From the Clouds to the Trenches

Learning to Manage the Marketplace

Eren Manavoglu, Partner Scientist

Microsoft Advertising, AI & Research

From the Clouds to the Trenches

Or How I Learned to Stop Worrying and Love Counterfactuals

Eren Manavoglu, Partner Scientist

Microsoft Advertising, AI & Research

Overview

Very Brief Intro to Search Advertising

Marketplace Objective

Marketplace Optimization

Understanding the Marketplace

I Very Brief Intro to Search Advertising

II Marketplace Objective

III Marketplace Optimization

IIII Understanding the Marketplace

Search Advertising

Text Ads

Product Ads

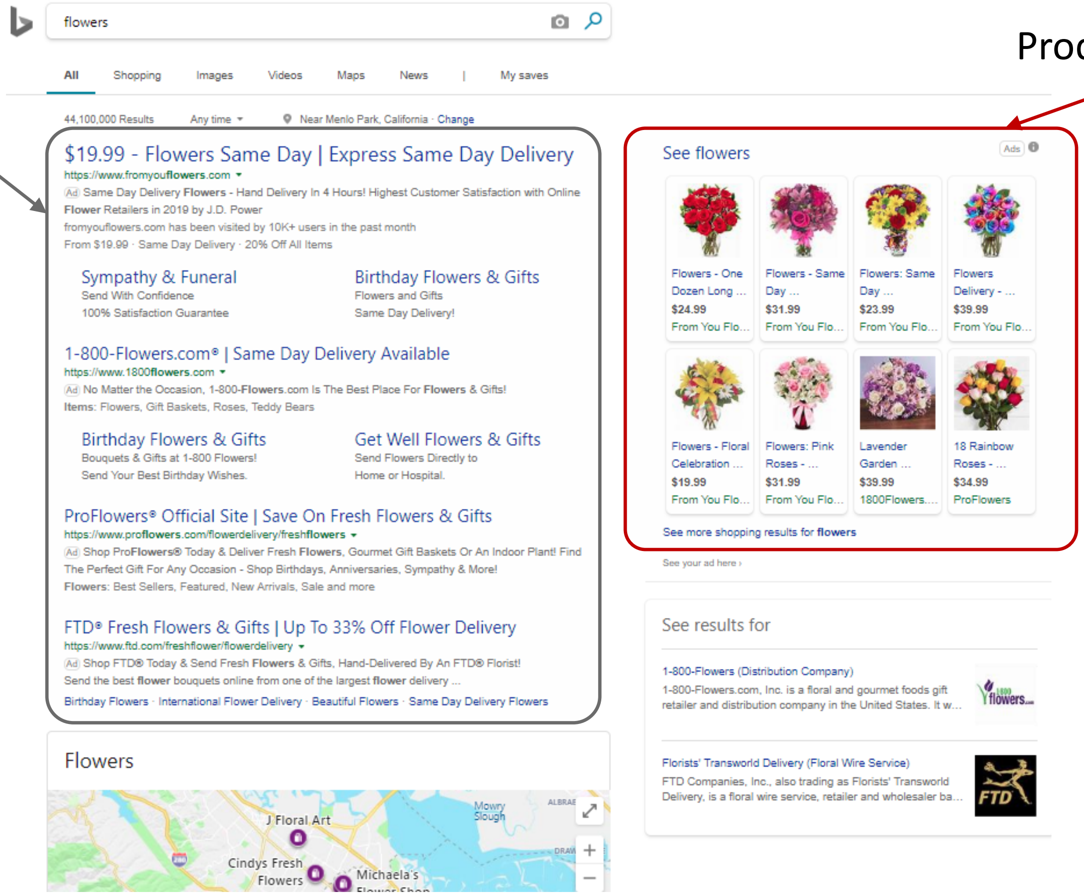

- Advertisers are charged per click
  - Ad platforms typically provide features to optimize for other targets
- Ads can have "decorations", making slot sizes variable
  - Decorations can be advertiser provided or generated by the platform
- Different ad products coexist on the same page
  - E.g. Text Ads and Product Ads can compete for the same slots

