Part II: Bidding, Dynamics and Competition

Jon Feldman
S. Muthukrishnan

Campaign Optimization

Budget Optimization (BO): Simple

- Input:
  - Set of keywords and a budget.
  - For each keyword, (clicks, cost) pair.
    - Same auction all day, same competitors, bids.
- Model:
  - Take the keyword or leave it, binary decision.
  - Maximize the number of clicks, subject to the budget.
- Output:
  - Subset of keywords.

BO: Simple

- Well-known Knapsack problem.
  - Each KW is an item, cost = weight, clicks = value. Total budget = weight knapsack can carry.
- NP hard in general.
- Algorithm:
  - Repeatedly take item largest value/weight (clicks/cost), or lowest cost per click. Last item will be fractional. Provably optimal.
  - Undergrad algorithms: Sort by density=clicks/cost and be greedy.

BO: Multiple Slots

- Input:
  - For each keyword, multiple (clicks, cost) pairs.
- Generalized Knapsack:
  - Same item can be picked in different combinations.
  - NP hard in general.
  - Discrete problem solvable by Dynamic Programming. Pseudo-polynomial time.

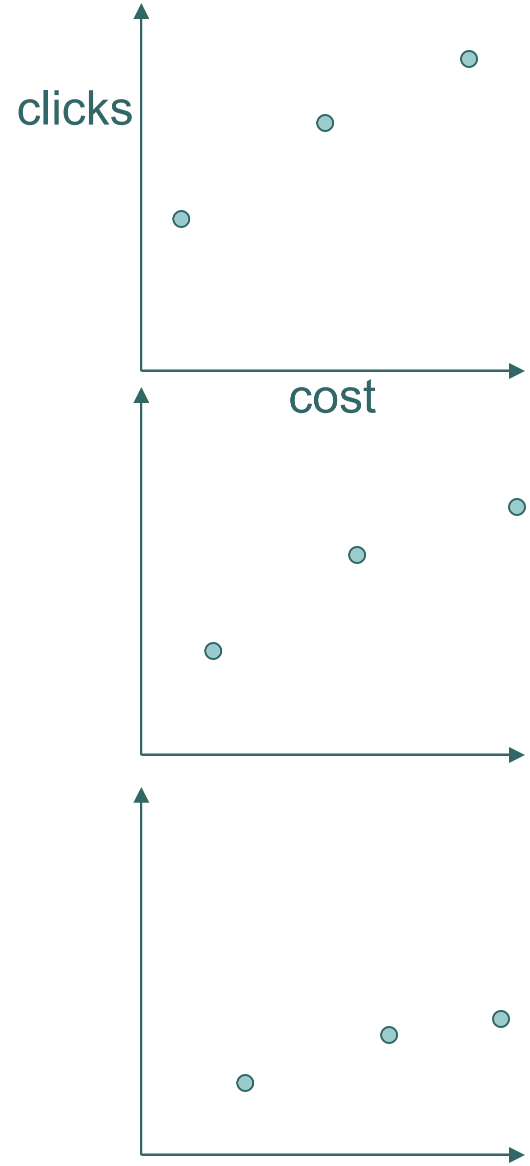

Multiple Slots BO: Some Observations

- Convex Hull. Taking convex combination will dominate other points.

