Analysis and prediction of daily orderlines volume at Bausch + Lomb ELC

Research Paper Business Analytics
January 26, 2014

Author:
Bas de Jong

Analysis and prediction of daily orderlines volume at Bausch + Lomb ELC

Research Paper Business Analytics
January 26, 2014

Author:
Bas de Jong

Supervisors:
Dr. René Bekker & Dr. Fetsje Bijma

VU University Amsterdam
Faculty of Sciences
De Boelelaan 1081
1081 HV Amsterdam
The Netherlands

Bausch + Lomb
European Logistics Center
Koolhovenlaan 110
1119 NH Schiphol-Rijk
The Netherlands

Preface

This research paper is part of the Master of the study Business Analytics at the VU University Amsterdam. As a dual student I have opted to do this research at the end of my first year of the Master. The objective of the paper is to study a problem with mathematical and computer science aspects that also has practical value for a company. The company in this case is Bausch + Lomb, my employer for 4 years running.

The assignment was formulated by Dean van der Zanden, groupleader Process Control at Bausch + Lomb. He and his colleagues are confronted with variability in order volume and the consequences of this variation on a daily basis. He wondered whether the volume could be predicted to a certain degree. This could help the company in planning the right number of resources without over- or understaffing. I would like to thank him for defining and setting up the problem and his feedback during the process.
From the VU I would like to thank dr. Fetsje Bijma and dr. René Bekker for their guidance and knowledge. Although it is quite unusual to have two supervisors from the VU for the BA Research Paper, I am thankful for their offer to help me. Being offered two points of view was always interesting as it kept me sharp, but it was also challenging to combine these into a single direction for the paper. In my opinion, it has only improved the research and the result.

The intended audience for this report is primarily operational management (groupleaders and the operations manager) at Bausch + Lomb, as well as students and teachers interested in (demand) modeling. In order to keep this report readable for people with little or no statistical background all theory, formulas and tests are discussed in separate sections within the text. These sections can be skipped or read back at any time.

Amsterdam, January 2014

Summary

The objective of this paper is to investigate whether the future daily orderline volume at the Bausch + Lomb European Logistics Center (ELC) at Schiphol can be modeled to a certain degree of accuracy, so that the results of this model can be used to improve the staff planning. This should help to avoid under- and overstaffing as much as possible. In order to meet this objective, influence factors were determined, trends were identified and a model was built and validated.

The orderlines volume received per day for the period [CONFIDENTIAL] is used for this paper. Season effects, large national or European events and the effect of day of the week were studied. More complex factors such as weather conditions or opening hours of stores are not considered in this paper.

The following influence factors were observed and studied in chapter 2:

- [CONFIDENTIAL]
- National holidays and the workday after a national holiday in large European markets have a significant impact on the volumes, as do national holidays in the smaller European markets.
- Large European holidays such as Eastern, Ascension, Christmas and New Year have a big impact on the orderlines volume. The volumes at or around these days are very low and are therefore separately modeled from the national holidays.
    - Although only 4 observations are available for Eastern and Ascension, the volume pattern is obvious and useful for staff planning. More data needs to be collected on these events to improve the accuracy of the model.
    - For Christmas and New Year fewer observations are available. For certain days only 1 or 2 observations were available, and this is not enough to draw reliable conclusions from. The period around Christmas and New Year was left out.
- The volumes differ significantly per weekday. Mondays and Tuesdays have the highest volumes in the week, with a steady decline of volume to Friday.
- The Bausch + Lomb monthly financial close has a negative impact on the orderlines volume in the last week before the close. Although the drop in volume is small, it is deemed significant while the effect of the quarter close is not.

Based on the observation from chapter 2, three types of models are built in chapter 3 using R statistical software: a linear, Huber linear and a log-linear model. The latter two are considered because of the outliers present in the linear model. The residuals in the linear model are not normally distributed and it is checked whether this is improved using another type of model. The variables for all three models are nearly the same, with 1 or 2 minor differences. The models are tested and compared with each other using leave-one-out cross validation. Based on this, the Huber linear model is the best performing model as the sum of squares for this model is the lowest by a small margin compared to the linear model. Also, measured by the Bausch + Lomb performance criterion [CONFIDENTIAL] the Huber model is the best model (80.0%) although it does not meet the ambitious threshold of 95.0% from Bausch + Lomb for the model to be deemed useful in practice.

