to the agreement of two of the three methods, as well as the inability to draw conclusions from the PACF plot, it was decided that $ARMA(2,3)$ would be chosen as the best model for the concrete data, would be used for eventual forecasting of the time series.

Once again, $ARMA(0,0)$ did a very poor job of modeling the series, ranking dead last, with the highest AIC and highest RMSE of any model attempted. At the very least, this showed that the movement of the concrete index could not be chalked up to pure white noise.

The $ARMA(2,3)$ model created from the iron/steel data was

$$y_t = 0.12y_{t-1} - 0.80y_{t-2} + 0.21\epsilon_{t-1} + 0.84\epsilon_{t-2} + 0.33\epsilon_{t-3} - 0.0005$$

with a near-negligible intercept value of -0.0005. While not extremely large, the positive first AR term coefficient of 0.12 means that there is some continuity throughout the monthly changes of the concrete series. This tells a similar story to that of Figure 3.6, since the AR spectrum values did fall as the frequency increased, but only slightly, and the massive swings in the spectrum values made for an unconvincing trend. Nonetheless, a real, continuous trend in the time series could be used to more accurately interpret the past changes in the index value for concrete, and allow for a more accurate understanding of how concrete prices will move in the near future.

Due to its more moderate price index swings relative to the other two commodities, concrete saw model residuals that were smaller on average. The exception to this being during the late 1930s, which has outlying large residuals, mostly corresponding to similarly large index percentage changes in the same time period. This can be seen in the top of Figure 4.8, and can most likely be attributed to the global unrest and economic turmoil that lead to World War II in the following years. The small spread of the residuals is represented in the figure's bottom right plot, which depicts a bell-shaped distribution. There is very low variance in the residual values overall, but this is not especially concerning, given that the distribution is centered around 0, and takes the aforementioned bell shape that indicates normality. ACF could potentially be a concern, with the magnitude of multiple values

