Externalities

Sponsored Search Auctions with Markovian Users
(Aggarwal, Feldman, Muthukrishnan, Pal)

Externalities in Keyword Auctions: an Empirical and Theoretical Assessment
(Gomez, Immorlica, Markakis)

Chonnikarn (Fern) Jira
& Pichayut (Petch) Jirapinyo

Sponsored Search Auctions: Recap

- Advertisers ranked and chosen to positions by **score** = $b_i w_i$ (w_i = weight associated with advertiser i, b_i = advertiser i's bid)
- **Separable click-through rate**:
  - P(click ad i | view ad i) = p_i (advertiser-specific)
  - P(view ad i | slot j) = α_j (slot-specific)
  - Click-through rate = P(click ad i | slot j) = $p_i \alpha_j$
- Explored various scoring and pricing rules, assuming strategic behavior of bidders
  - Most studied: GSPs
- Measures of desirability: efficiency, revenue-maximization
- Also explored effect of expressiveness vs. simplification
- Assumed no ***externalities*** …

A New Model for User's Behavior

- Ordered Search
  - Users browse from top to bottom
  - Make clicking decisions slot by slot
  - Implicit assumption: first ad always read

Externalities

- **Idea:** value of acquiring slot in sponsored search list highly depends on who else is shown in the other sponsored positions
- **Two types:**
  - Position Externalities
    - Your iPod ad is under an Apple store ad
  - Information Externalities
    - Your diet pill ad is under an ad warning about the diet pill
- Depart from separable click-through rate models

Sponsored Search Auctions with Markovian Users

Aggarwal et al

Aggarwal et al: big picture

- Model for user behavior: Markovian user model
- Considers negative externalities from positions (no information externalities)
- Based on this user model
  - characterizes efficient assignment
  - gives algorithm to find it
  - efficient assignment not GSP, but has some desirable properties of GSP (allows for *intuitive bidding*) and can be made IC with VCG payment
- Focuses on how to allocate ads to slots efficiently/to max revenue *given the bids*: **no equilibrium analysis**

Markovian User Model

- Recall separable click-through rate model:

  P(click ad i | slot j) = $p_i \alpha_j$; where α_j decreasing in position

- **Motivation:**
  - Model process such that decreasing α_j arises naturally
  - Each ad effects user's clicking on that ad as well as looking at other ads
- **Formulation:** user as Markov Process
  - n bidders $ß = \{1, \ldots, n\}$, k positions
  - For each bidder, begin scanning ads from the top down. When position j is reached:

    p_i = probability that a user will click on ad i, given he *looks* at it

    q_i = probability that a user will look at the next ad in a list, given that he *looks* at ad i

Markovian User Model (2)

- Compare with separable click-through rate model:
  - $\alpha_j = \prod_{i' \epsilon A} q_{i'}$; where A is set of ads above position j
- **No longer separable**, but reduces to same expression as separable model in the first slot, where α_1 =1 and click-through rate = p_i
- **Position externality**: q_i introduces tradeoff between probability of clicking an add (p_i) and the ad's effect on slot below it (q_i)

Optimal Assignment

- **Expected Cost per Thousand:** $e_i = p_i b_i$ = value of an "impression" = how much bidder values user looking at ad *i*

- What is **"best/optimal assignment"** ?

Given bid b_i for ad i, find assignment of ads $(x_1, x_2, \ldots, x_k)$ with corresponding p_{xi} and q_{xi} such that

$$max \ \{e_{x1} + q_{x1} (e_{x2} + q_{x2}(e_{x3} + q_{x3}(\ldots +q_{x(k-1)}(e_{xk}))))\}$$

- **Interpretation:** If assume that bidders are truthful in reporting values as bids ($v_i = b_i$), then want to maximize overall expected value of assignment to bidders

Properties of Optimal Assignment (1)

- **Adjusted ecpm** (a-ecpm) = $e_i/(1\text{-}q_i)$

THEOREM 1. *In the most efficient assignment, the ads that are placed are sorted in decreasing order of adjusted ecpm* $a_i = e_i/(1 - q_i)$.

