Forecasting Power Consumption in Pennsylvania During the COVID-19 Pandemic: A SARIMAX Model with External COVID-19 and Unemployment Variables

Jackson Au
Southern Methodist University, jau@mail.smu.edu

Javier Saldaña Jr.
Southern Methodist University, saldanaj@mail.smu.edu

Ben Spanswick
Southern Methodist University, bspanswick@mail.smu.edu

John Santerre
Southern Methodist University, jsanterre@mail.smu.edu

Forecasting Power Consumption in Pennsylvania During the COVID-19 Pandemic: A SARIMAX Model with External COVID-19 and Unemployment Variables

Jackson Au[1], Javier Saldaña Jr.[1], Ben Spanswick[1,2], and John Santerre, PhD[1]

[1] Master of Science in Data Science, Southern Methodist University, Dallas TX 75275 USA {jau,jsaldana,bspanswick,jsanterre}@smu.edu
[2] PPL Electric Utilities. 2 N Ninth St, Allentown, PA 18101
https://www.pplelectric.com/

Abstract. In this paper, we present how electrical consumption can reveal insight into the novel COVID-19 pandemic spread. We analyze electrical power consumption provided by PPL Electric Utilities, Department of Labor's unemployment claims, and the COVID-19 cases/deaths for the State of Pennsylvania to study the impact of the pandemic on the infrastructure. Using a SARIMA model as our benchmark and we analyzed the use of a SARIMAX model to forecast the power consumption in Pennsylvania 14 days ahead. Our work quantifies and illuminates the effect that the strict legislation passed to minimize the spread of COVID-19 had on power consumption. Most importantly, this study helps drive a greater understanding into the hidden cost of a global pandemic such as COVID-19.

Keywords: COVID-19 · Forecasting · Energy consumption · Unemployment · Economy · Pennsylvania

1 Introduction

With over 15 million cases and over 500,000 deaths across the world[21], COVID-19 is unprecedented in recent historical context for rapid global impact. The cost of the pandemic transcends biology and has the ability to bring infrastructure and the economy to the brink of collapse at a global scale. The stress of the pandemic on the U.S. infrastructure has forced state-level governments to enact draconian measures and issue mandatory stay-home orders in an attempt to mitigate the spread of COVID-19. With the engine of capitalism now forced to stay-home, many businesses responded to the severe reduction in revenue by reducing their cost, the workforce. The preservation of the health infrastructure was now stressing the economy. Being able to alleviate the stress allows governments the ability to create legislation that will mitigate the damage caused by pandemics.

It is estimated the H1N1 influenza pandemic of 1918 claimed at least 50 million lives[6], which is 10 million more than the global casualties of the First

World War. Since then, the world has experienced 3 influenza pandemics. On November 17, 2019, Chinese government records suggest a 55-year-old man from the Hubei province was the first person to have contracted the COVID-19[1]. On December 27, 2020 that Dr. Zhang Jixian from the Hubei Provincial Hospital of Integrated Chinese and Western Medicine advised Chinese health authorities the coronavirus they were dealing with was a novel strain.

On January 21, 2020, the state of Washington reported its first confirmed case[2]. The following week, the World Health Organization declares a global public health emergency as a result of COVID-19 and the U.S. responds by banning entry to foreign nationals who had traveled to China within 14 days[3]. Nearly a month later, on February 24, 2020, the Dow Jones began its decline and dropped over 1,000 points as fears of COVID-19 grew. By the end of March, the travel ban had extended to a substantial number of countries overseas, major sporting organizations (NCAA, NHL, MLB, NBA, NASCAR, Kentucky Derby, etc.) had either postponed or cancelled their seasons/events, elections were delayed, schools were closed, and local/state governments began to issue stay-at-home orders[4].

As Americans began to stay home, companies were forced to adapt. Over 30,000 brands surveyed reported revenues increased by an average of 65%[5] between March 14 and April 17, 2020. However, other major retailers already struggling were unable to overcome the additional hardships. As of July 20, 2020, 98 companies with at least 500 employees have filed for bankruptcy[6]. The U.S. airlines were given $25 billion to stay afloat during the pandemic under the condition they cannot fire employees until after September 30. Shortly after, United Airlines warned its employees it will downsize starting October 1 and to expect layoffs[7]. The auto industry began to sell their inventory online while

---

[1] See the "Coronavirus: China's first confirmed Covid-19 case traced back to November 17" article for more information at `https://www.scmp.com`. Last accessed 14 June 2020.

[2] See the "Coronavirus U.S. case confirmation: 1st coronavirus case is in Seattle" article for more information at `https://www.nbcnews.com`. Last accessed 14 June 2020.

[3] See the "China criticizes U.S. border closure as coronavirus death toll rises" article for more information at `https://www.nbcnews.com`. Last accessed 14 June 2020.

[4] See the "Coronavirus timeline: Tracking the critical moments of COVID-19" article for more information at `https://www.nbcnews.com`. Last accessed 14 June 2020.

[5] See the "Retailers Selling Non-Essentials See Double and Triple-Digit Increases In Online Sales During COVID-19 Crisis" post for more information at `https://www.forbes.com`. Last accessed 14 June 2020.

[6] See the "Coronavirus Bankruptcy Tracker: These Major Companies Are Failing Amid The Shutdown" post for more information at `https://www.forbes.com`. Last accessed 20 July 2020.

[7] See the "United Airlines Plans To Downsize Thousands Of Workers After Taking Billions Of Dollars From The Government" article for more information at `https://www.forbes.com`. Last accessed 14 June 2020.

2

maintaining their personalized approach through virtual meetings[8]. The industries across the economy continued to adapt, the workforce was also forced to evolve with the pandemic. Work-from-home became a new norm and the growing pains of virtual meetings were amplified by the massive influx of new users into the virtual environments[9].

In the wake of the COVID-19 outbreak, the global economy has been greatly impacted. From February 19, 2020 through March 19, 2020, the Dow Jones Industrial lost 33% of its value as the reality of the COVID-19 pandemic's economic impact set in. Since then, the US has experienced an increase in unemployment claims far higher than ever before with over 20 million jobs lost in April 2020[10]. The infrastructure was stressed as social distancing became the new norm and more people continued to perform their jobs in a virtual environment from home. In March, Netflix, YouTube, Amazon, and Disney all announced they would reduce the video quality of their services as they saw a substantial increase in demand that was crowding the digital infrastructure of many countries. By the end of April, the majority of the state governments in the US had issued stay-home orders[11]. Yet, the US Energy Information Administration (EIA) was predicting energy consumption to decrease as a result of the economic slowdown and stay-home orders[14].

With over 300 million Americans under stay-at-home orders during the month of April, talks of containment began to shift to plans of re-opening the economy. Market sentiment shifted with the public talks of reopening and the stock market reflected that sentiment. The Dow Jones, SP 500, and Nasdaq all increased between 12% - 15% during the month of April. Yet, unemployment claims continued to soar with 20 million claims filed during the same month. By the end of April, the unemployment rate was over 14%[12] while the stock markets continued to climb. In the month of May, several states across the country either modified or let their stay-at-home orders expire as the re-opening process began. As the country slowly re-opened, the federal government considered a second stimulus injection in order to jump start the economy. The unprecedented federal spend-

---

[8] See the "Coronavirus has dealerships moving to online sales, and car buying may never be the same" article for more information at `https://www.usatoday.com`. Last accessed 14 June 2020.

[9] See the "COVID-19 Forces More People To Work From Home. How's It Going?" article for more information at `https://www.npr.org`. Last accessed 14 June 2020.

[10] See the "Economists forecast an additional 2.4 million Americans filed for unemployment insurance last week" article for more information at `https://www.businessinsider.com`. Last accessed 14 June 2020.

