12-2008

Comparative analysis of forecasting techniques with intermittent demand

John DeForest

University of Arkansas, Fayetteville

University of Arkansas

Department of Industrial Engineering

Comparative Analysis of Forecasting Techniques with Intermittent Demand

Honors Thesis

By: John J. DeForest

Advisor: Dr. Manuel D. Rossetti

Submitted 12/5/08

1.Introduction

The world and life are filled with uncertainty. Statistics, and more specifically, forecasting techniques allow us to quantify uncertainty and make decisions based on that information. Many find forecasting advantageous in areas such as predicting consumer demand, stock prices, terrorist attacks, epidemiology, etc. This thesis will focus on forecasting intermittent demand, which is the study of analyzing sporadic demand. One application includes an airplane manufacturer's sporadic or intermittent demand for spare parts in their distribution center. Since spare parts are not needed on any regular schedule, it is thought of as intermittent. Knowing the projected quantity for a given time period will allow manufacturers to plan accordingly with suppliers, transportation logisticians, procurement agents and the like. Also, for long lead time items this information is especially helpful to ensure a high service level for the customer. Imagine if a customer orders a specific part and the lead time is nine-hundred days, without forecasting the customer will have to wait the entire lead time. By using intermittent demand forecasting techniques, on hand stock levels can be planned to decrease the likelihood that a customer has to wait for the part.

There are many forecasting techniques that have been created, though few specifically for intermittent demand. Intermittent demand is characterized by a time series that has a large percentage of time periods with zero demand and for which the time between non-zero demands is highly variable. Like most of forecasting, future demand is defined to be unknown within the positive real domain. Moreover, future demand is defined as erratic, that is, the demand pattern has large variability relative to its mean [Silver, Pyke, Peterson, 2004]. Experiencing a demand spike will necessitate action from the manufacturer or supplier causing longer lead times and potentially unsatisfied customers. Intermittent demand is difficult to quantify and consequently model for many reasons. Chiefly, zero demand periods give no information as to the historical structure of demand through time. For instance, one cannot easily analyze trends, seasonality, mean demand and so forth for demand equal to zero. Philosophically without information as to the structure of an entity, no model can ever be constructed. This is evident in the quantification of intermittent demand. Another reason for difficulty includes the demand's erratic nature leading to spikes, which in statistics are often discounted and labeled as outliers. Hence, quantification reduces to a model of exceptions, which is not very satisfying. Two prominent techniques are: Revised Croston's method and

Box-Jenkins model. This thesis will compare the effectiveness of applying these two techniques to intermittent demand.

First the Box-Jenkins model will be introduced. Relationships among customer demand are like many things in this world, dependent. That is to say, the current demand or lack thereof is highly dependent on a previous one. Most of elementary statistics assumes that observations are independent of one another. Although this is not always true, it makes for a simpler model. Observational dependence necessitates a more complex model, but allows for better prediction. In this case future customer demand is assumed to be a function of the past demand and unaccountable error. In the experiment, the dependence within the data will be exploited using the Box-Jenkins model, in hopes of a better prediction.

Secondly, the Revised Croston's method statistically quantifies the future demand by estimating the probability of occurrence of demand and of the size of demand when it occurs [Croston, 1972]. This model has different assumptions from Box-Jenkins, including independent and identically distributed occurrences and homoskedasticity. Different assumptions are representative of different perspectives of reality, of which the modeler considers important. Conclusions from this paper should deduce which assumptions are most appropriate for such a phenomena as intermittent demand.

2. Motivation

Motivation for such a project lies within the human inability to tell the future and discerning whether or not the past dictates the future. If the past does influence the future, which I believe is true, then there is a window of intersection between the recent past and the near future as in *Figure 2.1* below.

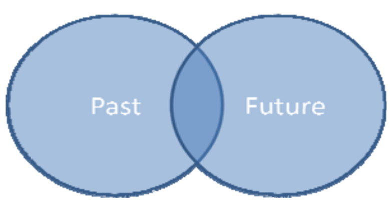

Figure 2.1: Venn Diagram of Forecasting Possibility

Practically, it is advantageous for humans to know the future in order to properly prepare for demand and help make better business decisions. Moreover, more philosophically, if we can perfectly predict the future, then the past dictates the future, whereby freedom as we know it does not exist. These implications of forecasting are indeed tremendous. I will conclude this section with an insightful passage from E.F. Schumacher's book *Small is Beautiful*.

> "When the Lord created the world and people to live in it – an enterprise which, according to modern science, took a very long time – I could well imagine that He reasoned with Himself as follows: "If I make everything predictable, these human beings, whom I have endowed with pretty good brains, will undoubtedly learn to predict everything, and they will thereupon have no motive to do anything at all, because they will recognize that the future is totally determined and cannot be influenced by human action. On the other hand, if I make everything unpredictable, they will gradually discover that there is no rational basis for any decision whatsoever and, as in the first case, they will thereupon have no motive to do anything at all. Neither the scheme would make sense. I must therefore create a mixture of the two. Let some things be predictable and let others be unpredictable. They will then, amongst many other things, have the very important task of finding out which is which."