Although the model will not be used on a daily basis to forecast orderlines volumes at the Bausch + Lomb ELC, valuable information was gathered for staff planning during European and national holidays and the summer period. The volume during these periods is constant from year to year and this fact can now be used to determine how many operators are required during these days.

The data available had too few observations to draw reliable conclusions for the yearly events such as European national holidays. It had too many observations for the weekly effect to react to

changes. To improve the model accuracy it is recommended to continue analyzing the yearly events and do a separate follow-up study focusing only on the weekly effects.

Table of Contents

1. Introduction

Companies would like to know future demand for their products so they can plan their resources to match this demand. The demand is often not known in advance, but the recent past can be an indication for the workload in the near future. This paper attempts to model the future workload measured in orderlines for the logistics center for Bausch + Lomb at Schiphol-Rijk using historical data.

1.1 Problem statement

The objective of this paper is to investigate whether the future daily orderline volume at the Bausch + Lomb European Logistics Center (ELC) can be modeled to a certain degree of accuracy, so that the results of this model can be used to improve the staff planning. To do so, either the workload should be known in advance or can be predicted with a particular degree of certainty.

To be able to answer this main question the research focuses on the following subjects:
- What factors influence the orderline volume from historical perspective?
- Which trends does the model need to take into account?
- How accurate is the model and can it be used to improve staff planning?

The first two questions are answered in chapter 2, while the models and their accuracy are presented in chapter 3. In chapter 4 the conclusions and recommendations as discussed.

1.2 Background

Bausch + Lomb is a worldwide operating company whose core business is the manufacturing and distribution of contact lenses, lens solutions, eye surgery products and eye medication. The company was founded 160 years ago in Rochester, New York in the USA and is responsible for numerous developments in the field of optics and eye health. The ELC at Schiphol-Rijk stores and distributes around 40.000 different items (SKU) to European customers, such as opticians, retailers and hospitals, but also to distributers and warehouses outside of Europe.

The order volume is not known in advance. The European sales volume mainly comes from independent and (small) chain opticians, as well as eye clinics and eye hospitals. Their order volume depends on the number of customers walking into their shops / clinics for contact lenses or surgical lens implants (IOL). Because of the high variation in shapes and diopter of the eyes, only the most popular items are kept on stock with the ELCs customers. To guarantee the short (maximum 48 hours) delivery times that the market demands (nearly) all items in the portfolio of Bausch + Lomb are kept on stock at the ELC. Orders that come in before a certain fixed time (country dependent) are required to be picked and shipped on the same day. For the past two years the ELC has shipped 99.4% of the orderlines in time (pick performance), with the target currently set at 99%.

The ELC wants to meet the daily order demand with the appropriate number of resources in such a way that under- and overstaffing are avoided as much as possible. Understaffing has a negative effect on the pick performance while overstaffing leads to inefficient use of the resources. The latter has been often the case at the ELC to make sure the target pick performance was met. For the past two years this performance is stable and operations management wants to know whether staffing levels can be optimized while maintaining the pick performance. The desire for appropriate staffing levels is also stimulated by the times of economic downturn when cost need to be kept at a minimum in order to stay competitive.

Overstaffing is not only bad from a cost perspective, also from a motivational perspective for the order pickers it is not good when the workload is low on a regular basis or for an extensive period of time during the day. In addition to reduction of cost also morale can be improved or retained when there is a better match between supply (resources) and demand. This fits perfectly in two of the 5 main mandates of the ELC: ‘Cost awareness’ and ‘Employee focus’.

The number of orderlines received on a day gives a good indication of the workload and thus the number of resources required. The reason the number of orderlines is selected as input variable rather than the number of received orders, is because there is a relatively constant relation between the number of orderlines and the number of picks (Figure 1.1). The number of picks is a good measure of workload because we know the (average) time required per pick. That makes orderlines indirectly a good measure of workload as well.

Figure 1.1: Daily ratio of picks per order and picks per orderline for period [CONFIDENTIAL][redacted]

The number of orderlines received has been logged since [CONFIDENTIAL] and provides a good basis for this analysis. Although products are separated by Business Units (Lens, Lens care & Surgical) and stored in different areas in the warehouse, with some variation between the Business Units in volumes, the total order volume per day over all Business Units is studied. This is because of the longer history available for the total volume, as the volumes per Business Unit have only been available since [CONFIDENTIAL].