Beyond Web Search

Shopping Vertical

Hotel Ads

Visually Similar Products

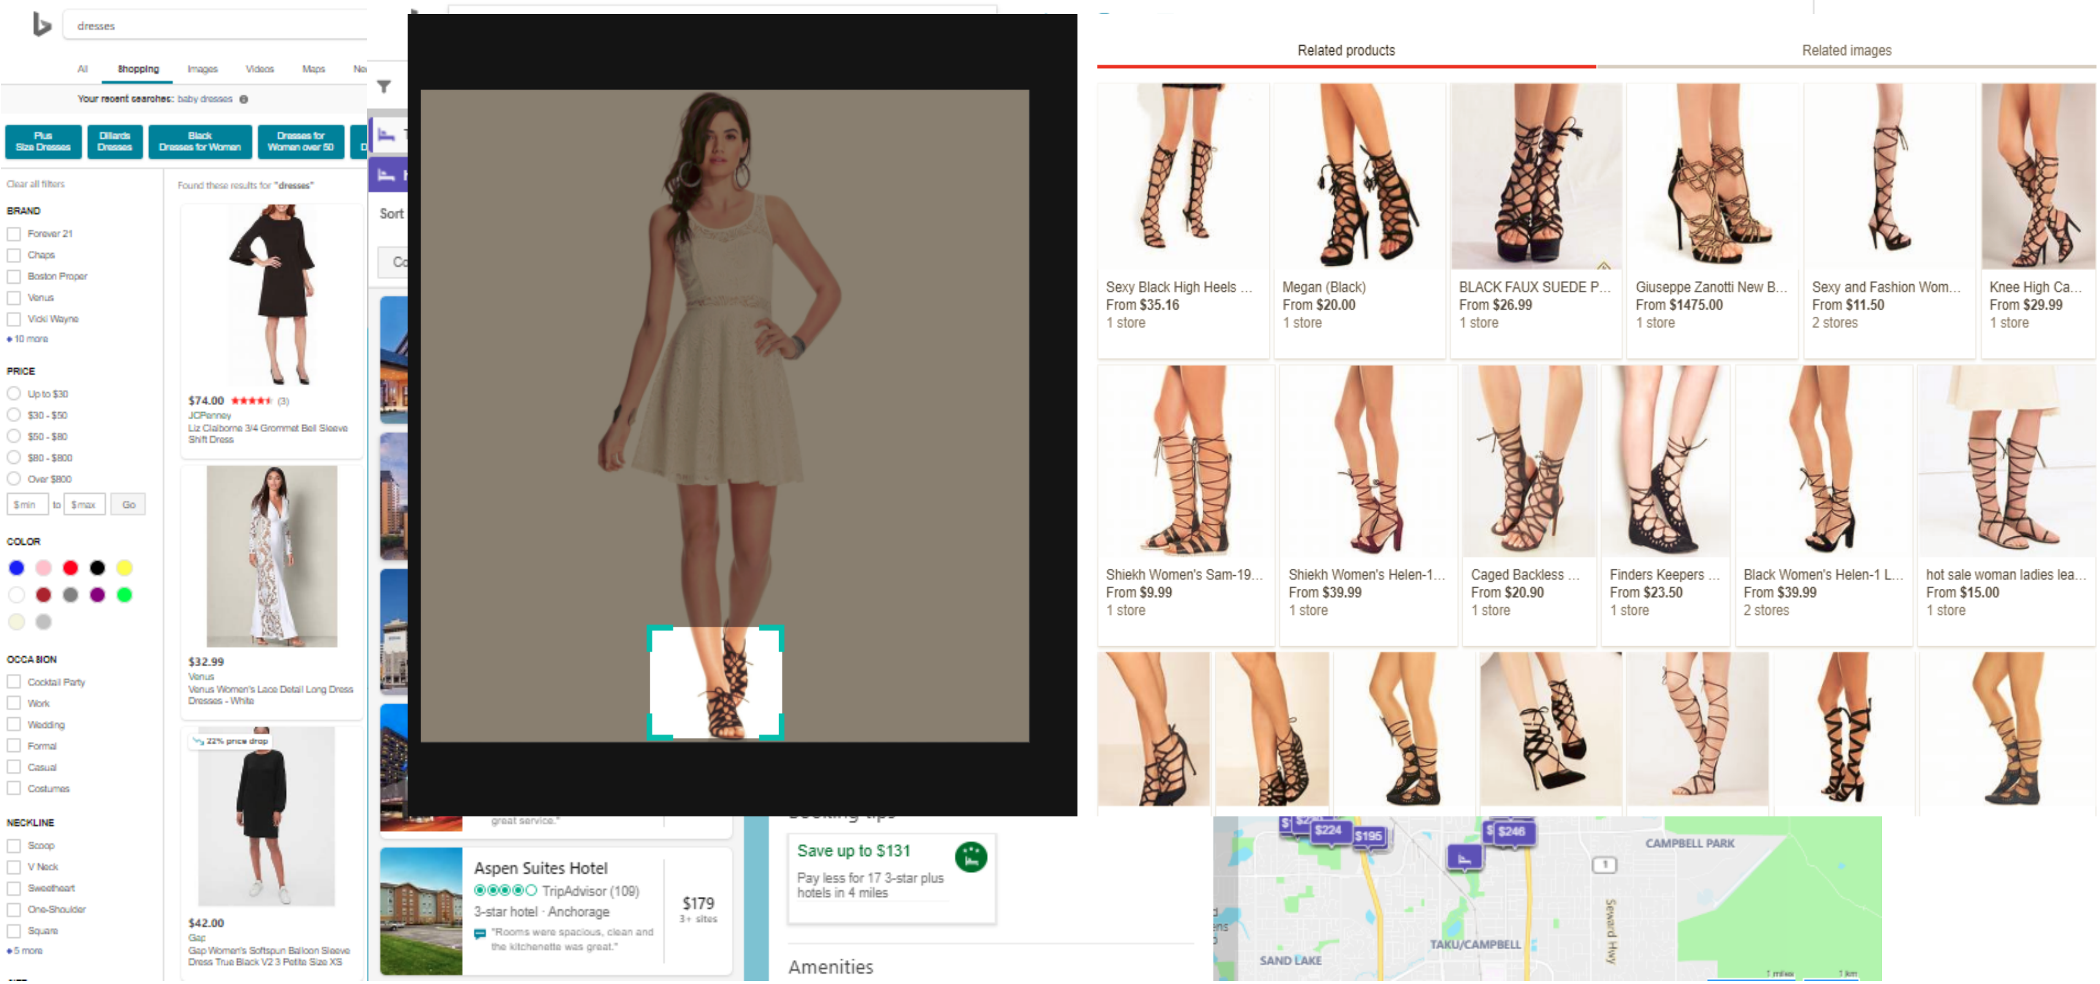

I Very Brief Intro to Search Advertising

II Marketplace Objective

III Marketplace Optimization

IIII Understanding the Marketplace

NEED TO PICK

Blue Sky

- Ultimate objective is to maximize *Long-Term* Revenue

$$Revenue = \#Users * Queries\ per\ User * Ads\ per\ Query * Clicks\ per\ Ad * Cost\ Per\ Click$$

Function of the user and the system

Function of the advertiser and the system

- Can we compute the long-term Revenue [think years]?
  - Need to estimate how our decisions would impact user activity and advertiser spend, over a long horizon.
    - E.g. how would showing more ads affect the user's search activity?
  - Not trivial to model the dependencies accurately
    - Reinforcement Learning provides a framework for a path forward

Down to Earth

- Assumption 1**:** satisfied users will engage more with the product
  - Short-term user satisfaction can be a proxy for long-term user activity
- **Assumption 2:** satisfied advertisers will increase spend
  - Short-term advertiser satisfaction can be a proxy for long-term advertiser spend
- Maximize all three short term metrics: Revenue, User Satisfaction and Advertiser Satisfaction
- Frequently formulated as:

$$\begin{array}{lll} \text{maximize} & (Short\ Term)\ Revenue & \\ s.t. & (Short\ Term)\ User\ Satisfaction & \geq K_u \\ & (Short\ Term)\ Advertiser\ Satisfaction & \geq K_a \end{array}$$

Trenches

- How do we measure user satisfaction?
  - User agnostic relevance metrics
  - or implicit user feedback (e.g. click through rate, short dwelltime click rate)
  - or a combination?