This quote captures the intimate relationship between life and forecasting. If Schumacher is correct in his assertion about some things in life cannot be predicted, then it necessarily follows that true randomness exists. Hence God himself could not predict the outcome of an unpredictable event. In any case, whatever the implications are for the development of forecasting, one thing is for certain, they are all thought-provoking, seemingly paradoxical and truly mysterious.

3. Background

Related statistics overview

Before discussing ARIMA forecasts, I will review linear regression and moving averages. First, linear regression is a model between a set of observations say $\{X_1,X_2,\ldots,X_n\}$ and their responses given by *y*. In the event that there is a linear relationship between the observations and the responses or a relationship that can be transformed into a linear one (maybe by means of *mathematical* transformations (e.g. natural logarithms), one can fit a model called linear regression that will determine an expected value and the respective errors of the model.

Second, the moving average is the arithmetic average with respect to time. So, if one is to forecast the future of the number of students demanding to take differential equations, then the prediction will just be the mean of the historical data for some time period, say five years. The next year, the forecast using a moving average will be the average of the five most recent observations. The simple moving average equation can be found below:

$$F_t = \frac{1}{N} * \sum_{i=t-N}^{t-1} D_i = \frac{1}{N}(D_{t-1} + D_{t-2} + \cdots + D_{t-N})$$

where F_t is the forecast for time *t*, D represents demand for that time period and N is the number of observations taken.

The ARIMA model combines linear regression and moving average to attain a robust model that better predicts the future than one using just regression or just moving averages. Picture lines in a two-dimensional plane moving with time, updated for each interval of time. These are the elements of the ARIMA model. Now for the model itself.

ARIMA model

ARIMA models are based on three components: the autoregressive model, the moving average process, and the autocorrelation function. First-order differencing will be used to combine these components into a cohesive model. Briefly, is a function of the sum between the mean and the autocorrelation function, which will be discussed further below (Box & Jenkins, 1976). The below equation states that the forecasted value is equal to the weighted average of the deviations from , where is observed value at time *t*. Also, the error, , is normally distributed with mean zero and variance σ^2. Initially, we will handle the autoregressive model given by the following:

where

The autoregressive model is similar to the regressive model discussed in the *Related statistics overview* above with the exception that the autoregressive model will update the regressive model each time more information is attained. Informally, updating is just an adjustment to the model which means in most cases either the slope or intercept or both will change slightly with new data. In addition to the autoregressive process, the moving average process is also necessary for ARIMA computation (Hamilton, 2004). The moving average process is given by:

$$\widetilde{z_t} = a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \cdots - \theta_q a_{t-q}$$

Above, is the difference of the errors at a specified observed time (Box & Jenkins, 1976). Moreover, the errors are weighted appropriately for the data being modeled. One can geometrically picture this equation as points above and below the x-axis in a two-dimensional plane representing error analogous to error distributed normally with mean zero. See *Figure 3.3* below:

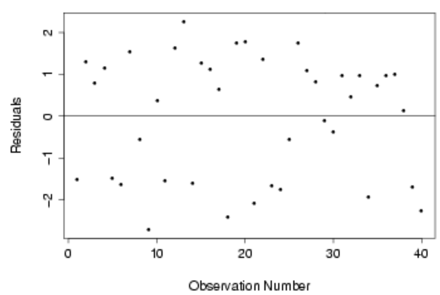

Figure 3.3: Residuals Graph with Approximately Zero Mean

When the two equations are integrated using first order differencing one gets:

$$\widetilde{z_t} = \varphi_1 \widetilde{z_{t-1}} + \varphi_2 \widetilde{z_{t-2}} + \cdots + \varphi_p \widetilde{z_{t-p}} + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - \cdots - \theta_q a_{t-q}$$

The above will be the model used for experimentation (Box & Jenkins, 1976).

Autocorrelation function

One important aspect of the ARIMA model is the autocorrelation function. As we will see, this function exploits the interdependencies among the data, which in turn yields a more realistic model. The function will be defined by:

$$\rho_k = corr(y_t, y_{t-k}) = \frac{E[(y_t - \mu)(y_{t-k} - \mu)]}{E[(y_t - \mu)(y_t - \mu)]}$$

Above reads that the dependency among data, ρ_k, is a function of two responses within a lag, *k* (Heiji, De Boer, Frances, Kloek, & Van Dijk, 2004). In the experiment, responses will be demand, which basically means that the demand series has statistical dependence. In general, one must realize the degree of dependency, denoted by *k*, one would like to model. For example, for positive correlation high values of the response *y* at time *t-1* will yield high values of *y* at time *t* and similarly for low values. It does not have to be the case that the degree of correlation is unity, for if one assumes that low values are followed by low values five time periods later then the degree (lag) of correlation will be five (Nahmian, 2005). For a graphical representation of autocorrelation process please refer to *Figure 3.4* below.