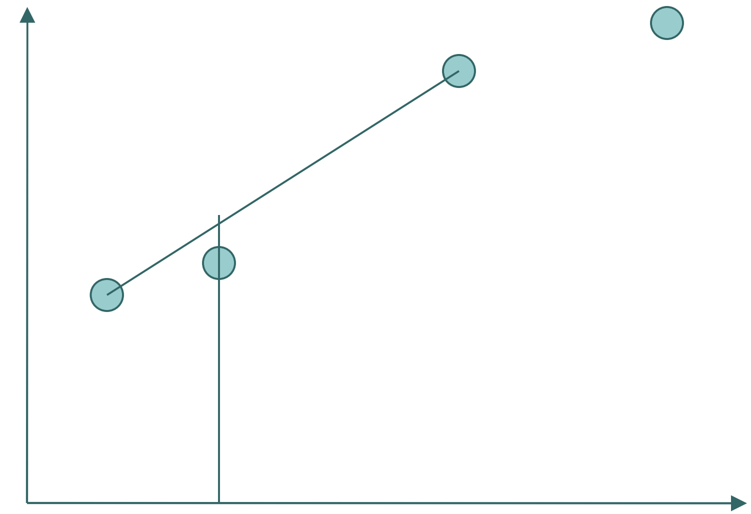

- Can treat each delta segment separately.

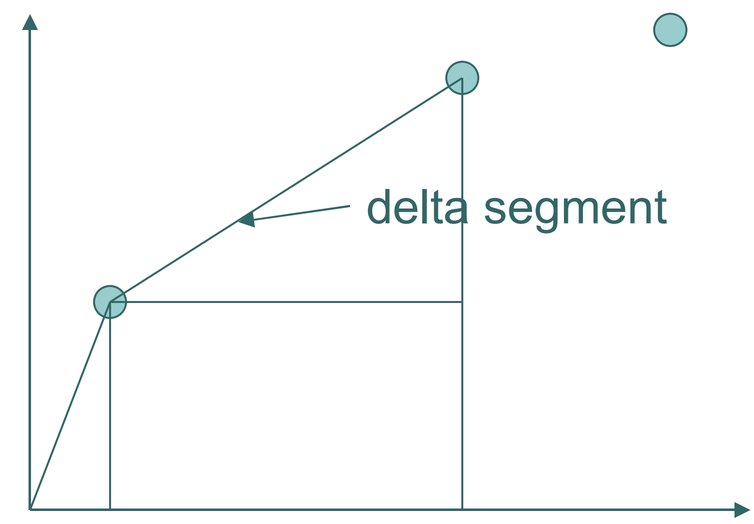

Multiple slots BO: Algorithm

- Consider each delta segment separately.
- Solve standard Knapsack as before.
  - Feasible since taken in order of decreasing clicks/cost.
  - Provably optimal.
- Message:
  - Algorithm produces x
  - Taking all delta segments (marginal) with

    cost-per-click ≤ x

    is the optimal solution.

Profit Optimization (PO)

- For each keyword (clicks, cost):
  profit = number of clicks * value – total cost.
- Profit Optimization: Maximize total profit.
- Take all profitable keywords. Optimal algorithm. No fractional issues.
- This algorithm targets marginal cpc = value.

PO with Budget

- Say budget B.
- Solve PO without B.
  - If spend < B, done.
  - Else, you will spend B. Then solve the BO problem given this B.

- [Homework] n KWs, k versions per KW. Preprocess them. Query is (V,B) or only V or only B. Solve BO or PO problems.
- Can be done in O(log (nk)) time. This data structure is landscapes.

XO: Optimizing X

- Conversion Optimization.
  - Given (conversions, cost), same algorithmics as above with cpc control knob.
- Maximize ROI = value/cost.
  - Get the 1 cheapest click!
- Improve ROI:
  - Bidding smartly
  - Improve the creative.
  - Change KW set,…

Target Positions

- Why?
- How?
  - Auction by auction.
  - Proxy bidding to average position target.
- BO/PO with Position Preference.
  - Simple: BO. Given budget B, for each KW, expected position < k.

Homework

- Given n keywords with k versions each find bids for keywords such that overall average CPC is at most x, and the number of clicks is maximized.
- Hint:
  - Algorithm will still proceed in increasing order of marginal CPCs.
- Formally,
  - Take increasing order of DeltaCost_i/DeltaClick_i.
  - Claim: sumDeltaCost_i/sumDeltaClick_i is also increasing. Hence stop when you get target average CPC.

XO Complicated

- 3 Examples:
  - Keyword Interaction
  - Stochastic Information
  - Broad Match

Keyword Interaction, BO Reexamined

- Keyword's interact.

