ISSN 2091-1521 (Print)
2091-153X (Online)
International Journal of Operational Research/Nepal - IJORN - 2019, Vol. 8, Issue 1

Optimizing Car Imports in Small Island States

Shiva Ramoudith[1], Patrick Hosein[1], Kevin Blake-Thomas[2]

[1] Department of Computing and Information Technology, The University of the West Indies, Trinidad
[2] Massy Technologies Ltd., Trinidad

Correspondence: Shiva Ramoudith, shivramx@gmail.com

Abstract

In most Small Island Developing States (SIDS), automobile dealers must import vehicles for sale on a periodic basis. This poses a problem for SIDS as the total demand is limited and the vehicle importation costs can be high since the only delivery option is by ship. We consider this problem for a single dealer in which the importation period is fixed and we use historical data (dates on which vehicles are bought, received and sold) to predict which brands/models should be purchased as well as the quantities of each to purchase using simple exponential smoothing. We demonstrate the potential of this method by investigating the time in which a vehicle spends on the lot as well as the probability of a vehicle being on the lot.

Keywords: inventory forecasting, purchase optimization, business intelligence, time series

1 Introduction

We consider the problem in which an automobile dealer must import cars for sale [2]. After investigating the procedures within the company, it was found that the orders are done on a periodic basis which is large enough so that a sufficient number of cars are imported to achieve bulk purchase benefits (such as supplier discounts, shipment costs, port costs, etc.). However, for such large periods one must be able to accurately predict the purchasing needs of customers. If too many cars are ordered, then several may remain on the lot for months, while if too few are ordered, then once those particular models run out customers will have to resort to competitors to make their purchases. In this paper, we develop a predictive model that ensures errors in over-ordering or under-ordering are sufficiently small. To measure the effectiveness of the model we propose two metrics, the probability of a vehicle existing on the lot and the average time spent by a vehicle on the lot. After determining the number of cars of each model that should be ordered, we may then be constricted by the budget. One must then determine which of these cars should be ordered to maximize profit. We also provide the solution to this optimization problem. Section 2 gives a brief overview of current literature. Sections 3 and 4 explain the proposed method and provide numerical results. Finally, Sections 5 and 6 give some discussion about the results and the conclusion.

2 Literature Review

This problem can be broadly classified as inventory forecasting which is a subset of time series and forecasting. A time series is a set of regular time-ordered observations of a quantitative characteristic of an individual or collective phenomenon taken at successive periods or points of time [8]. The goal of time series analysis is to make observations about the past and develop a model which describes the structure of the time series [1]. This model would then be used to make predictions concerning the future. A variety of forecasting models for time series exist, each having properties that make them suitable for different datasets. The size of the dataset greatly influences the type of forecasting model utilized [5].

Ostertagova and Ostertag [9] stated that exponential smoothing involves the unequal weighing of time series data (a procedure for continually revising a forecast in consideration of more recent experience, [6]). They mentioned that simple exponential smoothing is a powerful yet simple predictive technique for forecasting time series data in the short term where there is no trend or seasonal patterns. It was noted as very popular because it is simple (one parameter, alpha, needs to be tuned), computationally efficient and capable of change during forecasting. Simple exponential smoothing was therefore selected as the method of forecasting.

Lin [7], used sales, financial and economic indicators to predict new car sales in the United States of America. The following indicators were checked for their impact on vehicle sales: unemployment rate, the S & P stock index, disposable personal income, consumer price index, interest rates, and inflation rates. Data was available for every quarter from 1990 to 2013. Linear Regression and Time Series modeling were used to perform predictions on the data. They were compared based on the R-Square metric which concluded that the Time Series model was more accurate in forecasting. All indicators, with the exception of employment rates and the S & P stock index did play a significant role in the forecasting of car sales.

