Transactions on Machine
Learning and Data Mining
Vol. 10, No. 2 (2017) 57-65

ISSN: 1864-9734,
ISBN: 978-3-942952-54-5

ibai Publishing
www.ibai-publishing.org

Simulation based Approach for Special Discounting Threshold

Shashank Shekhar[1], Kaushik Shankar[1], Sushanta Kumar Mishra[1], Sam Pritchett[1], and Parag Chitalia[1]

[1]VMware Inc., Palo Alto, USA
{sshekhar, shankark, skumarmishra, spritchett, pchitalia}@vmware.com

Abstract. There are many existing solutions available that provide a recommendation to increase the threshold of the minimum order size to avail special discounts (SPF) by minimizing the loss in order conversion. The goal of such solutions is to maximize the earning and maximize quote to order conversion. However, these solutions are primarily based on the relationship between order size and quote to order conversion and utilize historical conversion rates at different price points (List Price-Discount). This paper illustrates the use of Single Item Single Period Newsvendor problem and Extreme Gradient Tree Boosting methods to simulate a case for increasing order thresholds and calculating the corresponding loss in order conversion. The boosting method was applied on the training data and simulated and cross-validated across three test data sets to test three different price change scenarios. We computed the expected quotes at the end of the period and expected conversion at the end of the period to maximize profit.

Keywords. Extreme Gradient Tree Boosting, SPF, List Price, B2B, Single Item Single Period Newsvendor problem, quotes, conversion

1 Introduction

VMware (VMW) is a virtualization, end-user computing and cloud company with annual revenues of USD 7 BB (as of 2016) and a market cap of USD 38 BB [1]. Total revenue is comprised of license revenue of $2,794 million and services revenue of $4,299 million. VMware sells products in the Software Defined Data Center (vSphere, VSAN, NSX for computing, storage & network virtualization respectively), end-user computing (Airwatch, Horizon, and Fusion/Workstation) and cloud. These are all sold to B2B customers.

The prices VMware (VMW) of products at VMW have been rarely changed over time [2]. However, sales reps could offer discounts via SPF flag which is a special discretionary discount which should ideally be given for large orders. The Pricing Business Unit was keen to figure out a way to get to optimal list prices at VMW. Sometimes order quotes are drafted using a Special Pricing Form (SPF). Typically, large order sizes qualify under SPF quotes. The sales reps provide a discount to the customer for such orders called as SPF discount. The paper focusses on building a case for increasing SPF order size thresholds from the current value to a higher value (50% increment/100% increment/150% increment). The goal of the business was to reduce costs by reducing the number of SPF orders, thereby maximizing profit and evaluating the potential risks associated with it.

We attempted to determine the optimum threshold for the Special Pricing orders by lowering the discount value. Here, the discount rate was used as a proxy for order size threshold. Lowering the discount will lead to a higher order price which results in a higher threshold. Instead of analyzing the quotes in a descriptive manner, here we focused on developing a predictive model and simulate what-if scenarios for different SPF thresholds.

The objectives of our research are:

- Analyze historical SPF data/price and discount trends/order conversions
- Identify significant drivers of SPF order conversion
- Compute expected quotes (demand) and expected conversion at the end of the period
- Build a model to determine optimum deal size in order to qualify for an SPF transaction.

2 Solution Framework

2.1 Strategy

A traditional discounting optimization problem would use order conversion based on discounts given on different deal size. But since the threshold for SPF discount was fixed, the data did not indicate the probability of order conversion for an increased threshold. The data was simulated to get a non-SPF discount for quotes which qualified for SPF discounts and then the probability of quote to order conversion was computed. The trend of SPF discounts was visualized using the quote to order matching method. This included product-specific weighted discounting patterns, GEO/Region patterns, seasonality analysis, reason code, etc. and also included assessing if List Prices are within acceptable ranges (e.g. are we using SPF more because of high list price?)