- **Proof sketch:** swap argument
  - Suppose not: consider a_i and $a_{i'}$ of ads in optimal assignment positions *j* and *j+1* such that $a_i < a_{i'}$
  - Compare contributions of positions j to n to efficiency for both orders of *i* and *i'* in slots *j* and *j+1*
  - Contradiction

Illustrative Example

- Theorem 1 tells *how to sort the ads* to select, but doesn't tell *which k ads* to select

Bidder	e_i	q_i	$a_i = e_i / (1-q_i)$
1	$1.00	0.75	4
2	$2.00	0.20	2.5
3	$0.85	0.80	4.25

- **Ranking by ecpm** doesn't yield optimal assignment
  - (2,1): efficiency = $2 + 0.2($1) = $2.20
- **Ranking by a-ecpm** doesn't yield optimal assignment
  - (3,1): efficiency = $0.85 + 0.8($1) = $1.65
- **Optimal Assignment** is (1,2): efficiency = $1 + 0.75($2) = $2.50

Properties of Optimal Assignment (2) Bidder dominance

- While having higher a-ecpm alone doesn't allow a bidder to dominate another, having both higher a-ecpm and ecpm does suffice

THEOREM 2. *For all bidders i in an optimal assignment, if some bidder j is not in the assignment, and $a_j \geq a_i$ and $e_j \geq e_i$, then we may substitute j for i, and the assignment is no worse.*

- **Intuition:** for a special case, if $a_i = e_i$ for all i, then assignment reduces to GSP

Properties of Optimal Assignment (3) Subset Substructure in Optimal Assignments

- Subset structure between optimal assignments to different numbers of slots
- OPT(C, j) = set of all optimal solutions for filling j positions with bidders from C
- Optimal solution S ϵ OPT(C, j) = set of agents assigned to slots 1-j

THEOREM 3. *Let $j \in \{1, \ldots, k\}$ be some number of positions, and let C be an arbitrary set of bidders. Then, for all $S \in OPT(C, j-1)$, there is some $S' \in OPT(C, j)$ where $S' \supset S$.*

- **Intuition:** for each additional number of slots, find a proper position to insert another bidder into sequence of bidders assigned

Properties of Optimal Assignment (4) Monotonicity of Position/Click Probs

- Optimal assignment supports **intuitive Bidding**
  - higher bids translate to higher positions and more clicks
  - Allows bidder to adjust bid intelligently without global knowledge of other bids

THEOREM 4. *As a bidder increases her bid (keeping all other bids fixed):*

(a) the probability of her receiving a click in the optimal solution does not decrease, and

(b) her position in the optimal solution does not go down.

Computing Optimal Assignment

1. **Optimal Assignment using Dynamic Programming**
   - Sort ads in decreasing order of a-ecpm: O(*n log n*)
   - *F(i,j)* = efficiency obtained (given you reach slot *j*) by filling slots (*j*…*k*) with bidders from set {*i*,…,*n*}

$$F(i,j) = max(F(i+1, j+1)\, q_i + e_i\, ,\, F(i+1, j))$$

   - Solving for *F(1,1)* yields optimal assignment: O(*nk*)
   - **Overall:** O(*n log n* + *nk*)

Computing Optimal Assignment (2)

2. Near-linear time algorithm (in place of dynamic programming)

- Oracle: for any $j, j' \epsilon ß$, return bidder $j \leq y \leq j'$ that maximizes $f(q_y, e_y)$ for an arbitrary linear function f : $O(log^2 n)$
- Algorithm constructs solution S_i ϵ OPT($ß$, i) for $i = 1,\ldots,k$. Final one S_k is the overall optimum. Make $O(k^2)$ calls to oracle
- **Overall:** $O(n\ logn + k^2 log^2 n)$

Theorem 5. *Consider the auction with n Markovian bidders and k slots. There is an optimal assignment which can be determined in* $O(n \log n + k^2 \log^2 n)$ *time.*

Using **VCG pricing** -> truthful mechanism for sponsored search with Markovian users

Discussion: limitations

- Statistical and machine learning problem: need to find a good way to find model parameters, especially q_i's
- Open problem: initialization for new users coming in
- q_i can depend on location
- Model extension for general layouts

Externalities in Keyword Auctions

Gomes, Immorlica, Markakis

Two Types of Externalities

- Position Externalities
  - User tires of the search
    - λ_j = Pr(Continue | View j + Don't Click j)
  - User finds what she looks for
    - γ_j = Pr(Continue | View j + Click j)
- Information Externalities
  - H= { j: link j received a click}
  - $F_j(H)$= Pr(Click j | View j + Click H)
  - $F_j(\{\emptyset\})$= Base-line Click-through rates

Example: Externalities

λ_j = Pr(Continue | View j + Don't Click j)