[11] See the "See Which States and Cities Have Told Residents to Stay at Home" dashboard for more information at `https://www.nytimes.com`. Last accessed 14 June 2020.

[12] For more information about the historical unemployment rate visit `https://data.bls.gov/timeseries/LNS14000000`. Last accessed 14 June 2020.

3

ing spurred by the pandemic is projected to leave the U.S. with a federal deficit around \$4.3 trillion[13].

With many states projecting the worst has passed, the debate has now focused on how fast to re-open the economy without compromising the health of its citizens. Regardless, COVID-19 came at a high price and continues to take human and business casualties. Pennsylvania issued its stay-at-home order on April 1 and extended it to June 4, 2020, while only applying it to specific counties allowed to re-open under the state mandated guidelines[14]. Like every other state in the country, Pennsylvania was not immune to the virus and endured its own struggles.

For this study, we partnered with PPL Electric Utilities to provide insights into the electrical power consumption behavior of the state of Pennsylvania. The goal of this paper is to circumvent the shortfalls of epidemiological data and leverage the strength and reliability of data gathering systems already in commercial use to gain insight into the indirect impact of COVID-19. By combining the electrical power consumption of Pennsylvania with external variables such as the new COVID-19 cases and unemployment claims filed with the U.S. Department of Labor, we uncover insights of the impact COVID-19 is having across the infrastructure in Pennsylvania. We developed and tested a seasonal autoregressive integrated moving average with external variables (SARIMAX) model, utilizing power consumption, unemployment, and COVID-19 data in an effort to forecast power consumption 14 days ahead in a time of crisis. We compared its performance to benchmark models such as an ARIMA and SARIMA model in order to determine the effectiveness of the external variables with respect to forecasting. We anticipated power consumption will decrease during a global pandemic as the non-essential enterprises shut down and offset the increase in residential consumption, which is expected from a state-wide stay-home order. Our hypothesis is that the SARIMAX model will perform better than the SARIMA and ARIMA models by having the unemployment and COVID-19 variables to further explain the variance in the short term. The impact of this study is economic and crucial as a means of preparedness for future catastrophic events.

---

[13] See the "Coronavirus, CARES And PPP Will Explode The Federal Deficit And Debt" article for more information at `https://www.forbes.com`. Last accessed 14 June 2020.

[14] See the "Pa. extends coronavirus stay-at-home order to early June, as 24 counties move into yellow reopening stage" article for more information at `https://www.pennlive.com`. Last accessed 14 June 2020.

4

2   Related Work

2.1   Forecasting with External Data

The use of external data to enhance forecasting models is a well documented area of academia. While univariate models may explain some of the variance, researchers and practitioners are looking for explanatory variables in the form of external data to explain the remaining variance. For the purpose of this study, we will define external data as any data that crosses organizational boundaries[10, 16]. However, integrating external data can prove to be unreliable and does not always provide the most useful models. Furthermore, the inclusion of additional variables into a time series model complicates the model further and increases the opportunity for error which may affect the performance of the forecast.

In a study by Adreoni and Postorino, they sought to compare an ARIMA (p,d,q) model and ARIMAX (p,d,q) (X) model (also commonly referred to as transfer function methods) in an effort to forecast demand for air travel[3]. They combined airport traffic data and air passenger survey data in order to see the foot traffic at a macro-scale and be able to gain insight into the characteristics of the passengers. They applied a univariate ARIMA model and multivariate ARIMAX model to air travel demand and found that the univariate ARIMA model performed slightly better than the ARIMAX model. However, they do mention that both models performed at a satisfactory level, they believe the univariate ARIMA model only performed well since the data was stable and within its normal boundaries. The additional explanatory variables in the multivariate ARIMAX model would have been able to adapt to the extreme events.

In 1984, McCleary and McDowall conducted a study to forecast Swedish population growth by testing a univariate ARIMA model and a multivariate AMIRAX model using external data[18]. The ARIMA model was modeled using a population change calculator which subtracted the birth rate and the death rate. In the ARIMAX model, harvest yields were included with the population change. While the models produced statistically significant results, they came with their share of problems. The ARIMAX model only utilized 2 independent variables despite the recommendation from Box and Jenkins to use a minimum of 3. In addition, the assumption of dependency is violated since harvest yields may have an impact on the population change. In the years following low harvest yields, marriage rates and legitimate fertility rates decrease as well. In addition, the analysis found that fertile women were more likely to emigrate to other countries during time period of low crop yields. They proposed a vector ARIMA model in order to address these concerns.

Assumption violations are not the only cause of uncertainty for multivariate ARIMAX/SARIMAX models. As the models become more sophisticated it does not equate to it being more useful. Durka and Pastorekova compared the performance of an ARIMA model and ARIMAX model when forecasting gross domestic product per capita and added unemployment claims as the external variable[13]. They found that the ARIMA model performed better than the ARIMAX model by producing a lower root mean squared error. Despite

5

the slightly superior performance from the ARIMA model, the ARIMAX model still accounted for over 92% of the variance. The performance issues could be attributed to the observation that the final 4 quarters of the data for unemployment saw a sharp increase while the gross domestic product per capita remained unfazed. This additional noise could be a contributing factor as to why the ARIMAX model would have underestimated the gross domestic product per capita in the study.

Another study by Arunraj, Ahrens, and Fernandes, they proposed a SARIMAX would perform better than a standard SARIMA model when attempting to forecast the daily retail sales of perishable foods[4]. Aside from daily sales figures, the authors believed the addition of holidays, seasonality, and price reduction would help explain the variance from the SARIMA model. By adding these additional explanatory variables, the authors were able to improve their adjusted r-squared from 38% to 61% and are supported by a reduction in the root mean squared error. The authors disclose that the forecasting model was only as useful as the quality of the data and relied heavily on the forecasting horizon and data availability.

2.2 Energy and Pandemics

While there are extensive published reports and white papers regarding the threat of a pandemic and the preparedness of the energy sector, we were unable to find any public forecast modeling studies which touch this relationship in an in-depth scale or provide any statistical analysis. As a result, we are confident to be one of the few, if any, to produce literature for this subject.

According to the International Energy Agency (IEA), global energy consumption has decreased by 3.8%, with the expectation to possibly reach 6-8%[2]. Before governments issued stay-home orders, most people started to stay home when their local/state governments declared a state of emergencies. The increase in work-from-home resulted in a small increase in power consumption, which was then offset by the global stay-home orders issued across the world. Major energy consumers such as public arenas, schools, and offices were all forced to close and appeared to create a decline in electricity consumption. Shortly thereafter, the price of oil sharply fell with the demand. The IEA is expecting an annual decline of 8% in carbon emissions.

The assumption that residential demand would simply fill the void from commercial demand has not been supported by the data. Residential power consumption began to increase across the United States between March 20, 2020 and March 30, 2020 and peak consumption shifted from 9:00 AM and 7:00 PM with a dip between 10:00 AM and 4:00 PM to a steady increase starting at 5:00 AM and peaking at 6:00 PM[22]. However, the increase in power consumption was also softened by the increase in demand for alternative energy sources such as clean and renewable energy solutions[22]. The shift to alternative energy sources is one that can have a direct financial impact on large utility companies that are not prepared to adapt to the change in demand. As renewable energy sources are

6

entering more homes and helping reduce the cost of energy, the hardest impact against utility companies would result as commercial consumers begin to shift.

2.3 Energy and Economy

On June 8, 2020, the National Bureau of Economic Research announced the United States' economy is officially in a recession and has been since March 2020[19]. While there is strong evidence the COVID-19 pandemic was a strong factor in the cause of the recession, we can explore the relationship between recessions and energy by looking to the Great Recession in 2008.