How (not) to Pick your Metrics

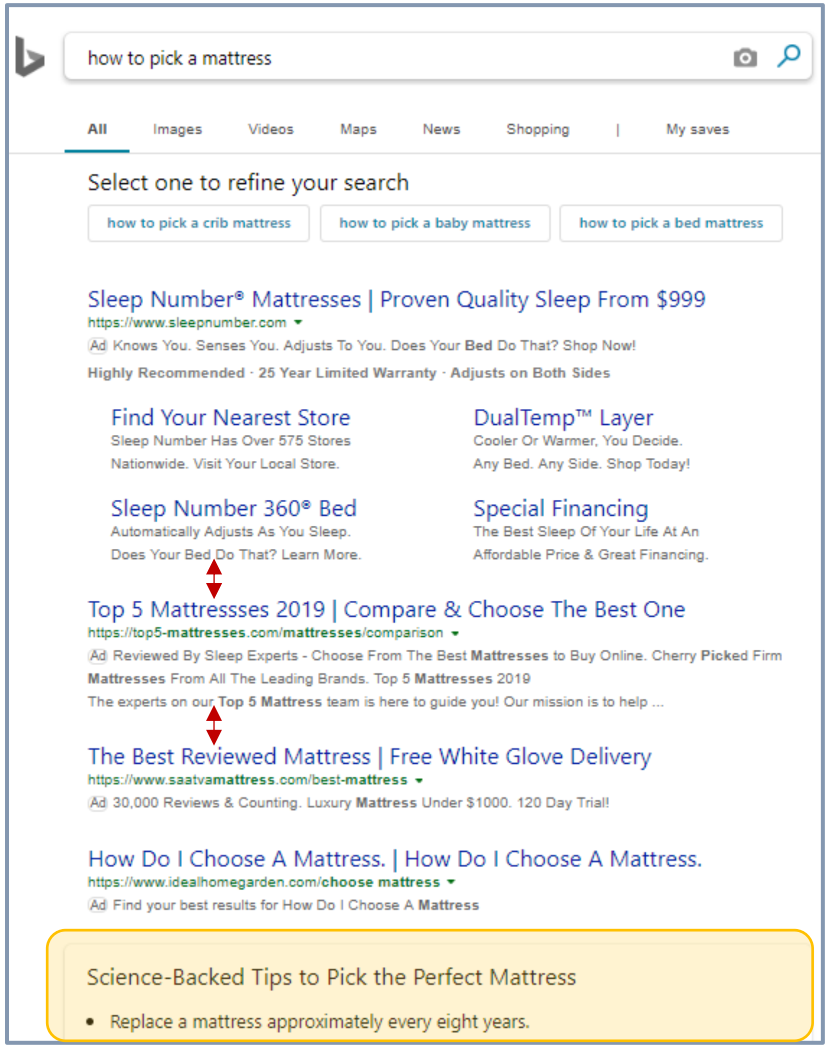

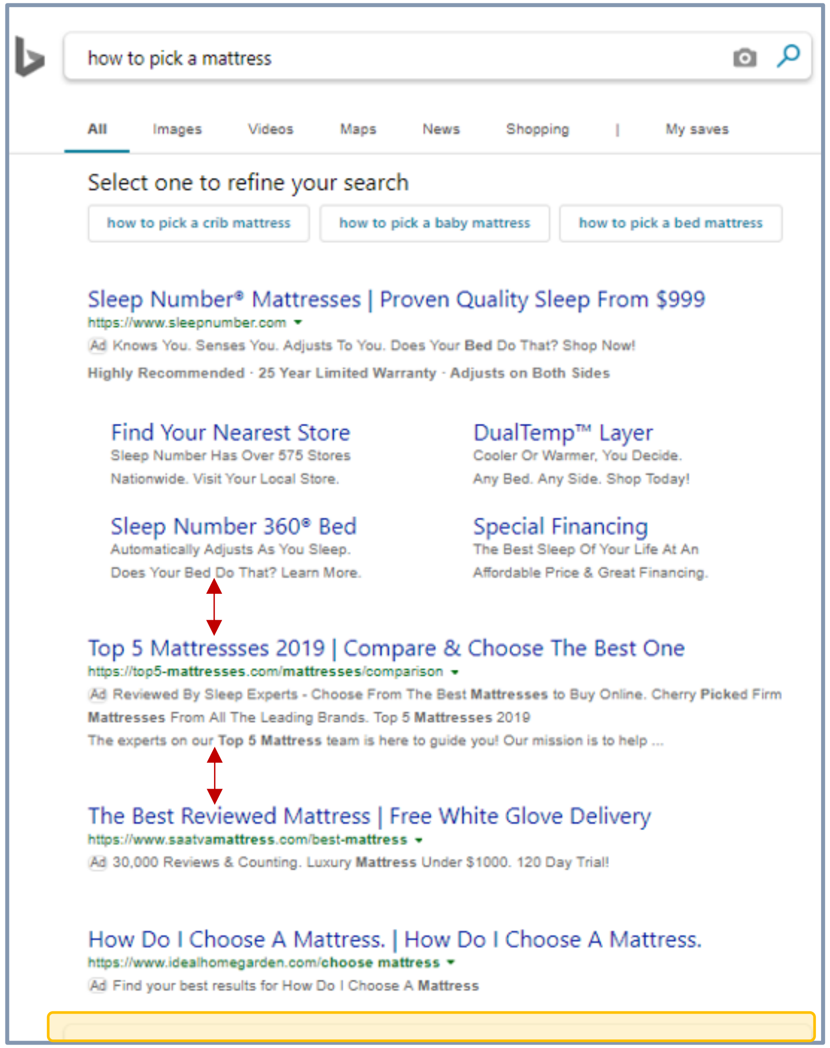

Adding whitespace does not change the relevance of ads

Pushing down all the other content typically improves click-through rates

Is the page with only ads visible *better* for the user?

Trenches

- How do we measure user satisfaction?
  - User agnostic relevance metrics
  - or implicit user feedback (e.g. click through rate, short dwelltime click rate, *space taken*)
  - or a combination?
- How do we measure advertiser satisfaction?
  - Long Dwelltime Click Through Rate, Conversion Rate, Quality of Match?

One Size Does Not Always Fit All

- User query: “a z office supplies”
- Ad keyword: “office supplies”

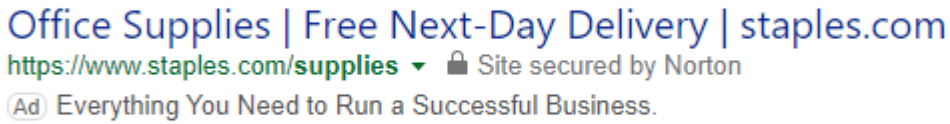

- Click and Dwelltime metrics are reasonable
  - No advertiser concern on performance as they measure it
- However Advertiser complains about the brand mismatch
  - Not a concern shared by other advertisers given the ads are performing

Trenches

- How do we measure user satisfaction?
  - User agnostic relevance metrics
  - or implicit user feedback (e.g. click through rate, short dwelltime click rate, *space taken*)
  - or a combination?
- How do we measure advertiser satisfaction?
  - Long Dwelltime Click Through Rate, Conversion Rate, Quality of Match?
- Single metric rarely captures all information
  - No need to artificially limit ourselves to using one metric alone

Trenches

- How do we evaluate our choice of metrics?
  - Run long-term experiments to measure the relation between the proposed proxies and long-term metrics?
    - Challenges:
      - Treatment dilution due to limitations in identifying users
      - Geo-based experiments can be tricky to analyze even with synthetic controls
        - Unexpected events can impact only one region
        - Advertisers target locations, may be tricky to separate advertiser and user response
  - Use user/advertiser complaints to verify your metric choices?

