Reprice with Confidence:

Dynamic Pricing with Robust Time-series Forecasting

BMW Group Team:
Dr. Erik Pfeiffenberger
Dr. Steffen Illig
Faculty Advisor:
Prof. Negin Golrezaei
PhD Mentor:
Jason Cheuk Nam Liang


Pan Asavamongkolkul


Elva He

Overview

Before we joined BMW Group, the "Advanced Analytics" team had a model to predict cars' demand elasticity with limited functionality; this model had moderate errors and could not capture seasonality. Our project goals are to improve this model and apply optimization to advise car prices across BMW headquarters' owned dealerships in Germany. **Our approach has 4 steps.**

1. **Data Cleaning Pipeline** →
2. **Time-series Elasticity Estimation** →
3. **Profit Optimization** →
4. **Interative Dashboard**

Data

Transactional Data: This data records sales orders for BMW headquarter-owned dealerships. Key variables include pricing information, customer types, and contract types.

Car Features: This data records characteristics of each car model e.g. fuel type, body style, production date.

Corporate KPI: This data is user-input KPI, used to set constraints for optimization. Currently, this includes sales target, and allowable price swing.

Data Challenge 1: There are low sales in model code level

	Amount
Series	~20
Development Code	~80
Model Code	~1000

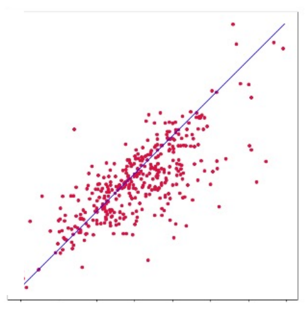

BMW desired that the model predicts in model code level. However, since there are around 1,000 model codes, the average monthly sales of each model code are low, leading to noisy features. To solve this problem, we assumed that the elasticity of the model code is equivalent to that of its corresponding development code (the assumption was accepted by BMW, at 17% error). **Therefore, the prediction model will predict at the development code level.**

Data Challenge 2: Price & Sales Relationship is Biased

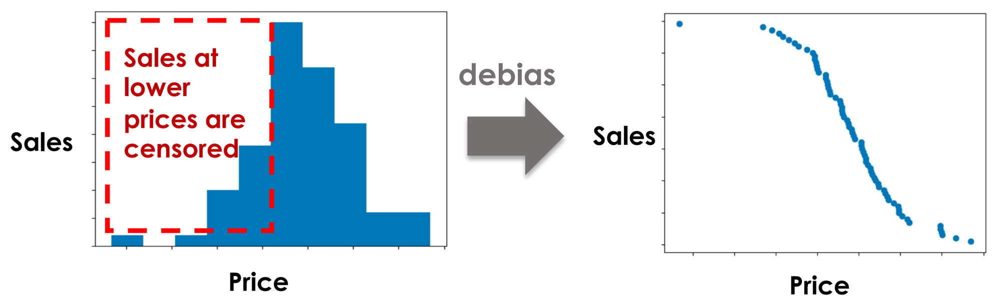

The left histogram shows that most sales are around the average price; it is counterintuitive since there should be more sales at a lower price. However, this is because not all price ranges were offered to every customer. **To un-bias this data, we used a cumulative sum of sales method** with an assumption that customers are always willing to purchase the same cars at the lower price.

Methodology

Elasticity Estimation

We aim to estimate the relationship between the sales and the discount of each car model. However, since each model has different scales of sales, we focus on the sale uplift instead. Tree-based models are selected as they can well capture the non-linearity in the data. **Our models' output can estimate the elasticity and accurately separate the "month trend effect" and "discount effect" to better understand the impact of discounts on sales.**