The Great Recession of 2008 saw many families lose homes, jobs, and financial security. Not surprisingly, electricity consumption plummeted during this time as households were looking to curb their spending and save on bills as much as possible. A case study in North China by Huiri Zhao, et. al. found evidence of a strong correlation between economic growth and power consumption during the 2007 economic boom and 2008 crash[25]. The recession impacted almost every facet of society and China as a whole saw a 14% decline in power consumption. However, the seasonal flux was only slightly affected by the recession, which would indicate the seasonality of the data remained regardless of the level of consumption.

Since electric and utility data is readily available, this can be a great resource for determining the economic health for a particular country. If the data indicates that the consumption of energy is deviating from the norm, in an extreme way, this could indicate a major event is occurring. Put another way, if disease data or defaulting loan data were not available, electric and utility data could have been great alternative sources for predicting whether a major event was happening.

2.4 Economy and Pandemics

Unlike the energy sector, there are a considerable number of economic models that are utilized during a pandemic in order to gain insight into the economy after the pandemic. Historical reports are a great data source into the behavior of economies during pandemics. In the 1918 Influenza Pandemic, the draft from World War I had crippled the industrial sector by taking most able-bodied men to the war[15]. However, the reduction in available workforce and increase in demand positioned the industry to see wage growth. Based on the supply and demand economic principles, a reduction in available workforce and increase in demand should result in a driver of wages. The war created an increase in workforce demand while the draft and high casualties from the pandemic only drove wages even higher[15]. Brainerd and Siegler studied the impact of the pandemic on the state income growth for the decade that followed the pandemic[8]. They found that the states with the highest mortality rates also experienced the highest state income growth by 1930. The studies suggest the economic impact of the pandemic was short-lived and far too costly.

While wage and income growth provide insight into the income being generated, it is far from a strong measure of the economy. The gross domestic product

7

is the preferred indicator of economic health. The recent H1N1 outbreak in 2009 raised similar concerns and researchers began to study the economic impact of the pandemic by focusing on the gross domestic product. An equilibrium economic model developed using United Kingdom economic data from 2004 found that the United Kingdom's gross domestic product would shrink between 0.5% - 4.3%[23]. According to the model, in low-fatality scenarios with early mitigation, the loss of gross domestic product could be kept under 1%. However, in high-fatality scenarios with extensive mitigation, the loss of gross domestic product would be between 2% and 3%, which is half of the losses suffered from the 2008 financial crisis.

The downside to using the gross domestic product as an indicator is its scarcity with regards to data instances. The U.S. Government calculates the gross domestic product on a quarterly and annual basis. At the end of the fourth quarter of 2019, the real gross domestic product had increased by 2.1% from the previous quarter. By the end of the first quarter of 2020, the real gross domestic product had decreased by an annual rate of 4.8%[5]. Considering two data points is insufficient for a strong analysis, we resorted to unemployment claims as an indicator of economic health.

3 Time Series Methodology

In this study, our primary objective is to develop a mathematical model which is capable of describing the relationship of peak power consumption utilizing secondary data sources. Time series is a stochastic process, which means that a series of random variables are indexed by a mathematical set. In our data sets, that mathematical index is the date range and the random variables are the combination of the data sets' respective variables. By building a benchmark autoregressive integrated moving average model as our control and multivariate seasonal autoregressive integrated moving average model as our test, we can measure the performance of the model with root mean squared error (RMSE). In this section, we will address the foundation of the SARIMAX model and methods to test for assumptions of the model.

3.1 Autoregressive Integrated Moving Average (ARIMA)

The ARIMA (p,d,q) model was developed by Box and Jenkins where p is the autoregressive term, q is the moving average term, and d is the non-seasonal difference[7]. ARIMA (p,d,q) model derives from an ARMA (p,q) model. An ARIMA model is an autoregressive integrated moving average model and contains dependency on previous observations and innovations terms[11] with a component of trend in the data. As a result, ARIMA models can defined as Equation 1[11, 9, 1]:

$$\phi_p(B)(1-B)^d Z_t = c + \theta_q(B)\varepsilon_t \qquad (1)$$

8

Where $\phi_p(B)$ is the autoregressive operator, B is the delay operator, $(1-B)^d$ is the differencing operator, Z_t is an instance at time t, c is the constant, $\theta_q(B)$ is the moving average operator, and ε_t is the residual error

3.2 Seasonal Autoregressive Integrated Moving Average (SARIMA)

Like an ARIMA (p,d,q) model, SARIMA also derives from an ARMA (p,q) model and contains the same components as an ARIMA (p,d,q) model with an additional element of seasonality [11]. In order to be able to model SARIMA models, the trend and seasonality must be removed through differencing. As a result, SARIMA models may be defined as Equation 2:

$$\phi_p(B)\Phi_P(B^S)(1-B)^d(1-B^S)^D Z_t = \theta_q(B)\Theta_Q(B^S)\varepsilon_t \qquad (2)$$

Where in addition to the definitions provided earlier, $\Phi_P(B)$ is the seasonal autoregressive operator, S is the length of the season, $(1-B)^D$ is the seasonal differencing operator, and $\Theta_Q(B)$ is the seasonal moving average operator.

3.3 Multivariate Seasonal Autoregressive Integrated Moving Average (SARIMAX)

The SARIMA (p,d,q) (P,D,Q) model is useful for univariate analysis but not ideal for our study. In order to incorporate the unemployment and COVID-19 data, we must consider a seasonal autoregressive integrated moving average with external variables (SARIMAX) (p,d,q) (P,D,Q) (X). Unlike SARIMA which only conducts univariate analysis, the SARIMAX was developed by Tiao and Box in 1981 and has the capability of incorporating regressors in order to understand what may appear to be outliers in the time series' behavior[24]. As such, the SARIMAX model can be defined as Equation 3:

$$Y_t = \beta_0 + \beta_1 X_{1,t} + \beta_2 X_{2,t} + ... + \beta_k X_{k,t} + \omega_t \qquad (3)$$

Where Y_t is the dependent variable at time t, $/beta_0, /beta_1, ..., /beta_k$ are the regression coefficients of k external variables, $X_{1,t}, X2, t, ..., Xk, t$ are the instances of k external variables in time t, and ω_t is the residual which could be calculated as Equation 4:

$$\omega_t = \frac{\theta_q(B)\Theta_Q(B^S)}{\phi_p(B)\Phi_P(B^S)(1-B)^d(1-B^S)^D}\varepsilon_t \qquad (4)$$

3.4 Assumptions

The principal assumption of any of the ARIMA models provided is that the model is stationary, which means the mean and the variance must be constant[7]. This assumption can be tested using the Dickey-Fuller test which tests whether a variable has a unit root[12]. The Dickey-Fuller test fits the model using ordinary

9

least squares but faces a serial correlation problem. In order to control for serial correlation, the augmented Dickey-Fuller test was developed can be defined as Equation 5:

$$\Delta y_t = \alpha + \beta y_{t-1} + \delta t + \sum_{j=1}^{k} \zeta_j \Delta y_{t-j} + e_t \qquad (5)$$

The null hypothesis of the Dickey-Fuller test is that the variable has a unit root or if beta = 0. If evidence of a unit root exists, the model must be modified further to achieve stationarity. However, there are well documented issues that have surfaced on the low power of the Dickey-Fuller test. Nelson and Plosser applied the Dickey-Fuller test to 14 annual economic U.S. time series data sets and in all but 1 they failed to reject the null hypothesis[20]. Even then the data has a clear deterministic trend, the test was unable to detect the trend. As a result, an alternative test with greater sensitivity was developed by Kwiatkowski–Phillips–Schmidt–Shin and it became known as the KPSS test[17]. The test checks for stationarity by searching for a unit root as seen in Equation 6:

$$x_t = r_t + \beta_t + \varepsilon_1 \qquad (6)$$

4 The Data

4.1 COVID-19

The COVID-19 data is sourced from the New York Times and used in their COVID-19 Tracker dashboard published on their website[15]. The New York Times gathered the reports from local and state health agencies in order to compiled their data set. The data is compiled and posted to their Github, which was then made accessible to the general public. There are three data sets, each with its own level of granularity. Considering our interest for this study was in Pennsylvania, we utilized the county-level data set. The data set contains daily reports from all states and their respective counties. There are disclaimers that are addressed in the data preparation step such as the inclusion of probable related death in the tally under specific days. The data set was updated daily and maintained by the New York Times.