Deleted:

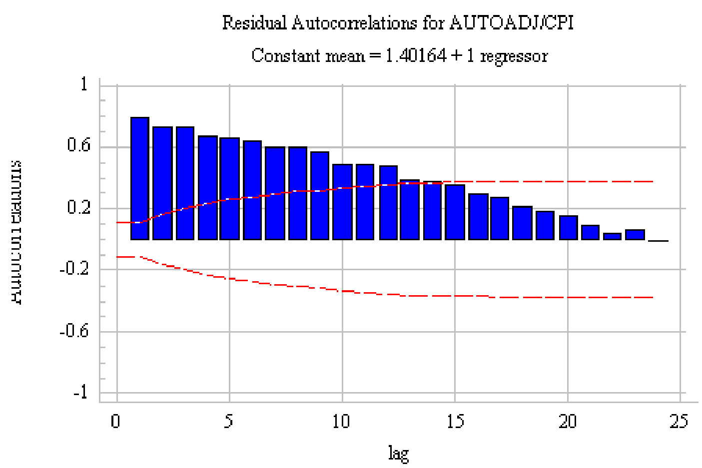

Figure 3.4: Example of Autocorrelation Function

Lastly, the autocorrelation function is very important because the product of the autoregressive function and the autocorrelation:

$\varphi(z) * \psi(z) = \theta(z)$, respectively.

From above we know that z is just the result of the ARIMA model after all the weights are applied to responses.

Revised Croston's Method

The Revised Croston's Method was developed by Croston and Syntetos. It models intermittent demand by modeling the inter-arrival interval and the size of demand realized and accommodating for randomness in the process. Like most forecasting methods, the Revised Croston's Method takes historical information to make a prediction for the next time period.

Croston's original method is based on the following ideas: If a transaction does not occur then: (1) transaction size estimates are not updated and (2) estimated value of the number of periods since the last transaction is equal to the previous one. Though, if a transaction occurs then updating occurs. The following constitutes Croston's forecasting method. F is the forecast for time t, $\hat{z}_t$ is the estimated average demand size in period t, z_t is the transaction size at time t, $\hat{n}_t$ is the estimate of number of periods till next transaction, n_t is the number of periods since last transaction, α_z is the smoothing constant associated with the average transaction size, α_n is the smoothing constant associated with the inter-demand period and α is the adjusted value for the forecast.

$$\hat{z}_t = \alpha_z z_t + (1 - \alpha_z)\hat{z}_{t-1}$$

$$\hat{n}_t = \alpha_n n_t + (1 - \alpha_n)\hat{n}_{t-1}$$

$$F_t = \frac{\hat{z}_t}{\hat{n}_t}(1 - \frac{\alpha}{2})$$

As seen in the Experimental Analysis section, $\alpha_z = \alpha_n$ = α, but later that constraint was relaxed. In the event that $\alpha_z \neq \alpha_n$, then we will assume the adjusted α = α_z. This is a valid assumption since the forecast is one of demand size and α_z is the smoothing constant associated with the size of demand.

The model, developed especially for intermittent demand, measures the time elapsed since the last non-zero observation to update the probability of demand occurrence. In order to elucidate Croston's theory a small example is give below with the following data set.

Croston Example

Period	Demand
1	1
2	0
3	2
4	0
5	1
6	0
7	1
8	2
9	1

Figure 3.5: Sample Demand Data

First, one must initialize the forecast using some technique. Croston does not give any methodology for initializing such forecasts, so we approached initialization by choosing the first three non-zero demands, summing them and dividing by the number of periods with non-zero demand. In the example, the first non-zero demands are {1, 2, 1} and they are in three non-zero periods. Hence, the initialization is:

$$\frac{1}{n}\sum_{i=1}^{n} D_{NZ}$$

$\overline{D}_{NZ}$ – average non-zero demand size over n

n – number of non-zero demands to initialize (e.g. 3)

Following the progression of the initialization, the result would be 4/3 = 1.33 yielding the estimate for the average transaction size. Consequently, all transactions sizes before period 5 will be zero. After initialization, one can find *n*, which is 2, because the last non-zero demand before 5 was at period 3, so 5-3 = 2. Now, one can find a forecast for the period 5+1, by using the equation and plugging in the variables yielding:

$$F_t = \frac{\hat{z}_t}{\hat{n}_t}\left(1 - \frac{\alpha}{2}\right) = \frac{1.3333}{2}\left(1 - \frac{0.3689}{2}\right) = 0.5437$$

Continuing in this fashion we can iteratively calculate transaction sizes and forecasts and ultimately the error. Below is a summarized table of the results assuming α = 0.3689.

Period	Demand	Average Transaction Size Estimated	Estimated Number of Periods till next Transaction	Forecast	Error
1	1	0.0000	0	0	0
2	0	0.0000	0	0	0
3	2	0.0000	0	0	0
4	0	0.0000	0	0	0
5	1	1.3333	2	0	0
6	0	1.3333	2.00000	0.54370	0.54370
7	1	1.2104	2.00000	0.54370	-0.45630
8	2	1.5017	1.63110	0.49356	-1.50644
9	1	1.3166	1.39829	0.75083	-0.24917

Figure 3.6: Forecasting Calculations Example

The example above differs from Croston's original model, which was later revised by Syntetos. Originally, Croston's method did not handle the statistical bias whereby increasing alpha values led to an overestimate of the average demand per period. This bias led to a forecast with larger error than desired. Syntetos corrected Croston's original model to accommodate for the bias and renamed the model to the Revised Croston's Method. By incorporating the alpha in the average demand per period, Syntetos was able to stabilize the estimate for the average demand per period. The method used in the experimental analysis below is the Revised Croston's Method.