There is considerable variation in the orderlines volume on a day-to-day basis and Operations Management at the ELC wants to know whether this variation can be predicted in a model so it can be used for short to medium term staff planning (1-4 weeks ahead).

4. Conclusions and recommendations

This chapter presents the conclusions and recommendations from the study performed in the chapters 1 through 3. First the conclusions will be discussed and later some recommendations will be presented for future research.

4.1 Conclusions

Based on the results, the Huber linear model performs best compared to the linear and the log-linear model. The Huber model put less emphasis on the heavy tailed distribution of the residuals. As a result, the Huber model performs better compared by both the sum of squares of the cross-validation and the Bausch + Lomb performance criterion [CONFIDENTIAL]. The model is with 80.0% by Bausch + Lomb standard not accurate enough as the minimum desired performance was 95%.

The fact that the criterion is not met does not make the model unusable. The model has provided some valuable insights that were previously not known or not quantified.

- **Effect of European and national holidays.** It was shown that the decrease in volume for European and national holidays often follow a certain pattern: due to timing of the order cut-off, a holiday not only has effect on the day itself but also on the next working day. On both days the orderline volume is lower than normal. Although the number of observations per holiday is too small to quantify these effect more precise (see recommendations), the effect does have a similar pattern from 2010 till 2013.
- **The effect of the summer period.** Similar to the holidays, the number of data points available is not enough to draw strong conclusions on the size and timing of the orderline volume trend during the summer, but it is clear enough to be useful for staff planning [CONFIDENTIAL]
- [CONFIDENTIAL]
- **Financial close.** The monthly or quarterly financial close has little impact on the orderline volumes. An increase in orderlines was expected as this effect was observed with non-European orders. There is no significant change for European orders in the last week before the quarter close. A small decrease in orderlines is observed in the last week before each monthly close. This was unexpected, but not considered to be strange.
- **Monthly trends.** No trends are observed when the volumes are summed per month. All effects can be explained on week level or on day level.

Although the model as a whole is not accurate enough according to the limit posed by Bausch + Lomb, the new knowledge on the holidays and the summer period can be used to adjust the staffing levels accordingly for these specific events.

4.2 Recommendations

The following recommendations are made for future research or for implementation of this model within Bausch + Lomb ELC:

- **Model accuracy.** The best performing model is still far from the Bausch + Lomb desired model performance. More complex model types, such as neural networks, could provide better results but will probably still not make the 95% threshold. The target set by Bausch + Lomb has proven to be very ambitious. Internally should be discussed how big the risks are when a lower accuracy is accepted and whether they are willing to accept these risks.

- [CONFIDENTIAL]
- **Improve logging of data.** No cause for the outliers in the data could be given, because no additional information was available. If the logging would be improved by adding comments to the daily volume on system, technical or other issues, reasons for the deviations could be found. Based on this information decisions can be made to include or exclude certain observations.
- **Continue analysis of yearly effects.** With over 3 years of data available, this was sufficient to quantify daily effects such as the weekday effect or long term effects such as the steady decline in volume. It did provide valuable insight into the effect of European holidays, such as Eastern or Christmas, and the summer period but not enough data points were available to quantify the effect with sufficient accuracy. It is recommended to continue gathering information on yearly effects so these effects can be better quantified.
- **Short term effects.** The effect of a weekday on the orderline volume was considered to be constant throughout the more than 3 years. It could very well be that these factor values change throughout time. For optimal performance of the model it is recommended that the scale of the daily and weekly effects are determined based on a shorter interval. A rolling year, or maybe shorter, could provide better results in predicting the orderline volume.

References

Below is an overview of the references used for the research of this paper. References to which specific parts of the literature below are used are included in the text.