γ_j = Pr(Continue | View j + Click j)

$F_j(H)$ = Pr(Click j | View j + Click H)

- Two ads: A_1, A_2
- Pr(click on A_1) = F_{A_1}
- Pr(click on A_2) = $(1 - F_{A_1})\lambda_{A_1} F_{A_2} + F_{A_1}\gamma_{A_1} F_{A_2}(\{A_1\})$

Representation of an Event

- A pair of tuples
  - First tuple = the ads displayed
  - Second tuple = the ads clicked
  - $\{j, k, \emptyset; k, \emptyset, \emptyset\}$= Event when *j* and *k* were displayed and only *k* was clicked.

Example: Representation

λ_j = Pr(Continue | View j + Don't Click j)

γ_j = Pr(Continue | View j + Click j)

$F_j(H)$ = Pr(Click j | View j + Click H)

Ordered Search Model:

$$\begin{aligned}\text{Prob}\left(\{j,k,l;j,\emptyset,\emptyset\}\right) &= F_j(1-\gamma_j) + \\ &\quad F_j\gamma_j(1-F_k(\{j\}))(1-\lambda_k) + \\ &\quad F_j\gamma_j(1-F_k(\{j\}))\lambda_k(1-F_l(\{j\}))\end{aligned}$$

Separable Model:

$$\text{Prob(advertiser } j \text{ gets a click on slot } k) = s^k \cdot f_j$$

s^k, slot specific CTR; f_j, advertiser specific CTR

Representation: all possible events

$$
\begin{aligned}
\text{Prob}\,(\{j,k,\emptyset;\emptyset,\emptyset,\emptyset\}) &= (1-F_j)(1-\lambda_j)+\\
&\quad (1-F_j)\lambda_j(1-F_k),\\
\text{Prob}\,(\{j,k,\emptyset;j,\emptyset,\emptyset\}) &= F_j(1-\gamma_j)+F_j\gamma_j(1-F_k(\{j\})),\\
\text{Prob}\,(\{j,k,\emptyset;k,\emptyset,\emptyset\}) &= (1-F_j)\gamma_j F_k,\\
\text{Prob}\,(\{j,k,\emptyset;j,k,\emptyset\}) &= F_j\gamma_j F_k(\{j\}),\\
\text{Prob}\,(\{j,k,l;k,\emptyset,\emptyset\}) &= (1-F_j)\lambda_j F_k(1-\gamma_k)+\\
&\quad (1-F_j)\lambda_j F_k\gamma_k(1-F_l(\{k\})),\\
\text{Prob}\,(\{j,k,l;l,\emptyset,\emptyset\}) &= (1-F_j)\lambda_j(1-F_k)\lambda_k F_l,\\
\text{Prob}\,(\{j,k,l;j,k,\emptyset\}) &= F_j\gamma_j F_k(\{j\})(1-\gamma_k)+\\
&\quad F_j\gamma_j F_k(\{j\}\gamma_k(1-F_l(\{j,k\})),\\
\text{Prob}\,(\{j,k,l;j,l,\emptyset\}) &= F_j\gamma_j(1-F_k(\{j\})\lambda_k F_l(\{j\}),\\
\text{Prob}\,(\{j,k,l;k,l,\emptyset\}) &= (1-F_j)\lambda_j F_k\gamma_k F_l(\{k\}),\\
\text{Prob}\,(\{j,k,l;j,k,l\}) &= F_j\gamma_j F_k(\{j\})\gamma_k F_l(\{j,k\}).
\end{aligned}
$$

Dataset: Observations

- CTRs depend on users' search objectives
  - Restricts experiments to a single search objective
- Requires variations on clicking history
  - Takes only events with at least 2 of 3 top advertisers displayed together
- The representation assumes top-to-bottom search for all users

Dataset

keyword	advertisers	# of obs.
ipod	(A): store.apple.com (B): cellphoneshop.net (C): nextag.com	8,398
diet pill	(A): pricesexposed.net (B): dietpillvalueguide.com (C): certiphene.com	4,652
avg antivirus	(A): Avg-Hq.com (B): avg-for-free.com (C): free-avg-download.com	1,336

Table 1: Keywords and Advertisers

slot	ipod	diet pill	antivirus
first	1,572 (74.08%)	640 (56.73%)	205 (43.15%)
second	524 (24.69%)	384 (34.04%)	259 (54.52%)
third	30 (1.41%)	104 (9.21%)	11 (2.31%)
total	2,122 (100%)	1,128 (100%)	475 (100%)