As of June 19, 2020, the data set contains over 236,000 data entries and ranges from January 21, 2020 through June 15, 2020 (as of this publishing). It has 6 columns, which are described in Table 1:

---

[15] See The New York Times for more information at `https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html`. Last accessed 14 June 2020.

10

Table 1. Column description of COVID-19 data.

Column	Description
Date	The date of entry for the cases and deaths
Country	The country in which the data is being reported
State	The state in which the data is being reported
FIPS	The Federal Information Processing Standard code
Cases	The cumulative number of cases on the given date
Deaths	The cumulative number of deaths on the given date

4.2 Transmission Peak Power Consumption

The transmission peak power consumption data was provided by PPL Electric Utilities[16]. Daily peak power consumption is a metric that assesses total power used on a given day across a bulk power network. The industry defines transmission network as the grid that moves wholesale power from generators to distributors, typically above 69kV in power. Peak power consumption can be measured on a variety of scales, with the peak power consumption of the United States electrical grid ranging from approximately 400 - 700 gigawatts (GW) at any given point of time. As much of the electrical grid is controlled by a network of Supervisory Control and Data Acquisition Control System Architecture (SCADA) devices, at any given time, utility networks can aggregate values and provide a snapshot of current electrical demand on their network. Once recorded this data is reported to the Public Utility Commission (PUC) as well as the Federal Energy Regulatory Commission (FERC) on a real time basis in order to ensure the reliability of the US electrical grid.

For the purpose of this study, a subset of the electrical grid were used. This subset represents approximately 1.4 million customers in Pennsylvania and captures the bulk electric needs of various Northeast US cities, such as New York, NY and Philadelphia, PA. The sample size for this data contains aggregated values of peak power consumption over the past five years and is represented in megawatt hours (mWh). The data set is also paired with the high, low, and average temperature for the given day across the grid. The data set has a total of 1896 entries and spans from March 19, 2015 to May 26, 2020. The structure of the data set is further illustrated by Table 2:

4.3 Unemployment Claims

The unemployment claims data was obtained from two government sources. In order to match the level of detail in our other data sets, we obtained daily unem-

---

[16] To learn more about PPL Electric Utilities, please visit `https://www.pplelectric.com`

11

Table 2. Column description of Transmission Peak Power Consumption data.

Column	Description
Date	The date of entry for the power consumption
Load	The aggregated peak load on the given day in megawatt hours (mWh)
Min	The minimum temperature in Fahrenheit
Avg	The average temperature in Fahrenheit
Max	The maximum temperature in Fahrenheit

ployment claims filed from the United States Department of Labor[17]. The data set contains entries starting in March 15, 2020 and running through June 4, 2020. In each entry, there was only the number of claims filed and the given day of the week. The data structure and description is illustrated by Table 3.We attempted to obtain daily unemployment claims going as far back as January 1, 2015 but could not find a reputable source. In order to account for this shortcoming, we opted for a blend of data sources between daily and weekly.

Table 3. Column description of Daily Unemployment Claims data.

Column	Description
Day	The day of the week
Date	The date of entry
Total New Claims	The number of claims filed on given date

The weekly number of unemployment claims filed were obtained from the United States Bureau of Labor Statistics[18], which spans from January 1, 2015 through June 1, 2020. The United States Bureau of Labor Statistics provides the unemployment totals for each county in Pennsylvania, which are then aggregated to a weekly number. The data set contains great insight into the ongoing active claims, and aggregated counts. Table 4 provides a column description of the data set in its raw format.

In order to be able to conduct thorough analysis, we decided to blend these data sources, while giving the daily filed claims priority over the weekly. More

---

[17] To learn more about the U.S. Department of Labor, please visit `https://www.dol.gov/`

[18] To learn more about the U.S. Bureau of Labor Statistics, please visit `https://www.bls.gov/`

12

Table 4. Column description of Weekly Unemployment Claims data.

Column	Description
State	Name of state
Filed Week Ended	The final date of that week
Initial Claims	The number of new claims filed for that week
Reflecting Week Ended	Date of previous week
Continued Claims	Rolling count of total active claims
Covered Unemployment	The total count of people who are employed
Insured Unemployment Rate	Estimated unemployment rate for that week

information about the blending process will be covered in the exploratory data analysis section.

5 Exploratory Data Analysis

5.1 COVID-19

The raw data for the New York Times COVID-19 data set was taken and filtered to show only instances for Pennsylvania. Even then, it had over 40,000 instances in its current structure. As a result, we created a pivot table and were left with 136 columns and over 100 date entries. The columns represented each county's respective cases or deaths at the given date. Each column had at least some null values and when checked, those null values came from the pivot table not being able to attribute values before cases and/or deaths started occurring in that county. As a result, we replaced all null values with 0, which would equate to the starting count for all counties with regard to the tally of cases and deaths.

In addition to reformatting the data, we also created new measures out of the current data. By transposing the entries for each county and aggregating them, we calculated the number of new cases and new deaths on a given date. As a result, the final structure of the data is demonstrated in Table 5:

Table 5. New Data Structure for COVID-19 Cases and Deaths.

Column	Description
Date	Date of record
New cases	Number of new cases reported compared to prior period
New deaths	Number of new deaths reported compared to prior period

By using these measures, we gained a greater understanding into the spread of the virus within the state of Pennsylvania. In Fig. 1, the number of new cases

13

exponentially grows and then begins to decline. The figure also illustrates new deaths followed a similar lagged behavior. The vertical dashed lines pertain to the start and end of the stay-home order that was issued by the governor in Pennsylvania. The order lasted for a period of 2 months and appears to have been successful at reducing the spread considerably.
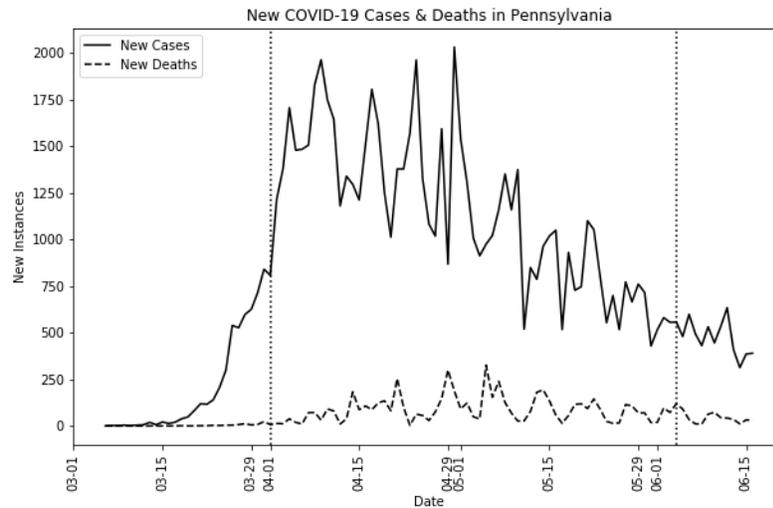

Fig. 1. The figure illustrates the number of new COVID-19 cases and deaths across the State of Pennsylvania.

For the purpose of this study, we only used new cases. We removed new deaths because of an interdependency issue among the variables. Most of the time, a death recorded is a prior case recorded which could amplify the predictions. The other columns only showed aggregations and did not provide any meaningful insight and were dropped as a result.