4. Experimental Analysis

To compare both methods a sufficiently large number of data sets were selected, 49, having the intermittent demand property. Both models will then take a series of data representing a stock keeping unit (SKU) and step through the data to make a forecasts based on the past. The difference between the forecasted and observed values will constitute the error and is the metric to assess which model is superior, if any in the case they are equal. A sampling of error using the paired t-test determines if the two error populations are statistically equivalent. For the purposes of the experiment the mean squared error (MSE) will be used to assess the error between the two forecasting techniques and is seen on the next page.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n} e_i^2 = \frac{1}{n}\sum_{i=1}^{n}(\hat{X}_{i+1} - X_i)^2$$

n-is the number of observations

e-is the forecast error

$\hat{X}_{i+1}$ – is the forecast for period $i+1$

X_i - observed demand

Upon receiving the result from the t-test, model superiority will be deduced and presented.

Given a set of random time series, the revised Croston's method was used to forecast the next step ahead forecast, when updating occurred. As seen above, the forecast depends on two parameters or "smoothing constants", one for the size of the demand and the second for the time between non-zero demands. These parameters were first constrained to equal each other and later that condition was relaxed. In order to evaluate both Croston and ARIMA, the best forecast that could possibly be produced from either was analyzed. Henceforth, the experimental analysis included the use of the parameters optimized such that it minimized the MSE. As seen below, it is no surprise that relaxing the equality of parameters resulted in a lower MSE.

After the revised Croston's method was optimized the time series were analyzed using an ARIMA model. The ARIMA model is such that its three parameters are discrete and relatively small, between zero and five, which shrinks the search space of the "optimization." Due to the complexity of ARIMA, an optimization was not created as in Croston's method but rather an educated guess at the best parameters was made. Even though this is not optimal, the findings of the experiment are interesting.

The data used in the experiment originated from the United State Navy and met the criteria for intermittency. Each data set, representing a SKU, had 127 periods where the SKU was recorded as either experiencing demand or not. Table 4.1 presents a select number of demand statistics that provide evidence towards the fact that indeed these series are intermittent. For example, the variance is greater than the mean, which is one indicator that the data could be intermittent. Also, a high probability for zero demand occurrence (0.565) was found. These indicators, as well as the other found in the appendix, are all signs of intermittent data.

Table 4.1 Average values of demand statistics

Mean	Variance	NonZero Mean	NonZero Variance	Interval b/t NonZero Demand Mean	Interval b/t NonZero Demand Variance	P-value for zero demand
0.7711	1.38	1.74	1.41	2.36	6.82	0.565

Histogram and summary statistics of this data can be found in the Appendix in Figure 8.1 to 8.7, found in the order listed in Table 4.1.

5. Results

One important note before reporting the results; when analyzing the intermittent demand series, one series in particular stood out and the time series plot is given in Figure 5.1.

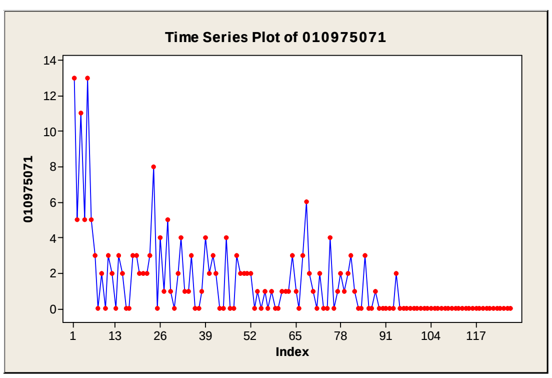

Figure 5.1: Plot of Possible Non-representative Intermittent Demand Item

Although the above may technically be defined as intermittent for its numerous zeros and erratic nature, I would disagree. Keeping in mind the entity of which the experiment analyzes, SKU, I argue that this SKU merely was discontinued and therefore should not be considered as intermittent. This is true, because in the long run, if one knows that a SKU is discontinued, then it will be easy to forecast the demand, namely zero. Therefore in an effort of completeness and soundness I will present the results first with the above time series and then without. As one will see later, the outcome does not change, but rather the confidence of the conclusions can be improved.

Now, a brief overview of each model's error is presented. After parameter optimization for Croston's model and a guess-and-check methodology for ARIMA parameters the MSE were computed and below a box plot is found depicting those MSE values for each forecasting technique.

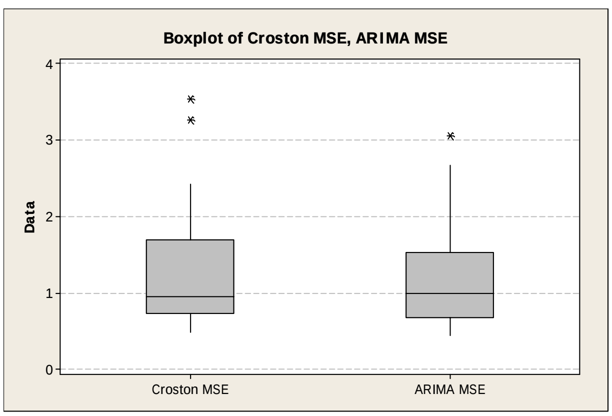

Figure 5.1: Boxplot Comparison of MSE between Croston and ARIMA

One can see that the ARIMA model has a smaller span between the 1st and 3rd quartile, though it is interesting to note that Croston had a lower median. Croston's method realized a median of 0.9529 and ARIMA 0.991. Also, notice the lower bound whiskers of the ARIMA model are barely beating out the Croston at 0.4523 compared to 0.4834. These box plots help us to understand the distributional properties of the MSE. A dot plot of the distribution of errors is presented in Figure 5.2. Without the questionable time series above the box plots change slightly, but Croston still has a lower median at 0.9328 compared to 0.975. As one can see, with negative differences favoring the ARIMA model, the dot plot has a higher frequency of negative numbers as compared to positive.