Table 3: Distribution of Clicks per Slot

slot	ipod	diet pill
first	(A): 6,460 (76.92%) (B): 1,864 (22.20%) (C): 74 (0.88%)	(A): 1,912 (41.10%) (B): 908 (19.52%) (C): 1,832 (39.38%)
second	(A): 1,438 (17.12%) (B): 5,826 (69.37%) (C): 1,134 (13.50%)	(A): 1,848 (39.72%) (B): 1,988 (42.73%) (C): 816 (17.54%)
third	(A): 26 (0.31%) (B): 22 (0.26%) (C): 950 (11.31%) (other): 7,400 (88.12%)	(A): 472 (10.15%) (B): 692 (14.88%) (C): 668 (14.36%) (other): 2,820 (60.62%)

	antivirus
first	(A): 1,233 (92.29%) (B): 71 (5.31%) (C): 32 (2.40%)
second	(A): 88 (6.59%) (B): 674 (50.45%) (C): 574 (42.96%)
third	(A): 9 (0.67%) (B): 21 (1.57%) (C): 355 (26.57%) (other): 951 (71.18%)

Table 2: Distribution of Advertisers per Slot

Estimation

- Estimating parameters using the maximum likelihood method, maximizing the probability of the sample:

$$\log L = \sum_n \log \left[\text{Prob} \left(\{ j_n, k_n, l_n; c_n^1, c_n^2, c_n^3 \} \right) \right]$$

Estimation Results

keyword	ipod	diet pill	antivirus
F_A	0.210 (0.005)	0.210 (0.008)	0.151 (0.010)
$F_A(\{B\})$	0.250 (0.038)	0.232 (0.032)	0.00 (0.074)
$F_A(\{C\})$	—	0.317 (0.065)	—
$F_A(\{B,C\})$	—	0.664 (0.075)	—
F_B	0.087 (0.006)	0.150 (0.009)	0.206 (0.038)
$F_B(\{A\})$	0.030 (0.022)	0.146 (0.034)	0.364 (0.050)
$F_B(\{C\})$	—	0.663 (0.080)	—
$F_B(\{A,C\})$	—	0.334 (0.083)	—
F_C	0.104 (0.012)	0.051 (0.004)	0.215 (0.042)
$F_C(\{A\})$	0.040 (0.032)	0.052 (0.017)	0.242 (0.042)
$F_C(\{B\})$	0.095 (0.032)	0.088 (0.029)	0.121 (0.889)
$F_C(\{A,B\})$	0.327 (0.190)	0.664 (0.089)	0.125 (0.699)
λ_A	0.676 (0.056)	0.760 (0.064)	1.0 (0.217)
λ_B	0.627 (0.042)	0.673 (0.057)	0.183 (0.049)
λ_C	1.00 (0.057)	0.579 (0.037)	0.424 (0.201)
γ_A	1.00 (0.777)	0.940 (0.195)	1.00 (0.231)
γ_B	1.00 (0.820)	1.00 (0.743)	0.686 (0.902)
γ_C	—	1.00 (0.892)	—

Table 4: Estimates of the Ordered Search Model

λ_j = Pr(Continue | View j + Don't Click j)

γ_j = Pr(Continue | View j + Click j)

$F_j(H)$ = Pr(Click j | View j + Click H)

Making Sense of the information externalities

- $F_j > F_j(\{k\})$:
  - k has a *negative externality* effect on j
- $F_j < F_j(\{k\})$:
  - k has *positive externality* effect on j
  - or evidence for the *selection* effect

$F_j(H) = \Pr(\text{Click } j \mid \text{View } j + \text{Click } H)$

Selection effect

- "the group of users that make at least one click may be fundamentally different from the total pool of users that perform searches on Microsoft Live."
- The selection effect:
  - $F_j < F_j(\{k\})$ - same as positive externalities
  - $\lambda_i < \gamma_i$

λ_j = Pr(Continue | View j + Don't Click j)

γ_j = Pr(Continue | View j + Click j)

$F_j(H)$ = Pr(Click j | View j + Click H)