5.2 Transmission Peak Power Consumption

The power consumption data set was imported as is in order to protect its integrity. The data set did not include any nulls and was only converted to a time series data set. As Fig. 2 illustrated, there were clear seasonality issues that had to be addressed. This was expected in power consumption since the standard household consumes less power during business days than they would be during the weekend. This became evident within the annual seasonality of the data which may be explained by the weather. While the data set came with partial temperature measures, it only covered a small portion of the data set in

14

2020. Because of its large missing data, we removed it from the data frame and only kept peak power consumption.
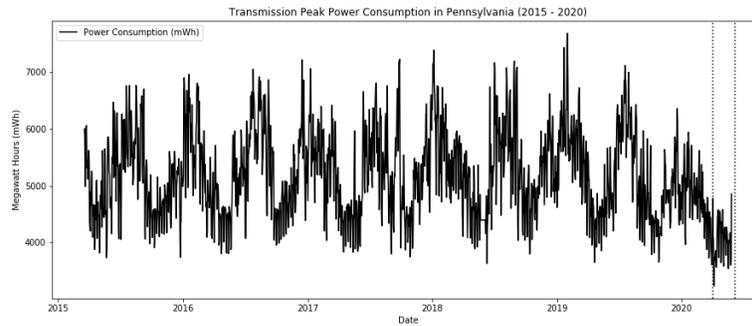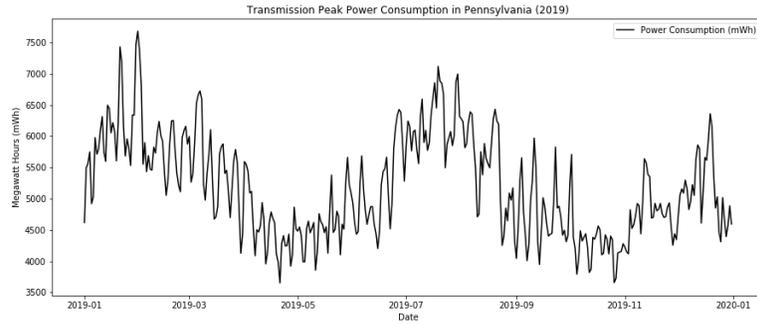

Fig. 2. The figure illustrates the transmission peak load power consumption across the subset power network from 2015 - 2020.

Regardless, Fig. 3 demonstrates the seasonality of the data set in an annual view from 2019. It is visible that power consumption increased during the height of winter and summer months. This is common across the United States as extreme temperatures force residents to cool off inside during the summer or warm up inside during the winter. By recognizing this trend, it would appear that COVID may have created an anomaly initially but that event quickly stabilized and the consumption proceeded with its seasonal path. In Fig. 4, the drastic impact COVID-19 had on the power consumption prior to the state-wide stay-home order is clearly illustrated. Yet, it appeared to level off towards the end. The transitions from Spring to Summer and Fall to Winter see the biggest dips in power consumption.
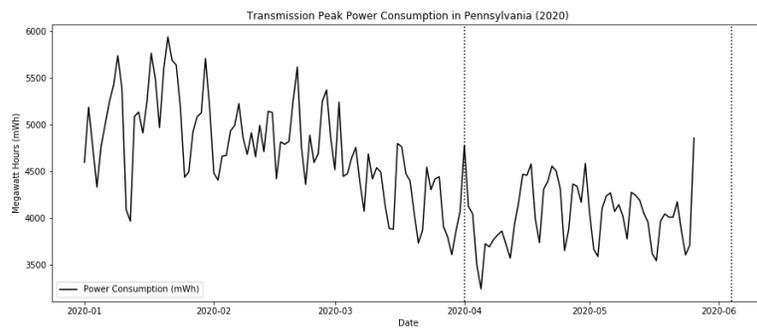
5.3 Unemployment Claims

Considering the data sources have different level of detail, we blended them together at the daily level. We extrapolated the weekly data of unemployment claims and filled in the same values for the days of the week when the report was made. We understand an alternative would have been to roll up the data sets to weekly levels in order to preserve the integrity of the data. We decided against that approach because of the scarcity of COVID-19 data. Rolling up the data to weekly would have only given less than 10 instances of COVID-19 data.

The final blended data for unemployment claims demonstrate a stable history that is greatly impacted by COVID-19. The average daily unemployment claims reported were 2,229 prior to COVID-19. Fig. 5 illustrates the large spike in unemployment claims once COVID-19 started having a local impact on business in Pennsylvania.

15

Fig. 3. The figure illustrates the transmission peak load power consumption across the subset power network 2019

Fig. 4. The figure illustrates the transmission peak load power consumption across the subset power network in the year 2020.

16
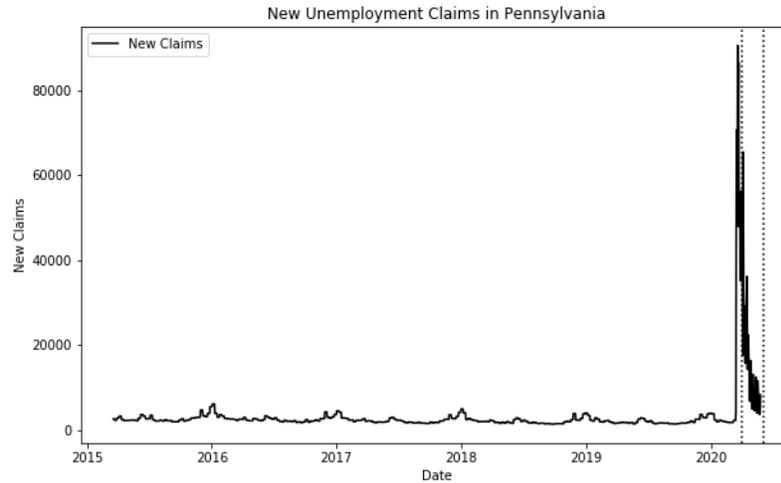
Fig. 5. The figure illustrates the number of unemployment claims filed in Pennsylvania (2015 - 2020).

Despite the market's reaction to COVID-19 in February when the stock market indexes began to plummet, the number of unemployment claims held steady until the free fall in the stock market finally hit a bottom. From that point on, as the stock prices and COVID-19 cases began to increase in Pennsylvania, so did the unemployment claims. In fact, Fig. 6 demonstrates the economy began to layoff employees well before the stay-home order was issued.

6   Time Series Analysis

In order to test the effectiveness of adding the external variables, we decided to use the standard ARMA model as our benchmark and compare the root mean squared error of the models. Once we determine which model produces the lowest root mean squared error, we then forecast the next 14 days using the ideal model. The data set was divided into three new sets. The training set contains data from March 19, 2015 through April 28, 2020. The testing set contains data from April 29, 2020 through May 12, 2020. The validation set contains data from May 13, 2020 through May 26, 2020. This allowed the model to capture a considerable amount of power consumption before and during the COVID-19 pandemic. Furthermore, we predicted the final two weeks of the data set and validate that prediction.

17
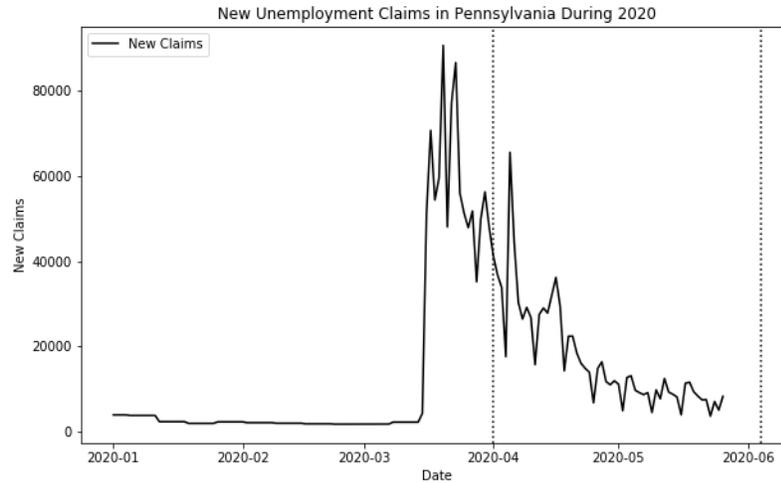
Fig. 6. The figure illustrates the number of unemployment claims filed in Pennsylvania 2020.