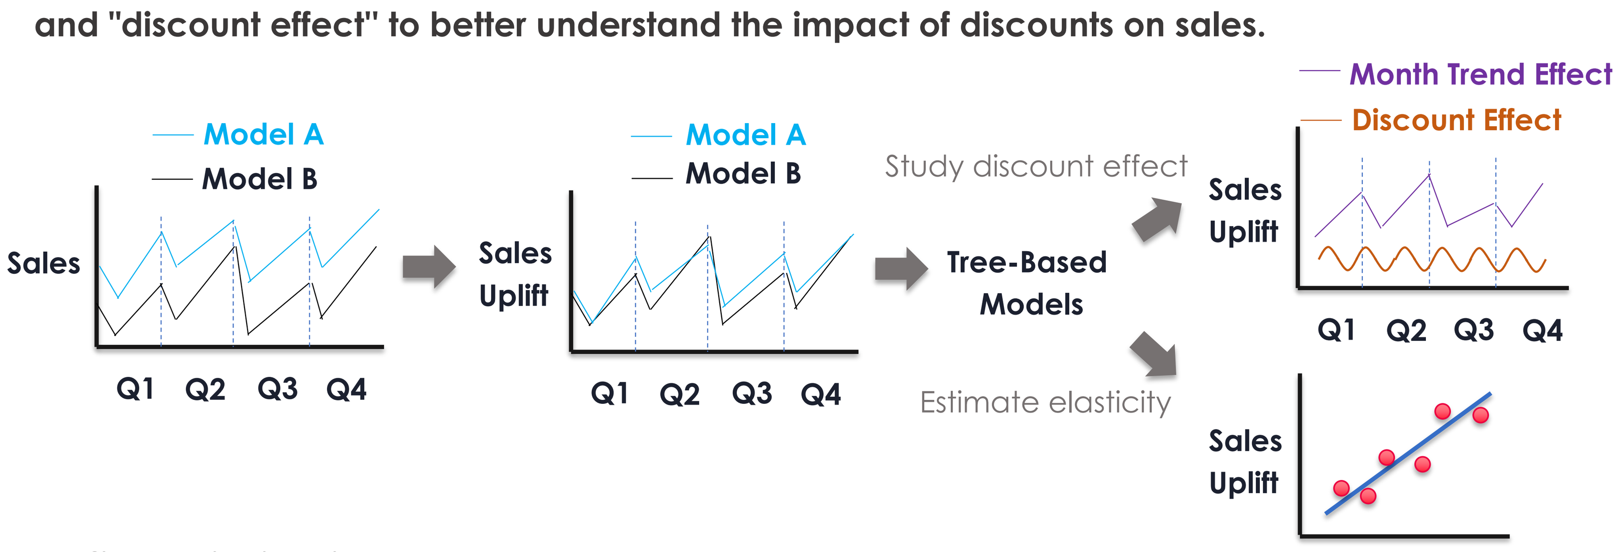

Profit Optimization

The profit optimization is framed as a quadratic programming problem that can be timely solved using open-source solvers (SCS).

Objective: Maximize Total Profit

$$\max \underbrace{\sum_{i=1}^{m}}_{\text{Number of Models}} \left[p_{i,opt} - p_{i,BE}\right] D_i(p_{i,opt})$$

Constraint 1: no abrupt change in price

$$0.9 \leq discount\ uplift_i(p_{i,opt}) \leq 1.1$$

Constraint 2: control sales share (KPI)

$$\underbrace{L_i}_{\text{Lower bound (\%) sale share for model i}} * \sum_{i=1}^{m} D_i(p_{i,opt}) \leq D_i(p_{i,opt}) \leq \underbrace{U_i}_{\text{Upper bound (\%) sale share for model i}} * \sum_{i=1}^{m} D_i(p_{i,opt})$$

Variables:
$P_{i,opt}$: optimal price for model i
$P_{i,BE}$: breakeven price for model i
$D_i(P_{i,opt})$: predicted demand for car model i at the optimal price

References

- Ferreira, Kris & Lee, Bin & Simchi-levi, David. (2015). Analytics for an Online Retailer: Demand Forecasting and Price Optimization. Manufacturing & Service Operations Management.
- Bi, W., & Liu, M. (2014). Product demand forecasting and dynamic Pricing Considering consumers' mental accounting AND Peak-End Reference Effects. Journal of Applied Mathematics, 2014

Results

Model	Testing Period		
	R^2	MAE	SMAPE
Regression Tree	0.85	13.3	22.6%
Random Forest	0.87	12.3	22.0%
Light GBM	0.90	11.4	**21.5%**

Important Key Features

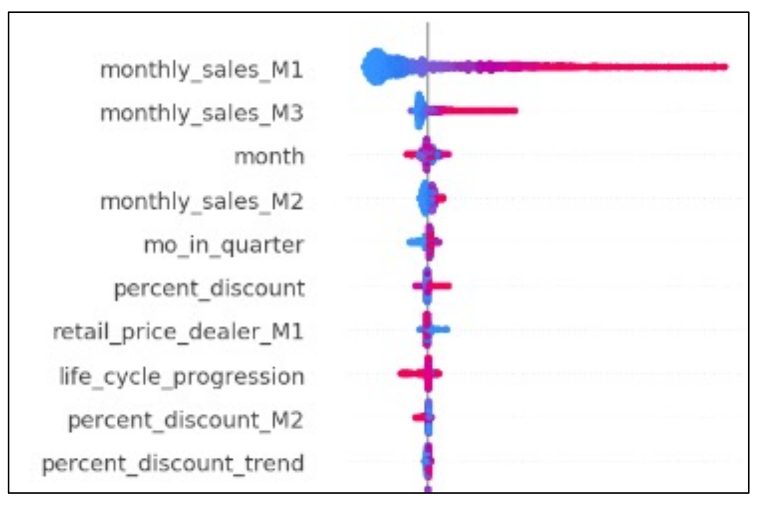

Our models reduce the error from the baseline by 14% (the baseline has SMAPE of 25%).