I

Very Brief Intro to Search Advertising

II

Marketplace Objective

III

Marketplace Optimization A Counterfactual Story

IIII

Understanding the Marketplace

How to Allocate Ads

- Rank and allocate ads to optimize the objective:

$$\begin{aligned} \text{maximize} \quad & (Short\ Term)\ Revenue \\ s.t. \quad & (Short\ Term)\ User\ Satisfaction \geq K_u \\ & (Short\ Term)\ Advertiser\ Satisfaction \geq K_a \end{aligned}$$

- Can be solved via the Lagrangian Relaxation:

$$maximize \quad Revenue + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

- Price is determined only after allocation. Replace Revenue ($p(click) * price$) with Welfare ($p(click) * bid$)

$$objective\ function = Welfare + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

How to Allocate Ads

A Per Slot Greedy Allocation Algorithm

$$objective\ function\ (rankscore, rs) = Welfare + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

Generalized Second Price

rs_1

rs_2

rs_3

Need probability of click, user satisfaction and advertiser satisfaction for that slot. E.g. $p(click|slot = i)$

Pricing

smallest bid b' such that $rs_1(b') \geq rs_2$

How to Allocate Ads

A Per Slot Greedy Allocation Algorithm

$$objective\ function\ (rankscore, rs) = Welfare + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

Generalized Second Price

rs_1

rs_2

rs_3

Need probability of click, user satisfaction and advertiser satisfaction for that slot. E.g. $p(click|slot = i)$

Pricing

smallest bid b' such that $rs_1(b') \geq rs_2$

What if slot size is variable?

rs_1

rs_2

rs_3

Condition on size too?
$p(click|slot = i, size = x)$
or

How to Allocate Ads

A Per Slot Greedy Allocation Algorithm

$$objective\ function\ (rankscore, rs) = Welfare + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

Generalized Second Price

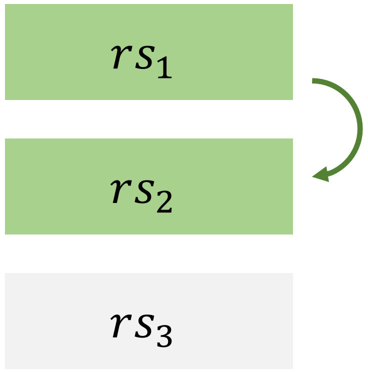

Need probability of click, user satisfaction and advertiser satisfaction for that slot. E.g. $p(click|slot = i)$

Pricing

smallest bid b' such that $rs_1(b') \geq rs_2$

What if slot size is variable?

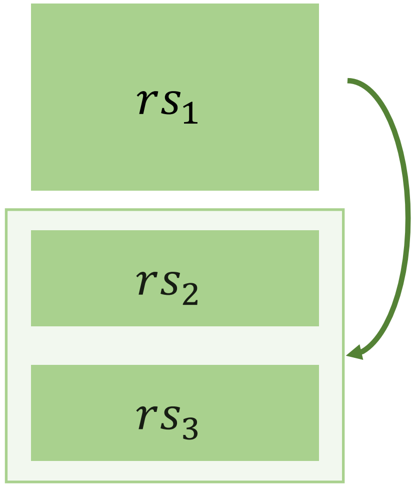

Use a coalition of ads to compete with larger ads

Back to the Objective Function

- Need to compute the $\lambda's$

  $$Welfare + \lambda_u User\ Satisfaction + \lambda_a Advertiser\ Satisfaction$$

- λ_u and λ_a can be interpreted as shadow prices:
  - $\boldsymbol{\lambda_u}$ is the cost of degrading user satisfaction by one unit
  - $\boldsymbol{\lambda_a}$ is the cost of degrading advertiser satisfaction by one unit
- Estimate using long-term experiments
  - Requires high accuracy. Small differences in the estimate may result in large differences in the outcome.
- Tune $\lambda's$ to meet business constraints and maximize the objective

How to Tune λ's

- If we could estimate the outcome of setting $\lambda's$ to any value, we could find the values that maximize the objective

λ_u	λ_a	Revenue	Long Dwelltime Click Yield	Conversion Rate
1	1	120	0.080	0.010
1	10	118	0.080	0.020
1	20	116	0.090	0.030
…	…	…	…	…
100	50	110	0.120	0.025
100	100	105	0.130	0.027

Maximum Revenue
s.t.
Long Dwelltime Click Yield > 0.11
Conversion Rate > 0.022

How to Tune λ's

- If we could estimate the outcome of setting $\lambda's$ to any value, we could find the values that maximize the objective.
- How to estimate the outcome of different $\lambda's$?