6.1 Model Identification

In the ARMA benchmark, we only utilized the power consumption data as a univariate analysis. The power consumption data shows clear signs of seasonality but questionable signs of stationarity (see Fig. 7). The power consumption peaks goes through a cycle on a weekly and seasonal basis. There is greater power consumption during the weekends when compared to weekdays. In addition, we see an overall increase in power consumption during summer and winter months as people use the heaters or air conditioning units. The lags in Fig. 7 also demonstrate the seasonality and the questionable stationary behavior. We can see a slow dampening pattern in the sample autocorrelation frequency, which may be misleading because of the size of the cycles. In order address the question of stationarity, a Dickey-Fuller test and KPSS test was performed to determine if there is a unit root in the data set.

6.2 Parameter Estimation

We addressed the current violation of stationarity by conducting the Dickey-Fuller test and the KPSS test. The Dickey-Fuller test indicates the data is stationary at a max of lag 12 and produces a p-value less than 0.01. However, the KPSS produces a test statistic of 0.67, which is less than the critical value for a significance level 1% but greater than the significance value for significance level of 2.5% at 12 lags as well. By differencing the data by 1, we are then able to produce a test statistic of 0.0056, which is well below the critical values

18
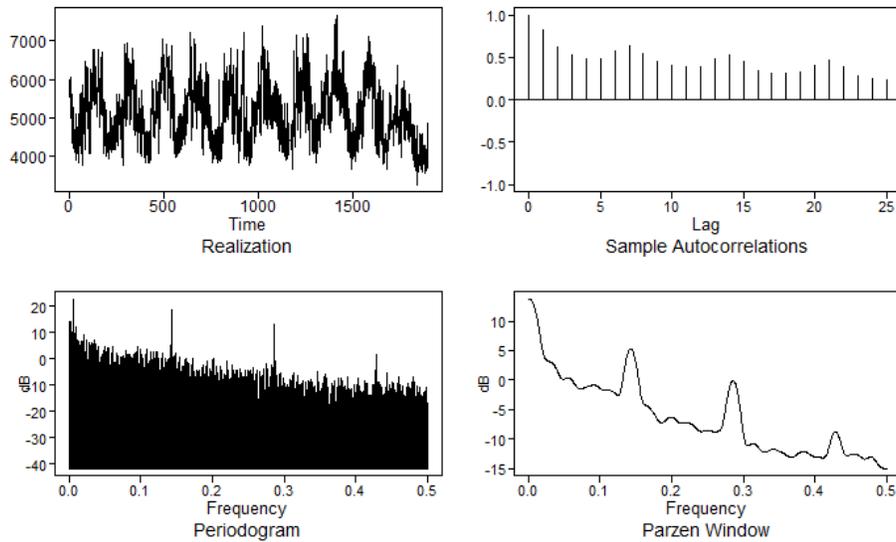
Fig. 7. The figure contains four plots, which demonstrate the questionable stationary behavior and seasonality in the raw data.

for significance levels less than 10%. After differencing the data, the KPSS and Dickey-Fuller test produce p-values under 0.01.

Differencing the data removed the trend from the series but failed to remove the seasonality. The sample autocorrelations from the newly differenced series showed evidence of oscillating behavior (see Fig. 8) after every 7 days. This was considered supporting evidence to the industry standard regarding the daily power consumption cycle most individuals engage in when consumption peaks in the weekends and drops during the week. The lag at every 7 days shows evidence of the seasonality. We addressed this by transforming the data and by removing the seasonality from the data. When the data was transformed by 7 (s = 7), there was minimal evidence of seasonality left in the series (see Fig. 9).

ARIMA We performed a stepwise parameter search for the lowest Akaike information criterion (AIC) using the differenced data with the seasonality still present (d = 1) to determine the autoregressive order of p and moving average order of q in the model. The output from the parameter search indicated the power consumption data was an ARIMA (5,1,3) model. The AIC value of the model order was 27664.

SARIMA/SARIMAX We performed a parameter search as well to determine the autoregressive order of p and the moving average order of q with the trend (d = 1) and seasonality (s = 7) removed Using the lowest Akaike information

19
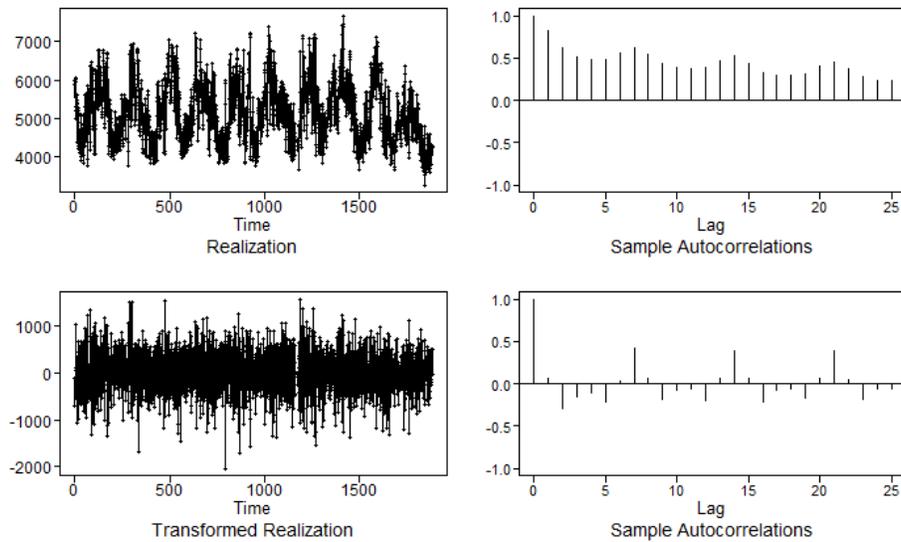
Fig. 8. The figure contains realization and sample autocorrelations before and after the data was differenced by one (d = 1).
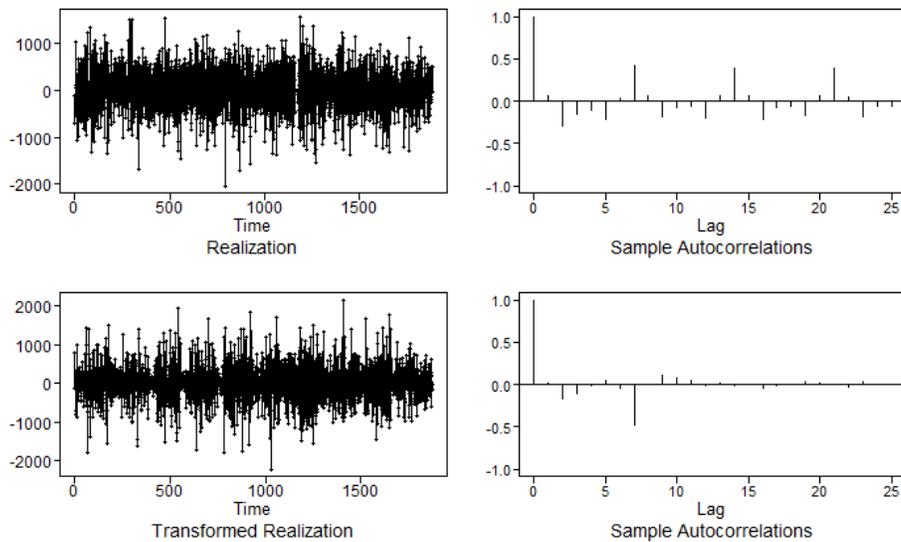
Fig. 9. The figure contains realizations and sample autocorrelations of the data before and after the seasonal transformation (s = 7)

20

criterion (AIC). The output from the parameter search indicates the ideal model order of p and order of q are 1, 1. Based on the AIC, the model order used for the seasonal model was SARIMA (2, 1, 1)(1, 1, 1, 7), where p = 2, d = 1, q = 1, P = 1, D = 1, Q = 1, and S = 7. The AIC generated from the order was 27232.