Note : Baseline model is BMW's past models

Actual Sales	Optimal Sales	Actual Revenue	Optimal Revenue	Actual Profit	Optimal Profit
Baseline	-11%	Baseline	-9%	Baseline	+7%

The optimal increase in profit based on the given KPI is 7%. We estimate the elasticity as a linear function (demand = m · price + c). Therefore, the **robustness test** was done on the slope and coefficient with a range of 0-10% error. The results show that even with the worst-case scenario, the optimal profit is still positive (2%).

Profit Increase

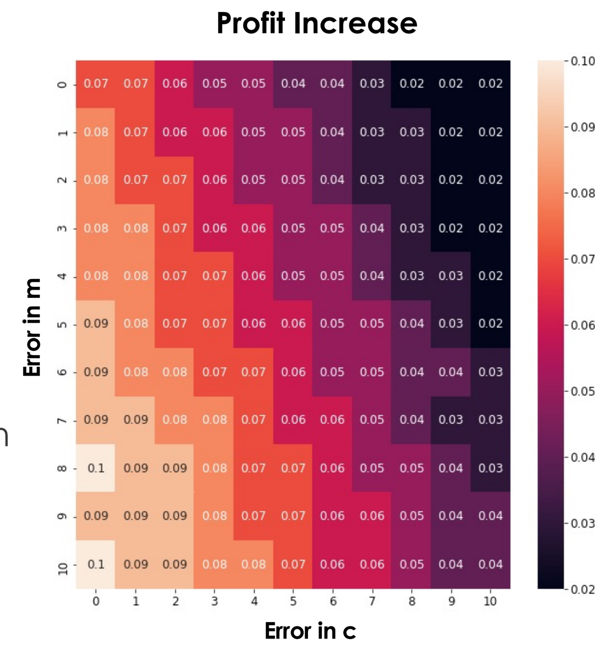

7% increase in profit

- Lower bound: 2% increase
- Upper bound: 10% increase

Timeline

Scope definition	February – April
Data collection / cleaning	April – July
Elasticity estimation	May – July
Front-end (Tableau) development	June – August
Profit optimization	July – August
Handover	August
Part-time	February – April
Full time	July – August

February · March · April · May · June · July · August

Deliverables

We provide BMW with **easy-to-modify Python scripts** and **interactive Tableau dashboard**.

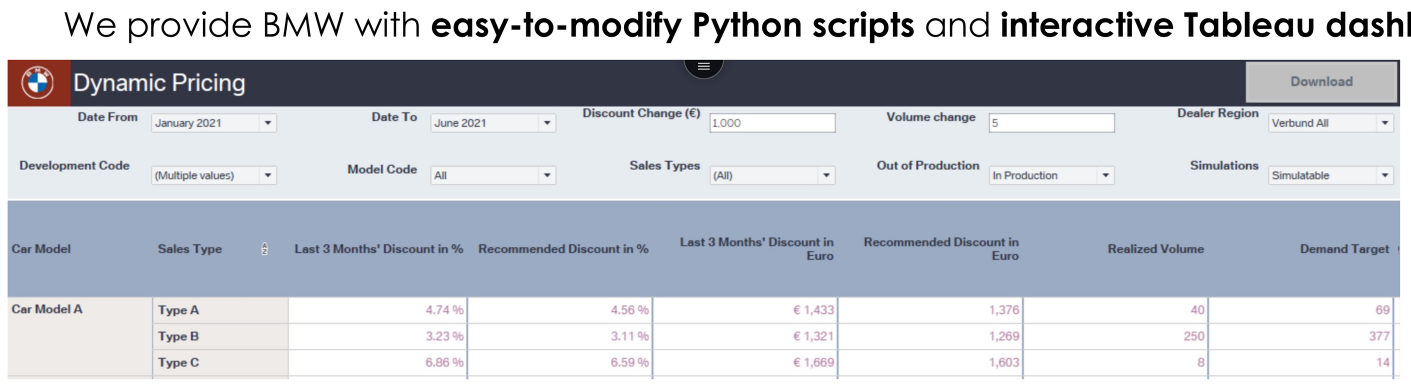

User-input

Optimized results

Note: The example here is for illustrative purposes only and do not represent actual numbers

Next Steps

With proper data, this project can also be extended to independent BMW dealerships or other countries' markets.

Our contribution to the BMW Group Advanced Analytics team was pioneering work using optimization on sales/marketing use-case. We showed the impact of this project, which will pave the way for more analytics works, especially in optimization in the future.