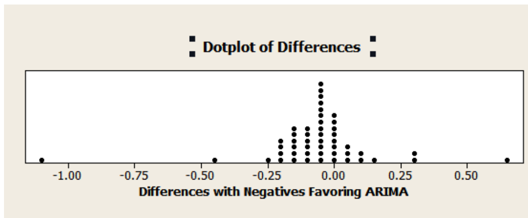

Figure 5.2: Dot Plot of Error between Forecasting Techniques

One important aspect in forecasting using either Croston's Method or ARIMA is their respective parameters. Parameter selection is key to a statistically good forecast. Since the ARIMA model takes only discrete values for parameters, the mode of the forecasting method was taken and reported below. For the complete list of the ARIMA parameters used in the experiment, please refer to Table 8.1 in the Appendix.

Table 5.1: Mode of ARIMA Parameters

	Autoregressive	**Difference**	**Moving Average**
Mode of ARIMA Parameters	**1 and 4**	**1**	**1**

Hence, if one must unknowingly guess which parameters to use in the ARIMA model, one should use 1,1,1 or 4,1,1. For Croston the parameters are different, in that they are real numbers with a range between 0 and 1 and summary statistics are found in Figure 5.2. One can almost certainly conclude that a forecaster can realize better forecasts if one does not assume equality between alphas associated with average transaction size and time between non-zero demand. This statement was verified using a paired t-test, which found an extreme statistical difference with a test statistic equal to 4.97. Moreover, one can deductively prove that two possibly unequal alphas will always be equal to or better than two equal alphas, when optimizing MSE in the manner in which it was performed in the analysis.

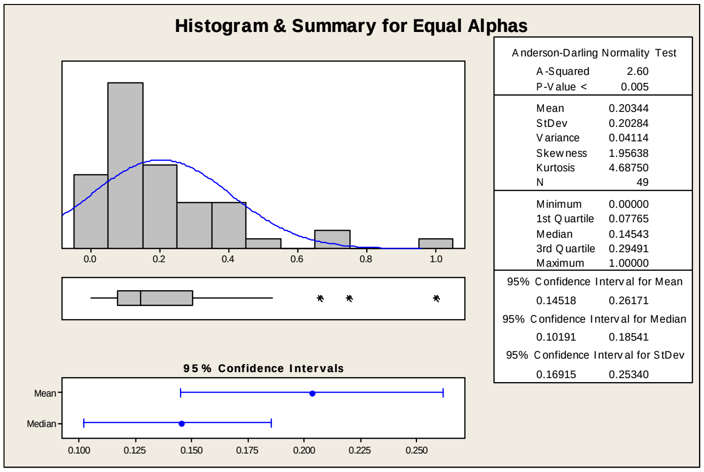

Figure 5.2: Descriptive Statistics for Croston's Method with Equal Alphas

Now, relaxing the constraint with equal alpha the histograms yields Figures 5.3 and 5.4.

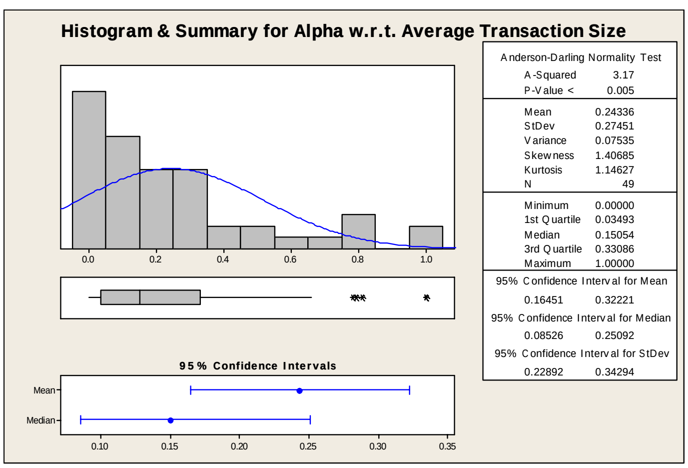

Figure 5.3: Descriptive Statistics for Croston's Method for Transaction Size Smoothing Constant