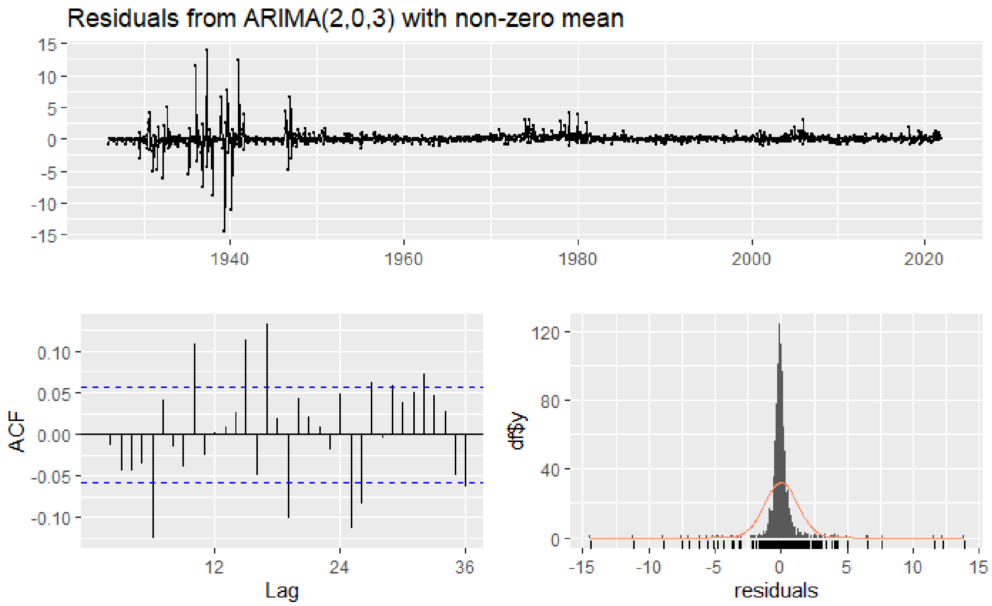

Figure 4.8: Concrete ARMA Model Residuals

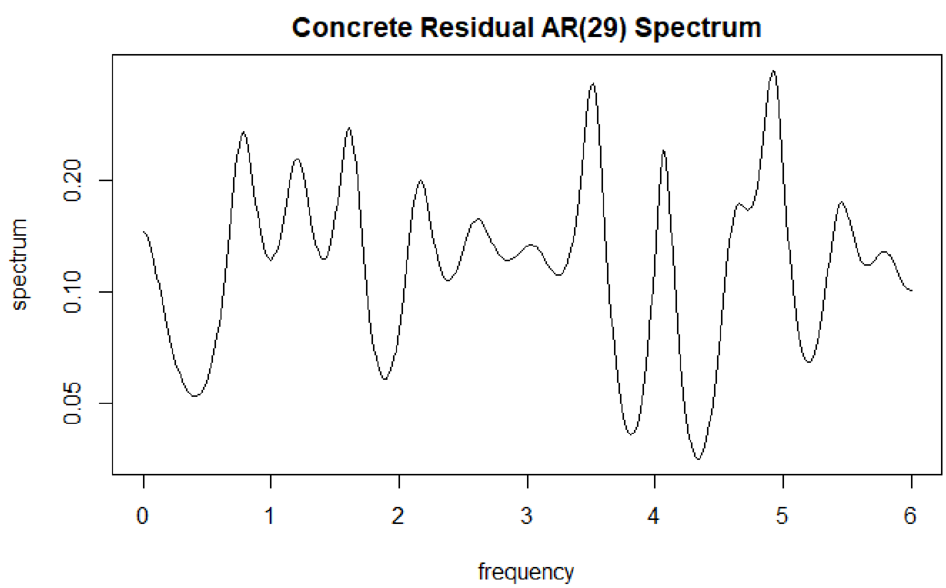

Figure 4.9: Concrete Model Residual AR Spectrum

standing out among the rest. However, there is no apparent pattern across lags, and the lack of any overall trend makes any concerns regarding a specific lag or two somewhat of a reach when viewing the residual data as a whole.

Similar to iron/steel, a large AR spectrum order of 29 was picked for the concrete data before it was used for plotting. The periodogram, seen in Figure 4.9, does not have any highs or lows on a different scale than the rest, and does not have an overall trend of increase or decrease. Thus lacking any indication of either cyclicity or point-to-point continuity, it can be assumed that the $ARMA(2,3)$ chosen to model the concrete data is indeed a proper fit.

Also supporting this claim is the model's residual RMSE of 1.25, the lowest among the three optimal commodity models produced. This low value is somewhat unsurprising, given the relatively small residual values seen in the last 70 or so years of the series, but nonetheless is a great indicator of proper model fit, since an average percentage change prediction error of 1.25% is a result that many experts in the construction industry would likely wish they could attain on a monthly basis.

CHAPTER 5

Model Forecasting

In order to make proper use of the attained best models from the prior ARMA modeling, future forecasting of each producer price index time series was performed. Each model was used to forecast 8 months of index results, beginning with January 2022, as it directly follows the final point of the dataset in December 2021. The 8 month cutoff was determined based on common forecasting practices in the construction industry. Many well respected leading indicators of the state of the construction industry, including the American Institute of Architects' Architecture Billings Index or the Associated Builders and Contractors' Construction Backlog Indicator, tend to forecast around 8 months past present day on average [3][1]. Predictions any further out are considered to be too unreliable to place trust in, due to the natural volatility and general rate of change inherent in construction.

5.1 Lumber

After obtaining the optimal $ARMA(4, 2)$ model seen in Section 4.2, the necessary monthly averaged percentage change values were plugged into the model equation to obtain the forecast values. Along with the forecast numbers for the index itself, 80% and 95% confidence intervals were calculated for each month of the forecast. To make plotting possible, the monthly averages were removed from the percentage changes, which were then converted back into raw index numbers whose scale corresponded to that of the original time series data. For visualization and scaling reasons, only the most recent 12 years of the index are plotted alongside the forecasts in Figure 5.1.

