The Value of a Dollar:
Optimizing Bidding Strategies for Suppliers


MBAn Team: Rocky Xie, Jessie (Jiesi) Zhou

MANAGEMENT BUSINESS ANALYTICS | wayfair

Faculty Advisor: Prof. Georgia Perakis
PhD Mentor: Manuel Moran Pelaez
Wayfair Sponsor: Rajesh Bawa, Stefan Poninghaus

Problem Statement

*How can we give **visibility** and **insights** to Wayfair **suppliers** to help them launch better advertising campaigns, while balancing **Wayfair** revenue and **customer** experience?*

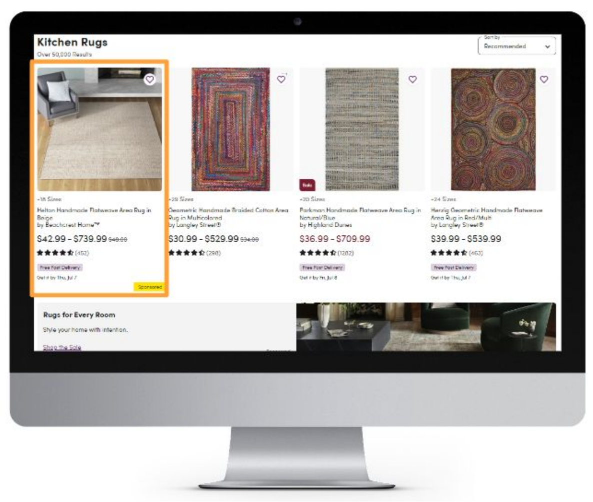

Context

- Suppliers **bid** for their products to be advertised on Wayfair.
- Suppliers are charged for each **click** on their advertised products.
- The final product **placement** is determined by bid and product quality, which influences the product's **performance**.
- Currently, Wayfair does not provide suppliers informed **recommendation** on advertising setups to reach intended outcomes.

Supplier A:
I want to sell as much as I can…

Inside the Brain of suppliers
- **How much** should I bid?
- **What products** should I advertise?
- **What budget** do I need to reach my goal?

Supplier B:
I want to improve our brand awareness…

Our Proposed Approach

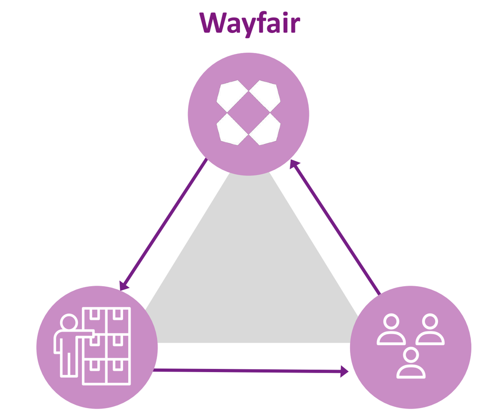

A bid recommendation engine powered by a mixed-integer optimization model to maximize target outcomes while satisfying Wayfair constraints and balancing customer experience.

Dataset

Scope

~125K rows of daily product level data (May 15th - June 15th 2022) from area rugs on browse pages on Wayfair.com

Advertised Products	Auction Outcome	Ads Performance
Product ID	Date	Return on Ad Spend
Price	Page Number	Conversion Rate
Product Quality Score	Slot ID	Clicks
Bid	Cost per Click	Impressions