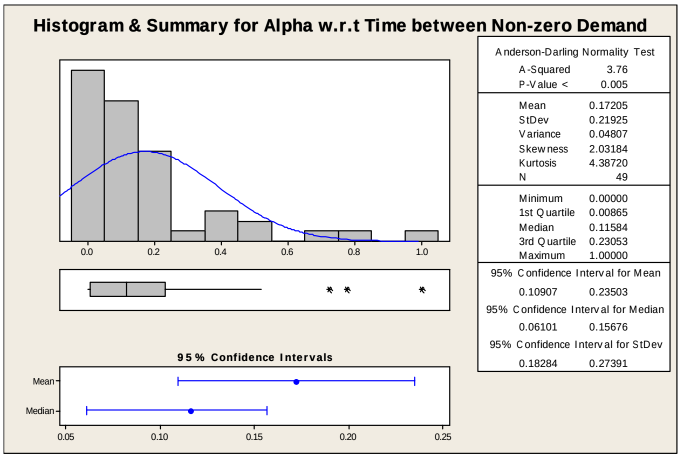

Figure 5.4: Descriptive Statistics for Croston's Method for Time between Non-zero Demand Smoothing Constant

Although the alpha used to calculate the expected value of the time between non-zero demands fell between the recommended bounds, namely 0.1 and 0.2, the alpha used to estimate the expected transaction size estimate did not. Rather, optimized values were on average 21% higher than the maximum recommended bounds. Moreover, the standard deviation of the alpha covered a large area relative to their domain. The user should be cautious in choosing the value for alpha.

After comparing the two forecasting techniques, ARIMA and the revised Croston's method, a t-test found a significance difference with an alpha level of 0.05. The interval was (0.0001, 0.1261), which does not include zero and therefore there exists a significant statistical difference between the two methods. Without the questionable intermittent series, which was forecasted with a much greater accuracy with the revised Croston's method, is the evidence is even more in favor of ARIMA over Croston's method. When the time series was taken out and the t-test revealed that with 99% confidence the ARIMA model is superior to the Croston with an interval of (0.0018, 0.1536).

Lastly, in order to validate the t-test assumptions, plots of the MSE differences were graphed to check for independence, trend and normality and can be seen in Figures 5.5, 5.6 and 5.7, respectively.

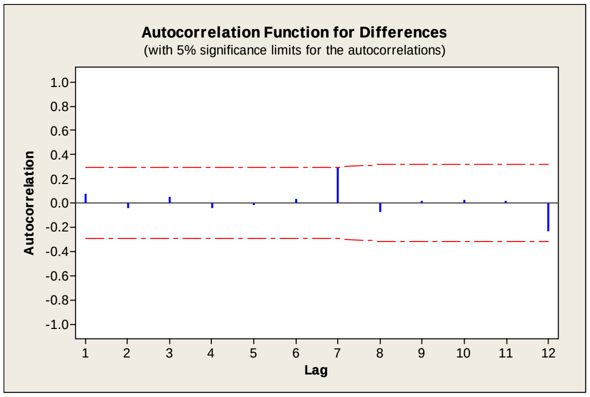

Figure 5.5: Autocorrelation Graph of MSE Differences

An autocorrelation such as this suggests that the data is independent and therefore validates the corresponding t-test assumption. As for Figure 5.6, with the exception of the one outlier found at index 34, the time series plot seems to have little to no trend in data. Likewise, the check for normality seems to be reasonably linear, although somewhat off in the tails of the distribution. From these plots, the independent and identically distributed normality assumptions appear reasonable. In addition, with a sample size of 49, it seems reasonable to conclude that the central limit theorem can be applied.

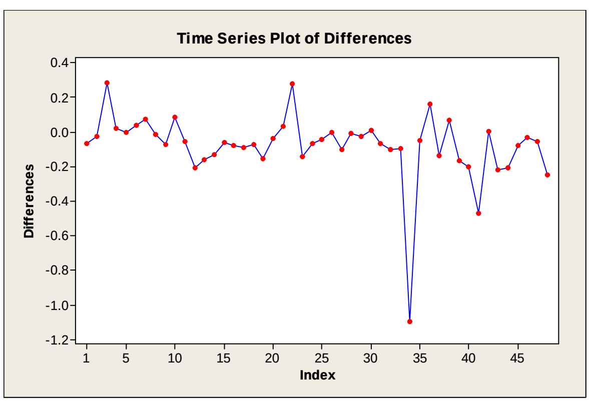

Figure 5.6: Plot of MSE Differences

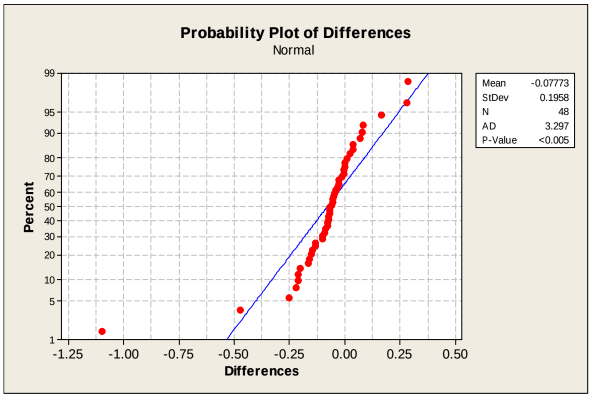

Figure 5.7: Plot of MSE Differences

6. Conclusion and Future Work

An experiment was performed where 49 SKUs were analyzed. ARIMA models were used to analyze the data and predict the future using a forecast, which yielded an approximate customer demand for the next business day. An appropriate conclusion to this preliminary research is that ARIMA was a superior forecasting model for intermittent demand; however, further research is needed. In both models, Croston and ARIMA, one still faces the problem of estimating the parameters.