Inclusion of External Variables The SARIMAX model inherited the same orders from SARIMAX. With the train data set ranging from March 2015 through May 13, 2020, the external variables were also split under these same parameters. Although the SARIMAX model is robust against dependency violations, there is an interdependence concern that is worth mentioning for the study. We are aware unemployment claims, new COVID-19 cases, and new COVID-19 deaths are strongly correlated but are confident the SARIMAX model is able to produce a useful model nonetheless.

7 Results

We used the root mean squared error (RMSE) as the deciding metric to diagnose the fitness of the models. Once the most useful model was identified, that model was used to forecast 14 days out from the latest date, which created a power consumption prediction from May 27, 2020 through June 9, 2020.

7.1 Diagnosis of Fitness

Using the orders for ARIMA and SARIMA, we fit the models in order to test their fitness. The models were tested on peak load power consumption from May 13, 2020 through May 26, 2020. It was evident from the output that the SARIMAX model appeared to have the greatest error (see Fig. 10). This was supported by its high RMSE and poor r^2 metric (see Table 6). From the results, it became apparent that the SARIMA model performed better than the ARIMA and the SARIMAX models.

Table 6. Model Performance Measures

Model	RMSE	r^2	MAE
ARIMA (5,1,3)	243.13	0.087	196.11
SARIMA (2,1,1)(1,1,1,7)	170.52	0.551	151.64
SARIMAX (2,1,1)(1,1,1,7)	293.34	-0.331	234.33

The performance improvement from the ARIMA to the SARIMA model is expected. Earlier we identified the seasonal element in the data and by incorporating it into the model, it clearly improved its performance. However, the

21

addition of the unemployment data and COVID-19 cases appears to have worsened the model's performance. Nonetheless, the SARIMA model is selected and used to forecast the final two weeks of the power consumption.
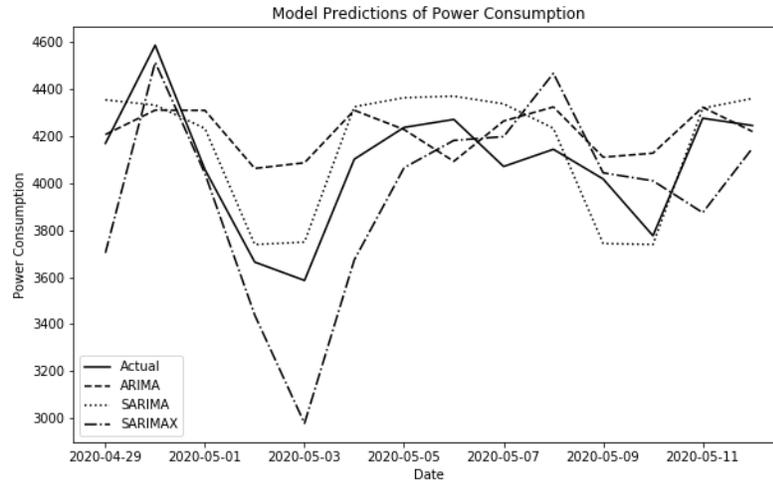
Fig. 10. The figure contains predictions of peak power consumption by the ARIMA, SARIMA, and SARIMAX for the test period

7.2   Forecasting and Validation

The SARIMA model was used to forecast the peak power consumption from May 13, 2020 through May 26, 2020. During this time period, the pandemic was still present, stay-home orders were still in effect, and unemployment claims were still setting all-time high records, which we were arguing have a impact on the consumption. Yet, the weekly seasonality was still intact and continues to fluctuate through the week. Forecasting through the SARIMA model, we can see the model predicts higher than the output trending downward with a slight spike at the end (see Fig. 11). Compared to the previous test, the SARIMA model generated a much higher RMSE of 308 and r^2 of 0.581. The increase in both RMSE and r^2 can be attributed to the slight downward trend exhibited by the true realizations. The spike at the end of the realization then leads to an increase in the r^2 value but a decrease in the RMSE value.

22
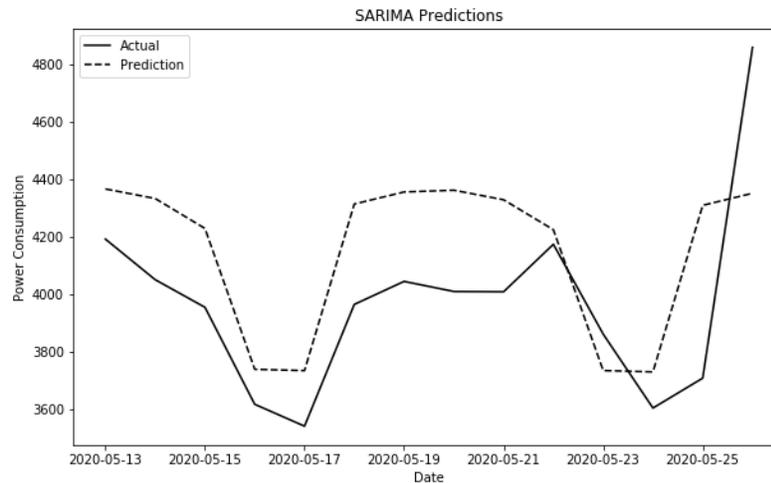
Fig. 11. The figure contains predictions of peak power consumption by the ARIMA, SARIMA, and SARIMAX for the test period

8 Conclusion

Based on this study, we felt there was sufficient evidence that the SARIMA model outperformed the SARIMAX model. The findings in this study were critical since they provided insight into the behavior of consumers in the State of Pennsylvania. The increase in COVID-19 cases and unemployment claims don't appear to be great indicators of power consumption. In fact, when added into the model they increase the noise and only worsened its performance. There are certainly other means that we believe would help improve the model. By adding standard industry variables such as weekend indicators and holiday indicators, utility companies are able to model the power consumption using stable external variables. However, those variables alone extend before and after times of crises. While the data from the crisis itself may not be as influential as we original believed, there is another element that could have a greater impact on power consumption. In this instance, we believe using the calendar variables and categorically mapping when the state-wide stay-home order and public state of emergency declaration were initiated would serve as a greater predictor.

While the COVID-19 pandemic was having a catastrophic impact all over the world, the United States' was catapulted to the top of the list in respect to new cases. While some Pennsylvanians took precautions to mitigate the impact, others proceeded with their everyday lives. When the State of Pennsylvania declared a state of emergency, its residents began to stay home on a voluntary basis. The commercial industries felt the decrease in traffic and responded by downsizing in an effort to get ahead of the coming policy. A few weeks after,

23

Pennsylvania declared a stay-home order and forced the remaining residents to stay home. The 2 weeks between the state of emergency and stay-home order is when the power consumption saw its initial spike. Once the stay-home order was put in place, the decrease in commercial consumption caused a downward trend in power consumption throughout. While the fear of the COVID-19 pandemic drove the stock market to a significant crash in March, it was the reaction from the local governments that had the true impact on the power utility grid.