Wang et. al. [10] proposed a solution to optimize the allocation of car stock in various car ports in Singapore for a car rental company. Forecasting was used to determine the future demand. Again this problem is also concerned with inventory forecasting but in this scenario, reallocating vehicles is done instead of purchasing them. The authors made an important point that aggregated forecasts are typically more accurate than dis-aggregated forecasts. This study utilized several forecasting techniques such as selective moving averages, Holt's model, and Winter's model for each time period. The most accurate method was then used. The rest of the paper presents an approach to replenish inventories and maintain an optimal number of vehicles in each car port. Mention is made to having a safety stock for each car port. The authors chose to monitor the shortage of cars and results conclude that they were able to achieve a service level of 93.1%. The only information that was given concerning the test data was that they initially started of with 10 vehicles in each port and then used 1000 sets of car rental demand data from hypothetical and real-life data.

3 Methods

3.1 Predicting the Number of Cars to be Imported

Let us first focus on a single brand/model and determine how many should be ordered. We assume that when the ith order for this model was made, P_i cars were purchased. When this order is placed it may take several months for the order to be processed, shipped and then cleared from the arrival port. The time taken between ordering and arrival of cars on the lot is denoted as T. When this shipment of cars arrives on the sales lot there may be

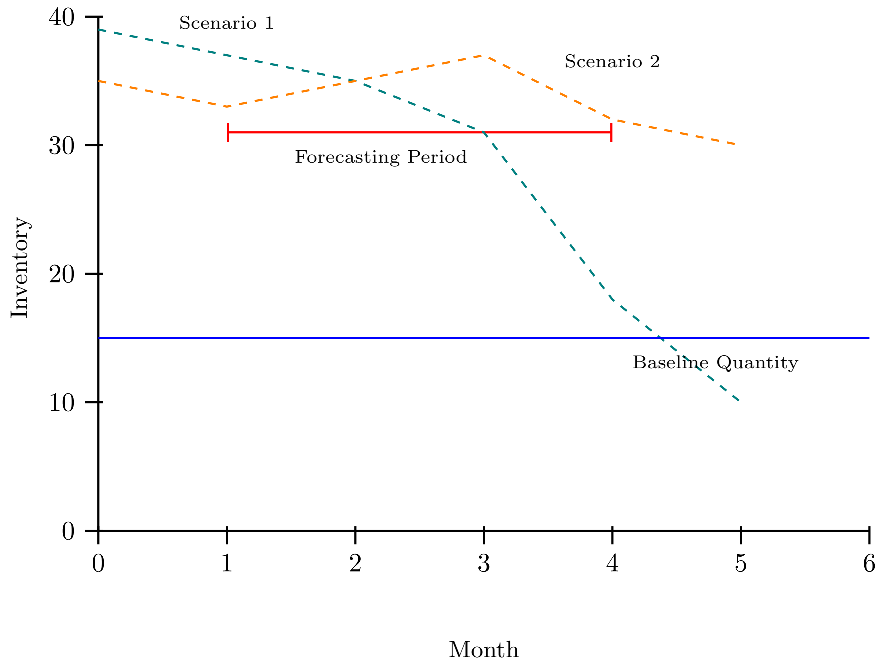

Figure 1: Prediction of quantity in stock for a single time period.

cars of this type from a previous shipment on the lot or the cars may have been sold out before the new shipment. If there are leftover cars we denote the number of these by R_i.

This model is illustrated in Figure 1. Orders are made every month and as mentioned above, they arrive T months in the future. At the beginning of month 0, cars arrive from the last placed order. At the end of month 0, the quantity in stock is the lowest and we have the remaining number of cars on the lot. We then forecast T months in advance to determine how many cars would remain at the end of the time period. For each month that we forecast, we considered the orders that were made previously as well as our estimate of the sales for each month.

We use the following approach to this problem. First we keep an exponentially smoothed estimate of the number of cars sold per month. If we denote the estimate for month i by $\bar{S}_i$ and the actual number of sales in month i by S_i then

$$\bar{S}_i = \alpha S_i + (1 - \alpha)\bar{S}_{i-1} \tag{1}$$

where $0 < \alpha < 1$ is the smoothing factor.