Faraway, Julian J., '*Linear models with R*', Chapman & Hall, 2005

Faraway, Julian J., *'Extending the Linear Model with R'*, Chapman & Hall, 2006

Appendix

The data analysis for this research paper is done using Microsoft Excel 2007 and R statistical software version 2.15.2. Below is the code displayed for conducting the analysis in R.

```
summary(LinesData$Lines)

# Plot key figures
par(mfrow=c(2,2))
boxplot(LinesData$Lines, main = "Boxplot")
hist(LinesData$Lines, xlab = "Orderlines received", ylab = "Count of observations", main = "Histogram")
symplot(LinesData$Lines)
qqnorm(LinesData$Lines)
shapiro.test(LinesData$Lines)        #W = 0.9933  p = 0.0005 Not normal distributed

#-----------------------------------------------
# Averages and boxplot per weekday
#-----------------------------------------------
LinesMon = subset(LinesData, LinesNoHol[,5] == 1)
LinesTue = subset(LinesData, LinesNoHol[,5] == 2)
LinesWed = subset(LinesData, LinesNoHol[,5] == 3)
LinesThu = subset(LinesData, LinesNoHol[,5] == 4)
LinesFri = subset(LinesData, LinesNoHol[,5] == 5)

summary(LinesMon$Lines)
summary(LinesTue$Lines)
summary(LinesWed$Lines)
summary(LinesThu$Lines)
summary(LinesFri$Lines)

wkday = c("Monday", "Tuesday", "Wednesday", "Thursday", "Friday")
boxplot(Lines ~ DayNum, data = dfh, names = wkday, main = "Boxplot of orderlines per weekday", xlab =
"Weekday", ylab = "Orderlines")

#---------------------------------------------
# Building a linear model
#---------------------------------------------
dfh = data.frame(LinesData)
dfh$RecDate <- as.Date(dfh$RecDate, "%d/%m/%Y")
attach(dfh)

baselm <- lm(Lines ~ 1, data = dfh)
weeklm <- lm(Lines ~ LTW, data = dfh)
wkdaylm <- lm(Lines ~ LTW + factor(DayNum), data = dfh)
hollm <- lm(Lines ~ LTW + factor(DayNum) + LHC + SHC + LDAC + SDAC, data = dfh)

anova(baselm, weeklm, test = "F")

summary(hollm)
sd(weeklm$residuals)
qqnorm(baselm$residuals)
shapiro.test(baselm$residuals)

plot(dfh$RecDate, wkdaylm$residuals, type ="p", xlab = "Date", ylab = "Residuals", main = "Residuals from long
term + weekday model")
```

```
val.daag <- CVlm(df = dfh, m = 25, form.lm = formula(Lines ~ LTW + YW + factor(VW) + factor(DayNum) +
factor(Holiday) + LHC + SHC + LDAC + SDAC))
val.glm <- cv.glm(dfh, comp_model, K = 10)

#--------------------------------------
# linear model without Xmas period
#--------------------------------------
dfx = data.frame(Data.exXmas)
dfx$RecDate <- as.Date(dfx$RecDate, "%d/%m/%Y")
attach(dfx)

model_base = lm(Lines ~ LTW + factor(DayNum) + LHC + SHC + LDAC, data = dfx)
model_exXmas = lm(Lines ~ LTW + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC + SDAC, data = dfx)
model_comp = lm(Lines ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx)
lin.mod = glm(Lines ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx, family = "gaussian")
summary(model_comp)

lin.res = model_comp$residuals

model_cov = lm(Lines ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose + factor(DayNum)*factor(VW), data = dfx)
summary(model_cov)
anova(model_cov)

plot(dfx$RecDate, residuals(model_vac), xlab = "Date", ylab = "Residual", main = "Residuals from model")

shapiro.test(model_comp$residuals)
sum(abs(model_comp$residuals) <= 2000) / length(dfx$Lines)

anova(model_exXmas, model_base, test = "F")

par(mfrow=c(1,1))
plot(dfx$RecDate, dfx$Lines, xlab = "Date", ylab = "Orderlines", main = "Plot of actual and modelled received
orderlines")
lines(dfx$RecDate, lin.mod$fitted.values)

par(mfrow=c(1,2))
plot(dfx$RecDate, residuals(model_comp), xlab = "Date", ylab = "Residual", main = "Residuals from model")
boxplot(model_comp$residuals)
hist(residuals(model_comp), xlab = "Residuals of lin model", main = "Histogram of residuals")
qqnorm(model_comp$residuals, main = "Normal Q-plot for residuals")
qqt(model_comp$residuals, df=4, main = "T-4 Q-plot for residuals")

par(mfrow=c(1,1))