8.1 Future Study

While results generated failed to support our hypothesis, there is still a considerable amount of future work mandated by this subject. First and foremost would be the exploration of the stay-home order and/or the declaration of the state of emergency. The emergency declaration curbed behavior of cautious residents while the stay-home order forced all others to remain home. Other external variables to consider would be the distinction between commercial and residential power consumption and how both differed in a similar analysis. Being able to separate the residential and commercial power consumption will remove the offset and clarify some of the aggregated noise currently present in the data. In addition, expanding this study to other geographic areas would also help to provide a useful model that could withstand a generic application across a broader region. Although outside of the immediate scope of this study, we found that incorporating the weather, weekend/weekday, and holiday data as external variables helped the SARIMAX model outperform the SARIMA model. We believed this to be the case because those variables help reinforce the current seasonality in the data itself and help mitigate the black swan event created by the COVID-19 pandemic.

8.2 Ethical Considerations

Considering the data that was used and the model implications, we are well aware of the ethical responsibilities that stem from a study like the one presented. The COVID-19 data is sourced from various self-report government entities. This creates the potential to taint the data through reporting bias since the data is sourced from agencies that self-reporting. The Center for Disease Control has stopped reporting COVID-19 data[19] and the Trump Administration has discouraged hospitals from making COVID-19 data publicly available. Actions and censorship as such jeopardize the integrity of the data and, as a result, the integrity of the model.

Most importantly, the model's output has the potential to provide misguided insights into the power consumption market. The output of the model could

---

[19] See the "Disappearance of covid-19 data from CDC website spurs outcry" article for more information at `https://www.washingtonpost.com`. Last accessed 6 August 2020.

24

result in a reallocation of resources, which could cause more harm than good. Improper aggregation or manipulation of the power consumption data coupled with incomplete COVID-19 data creates opportunities in which the model has adverse effects than it was intended. In order to ensure we hedge against these issues, we sourced COVID-19 data from a reputable source and worked closely with our business partner to ensure the power consumption data was handled properly.

9   Acknowledgement

References

[1]   L. Aburto and R. Weber. "Improved supply chain management based on hybrid demand forecasts". In: *Applied Soft Computing* 7 (1 2007), pp. 136–144.

[2]   International Energy Agency. *Global Energy Review 2020: The impacts of the Covid-19 crisis on global energy demand and CO2 emissions*. Paris, Apr. 2020. URL: `https://www.iea.org/reports/global-energy-review-2020` (visited on 06/14/2020).

[3]   Alberto Andreoni and Maria Nadia Postorino. "A Multivariate ARIMA Model to Forecast Air Transport Demand". In: *Association for European Transport and Contributors* (2006).

[4]   Nari Arunraj, Diane Ahrens, and Michael Fernandes. "Application of SARIMAX Model to Forecast Daily Sales in Food Retail Industry". In: *International Journal of Operations Research and Information Systems* (7 2016), pp. 1–21. DOI: `10.4018/IJORIS.2016040101`.

[5]   BEA. *Gross Domestic Product, 1st Quarter 2020 (Advance Estimate)*. Bureau of Economic Analysis, Apr. 2020. URL: `https://www.bea.gov/news/2020/gross-domestic-product-1st-quarter-2020-advance-estimate` (visited on 06/14/2020).

[6]   Jessica A. Belser and Terrence M. Tumpey. "The 1918 flu, 100 years later". In: *Science* 359 (6373 Jan. 19, 2018), p. 255. DOI: `10.1126/science.aas9565`. URL: `https://science.sciencemag.org/content/359/6373/255` (visited on 03/04/2020).

[7]   George E. P. Box and Gwilym M. Jenkins. *Time Series. Analysis: Forecasting and Control*. San Francisco: Holden-Day, 1976.

[8]   Elizabeth Brainerd and Mark Siegler. "The Economic Effects of the 1918 Influenza Epidemic". In: *C.E.P.R. Discussion Papers* 3791 (Feb. 2003). URL: `https://cepr.org/active/publications/discussion_papers/dp.php?dpno=3791`.

25

[9] M. Cools, E. Moons, and G. Wets. "Investigating variablility in daily traffic counts using ARIMAX and SARIMA (X) models: assessing impact of holidays on two divergent site location". In: *Journal of the Transportation Research Board* (2136 2009), pp. 57–66. DOI: `10.3141/2136-07`.

[10] Thomas H. Davenport. "Competing on Analytics". In: *Harvard Business Review* (Jan. 2006). URL: `https://hbr.org/2006/01/competing-on-analytics` (visited on 06/09/2020).

[11] Marcel Dettling. *Applied Time Series Analysis*. Zürich, Switzerland: Zurich University of Applied Sciences, 2008.

[12] D. A. Dickey and W. A. Fuller. "Distribution of the estimators for autoregressive time series with a unit root". In: *Journal of the American Statistical Association* 74 (1979), pp. 427–443.

[13] *ARIMA vs. ARIMAX – which approach is better to analyze and forecast macroeconomic time series?* 30th International conference on Mathematical Methods in Economics (Silesian University, Sept. 11, 2012). Silesian University in Opava, School of Business Administration in Karviná, 2012, pp. 136–140. ISBN: 978-80-7248-779-0.

[14] EIA. *Short-Term Energy Outlook (STEO)*. June 2020. URL: `https://www.eia.gov/outlooks/steo/pdf/steo_full.pdf` (visited on 06/15/2020).

[15] Thomas A. Garrett. "Pandemic Economics: The 1918 Influenza and Its Modern-Day Implications". In: *Federal Reserve Bank of St. Louis Review* 90 (2 Mar. 2008), pp. 75–93.

[16] Angela Hoskisson and Gloria Phillips-Wren. "An analytical journey towards big data". In: *Journal of Decision Systems* 24 (1 2015), pp. 87–102. DOI: `10.1080/12460125.2015.994333`.

[17] Denis Kwiatkowski et al. "Testing the null hypothesis of stationarity against the alternative of a unit root". In: *Journal of Econometrics* 54 (1992), pp. 159–178.

[18] Richard McCleary and David McDowall. "A Time Series Approach to Causal Modeling: Swedish Population Growth, 1750-1849". In: *Political Methodology* 10.3 (1984), pp. 357–375.

[19] NBER. *Determination of the February 2020 Peak in US Economic Activity*. June 8, 2020. URL: `https://www.nber.org/cycles/june2020.html` (visited on 06/14/2020).

[20] C. R. Nelson and C. I. Plosser. "Trends versus random walks in macroeconomic time series: Some evidence and implications". In: *Journal of Monetary Economics* 10 (1982), pp. 139–162.

[21] World Health Organization. *Coronavirus disease (COVID-19) Situation Report*. 145. June 13, 2020. URL: `https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200613-covid-19-sitrep-145.pdf` (visited on 06/13/2020).

[22] Sense. *Home Energy Consumption Trends as Covid-19 Triggered Shutdowns Across the Country*. Apr. 2020. URL: `https://drive.google.com/file/d/13nNc_Q6In-5espEDExurPtWMzh1hnjTE/view` (visited on 06/14/2020).

26

[23] Richard D. Smith et al. "The economy-wide impact of pandemic influenza on the UK: a computable general equilibrium modelling experiment". In: *C.E.P.R. Discussion Papers* 3791 (Feb. 2003). URL: https://cepr.org/active/publications/discussion_papers/dp.php?dpno=3791.

[24] G. C. Tiao and G. E. P. Box. "Modeling Multiple Times Series with Applications". In: *Journal of the American Statistical Association* 76.376 (Dec. 1981), pp. 802–816.

[25] Huiru Zhao et al. "The Impact of Financial Crisis on Electricity Demand: A Case Study of North China". In: *Energies* 9 (Mar. 2016), pp. 1–13. DOI: 10.3390/en9040250.

27