We use the most recent value of $\bar{S}_i$ to determine how many vehicles would be sold after T months in the future. Consider month i and assume we want to estimate the number of cars P_{i+n}, on the lot in n months time then

$$P_{i+n} = L_i - \bar{S}_i n + A(i, i + n) \tag{2}$$

where L_i represents the number of cars presently on the lot, and $A(i, i + n)$ represents the number of cars that were previously ordered and expected to arrive between the present and

n months in the future. This means that the remaining number of vehicles on the lot for each forecast month is equivalent to the previous quantity in stock plus the vehicles that came onto the lot from previous purchases minus the forecast sales.

We compare this predicted quantity, P_{i+n}, and order sufficient cars in the present month so that when they arrive in n months time the number on the lot will be approximately $\tilde{L}$ which is some minimum number of cars we try to maintain. If $P_{i+n} \geq \tilde{L}$ then no additional cars are ordered. If $P_{i+n} < \tilde{L}$ then $\tilde{L} - P_{i+n}$ cars are ordered.

3.2 Performance Metrics

As mentioned previously, we proposed two metrics for evaluation of the model. The first metric is the average time that a car spends on the lot (i.e. the difference in time between sale of the car and its arrival on the lot). Naturally, we would like this to be as small as possible. The second metric is the probability that a customer can find the car they desire on the lot. In this case we want this probability to be as close to one as possible. We compute this as follows. For each model we count that number of months that at least one of the model was available on the lot and divide this by the total number of months considered. We simulated what would happen if ordering was done with the proposed approach and compare results with what actually happened. In order to have a fair comparison of the present mode of operation with the proposed mode of operation we assume that the same total number of cars of each model is ordered in both cases. Our method therefore rearranges purchases to reduce the average time spent on lot while attempting to maximize the probability of a vehicle existing on the lot.

3.3 Profit Optimization

Once we estimate the number of cars of each model that should be ordered we next have to determine if the budget allows for such an order and, if not, what models should in fact be ordered. Let $O_j(i)$ denote the computed number of cars of model i for the jth order. Let x_i denote the number of cars of model i that should be ordered, let p_i denote the profit per car for model i, let c_i be the unit cost to the dealer for model i and let B denote the total budget. Since we want to optimize the total profit given the allocated budget then the optimization problem can be stated as:

$$\max_{\vec{x} \in \mathbf{Z}^N} P(\vec{x}) \equiv \sum_{i=1}^{N} x_i p_i \tag{3}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} x_i c_i = B$$

$$\text{with} \quad 0 \leq x_i \leq O_j(i) \quad \forall i$$

where we assume N types of models. Note that we have assumed that the budget constraint is binding (i.e. that the total cost is exactly equal to the budget) since if this were not the case then additional cars could be purchased resulting in additional profit.

This problem can be solved using Lagrange Multiplier methods by introducing $\lambda, \vec{\mu}$ and $\vec{\gamma}$. The Lagrangian is given by,

$$\mathcal{L}(\vec{x}, \lambda, \vec{\mu}, \vec{\gamma}) = \lambda B + \sum_{i=1}^{N} x_i p_i - \lambda x_i c_i + \mu_i x_i + \gamma_i (O_j(i) - x_i) \tag{4}$$

$$\text{s.t } \mu_i x_i = 0, \;\; \gamma_i (O_j(i) - x_i) = 0, \;\; \mu_i \geq 0, \;\; \gamma_i \geq 0 \quad \forall i$$

Taking partial derivatives and setting to zero we obtain:

$$\frac{\partial \mathcal{L}}{\partial x_i} = p_i - \lambda c_i + \mu_i - \gamma_i = 0 \tag{5}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = B - \sum_{i=1}^{N} c_i x_i = 0 \tag{6}$$

Let us consider the various cases:

$$\begin{aligned}
\mu_i > 0, \; \gamma_i > 0 \quad &\Rightarrow \quad x_i = 0, \;\; x_i = O_j(i) \;\; \text{not possible} \\
\mu_i > 0, \; \gamma_i = 0 \quad &\Rightarrow \quad x_i = 0, \;\; \lambda > \frac{p_i}{c_i} \\
\mu_i = 0, \; \gamma_i > 0 \quad &\Rightarrow \quad x_i = O_j(i), \;\; \lambda < \frac{p_i}{c_i} \\
\mu_i = 0, \; \gamma_i = 0 \quad &\Rightarrow \quad 0 \leq x_i \leq O_j(i), \;\; \lambda = \frac{p_i}{c_i}
\end{aligned}$$

Therefore, we just need to find $\lambda = \lambda^*$ that satisfies these conditions. This can be accomplished by starting with the model with the largest ratio $\frac{p_i}{c_i}$ and purchasing the maximum number of cars of that model. We then go to the model with the next largest ratio and repeat. We keep track of the total cost and stop when this total reaches B.