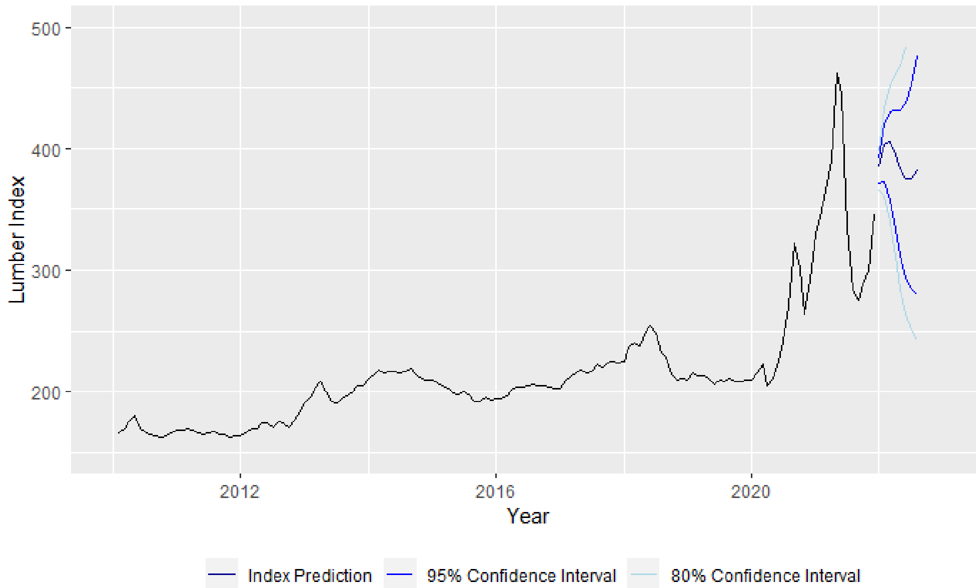

Figure 5.1: 8-Month Lumber Forecasts

Of the three commodity indices analyzed, lumber has by far the most variability in recent times, both in terms of direction and scale. Thus, it should not be surprising that the direction of the model forecast change frequently. After predicting an increase in the lumber index in early 2022, the model forecasts decrease in mid-2022 before beginning to rise again near the end of the year. Similarly, the large ranges of the confidence intervals seem unrealistic when compared to the data before 2020, but when viewing the scale of the data during the COVID-19 pandemic, tighter intervals would seem not only optimistic, but overly confident.

Only time will tell the full accuracy of the model's predictions, as much of the 8-month forecast period has yet to occur in reality. However, early signs are positive for the $ARMA(4,2)$ model. Lumber prices soared in the early months of the year, largely due to Russia's invasion of Ukraine and the change in transportation costs associated with it globally [13]. But since those early-year highs, prices have fallen alongside the recent rise in mortgage rates, and experts predict this trend to continue through the middle of the year, much like the trend seen in the ARMA forecasts.

5.2 Iron/Steel

Using the $ARMA(3, 4)$ model created in Section 4.3, forecast values were determined in a similar process to that of lumber. 80% and 95% confidence intervals were calculated for the iron/steel series forecasts as well. Unlike lumber, there was no monthly averaging performed for the iron/steel data, and therefore no correction was needed for the forecast numbers. The percentage changes were, nonetheless, converted back into a raw numerical index form to enable plotting of the forecasts alongside the original index data. For visualization and scaling reasons, only the most recent 12 years of the index are plotted alongside the forecasts in Figure 5.2.