SPF requests are useful indicators to infer the customer's assignment of value to a product. SPF flag is a discount that can be given by a sales representative on request, mostly given on large order sizes. An increase in the threshold of minimum order size for SPF discount can potentially lead into non-conversion of the order and hence can result in revenue loss for VMware, Inc. The intent of the project was to determine optimum threshold such that the increase in overall price offsets the revenue loss due to a decrease in conversion, thereby maximizing the overall profit. The following detailed steps were planned:

- Determine magnitude of SPF orders and their impact on overall revenue
- Visualize the trends of SPFs by type using a quote to order matching method
- Compute expected deals won/lost using Single Item Single Period Newsvendor problem [3]
- Identify drivers of discounting behaviors
- Recommend policy changes supported by findings

2.2 Solution Approach

The objective was to determine the magnitude of increment in minimum order size for SPF discounting which results in a minimum loss in conversion. There was a need to create a base set of research data (Customer -> Approved Quotes -> Order) which needed to be standardized and approved to be used for modeling. Exploratory data analyses were done to depict a relationship between the amount of discount and actual bookings. The analyses showed GEO level differences in behavior patterns. Clustering was done on the behavior patterns and the clusters were aligned to decision tree [4] output to determine the drivers of outliers.

The estimation of expected quotes and the conversion gave us an indication of the proportion of win/loss in different scenarios. The modeling was done to estimate the proportion of win/loss for different order size threshold. The final recommendation was the new threshold based on the conversion probability for each evaluated threshold.

3 Modeling

3.1 Dataset Creation

The data matrix for the model was aggregated at an order-SKU level using Greenplum and Hadoop. This is because our entity is a combination of an order and an SKU. Specific business level judgments were applied like:

- Dataset for a time period of FY14Q1 to FY16Q4
- Specific products like vCloud Air that had a subscription model/enterprise license agreements (ELAs) that have pre-determined prices transactions were removed while creating this dataset
- Given large amounts of data, the data was picked from Hadoop, where we took the last 5 years of dataset related to the site made available by IT in a flat tabular structure at order-SKU level.
- The continuous variables were subjected to coarse classing and fine classing depending on the nature of variable to arrive at bins which can be used as an input in the model.

- Opportunity data was added to the dataset to include the following:
  - Time from Opportunity to Quote
  - Time from Quote to Conversion
  - Outliers where a quote is generated after the Opportunity Close date
- The key independent variables were Quote Value, Total Deal Score, Product platform, SPF deal size, industry vertical, sales classification etc.
- The dependent variables were a booking flag which indicated quote to order conversion

3.2 Estimation of expected profit based on Deals Won/Lost

We used part of Single Item Single Period Newsvendor problem to estimate the numbers of deals won/lost at the end of a period to give an indication of a net change in revenue.

Let;

p = List Price per unit

i = Invoice price per unit

r = rebate/discounts other than SPF per unit per period ($ per unit per period)

s = SPF discount in $ per unit

k = Incremental revenue due to change in threshold deal size for SPF discount

Q = quantity sold (no. of units)

x = demand per period (No. of quotes per period); stochastic with F(x) as the cumulative distribution of total demand and f(x) as the probability distribution function

π = Incremental profit

Let μ = E[x]be the expected demand per period

Let I(Q) = Expected deals won at the end of the period

Let B(Q) = Expected deals lost at the end of the period

$$\mathrm{I(Q)} = \int_0^Q (\mathrm{x}-\mathrm{Q})\mathrm{f(x)dx}\ \ (continuous\ demand);\ \mathrm{I(Q)} = \sum_0^Q (\mathrm{x}-\mathrm{Q})\mathrm{f(x)}(discrete\ demand)$$

$$\mathrm{B(Q)} = \int_Q^\infty (\mathrm{x}-\mathrm{Q})\mathrm{f(x)dx}\ \ (continuous\ demand);\ \mathrm{B(Q)} = \sum_{Q+1}^\infty (\mathrm{x}-\mathrm{Q})\mathrm{f(x)}(discrete\ demand)$$