Of course, forecasting is elementary when one knows *a priori* what the parameters should be. The challenge in everyday life is finding those magical parameters that minimize error. In any case, one can conclude that in the case of this research that the ARIMA model was competitive with Croston's method for intermittent demand. One potential reason as to why ARIMA turned out on top is due to its ability to account for the interdependence among observed data. Lastly, upon completing this experiment, I have provided comments on forecasting and the decision making in inventory system for future areas of study.

After reading Croston's and Syntetos' research paper I was initially satisfied with their model, but it quickly subsided. Syntetos revised Croston's method to reflect the corrected error and ended with an unbiased equation. Even though the calculations are now unbiased, it does not in most decision-making-processes make any difference because it does not discriminate between different kinds of error.

To expand Croston's method I would like to propose two ideas that will make the model even more robust than it currently is, namely discriminate against types of errors and calculate anticipated demand. I will qualify both, but first the error. When analyzing inventory, we should not be so quick to assume that an error in the positive direction is equal to an error in the negative direction, because frankly they do not have the same consequences. The MSE ultimately provides one with smoothing constants that take the absolute value of the error in the calculations. The reason that this is unacceptable is because engineers will make different decisions based on different errors.

Allow me to elucidate. If the service level is of the utmost importance then one should likely choose the forecasting technique that overestimates the demand and therefore sufficiently supplies the client. Though, if cost is the most important then underestimating demand would be the technique one should choose. But, in both policies, service level and cost, they still assume the errors are the same. Which is untrue! To accurately portray error in light of decision-making-processes it should be translated into a monetary value. The inventory policy that minimizes cost will inherently do this, but the service level policy will not. Accordingly, the error should be transposed to a monetary figure that would reflect the penalty for not fulfilling demand. Obviously, this will be a hard number to compute but will more accurately represent reality.

Second and lastly, I believe one is losing precious information when not computing anything after a zero demand. It is as if we are not updating what needs to be updated, namely the average number of time intervals between intermittent demands. For example, say that the average number of time intervals between demands is 5.6 and at present we have not seen demand in 11 time intervals. Then, rather than doing nothing we should be expecting demand relatively soon (depending on variance). So, the probability that we will experience demand should be on the rise. This should be demonstrated in the model, but is not currently. That is to say, rather than doing nothing we should calculate the probability of demand at any time and when it reaches a certain decision threshold, we should order products.

7. References

Box, G. E., & Jenkins, G.M. (1976) .*Time series analysis forecasting & control.* Oakland, CA: Holden-Day.

Croston, J. D. *Forecasting and Stock Control for Intermittent Demand.* Operations Research Quaterly 23 (1972): 289-303.

Hamilton, J. D. (1994). *Time series analysis.* Princeton, NJ: Princeton University Press.

Heiji, C., De Boer, P., Frances, P. H., Kloek, T., & Van Dijk, H. K., (2004). *Econometric methods with applications in business and economics.* NY, NY: Oxford University Press.

Nahmias, S (2005). *Production and operations analysis.* 5th ed. New York, NY: McGraw Hill.

Schumacher, E. F. Small is Beautiful. Hartley & Marks, 1999.

Syntetos, A. A., and J. E. Boylan. *On the bias of intermittent demand.* International journal of production economics 71 (2001): 457-66.