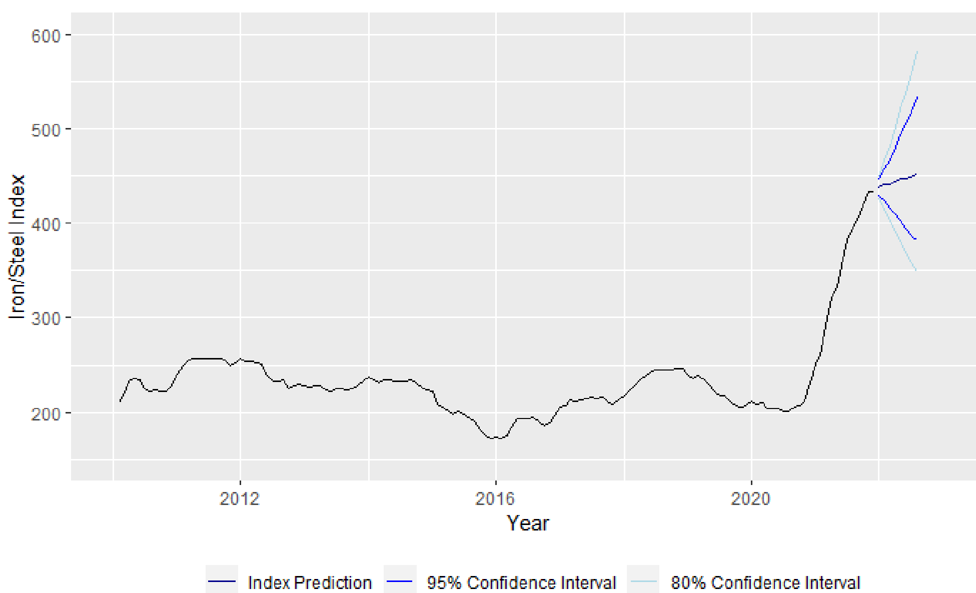

Figure 5.2: 8-Month Iron/Steel Forecasts

Overall, iron/steel has less volatility than seen previously with lumber. But the recent changes in the index have been on a similarly unprecedented scale, with values more than doubling over the course of 2021. Such uncertainty can be seen in the wide spread of the confidence intervals. While a 95% confidence interval having a range of over 100 index points at the 8th forecast month may seem huge, given that the final 8 months of the time series has a rise of around 100 points, the confidence interval is actually quite reasonable. While the scale of the iron/steel index value changes is somewhat variable, the direction of the index's changes is quite consistent. The forecast predictions reflect this consistency, as the 8-month period sees a steady, moderate increase in the prices of iron/steel.

Unlike the model projections, the price of steel in 2022 has been anything but steady. After moderate declines in January and February, massive increases have occurred in the months of March and April [19]. As Russia and Ukraine are two of the world's top ten iron-producing countries by volume, the escalation of their conflict and the worldwide trade sanctions that have come along with it are expected to cause a continuation of the increases seen only recently. While the iron/steel model forecasts have thus performed poorly up to present day (and are expected to for the rest of the 8 months in question), a situation with such an enormous global impact as that of the Russian invasion of Ukraine could not have been accurately predicted by even the most knowledgeable construction experts at the end of 2021, let alone by an ARMA model having to make 8 months of index predictions. Therefore, any inaccuracy in the model's forecasts should realistically be taken with a grain of salt, and a bit of perspective.

5.3 Concrete

Forecast values were also created for the concrete series, using the $ARMA(2,3)$ model created in Section 4.4, along with the corresponding 80% and 95% confidence intervals. There was monthly averaging done previously for concrete, so the first step in reformatting the forecasts was removing the respective monthly averages from each percentage change. Then, they were converted to the raw index form and then plotted with the original index data in the same figure. For visualization and scaling reasons, only the most recent 12 years of the index are plotted alongside the forecasts in Figure 5.3.