It can be shown that

B'(x) = F(x)-1

I'(x) = F(x)

For discrete demand it can be shown that $\Delta\{\sum_{Q+1}^{\infty}(x-Q)f(x)\}$ =F(Q)-1 where $\Delta\{f(r)\} = f(r) - f(r-1)$

Expected Incremental Profit = E[π] = Expected Revenue – Standard Rebate/Discount – SPF Discount + Incremental Revenue due to change in SPF threshold

E[π(Q)] = E[min(Q,x)]*p -r*I(Q) -s*I(Q) + K*I(Q)

It can be shown that maximizing expected profit is same as minimizing the number of deals lost (in-turn maximizing the number of deals won) due to increase in threshold deal size for SPF discount.

3.3 Model for Quote to Order Conversion

For the model, the target variable is a booking flag that indicates whether a quote is converted given that the quote is approved. It makes sense to only consider approved quotes in this regard, as a quote has nil probability of being converted if it is not approved by VMWare. We also confined ourselves to a very granular level of data to provide robustness to our model. Below we described the filters we applied to our training data:

- We considered only the quotes located in AMER region.
- We only considered quotes that contain at least one licensed product.
- We only considered a standard form of SPF and excluded quotes where blanket or recurring SPF were applied.
- We also removed all the quotes containing AirWatch products (MOBILE-EMM-EMM).

Our modeling approach is driven by the idea to explore the fact that had we offered a non-SPF discount on the same set of quotes, then what would possibly have happened in terms of probability of conversion.

Assuming that the current order size threshold for SPF discount is $X, we derived five deal size buckets, which are:

- < $ X
- $ X – $ 1.5X
- $ 1.5X - $ 2X
- $ 2X - $ 2.5X
- > $ 2.5X

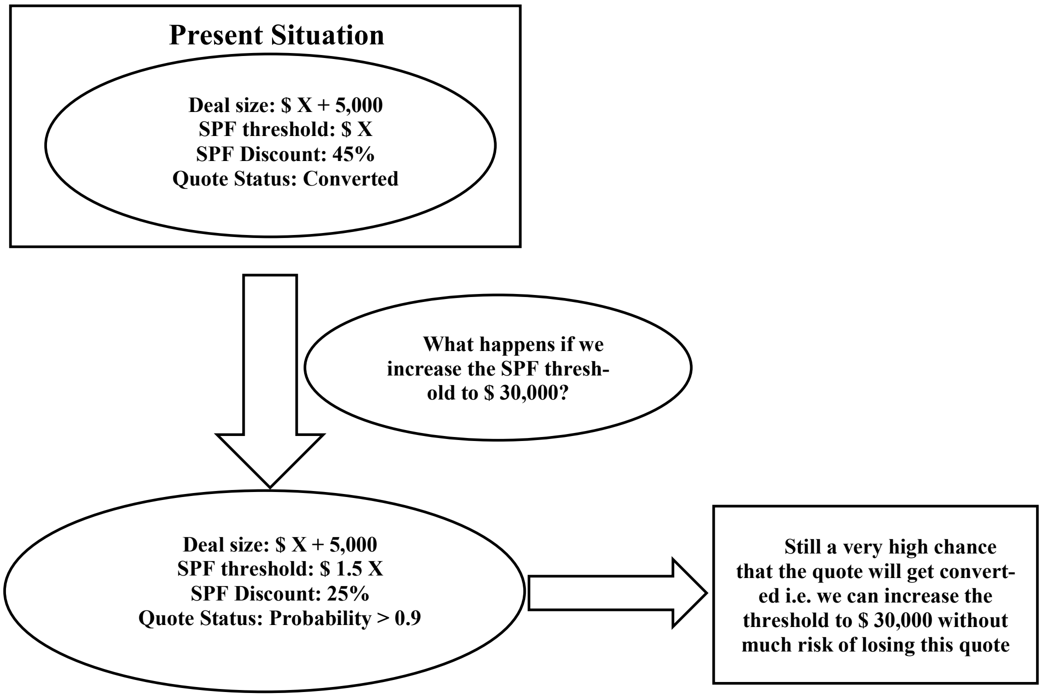

Fig. 1. Current and Proposed Scenarios

For each of the above deal size buckets we obtained the historical discounts given in quotes belonging to both SPF and non-SPF, **w**e trained an Extreme Gradient Boosting (xgboost) [5] model of the filtered data. We scored our models based on a 4-fold cross-validation and selection evaluation criteria as the area under the ROC curve (AUC) [6]. We considered different hyper-parameter combinations and chose the best model offering the best out-of-fold estimate (around 0.7).
For the test data, we created three different data sets, by keeping all other variables as is in the train data but changing the discount to non-SPF average discount in the respective buckets.

Table 1. Test Dataset 1: SPF threshold $ 1.5X (X = Current SPF threshold)

Deal Size	Discount Type
Less than $ X	Non-SPF
Between $ X and $ 1.5X	Non-SPF
Between $ 1.5X and $ 2X	SPF
Between $ 2X and $ 2.5X	SPF
Greater than $ 2.5X	SPF

Table 2. Test Dataset 1: SPF threshold $ 2X (X = Current SPF threshold)

Deal Size	Discount Type
Less than $ X	Non-SPF
Between $ X and $ 1.5X	Non-SPF
Between $ 1.5X and $ 2X	Non-SPF
Between $ 2X and $ 2.5X	SPF
Greater than $ 2.5X	SPF

Table 3. Test Dataset 1: SPF threshold $ 2.5X (X = Current SPF threshold)

Deal Size	Discount Type
Less than $ X	Non-SPF
Between $ X and $ 1.5X	Non-SPF
Between $ 1.5X and $ 2X	Non-SPF
Between $ 2X and $ 2.5X	Non-SPF
Greater than $ 2.5X	SPF

For each of the test sets, we obtained the probability of conversion in each of the buckets. To compare the simulated situations in test sets and the present discounting situation, we evaluated the mean conversion rate in each of the above five buckets result and these are shown in the table**s** below:

Table 4. Mock-up representation of Deals Lost for Deal Size threshold $ 1.5X (X = Current SPF threshold)

Deal Size Bucket`	Total Deals	Percentage Conversion Before	Percentage Conversion After	Deals Lost
≥ $ 2.5X	13642	0.201%	0.201%	0
$ 2X - $ 2.5X	1535	0.244%	0.244%	0
$ 1.5X - $ 2X	2034	0.268%	0.268%	0
$ X - $ 1.5 X	2393	0.276%	0.273%	7
< $ X	2166	0.31%	0.31%	0

The model indicated a significantly low decrement in order conversion for a 50% increase in the deal size threshold. The deal size bucket \$ X - \$ 1.5X is the only bucket affected by the proposed increase in the deal size threshold.

A similar evaluation was done for scenarios where threshold deal size was increased to \$ 2X and \$ 2.5X. The difference between the incremental revenue generated by the won deals and the revenue lost due to lost deals was computed.

Table 5. Mock-up representation of Deals Lost for Deal Size threshold \$ 2X (X = Current SPF threshold)

Deal Size Bucket`	Total Deals	Percentage Conversion Before	Percentage Conversion After	Deals Lost
≥ \$ 2.5X	13642	0.201%	0.201%	0
\$ 2X - \$ 2.5X	1535	0.244%	0.244%	0
\$ 1.5X - \$ 2X	2034	0.268%	0.208%	121
\$ X - \$ 1.5 X	2393	0.276%	0.273%	7
< \$ X	2166	0.31%	0.31%	0

Table 6. Mock-up representation of Deals Lost for Deal Size threshold \$ 2.5X (X = Current SPF threshold)