4 Results

4.1 Approach

The dataset was obtained from an automotive retailer. It spanned a total of 18 months and comprised over 50 vehicles. The five top selling vehicles were identified and used for testing. For the sake of anonymity, they will be referred to as vehicles A, B, C, D and E. The following information was extracted from each vehicle: ship date, entry date, sold date, cost and profit. This was then aggregated and the following values were obtained for each vehicle model:

- T - The average time (months) that a vehicle took to be shipped.
- T_h - The average time (days) on lot spent by vehicles using the present approach.
- P_h - The probability of a car being available on the lot for the present approach.

Afterwards, we utilized the method mentioned in Section 3.1 and forecast the number of vehicles required for one time period in advance. We then appended this result to the historical data and computed the following:

- T_f - The average time (days) on lot spent by vehicles using the proposed approach.
- P_f - The probability of a car being available on the lot for the proposed approach.

It should be noted that the method for computing T_h, T_f, P_h and P_f can be found in Section 3.2. When performing forecasting, we selected a value for α in conjunction with a value for $\tilde{L}$.

Now, after having the predicted number of vehicles for each model, we brought the budget limitation into the problem and utilized the method in Section 3.3 to obtain the number of vehicles that could actually be bought.

4.2 Numerical Results

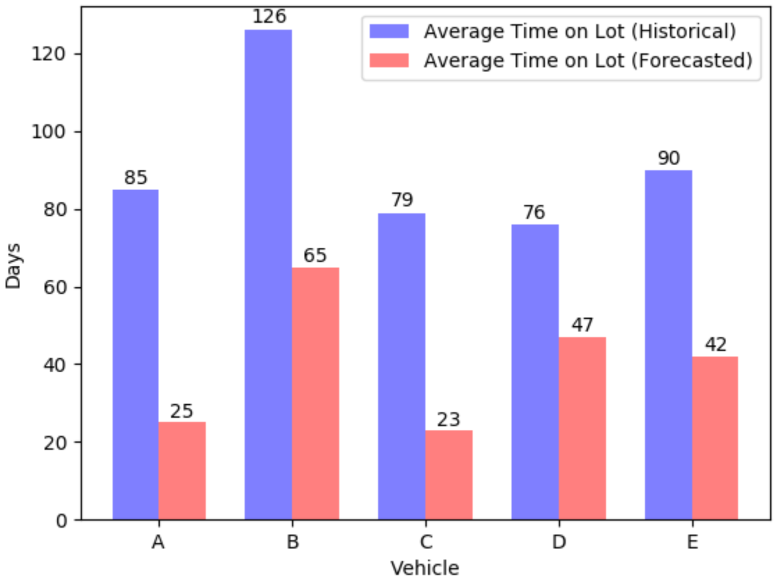

Figure 2: Average Time on Lot ($\tilde{L} = 30$)