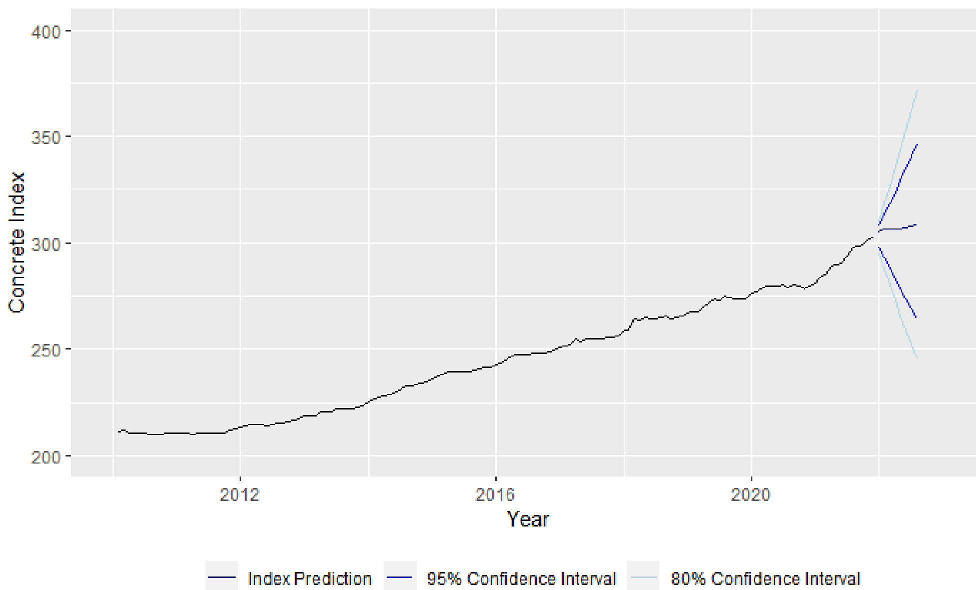

Figure 5.3: 8-Month Concrete Forecasts

Unlike either of the other two commodities, concrete has relatively stable changes in index values over the past decade, in terms of both scale and direction. The index increased slightly more than average during the COVID-19 pandemic, but not at the unprecedented levels of lumber or iron/steel. Because of this, the model forecasts similarly tame increases over the 8 months following the time series. The confidence intervals are wide for such a steady index, but the different rates of change seen near the end of the series in 2021 could be leading to this kind of spread.

Similar to iron/steel, the concrete model has thus far struggled to predict the way that the U.S. construction industry would change due to the war in Ukraine. Although the impact has been indirect, concrete prices have increased faster than anticipated by the model in the first few months of 2022. Quickly rising fuel prices have caused the transportation costs of ready-mix concrete, a commodity heavily dependent on local transportation within regions, to increase dramatically [15]. It should be mentioned, however, that although concrete prices have outpaced the model forecasts in reality, the error is not overly large, and the overall trend of increase throughout 2022 has so far been correct.

CHAPTER 6

Conclusion

In Chapter 1, The general structure of construction commodity time-series statistics was described, along with the motivations behind performing such an analysis. In Chapter 2, the data sourcing was explained and the data was then reformatted into stationary series in order to comply with the assumptions needed for later modeling. Chapter 3 saw the time series data assessed both visually and numerically, with seasonality plots showing that there was nothing to worry about in that regard. The AR spectrum for each of the three commodities was displayed on periodograms, which indicated potential cyclicity in the lumber and concrete series. Subtraction of the 12 monthly averages was performed on the two series in question, which controlled for any cycles present, and thus had the data meet the assumption needed to model a time series. Next in Chapter 4 the ARMA modeling method was explained, including how it is actually a combination of two separate model types, AR and MA, used to make time series forecasts. Various model selection methods were then explained, before all being applied to 25 separate ARMA models for each of the three commodities. Individually, the best ARMA models was selected for lumber, iron/steel, and concrete, based on a consensus optimal choice using the aforementioned model selection methods. In Chapter 5, the best model for each commodity was taken and used to make 8-month forecasts, beginning with the month directly following the end of the data series.

6.1 Final Conclusions

The lumber model has the highest point-to-point directional change continuity of the three commodities, based of the model coefficients. Its significantly positive first AR term indicates that the next observation in the series is more likely to move in the same direction that the previous observation did. It is possibly due to this that the model has so far been accurately predict the trend of lumber prices in 2022, including the early year rise that has been followed by a dip in the Spring.

The iron/steel model had much more trouble identifying small trends in the data, with next to no directional continuity being found between successive points. However, the steady, moderate rise seen over the course of long intervals of the index was modeled quite well, solidified by the fact that the $ARMA(0,0)$ white noise model performed terribly when modeling the time series. This trend showed up in the forecasts for iron/steel, whose slow, consistent increase has not aligned particularly well with the surging steel prices in reality caused by the Russian invasion of Ukraine, but does align with the general shape of the series.

Concrete was perhaps the easiest commodity to model, and the hardest to forecast, due to the prices of concrete in 2022 changing rapidly, in a market with a precedent of low volatility over the past few decades. The model forecasted the first 8 months of 2022 in a matter similar to that of the index before the COVID-19 pandemic: a consistent, mild rise with minimal variability in the magnitude of the monthly changes. While this may have turned out to be an accurate prediction in a peaceful economy, the spike in transportation costs due to the conflict in Ukraine have seen concrete prices rise much faster than their historical precedent for most of the Winter and Spring. While in practice, the concrete ARMA model has not done a very good job in forecasting the commodity index, it must be said that a variable the size of a multinational conflict affecting world trade could never had been accounted for without some level of precedent in the data.

6.2 Further Research

In performing the analysis of lumber, iron/steel, and concrete, three of the most vital construction commodities on the industry were modeled and forecasted. But there is much more to a project than specific materials, with aspects like the amount of contractors bidding for the right to a job, the expected inflation over the course of a build and the price escalation that comes along with it, and the local markets for manual and skilled labor where the construction is occurring. Any of these could in of themselves be a subject for a full time series analysis, either individually, or in tandem with commodity prices.