Exploratory Data Analysis

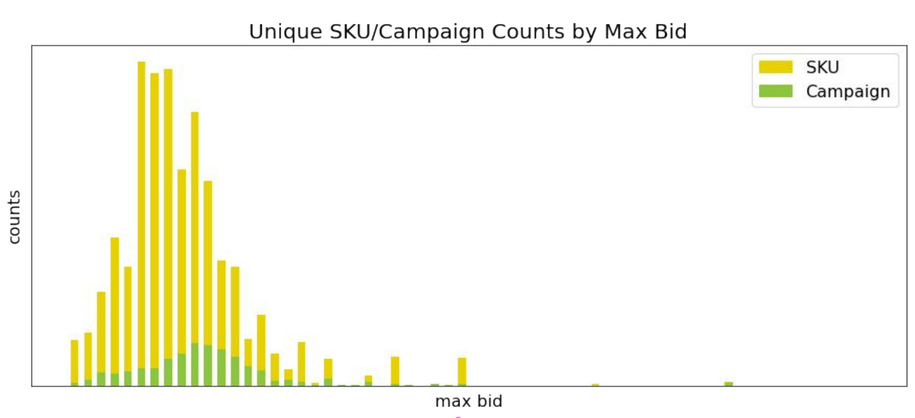

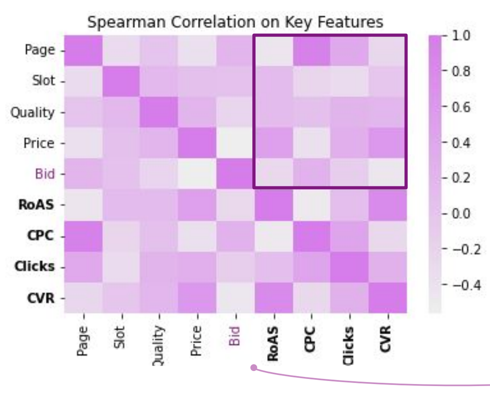

1. The majority of suppliers **bid below** certain dollar amounts with a few suppliers bidding excessively.
2. There is a **lack of clear relations** between explanatory and outcome variables.

3. After aggregating by bid, some **weak trends** emerge in outcome variables.

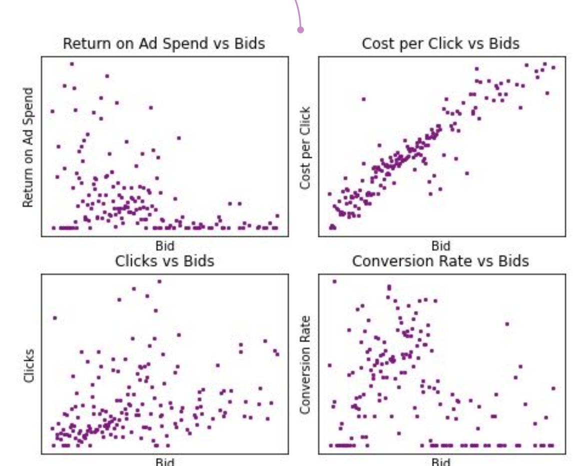

Explanatory variables
Page & Slot: the position of a product ad on the Wayfair page
Quality: a score that measures the product's quality
Price: a product's price
Bid: a supplier's maximum willingness to pay for a click on their ads

Outcome Variables
RoAS (return on ad spend): ratio of ad generated earnings to spending
CVR (conversion rate): ratio of number of orders from ad to number of clicks
Clicks: number of clicks the ad receives
CPC (cost per click): what Wayfair charges the supplier per customer click

Methodology

Reward Estimation

- Assume that our outcome variables are a function of **bids** and some unmeasurable noise, namely:

$$Metric = f(bids) + \varepsilon,\ where\ E[\varepsilon] = 0$$

- The function f is unknown; our objective, to **estimate rewards**, is equivalent to finding this function f.
- We propose to find f by **minimizing the difference** between the observed data and our fitted curves.
- We then add robustness by finding best fit curves from hundreds of **stratified bootstrap samples** and then using confidence intervals and averaging to estimate rewards.
- This procedure also allows reward estimation for **individual** SKUs by first finding similar products using k-NN.