Estimation Results

keyword	ipod	diet pill	antivirus
F_A	0.210 (0.005)	0.210 (0.008)	0.151 (0.010)
$F_A(\{B\})$	0.250 (0.038)	0.232 (0.032)	0.00 (0.074)
$F_A(\{C\})$	—	0.317 (0.065)	—
$F_A(\{B,C\})$	—	0.664 (0.075)	—
F_B	0.087 (0.006)	0.150 (0.009)	0.206 (0.038)
$F_B(\{A\})$	0.030 (0.022)	0.146 (0.034)	0.364 (0.050)
$F_B(\{C\})$	—	0.663 (0.080)	—
$F_B(\{A,C\})$	—	0.334 (0.083)	—
F_C	0.104 (0.012)	0.051 (0.004)	0.215 (0.042)
$F_C(\{A\})$	0.040 (0.032)	0.052 (0.017)	0.242 (0.042)
$F_C(\{B\})$	0.095 (0.032)	0.088 (0.029)	0.121 (0.889)
$F_C(\{A,B\})$	0.327 (0.190)	0.664 (0.089)	0.125 (0.699)
λ_A	0.676 (0.056)	0.760 (0.064)	1.0 (0.217)
λ_B	0.627 (0.042)	0.673 (0.057)	0.183 (0.049)
λ_C	1.00 (0.057)	0.579 (0.037)	0.424 (0.201)
γ_A	1.00 (0.777)	0.940 (0.195)	1.00 (0.231)
γ_B	1.00 (0.820)	1.00 (0.743)	0.686 (0.902)
γ_C	—	1.00 (0.892)	—

Table 4: Estimates of the Ordered Search Model

λ_j = Pr(Continue | View j + Don't Click j)

γ_j = Pr(Continue | View j + Click j)

$F_j(H)$ = Pr(Click j | View j + Click H)

Model Validation

Ordered Search Model vs. Separable CTR Model

keyword	ipod	diet pill	antivirus
F_A	0.210 (0.005)	0.210 (0.008)	0.151 (0.010)
F_B	0.087 (0.006)	0.150 (0.009)	0.206 (0.038)
F_C	0.104 (0.012)	0.051 (0.004)	0.215 (0.042)

Estimates of the Ordered Search Model

keyword	ipod	diet pill	antivirus
f_A	0.216 (0.005)	0.205 (0.008)	0.144 (0.009)
f_B	0.085 (0.004)	0.164 (0.009)	0.256 (0.036)
f_C	0.107 (0.011)	0.057 (0.004)	0.253 (0.038)
s^1	1.00 —	1.00 —	1.00 —
s^2	0.676 (0.036)	0.671 (0.037)	0.961 (0.136)
s^3	0.400 (0.072)	0.699 (0.056)	0.144 (0.043)

Table 5: Estimates of the Separable CTR Model

Model Validation (2)

Prob.	*ipod* realized CTR	*ipod* ordered	*ipod* separable	*diet pills* realized CTR	*diet pills* ordered	*diet pills* separable	*avg antivirus* realized CTR	*avg antivirus* ordered	*avg antivirus* separable
A	0.21	0.21★	0.22	0.23	0.21★	0.21	0.15	0.15★	0.14
B	0.09	0.09★	0.09	0.15	0.15★	0.16	0.17	0.21★	0.26
C	0.05	0.10★	0.11	0.03	0.05★	0.06	0.22	0.22★	0.25
$A\|B$	0.26	0.25★	0.15	0.22	0.23★	0.14	0.00	0.00★	0.14
$A\| \sim B$	0.13	0.13★	0.15	0.10	0.14★	0.14	0.04	0.03★	0.14
$A\|C$	0.00	0.00★	0.15	0.36	0.32★	0.14	0.00	0.00★	0.14
$A\| \sim C$	0.33	0.21★	0.15	0.11	0.12★	0.14	0.06	0.06★	0.14
$B\|A$	0.04	0.03★	0.06	0.21	0.14★	0.11	0.04	0.03★	0.06
$B\| \sim A$	0.06	0.06★	0.06	0.10	0.11★	0.11	0.06	0.06★	0.06
$B\|C$	0.00	0.00★	0.06	0.60	0.66★	0.11	0.00	0.00★	0.06
$B\| \sim C$	0.05	0.09	0.06★	0.10	0.09★	0.11	0.05	0.09	0.06★
$C\|A$	0.05	0.04★	0.07	0.07	0.05★	0.04	0.05	0.04★	0.07
$C\| \sim A$	0.07	0.07★	0.07	0.06	0.04★	0.04	0.07	0.07★	0.07
$C\|B$	0.19	0.09★	0.07	0.11	0.09★	0.04	0.19	0.09★	0.07
$C\| \sim B$	0.08	0.07★	0.07	0.07	0.03	0.04★	0.08	0.07★	0.07