Output of λ_u, λ_a values that were used to serve the request online

User Click

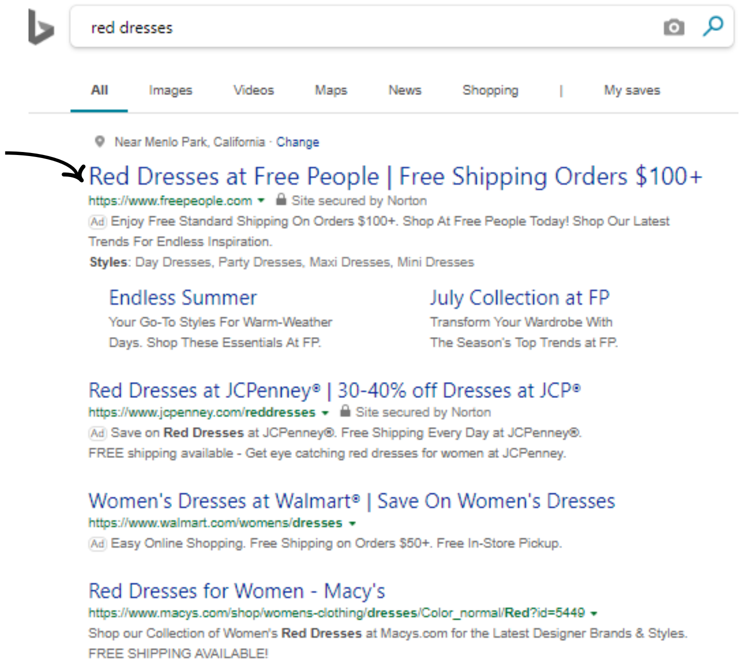

Output of new λ_u, λ_a values that were not observed online

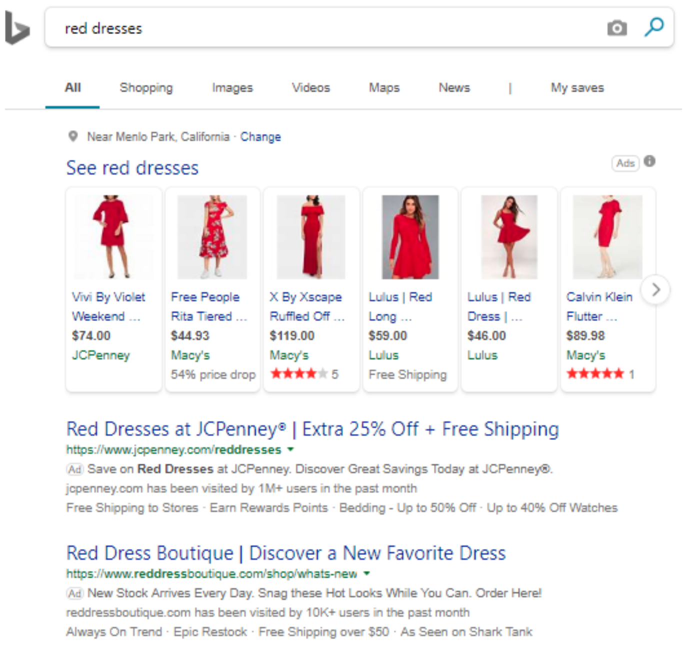

How would the user respond?

How to Tune λ's

- If we could estimate the outcome of setting $\lambda's$ to any value, we could find the values that maximize the objective.
- How to estimate the outcome of different $\lambda's$?
  - Simulate the output of the system
    - Requires the ability to replay the end-to-end stack offline
    - Comprehensive logging is critical for high fidelity simulations

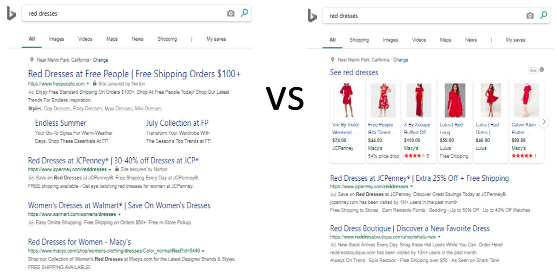

  - Simulate the user response
    - Requires estimating counterfactual probabilities
    - User model needs to be accurate for rarely seen ad slates as well

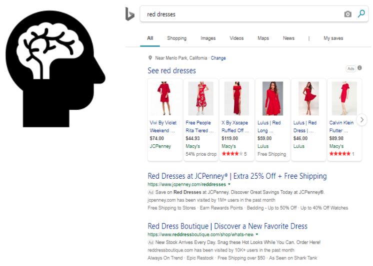

Counterfactual Click Modeling

- Goal: Estimate the *p(click)* for counterfactual allocations
- Model Inputs:
  - Query logs with click/no-click information
  - Post-allocation information
    - Ad position, ad size, other ads, page layout
    - Not available for the online models
- Model Output:

$$p(\boldsymbol{click}|query, user, ad, \boldsymbol{do}(\boldsymbol{allocation}))$$

- Need to handle *biases* that exist in *observational data*
  - E.g. Utilize Exploration and Propensity Scoring