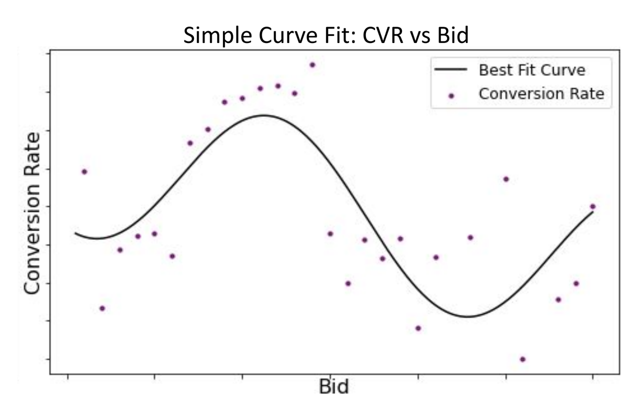

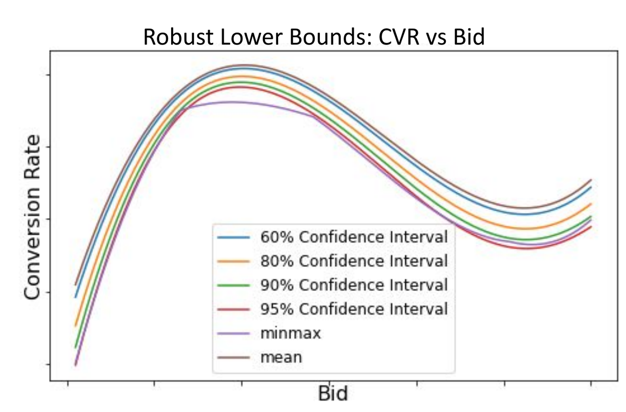

Optimization

- With the trade-offs between **interpretability**, **feasibility** and **flexibility**, we propose two optimization models to solve for optimal bids given estimated rewards.
- The main objective of optimization is to **maximize** our targeted **rewards** (e.g., conversion rate) while ensuring other metrics (e.g. budget) still meet **expectations**.
- We further consider **robust** rewards to stabilize our recommendations and deal with the large amount of uncertainty in our data.

MOO (Multi-Objective Optimization) Formulation

$$\max_{\boldsymbol{x}} \sum_{i\in P}\Big[\omega_1\mathbf{S}(x_i)+\omega_2\mathbf{CV}(x_i)+\omega_3\mathbf{CL}(x_i)\Big]x_i$$
$$s.t. \quad \sum_{i\in P} x_i = 1$$
$$\sum_{i\in P} c_i x_i \le b$$
$$\sum_{i\in P} cpc_i x_i \le \Omega$$
$$x_i \in \{0,1\}$$

Robustness:

$$\max_{\boldsymbol{x}} \mathbf{S}(x_i) \quad \Rightarrow \quad \max_{\boldsymbol{x}} \min_{\mathbf{S}\in\mathcal{P}_\mathbf{S}} \mathbf{S}(x_i)$$

Results

Bids Recommendation

Use Case 1
High Intent

Suppliers are **short-term** focused, mainly on immediate conversion. Recommended bids will tend to have higher **conversion rates**.

Use Case 2
Brand Awareness

Suppliers are **long-term** focused and concerned more about ads impression. Recommended bids will tend to have higher **clicks**.