Table 6: Model Validation

- J | K = 'user clicks on advertiser *J* occupying the second slot given that she clicked on ad *K* occupying the first slot'
- J | ~K = '... given that she DID NOT click on ad *K* occupying the first slot'

Equilibrium Analysis

- **GSP + Scoring Rule**
- Ads ranked by $s_{A_j} = w_{A_j} \cdot b_{A_j}$
  - b_{Aj} : A_j's bid
  - w_{Aj} : a weight (= F_{Aj} in the case of rank-by-revenue)
- GSP: payment equal to the smallest bid he could have submitted that would have allowed him to maintain his position:

$$p_{A_j} \cdot w_{A_j} = b_{A_{j+1}} \cdot w_{j+1} \quad \text{which gives} \quad p_{A_j} = \frac{b_{A_{j+1}} \cdot w_{A_{j+1}}}{w_{A_j}}$$

- Assumption: $F_{A_j} = F_{A_j}(H)$
- Can we design a scoring rule to maximize revenue?

Equilibrium Analysis (2)

Lemma 2. *Consider the GSP with scoring rule* w_{A_j}, *selling* K *slots to* $N > K$ *advertisers. Let advertisers* $A_1, ..., A_N$ *be the efficient assignees of slots* 1 *to* N *and assume advertisers submit bids according to:*

$$b_{A_j} = (1 - c_{A_j}) \frac{w_{A_{j-1}}}{w_{A_j}} v_{A_{j-1}} + c_{A_j} \frac{w_{A_{j+1}}}{w_{A_j}} b_{A_{j+1}}$$

$$\text{for } j \in \{2, ..., K\}, b_{A_{K+1}} = \frac{w_{A_K}}{w_{A_{K+1}}} v_{A_K}, \; b_{A_1} > b_{A_2} \qquad (2)$$

$$\text{and} \quad b_{A_j} < b_{A_{K+1}} \text{ for } j > K+1. \qquad (3)$$

If this bid profile constitutes a Nash equilibrium, than it maximizes the search engine's revenue among all pure strategy complete information Nash equilibria. We call it the greedy bid profile.

Equilibrium Analysis (3)

Proposition 1. *Consider the GSP with scoring rule w_{A_j}, selling K slots to $N > K$ advertisers. The greedy bid profile constitutes a complete information Nash equilibrium for all valuations and search parameters $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$ if and only if $w_{A_j} = \frac{F_{A_j}}{1-c_{A_j}}$ (up to a multiplicative constant). In this case, the equilibrium allocation is efficient and the search engines's revenue is maximal.*

$c_{A_k} = F_{A_k}\gamma_{A_k} + (1 - F_{A_k})\lambda_{A_k}$ = continuation probabilities

$a_i = e_i/(1 - q_i)$ = adjusted-ecpm (from the other paper)

Proposition 2. *Consider the GSP selling K slots to $N > K$ advertisers. There is no scoring rule w_{A_j} which depends solely on advertiser A_j's search parameters $(F_{A_j}, \gamma_{A_j}, \lambda_{A_j})$ that implements an efficient equilibrium with VCG payments for all valuations and search parameters $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$.*

Discussion

- More parameters, overfitting
- Does F really capture the information externalities?
  - The selection effect
  - Position externalities overlap with F
- Tractability
- Small dataset

Recap

	Separable CTR Model	Markovian User Model	Ordered Search Model
Advertiser-specific parameters	p_i = P(click ad i | view ad i)	p_i = P(click ad i | view ad i)	$F_i(H)$ = P(click | view ad i ∧ click H) $F_i(\{\emptyset\}) = p_i$
Position-related parameters	α_j = P(view ad i | slot j)	q_i = P(cont | view ad i) $q_1 \times \ldots \times q_{j-1}$ = P(view ad i | slot j)	λ_i = P(cont | view + no click i) γ_i = P(cont | view + click i) $q_i = (1-F_i(H))\,\lambda_i + F_i(H)\,\gamma_i$

Interesting questions...

- Top-down assumption
- Fairness for individual advertisers
  - Is it really a negative externality?
- Trade-off: insights gained vs. increased complexity
- Future work: Tractability? More work on segregating user types? Comparison of revenues?