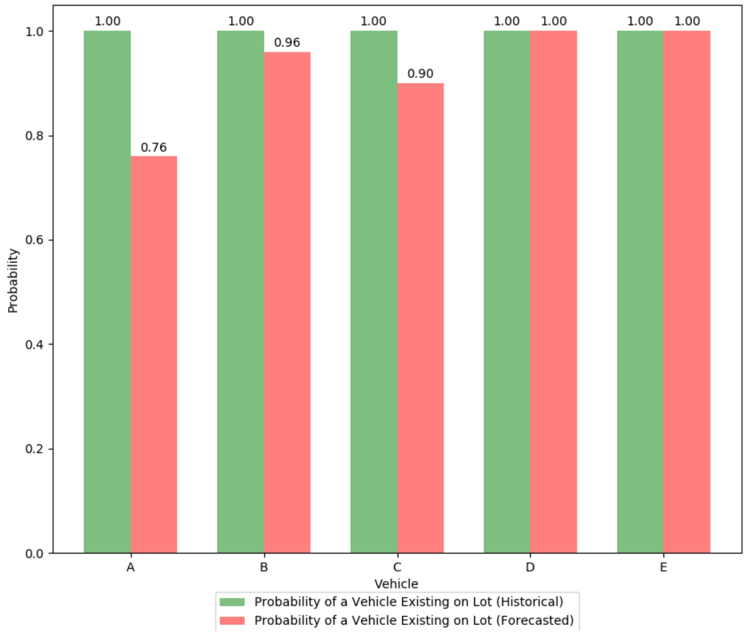

Figure 3: Probability of a Vehicle Available on Lot (with $\tilde{L}$ fixed at 30)

5 Discussion

5.1 Order Prediction

Current work in this area mentions the use of advanced forecasting techniques such as Artificial Neural Networks, Decision Trees, ARIMA and Support Vector Machines. We decided to use simple exponential smoothing to forecast vehicle sales because of additional constraints that need to be added for a practical deployment. Some additional reasons are:

- The time period in which our dataset spans (18 months) is very small compared to those mentioned in the related works section. The time period in the data used by [7], and [3] ranged a minimum of 20 years. Methods such as ARIMA, Linear Regression and Triple exponential smoothing require data spanning multiple years in order to better realize the trend and seasonality of the data, which is crucial in making forecasts. As a result, these methods are not feasible given our time range.
- We were only required to forecast purchases for one time period in advance. Therefore, the use of the complex methodologies to forecast sales for only one month in advance would be considered unnecessary. Furthermore, other complex forecasting methods require that multiple parameters be tuned (which decreases robustness).
- Our dataset comprised sales, purchase and vehicle information. In the case of [7], other

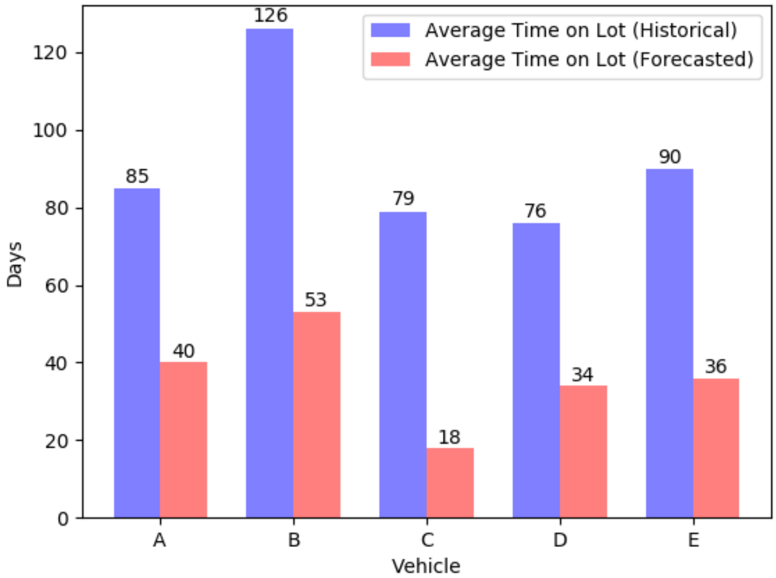

Figure 4: Average Time on Lot (using optimal $\tilde{L}$ for each model)

important features that could have correlation to the vehicles sales such as GDP, per-capita income and population density were used for improved accuracy in forecasting. We had no additional data so utilization of the decision tree model mentioned by [4] was not possible.

It should be mentioned that we ran tests initially with various alpha values ($0.1 \leq \alpha \leq 0.9$). We determined that $\alpha = 0.7$ provided reasonable results for all models. Initially, the value of $\tilde{L}$ was set to 30 for all vehicles. Results in Figures 2 and 3 indicate that this value was not optimal for all cases and hence we decided to optimize this parameter independently for each model. For the 5 models considered, we were able to decrease the average time spent on lot by 55 days whilst maintaining the probability of a vehicle existing on the lot at one. The optimal values of $\tilde{L}$ were 105, 42, 39, 26 and 30 for vehicles A through E respectively (refer to Figures 4 and 5).