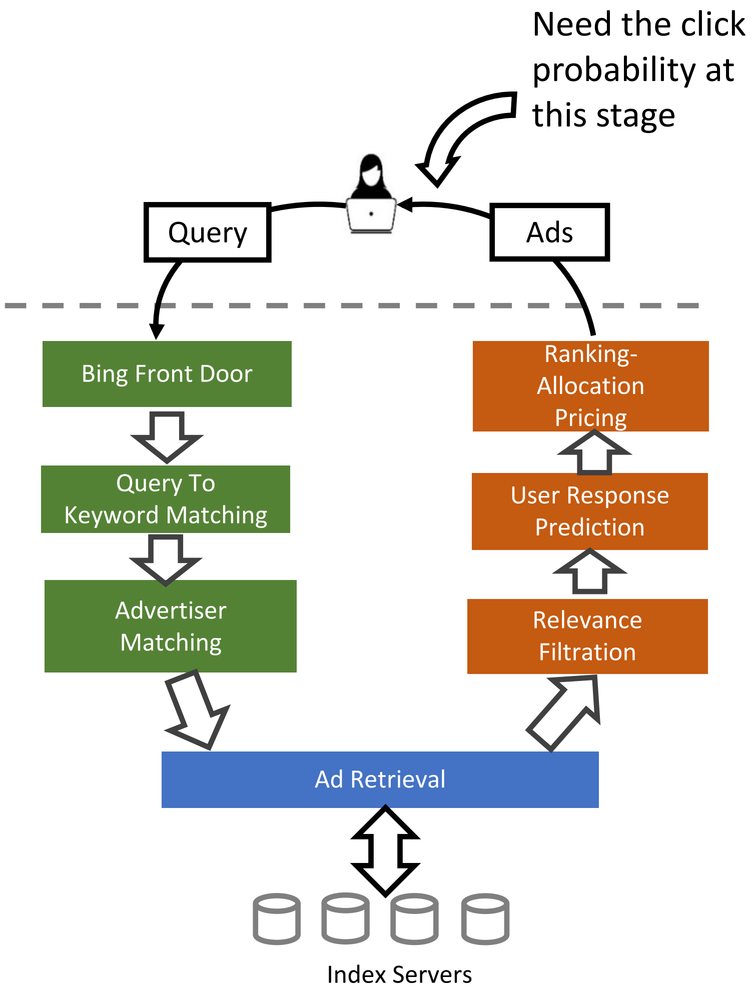

Alternative to Simulation

- Disadvantages of simulation:
  - Requires replaying the end-to-end stack, can take time
  - Simulating the user response accurately may be challenging
- Idea: Explore different values of $\lambda' s$ at run time
  - For some portion of real traffic, sample λ values from a distribution
  - Use importance sampling to compute estimated metrics

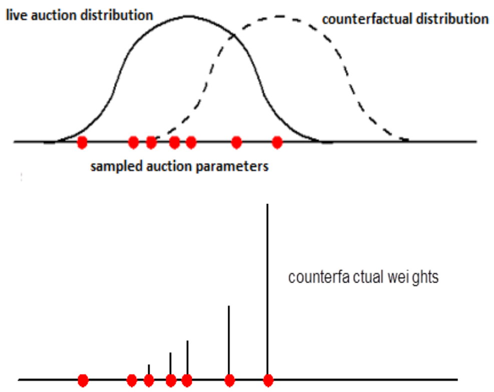

$$E_p[f(x)] = \int f(x)\frac{p(x)}{q(x)}q(x)dx = E_q\left[f(x)\frac{p(x)}{q(x)}\right]$$

- Disadvantages:
  - Randomization has short-term cost (can be reduced by joint sampling)
  - Confidence intervals widen as we increase the variance in exploration.

Very Brief Intro to Search Advertising

Marketplace Objective

Marketplace Optimization

Understanding the Marketplace

Explaining KPI Movements

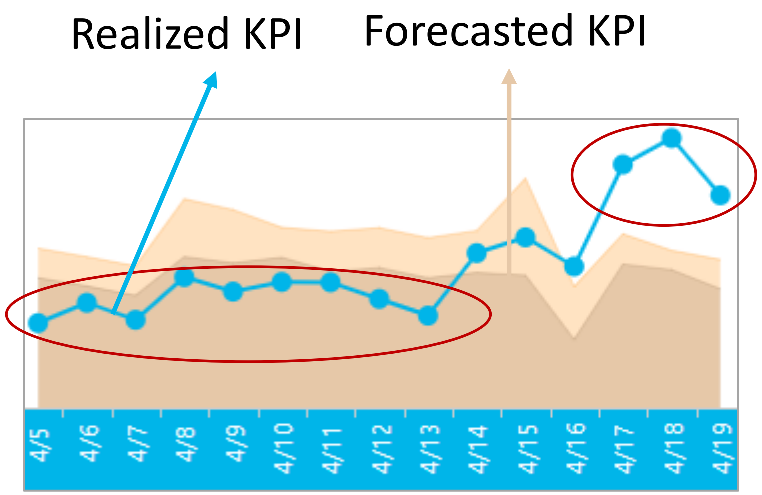

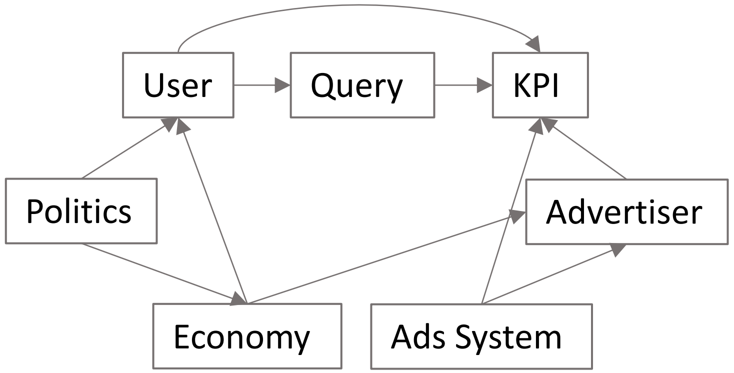

What caused the KPI to diverge from the forecast?

- Did supply (user/query) change?
  - Unexpected news events can cause major changes in query distributions.
- Did demand (advertisers) change?
  - Advertisers unseasonably increasing or decreasing their budget can result in unexpected KPI movements.
- Did the system change?
  - Did we introduce a bug?
- Some factors like economy might cause both the supply and the demand to change.