Recommended Bids with Robust Curve Fits

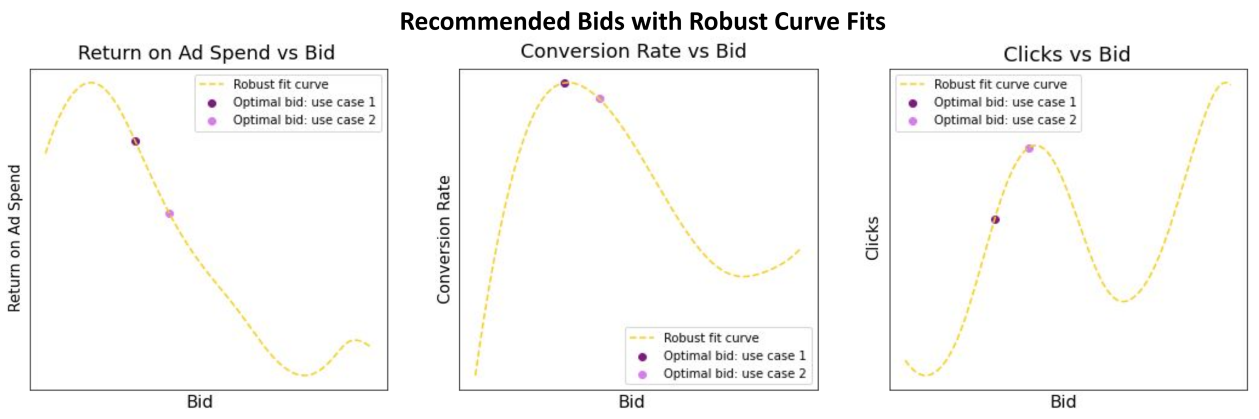

Monte Carlo Simulation

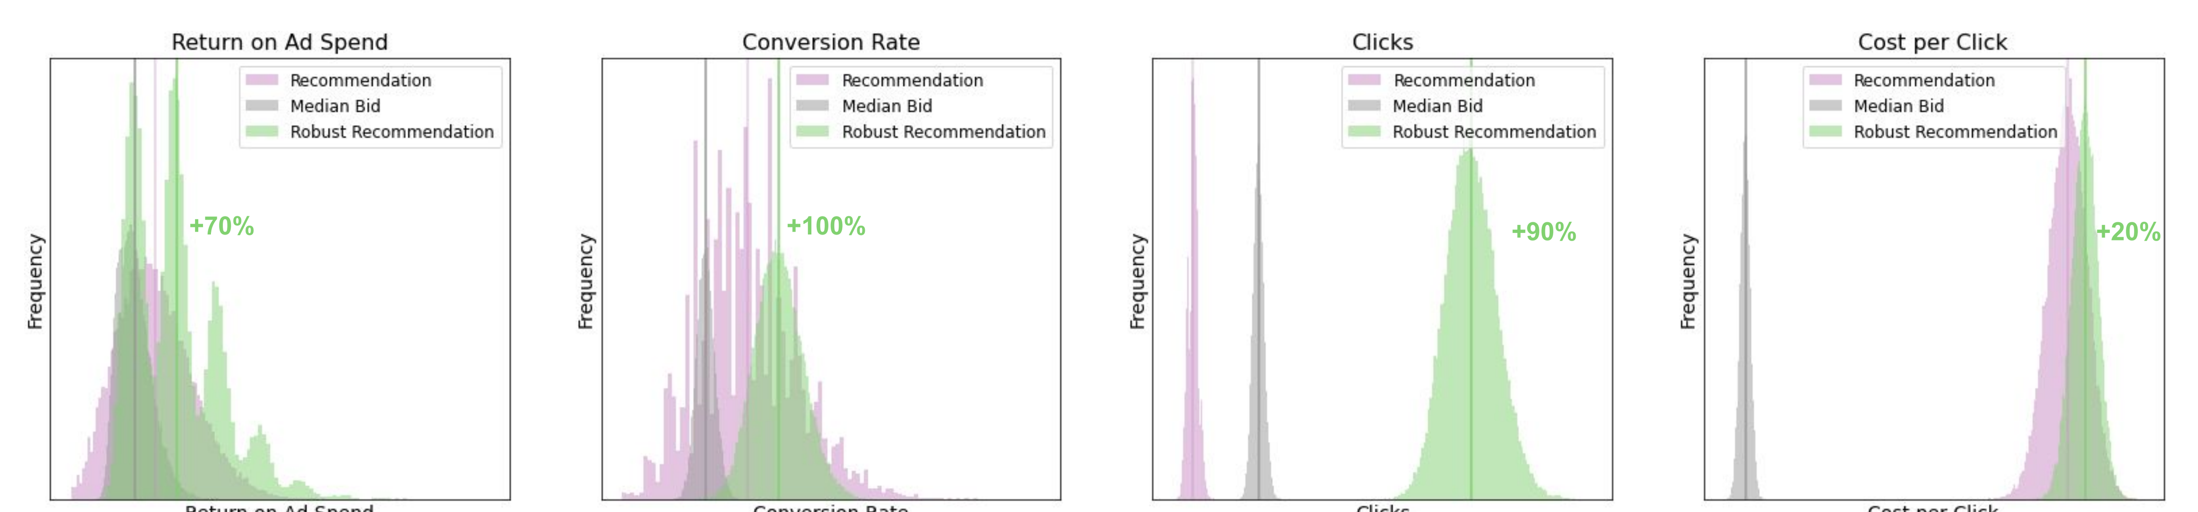

In the case above, we imagine on <u>June 15th</u>, an area rugs supplier with a focus on <u>immediate conversion</u>, wants to advertise her product. We provide a bid recommendation training on previous month's data, then simulate, using data from the following month, the outcomes.

Following our robust recommendation, we project an increase in conversion of roughly 2x over if she puts in the median bid of the previous month (current Wayfair recommendation).

Impacts

Improve supplier
SATISFACTION
by helping launch better informed and goal-oriented advertising campaigns

Projected Wayfair monthly revenue increase
$350K
in Area Rugs alone by more competitive bids and increased product sales

A generalizable data-driven
FRAMEWORK
for supplier advertising recommendations

Future Considerations

Improvement of dataset building

The quality of our recommendations is largely determined by the quality of our rewards estimation. Our current approach, although a solid start, could benefit from more sophisticated ML models or causal analysis on the effects of bid on outcome metrics.

Expansion on optimization

Our framework is easily generalizable to consider more use cases (we currently have two), or additional levers (we currently only use bid). Furthermore, our pipeline can be scaled up by increasing number of parallel processes for more robust curve fitting.

Pilot testing with suppliers

The next immediate step should be to run a pilot program with some suppliers to collect experimental data and concretely measure preliminary impacts.

Auto-bidding

Optimal bids from our recommendation engine can serve as the input to automatically adjust supplier bids to better reach supplier goals.