Some of the related works [11], [4] in this area originate from larger nations such as China and the USA, some of whom are manufacturers of vehicles. In these countries, the shipment times (from ordering to arriving on the lot) are considerably less when compared to small island developing states. Additionally, the issue of vehicle saturation is far less in these economies. Our problem is therefore unique as these factors prevent us from simply using forecasted sales data to directly impact purchases as seen with the current literature.

Lastly, mention should be made of [10] in particular, where a metric (service level) similar to ours (the probability of a vehicle existing on the lot) was proposed to evaluate a forecasting model. They are similar in the sense that they both monitor the probability of a customer

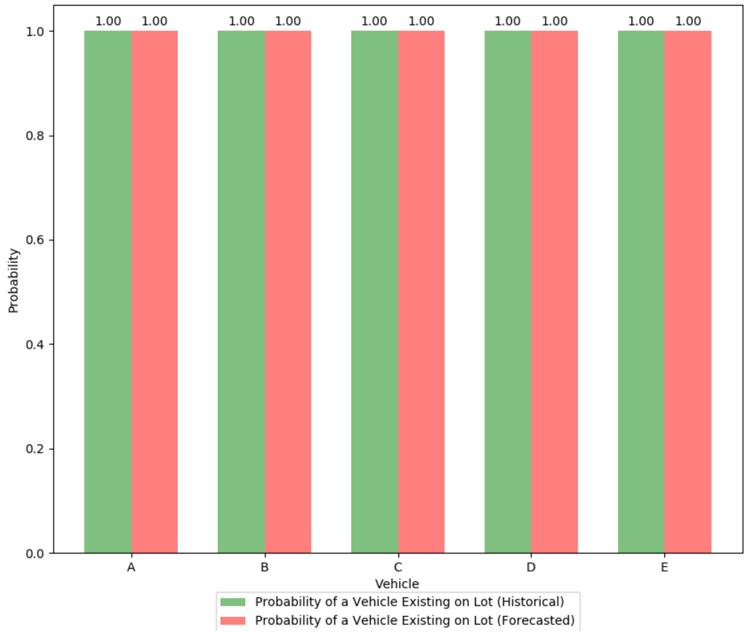

Figure 5: Probability of a Vehicle Available on Lot (using optimal $\tilde{L}$ for each model))

Vehicle	Quantity to Purchase	Profit (TT Dollar)	Cost (TT Dollar)
X	30	54781.66	115003.70
Y	15	44342.23	183313.21
Z	35	43499.39	153464.31

Table 1: Information for vehicles to be purchased

being served (cars exist on lot so a customer can purchase or rent). Forecasting the number of vehicles to reallocate at each rental site is similar to recommending the purchases that should be made in our scenario. In addition, the safety stock mentioned is also similar to the baseline quantity we proposed. [10] were able to achieve a service level of 93.1% while we achieved 100%.

5.2 Purchase Optimization

There was no insight given as to whether or not the automotive retailer was constrained by a budget for monthly purchasing so we decided to present a scenario to illustrate our approach. Let us consider a budget of six million dollars and the vehicle information provided in Table 1. Given the fact that the budget is restricted to six million dollars, we cannot purchase vehicles X, Y and Z in their desired quantities. We compare a trivial method (highest profit first) with the optimized approach.

First we demonstrate the trivial method. After looking at the profits made by each vehicle, we should purchase 30 of vehicle X first since we would make the most profit from selling this vehicle. Afterwards we should purchase 13 of vehicle Y as it has higher profit than vehicle Z. At this point, we do have sufficient capital remaining (166817.27 dollars) to purchase one unit of vehicle Z. With this method, our resultant profit is 2263398.18 dollars. We use the data in table 1 together with the approach previously outlined:

1. Compute the profit/cost ratio for all vehicles. They are 0.476, 0.242 and 0.283 for vehicles X, Y and Z respectively.
2. Select the vehicle with the highest profit/cost ratio that has not been purchased yet.
3. Once we can purchase the desired quantity with the remainder of the budget, we move onto the next vehicle (start at step 2).
4. If we cannot purchase the desired vehicles, we need to reduce the quantity so that we would be within range of the budget.
5. We repeat steps 2-4 until our budget is totally used.