`Deal Size Bucket	Total Deals	Percentage Conversion Before	Percentage Conversion After	Deals Lost
≥ \$ 2.5X	13642	0.201%	0.201%	0
\$ 2X - \$ 2.5X	1535	0.244%	0.188%	80
\$ 1.5X - \$ 2X	2034	0.268%	0.208%	121
\$ X - \$ 1.5 X	2393	0.276%	0.273%	7
< \$ X	2166	0.31%	0.31%	0

3.3 Determining Relative Importance of drivers of Discount

The relative importance of the factors influencing SPF discounts was determined by running Gradient Boost Method [7]. The target variable for finding the relative importance of driver was Perpetual Blended Discount. The measures are based on the number of times a variable is selected for splitting, weighted by the squared improvement to the model as a result of each split, and averaged over all trees. Here we fitted multiple trees in each stage in a one-vs-all way.

4 Conclusion

The solution provided for optimizing the order size threshold for SPF discounting VMware Inc. is unique and distinctive because the thresholds have been constant since inception and there has been hardly any instance of changing the threshold. The estimation of expected incremental profit based on deals won/lost after changing the SPF threshold is industry-first and there is no existing solution which does that. Also, the implemented approach was very different from standard optimization methodolo-

gies where historical changes are utilized to arrive at recommended price changes. From an algorithmic perspective, the approach is unique and there is no such implementation that utilizes simulation and boosting technique to arrive at the threshold which maximizes profitability.
The recommendation to increase the order size threshold for ~60% of the SKUs was given to the business. It was also recommended to increase the threshold in a phased manner and actively track the quote to order conversion. SPF usage is an indicator of customers' perception of extended list price and hence the recommendation was to make the change independent of other changes and to keep the standard rebate/discount constant as earlier. This will enable us to directly measure customers' perception of value. The recommendations from this project will help maximize profit for VMware, Inc. and also actively track price perception.

References

1. VMware Inc. Annual Report, http://ir.vmware.com/annuals.cfm
2. Rama K., Shekhar S., Kiran J., Rau R., Pritchett S., Bhandari A., Chitalia P.: List Price Optimization Using Customized Decision Trees. In: Perner P. (ed.) Machine Learning and Data Mining in Pattern Recognition. Lecture Notes in Computer Science, vol. 9729, pp. 88-97, Springer, Cham (2016).
3. Gallego G., Moon I.: The distribution free newsboy problem: review and extensions. Journal of the Operational Research Society 44, 825- 34 (1993).
4. Quinlan J.R.: Induction of Decision Trees. Machine Learning 1, 81-106, Kluwer Academic Publishers (1986).
5. Chen T., Guestrin C.: XGBoost: A Scalable Tree Boosting System. https://arxiv.org/abs/1603.02754 (2016).
6. Robin X., Turck N., Hainard A., Tiberti N., Lisacek F., Sanchez J.C., Mülle M.: pROC: an open-source package for R and S+ to analyze and compare ROC curves. BMC Bioinformatics 12, 77 (2011).
7. Friedman J.H.: Stochastic gradient boosting. Computational Statistics & Data Analysis 38(4), 367–378 (2002).
8. Dowle M., Srinivasan A., Short T., Lianoglou S. with contributions from Saporta R., Antonyan E.: https://cran.r-project.org/web/packages/data.table/data.table.pdf
9. Hastie T., Tibshirani R., Friedman J.H.: The elements of statistical learning: Data mining, inference, and prediction. Springer Verlag, New York (2001).
10. Breiman L., Friedman J.H., Olshen R.A., Stone C.J.: Classification and Regression trees. Wadsworth & Brooks/Cole Advanced Books & Software. Monterey, CA (1984).
11. Cohen J., Cohen P., West S.G., Aiken, L.S.: Applied multiple regression/correlation analysis for the behavioral sciences. 2nd edn. Hillsdale, NJ: Lawrence Erlbaum Associates (2003).