- World is more complex.
  - Competitors drop in and out.
  - Multipliers change, traffic prediction is hard, …
- Landscape functions are now complicated.

Strategy: BO with keyword interaction

Let C be the number of clicks obtained by an Omniscent bidder.

- there exists a bid b such that

  $$\text{clicks}(\text{uniform}(b)) \geq C/2.$$

- There exists a distribution d over two bids such that

  $$\text{clicks}(\text{uniform}(d)) \geq (1-1/e)\, C.$$

Better in practice and a very useful heuristic.

Feldman, Muthu, Pal, Stein. EC 07.

Proof Sketch

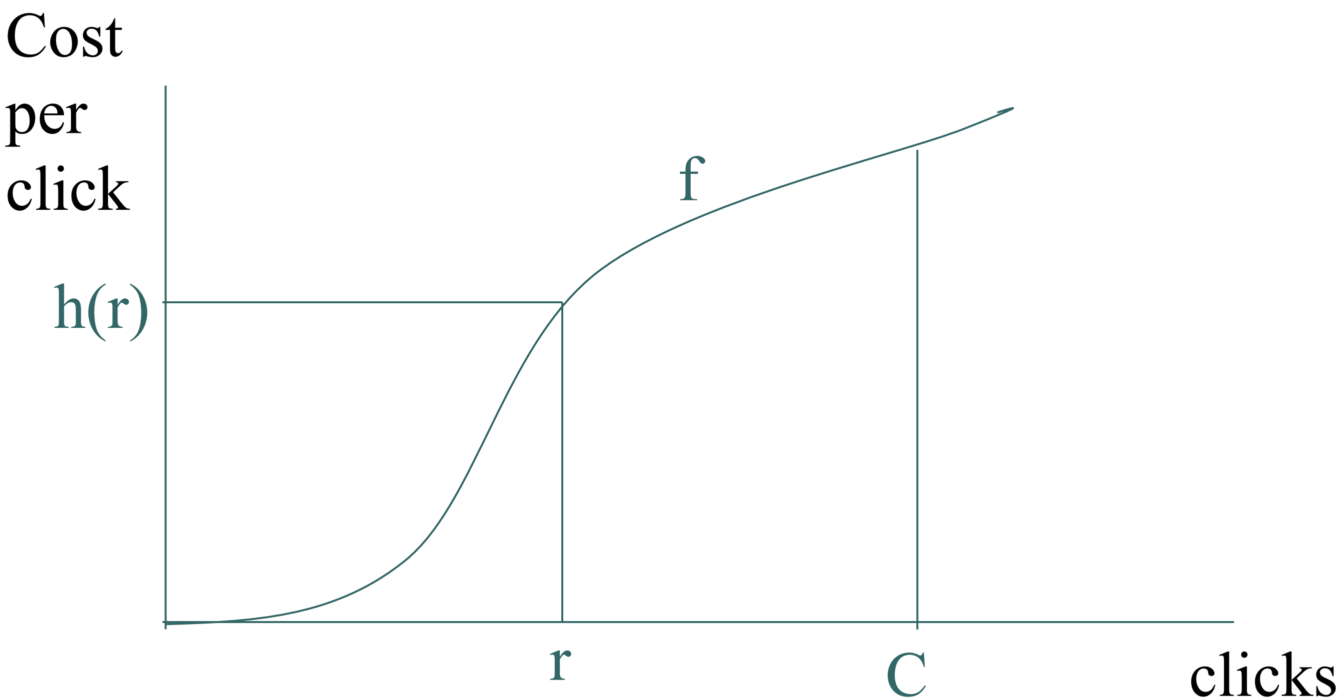

Bid h(r) on each query and

- get $\geq$ r clicks.
- spend $\leq$ r h(r).