8. Appendix

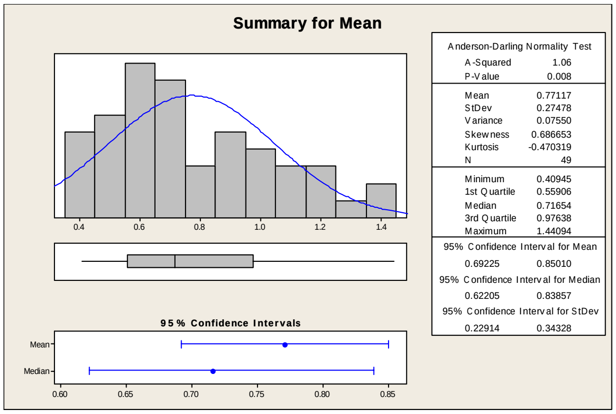

Figure 8.1: Descriptive Statistics of the Demand Data Mean

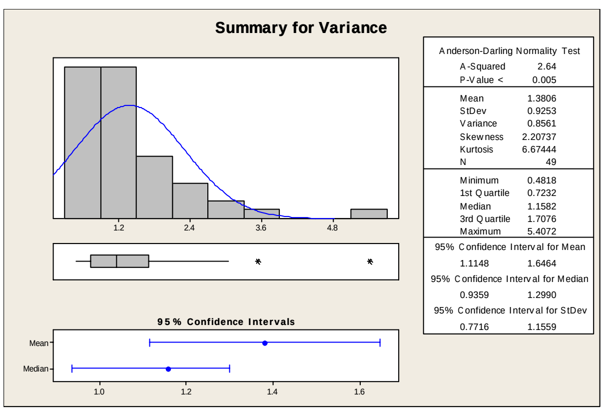

Figure 8.2: Statistics of the Demand Data Variance

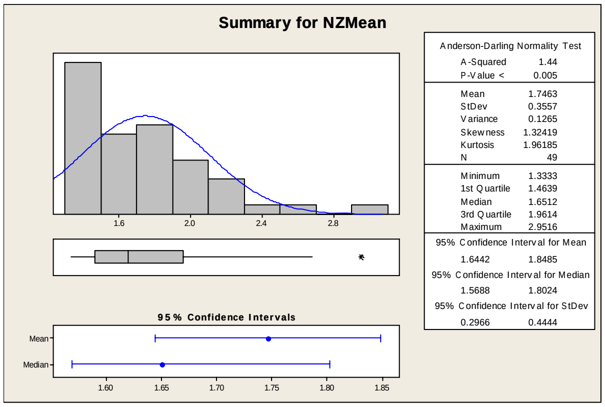

Figure 8.3: Statistics of the Demand Data Non-Zero Mean

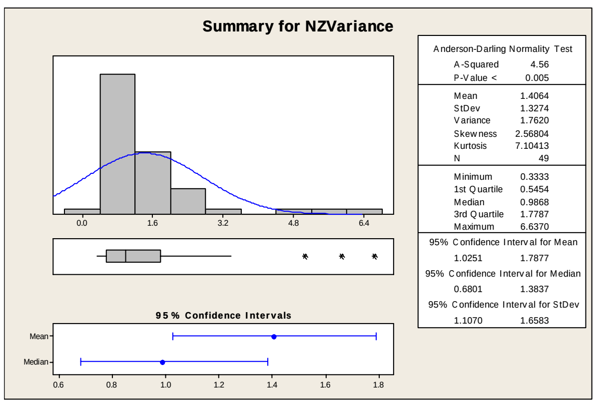

Figure 8.4: Statistics of the Demand Data Non-Zero Variance

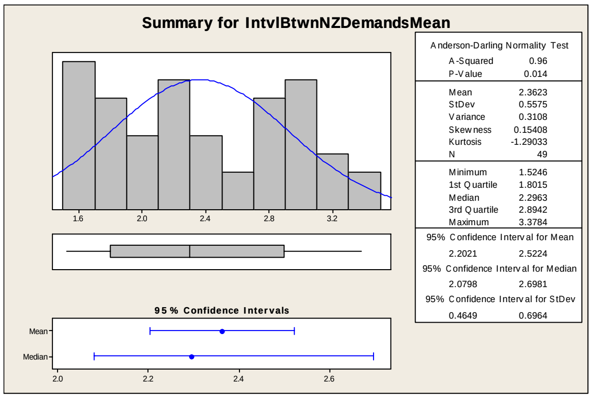

Figure 8.5: Descriptive Statistics of the Demand Data's Interval Between Non-Zero Demand Mean

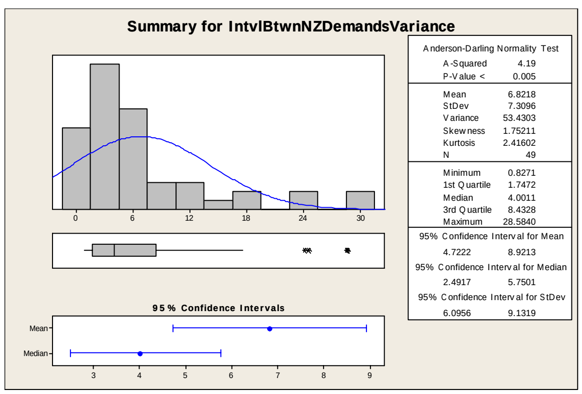

Figure 8.6: Descriptive Statistics of the Demand Data's Interval Between Non-Zero Demand Variance

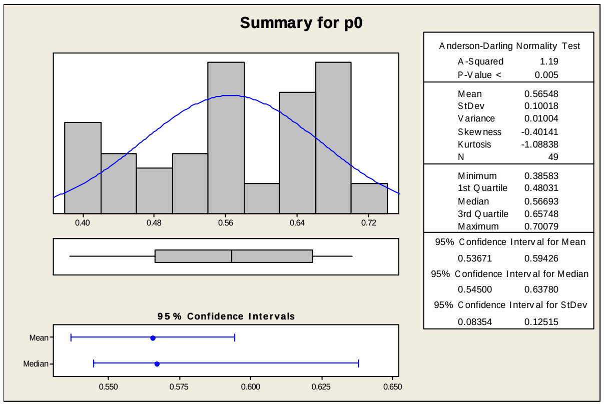

Figure 8.7: Descriptive Statistics of the Demand Data's Probability of Zero Demand

Table 8.1 List of ARIMA Parameters

Autoregressive	Difference	Moving Average
1	1	1
1	1	1
5	1	2
2	0	1
1	1	1
1	0	1
1	1	1
2	1	1
4	0	2
0	1	5
1	0	1
5	1	1
1	0	0
3	0	1
1	1	3
1	1	3
4	0	3
4	0	2
2	1	2
2	1	2
2	0	2
1	0	3
2	0	0
5	0	3
4	1	4
5	1	3
5	1	4
5	0	4
2	1	2
1	0	0
3	1	3
2	0	2
5	0	4
2	0	5
2	1	5
4	0	5
4	1	4
2	1	3
2	1	1
1	0	0
4	1	4
4	1	4
1	0	0
3	0	0
4	0	3
4	0	4
4	1	4
4	0	5
4	1	5
1	1	1