Supply As A Case Study

- Can we quantify how supply changes impact KPIs?
  - E.g. what's the contribution of supply differences in the analysis period vs the reference period on KPIs?
- Idea: Build a synthetic control for the reference period using only supply features

Analysis period searches

Query	User	Query Category	...	IsCommercial	Browser	KPI
flowers	user1	Flowers	...	Yes	Edge18	120
red dresses	user2	Apparel	...	Yes	Chrome65	75
windows update	user3	Computers	...	No	Edge18	10
...	...	...	...	...	...	...
maui hotels	userN	Travel	...	Yes	IE11	90

$$KPI_{predicted} = \sum_{q \in Queries} M_{KPI}(F(q))$$

$KPI_{Predicted}$: Estimate of KPI on the supply observed in analysis period with the fixed system and demand from reference period.

Trained model M_{KPI}

Reference period searches

Query	User	Query Category	...	IsCommercial	Browser	KPI
tarrifs war	user21	Politics	...	No	Chrome65	0
facebook	user21	Internet	...	No	Edge18	5
car insurance	user22	Finance	...	Yes	IE11	110
...	...	...	...	...	...	...
wood furniture	userM	Furniture	...	Yes	Edge18	90

$$\Delta KPI_{Predicted} = (KPI_{Predicted} - KPI_{Ref})/KPI_{Ref}$$

$\Delta KPI_{Predicted}$: Estimate for KPI change between reference and analysis period that is attributed to supply change.

Measuring Advertiser Response

- Typical questions about advertiser reponse:
  - If we were to improve {Conversion Rate, ROI} would the advertisers increase spend?
  - Would the increase (if any) be sufficient to overcome the first order revenue drop?
- Challenges:
  - Number of active advertisers is small, and spend per advertiser is very skewed
  - Not all advertisers have the same objective
  - Advertisers may not respond, or even if they do response times might be variable and long (quarters instead of within session)

Advertiser Experiments

- Approach:
  - Experiment on advertisers who are likely to respond
    - Assume a simple a-priori model for advertiser response (e.g. the more the prices change the likelier the advertisers respond)
    - Estimate first order effect of the treatment per advertiser using simulations
    - Pick the set of advertisers with maximum predictive power
  - Randomly assign the selected advertisers to treatment and control
    - Pairwise stratified randomization works better than IID
  - Find the optimal policy (i.e. which advertisers should get the treatment)
    - Build a better model of advertiser response (using the experiment data)
    - Use the new model to pick the advertisers that would react positively to the experiment

Optimal Policy Identification

- Inputs:
    - $\{x_j\}_{j \in J}$: Advertiser features, measured pre-experiment
    - $\{t_j\}_{j \in J}$: Treatment indicators
    - $\{y_j^{(t_j)}\}_{j \in J}$: Per sample reward, e.g. post-treatment spend per advertiser
- Task:
    - Find policy $\pi: X \to \{0,1\}$ which maximizes $R(\pi) = \frac{1}{n} \sum_{j \in J} y_j^{(\pi(x_j))}$
    - I.e. characterise advertisers for which $y_j^{(1)} > y_j^{(0)}$
- Model: Honest Random Forests
    - Use control and treatment advertisers to build a forest per advertiser using leave-advertiser-out
    - Effect on sample *j* estimated as the difference between treatment and control samples in the same node (after excluding *j*)

Full Population

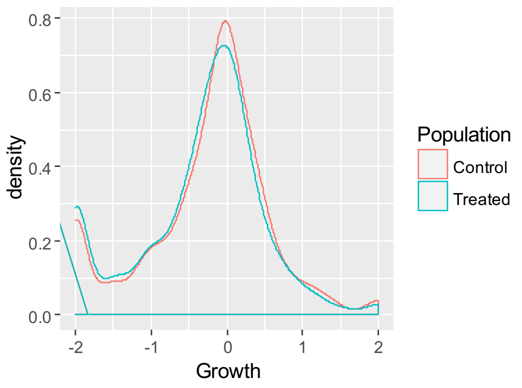

Effect on spend: -5.7% (-11.1% - -0.4%)

Selected Population

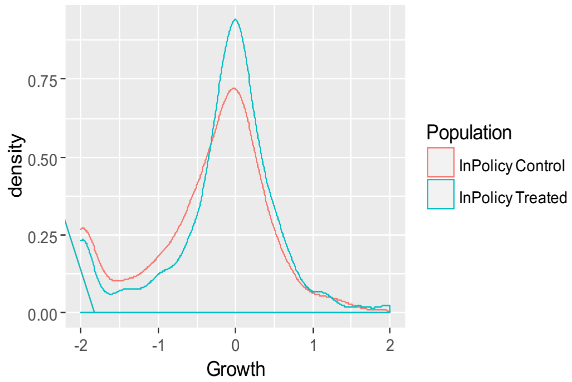

Effect on spend: 13.5% (8.1% - 18.8%)
3.9% (2.3% - 5.5%) increased spend

Very Brief Intro to Search Advertising

Marketplace Objective

Marketplace Optimization

Understanding the Marketplace

Looking Ahead

- Reinforcement Learning to directly optimize for the Long-Term
  - Already interesting work happening but under many assumptions
- Advertising Ecosystem is evolving
  - Advertisers are moving to AI/ML for everything: UX design, content generation, budget management and more
  - Modeling causal dependencies will be critical to react optimally
- User Interface is evolving
  - Definition of clicks or engagement needs to adapt