The various Producer Price Indices published by the U.S. Bureau of Labor Statistics, of which three are used in this analysis, are updated on a monthly basis. But many potential construction clients and advisors look for guidance on a more aggregate level such as quarterly or even yearly, especially on long-term megaproject builds that can take decades. To accomodate this, a separate analysis could be undertaken with data on a similarly aggregated scale, to identify possible trends that occur at a more macro level in terms of time.

While ARMA was the method of choice for modeling the indices in this analysis, there are a myriad of was to model a time series, including, but not limited to, exponential smoothing, generalized autoregressive conditional heteroskedasticity (GARCH), and neural network autoregression (NNETAR) [12]. In the future, some of these other methods could be applied to the same commodity index time series to compare forecasting accuracies and potentially find an optimal modeling method, or methods, for a specific set of data.

REFERENCES

[1] ABC CBI and CCI Methodology. https://www.abc.org/News-Media/News-Releases/categoryid/1061/Default. Accessed: 2022-04-24.

[2] Akaike information criterion. https://en.wikipedia.org/wiki/Akaike_information_criterion. Accessed: 2022-04-17.

[3] Architecture Billings Index (ABI). https://www.aia.org/resources/10046-architecture-billings-index-abi. Accessed: 2022-04-24.

[4] Producer Price Indexes. https://www.bls.gov/ppi/. Accessed: 2022-01-06.

[5] Producer Price Indexes Frequently Asked Questions (FAQs). hhttps://www.bls.gov/ppi/faqs/questions-and-answers.htm. Accessed: 2022-01-06.

[6] Root Mean Square Error (RMSE). https://c3.ai/glossary/data-science/root-mean-square-error-rmse/. Accessed: 2022-04-17.

[7] U.S. Producer Price Index. https://en.wikipedia.org/wiki/U.S._Producer_Price_Index. Accessed: 2021-04-16.

[8] Akaike's Information Criterion: Definition, Formulas. https://www.statisticshowto.com/akaikes-information-criterion/, 2015. Accessed: 2022-04-17.

[9] Richard Adaurhere, Innocent Musonda, and Chioma Okoro. *Construction Contingency Determination: A Review of Processes and Techniques*, pages 271–277. 01 2021.

[10] Jordan Berninger. Forecasting the Time Series of Apple Inc.'s Stock Price. Master's thesis, University of California, Los Angeles, 2018.

[11] George E. P. Box and Gwilym M. Jenkins. *Time Series Analysis, Forecasting and Control*. San Francisco: Holden-Day, 1970.

[12] Davide Burba. An overview of time series forecasting models. https://towardsdatascience.com/an-overview-of-time-series-forecasting-models-a2fa7a358fcb, 2019. Accessed: 2022-04-27.

[13] Will Daniel. Lumber prices have tumbled over 20% in March, but homebuyers shouldn't celebrate just yet. https://fortune.com/2022/03/24/housing-market-lumber-prices-fall-homebuyers-home-prices/, 2022. Accessed: 2022-04-24.

[14] IBM. Spectral Plots. https://www.ibm.com/docs/en/spss-statistics/28.0.0?topic=forecasting-spectral-plots. Accessed: 2022-04-16.

[15] Linesight. Construction Materials Prices Rose in Q4 2021 but Declines Expected in Q1 2022. https://www.forconstructionpros.com/business/press-release/22056153/linesight-construction-materials-prices-rose-in-q4-2021-but-declines-expected-in-q1 2022. Accessed: 2022-04-24.

[16] Charanraj Shetty. Time Series Models. https://towardsdatascience.com/time-series-models-d9266f8ac7b0, 2020. Accessed: 2022-04-17.

[17] Robert H. Shumway, David S. Stoffer. *Time Series Analysis and Its Applications*. Springer, 2016.

[18] Paul Teicholz. Forecasting final cost and budget of construction projects. *Journal of Computing in Civil Engineering - J COMPUT CIVIL ENG*, 7, 10 1993.

[19] Yoke Wong. Steel price forecast (2022): Market rally on supply woes. https://capital.com/steel-price-forecast, 2022. Accessed: 2022-04-24.