plot(dfx$RecDate, abs(residuals(model_exXmas)), xlab = "Date", ylab = "Abs Residual", main = "Abs Residuals
from model")
boxplot(residuals(model_exXmas), ylab = "Residuals", main = "Boxplot of residuals")

#--------------------------------------
# Huber lin model
#--------------------------------------
rlmod <- rlm(Lines ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx)
temp = summary(rlmod)$coefficients
```

```
write.table(temp, file = "Parameter values Huber.csv", sep = ",", eol = "\n", row.names = FALSE, col.names =
TRUE)

#Calculate p-values
pval = (1-pt(abs(temp[,3]),836))
write.table(pval, file = "P-values Huber.csv", sep = ",", eol = "\n", row.names = FALSE, col.names = TRUE)

qt(0.975, 836) #1.9628
qt(0.025, 836) #-1.9628

#B+L performance criterion
sum(abs(rlmod$residuals) <= 2000) / length(dfx$Lines)

#Plot residuals against time
par(mfrow=c(2,2))
plot(dfx$RecDate, residuals(rlmod), xlab = "Date", ylab = "Residual", main = "Residuals from model")

#Plot observations with Huber model
par(mfrow=c(1,1))
plot(dfx$RecDate, dfx$Lines, xlab = "Date", ylab = "Orderlines", main = "Plot of actual and modelled received
orderlines")
lines(dfx$RecDate, predict(rlmod, type="response"))

#--------------------------------------
# Poisson regression model
#--------------------------------------
poiss.mod <- glm(Lines ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx, family = "poisson")
temp = summary(poiss.mod)$coefficients
write.table(temp, file = "GLM_Poisson.csv", sep = ",", eol = "\n", row.names = FALSE, col.names = TRUE)

check.mod <- lm(log(Lines) ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx)
check1.mod <- lm(log(Lines) ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC
+ SDAC + MonthClose, data = dfx)
summary(check.mod)
temp = summary(check.mod)$coefficients
write.table(temp, file = "Log_model.csv", sep = ",", eol = "\n", row.names = FALSE, col.names = TRUE)

anova(check.mod, check1.mod, test = "F")

glm.res <- dfx$Lines - predict(poiss.mod, type="response")
mean(abs(glm.res))
sum(abs(glm.res) <= 2000) / length(dfx$Lines)

log.res <- dfx$Lines - exp(predict(check.mod, type="response"))
mean(abs(log.res))
sum(abs(log.res) <= 2000) / length(dfx$Lines)

plot(dfh$RecDate, glm.res, xlab = "Date", ylab = "Resdiuals")
hist(glm.res, main = "Histogram of residuals")
qqnorm(glm.res, main = "Normal Q-plot for residuals")

#Compare linear with log lin model outcome / deviation
boxplot(lin.res, glm.res, xlab = "Linear vs Log linear model", ylab = "Deviation in orderlines", main =
"Comparison of residuals")
```

```
#--------------------------------------
# log lin model
#--------------------------------------
log.mod <- glm(log(Lines) ~ LTW + YW + factor(VW) + factor(DayNum) + factor(Holiday) + LHC + SHC + LDAC +
MonthClose, data = dfx, family = "gaussian")
summary(model_mod)
log.mod$coefficients
hist(log.mod$residuals)
shapiro.test(log.mod$residuals)
qqnorm(log.mod$residuals)

plot(dfx$RecDate, dfx$Lines, xlab = "Date", ylab = "Orderlines", main = "Plot of actual and modelled received
orderlines")
lines(dfx$RecDate, exp(model_mod$fitted.values))

res.logmod <- sum((dfx$Lines - exp(log.mod$fitted.values))^2)
res1.logmod <- dfx$Lines - exp(log.mod$fitted.values)
hist(res.logmod)
shapiro.test(res.logmod)
sum(abs(res1.logmod) <= 2000) / nrow(dfx)

#--------------------------------------------------
# Validation
#--------------------------------------------------
require(boot)
#cost <- function(r, pi){mean((r-pi)^2)}

#Cross-validation of linear model
lm.cv <- cv.glm(dfx, lin.mod)
lm.cv$delta
sqrt(lm.cv$delta[1])

#Cross-validation of Huber model
hub <- fitted(rlmod)
hub.diag <- glm.diag(rlmod)
cv.err <- mean((dfx$Lines - hub)^2/(1 - hub.diag$h)^2)
sqrt(cv.err)

#Cross-validation of Log linear model
cost <- function(r, pi){mean((exp(r)-exp(pi))^2)}
log.cv <- cv.glm(dfx, log.mod, cost)
log.cv$delta
sqrt(log.cv$delta[1])
```