After performing the steps above, we end up with the optimal orders. Vehicle X has the highest profit/cost ratio of 0.438. We are able to purchase the quantity desired for vehicle X. Moving on, the vehicle Z has the next highest profit/cost ratio. We are unable to purchase our desired quantity so we purchase 16 vehicles instead. Remaining with a capital of 94460.04 dollars, we are unable to purchase any more of vehicles Z or Y. With this method, our resultant profit is 2339440.04 dollars.

5.3 Future Work & Concluding Remarks

In conclusion, we addressed the problem of improving the efficiency of automobile ordering under the specific conditions faced by small island states. In such countries automobile orders must be made in bulk and, in addition, the duration of time between ordering of vehicles and their arrival can be on the order of six or more months. Therefore accurate prediction of the needs of the country is essential when making decisions on purchase orders. We demonstrated a simple yet robust approach for doing this and illustrated its advantages using real data. Naturally further optimization is possible but we also need to address possible additional constraints in the ordering process. However, the point that data-based decision can result in significant cost savings has been illustrated in this particular case and we plan to eventually deploy this optimization approach for the dealer with whom we have collaborated.

Future work will include the following. Once data on indicators of the economy are obtained, we can incorporate them with the sales data and utilize more advanced forecasting methods mentioned earlier on. Once sale and purchase data spanning a wider time range are obtained, we can perform long term forecasting with the data using more advanced forecasting techniques (double and triple exponential smoothing). Furthermore, we can combine the sale and purchase data with service and part information for vehicles. We can then propose a method to determine what parts might be needed in the near future and keep stock of them in order to increase customer satisfaction as well as sales. With the introduction of electric and hybrid vehicles, gas counterparts would slowly see decreases in sales and would phase out eventually so forecasting when this could possibly occur would be beneficial for the automotive retailers.

References

[1] R. Adhikari and R. Agrawal. An introductory study on time series modeling and forecasting. *arXiv preprint arXiv:1302.6613*, 2013.

[2] M. Al-Noukari and W. Al-Hussan. Using data mining techniques for predicting future car market demand; dcx case study. In *Information and Communication Technologies: From Theory to Applications, 2008. ICTTA 2008. 3rd International Conference on*, pages 1–5. IEEE, 2008.

[3] J. Dargay, D. Gately, and M. Sommer. Vehicle ownership and income growth, worldwide: 1960-2030. *The Energy Journal*, pages 143–170, 2007.

[4] M. Hülsmann, D. Borscheid, C. M. Friedrich, and D. Reith. General sales forecast models for automobile markets and their analysis. *Trans. MLDM*, 5(2):65–86, 2012.

[5] R. J. Hyndman, A. V. Kostenko, et al. Minimum sample size requirements for seasonal forecasting models. *Foresight*, 6(Spring):12–15, 2007.

[6] P. S. Kalekar. Time series forecasting using holt-winters exponential smoothing. *Kanwal Rekhi School of Information Technology*, 4329008:1–13, 2004.

[7] Y. Lin. *Auto Car Sales Prediction: A Statistical Study Using Functional Data Analysis and Time Series*. PhD thesis, University of Michigan, 2015.

[8] OECD. *Data and Metadata Reporting and Presentation Handbook*. OECD, 2007.

[9] E. Ostertagova and O. Ostertag. Forecasting using simple exponential smoothing method. *Acta Electrotechnica et Informatica*, 12(3):62, 2012.

[10] H. Wang, R. Cheu, and D.-H. Lee. Logistical inventory approach in forecasting and relocating share-use vehicles. In *Advanced Computer Control (ICACC), 2010 2nd International Conference on*, volume 5, pages 314–318. IEEE, 2010.

[11] Z. Zhu, T. Chen, and T. Shen. A new algorithm about market demand prediction of automobile. *International Journal of Marketing Studies*, 6(4):100, 2014.