With some work, r clicks at cost rh(r)

Proof Sketch (uniform bid)

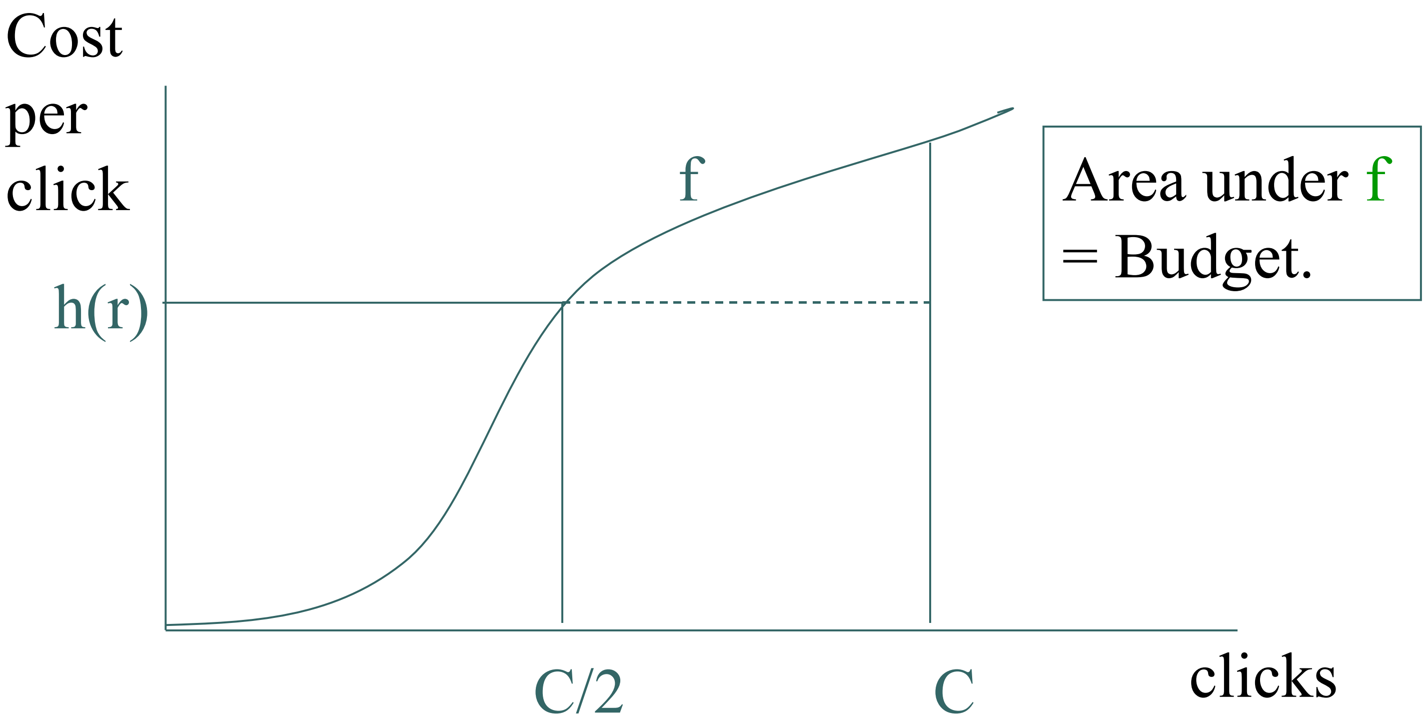

Bid $h(C/2)$ on each query and

- get $C/2$ clicks.
- spend $C/2\ h(C/2) \leq$ Budget

Analytical Puzzle

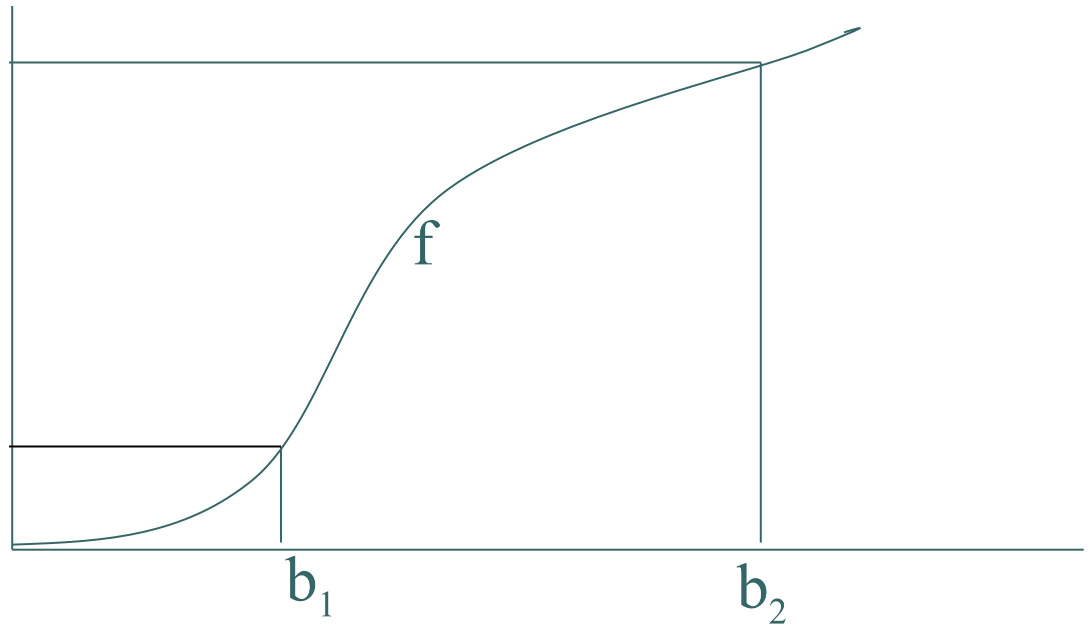

$$distribution: \quad \alpha_1 + \alpha_2 = 1$$

$$budget = \alpha_1 b_1 f(b_1) + \alpha_2 b_2 f(b_2)$$

$$\max \quad clicks = \alpha_1 b_1 + \alpha_2 b_2$$

PO with Keyword Interaction

- We can make up examples, so no profit approximation.
- Theorem: Say we can get profit P with value per click of V. Consider an uniform bidder with value eV/(e-1), gets profit at least P.
- Proof.
  - cl_o, co_o is what OPT gets and gives P_o.
  - Uniform theorm says there exists cl_u=(e-1)/e cl_o and co_u < co_opt.
  - Thus, if someone has value Ve/(e-1) then, profit_u= V e/(e-1) cl_u - co_u = v cl_o – co_o = profit_o.
- Open:
  - Position, Average CPC, etc. bidding when keywords have interaction.

Stochastic BO

- (click, cost) functions are random variables with dependencies.
- Three popular stochastic models:
  - Proportional
  - Independent
  - Scenario
- Variety of approximation algorithms known.

Muthu, Pal, Svitkina WINE07.

Stochastic BO: Scenario Model

- Each scenario gives (click, cost) distribution for keywords.
- There is a probability distribution over scenarios.
- Finding a bidding strategy to maximize expected clicks:
  - scaled by how much one overshoots the budget.
- Polylog approx, log hardness of approx.
- Technical key: “scaled” versions of combinatorial optimization problems.

Dasgupta, Muthu 09.

BO: Bidding Broad

- Advertisers have to choose how to bid Exact or Broad.
  - Because of impedance mismatch between user queries and bidding language for advertisers.
- Key technical difficulty in BO with broad match.
  - Bid on query/keyword q applies implicitly to keywords eg., q'.
  - While value from q may be large, value from q' may be even negative!

Bidding Broad

- Pick subset of queries to bid broad to maximize profit.
  - Polynomial time algorithms, even for budgeted versions.
- Bid on exact or broad on keywords to maximize profit.
  - Hard to even approximate (independent set).
  - O(1) approx if profit >>> cost.

Even-Dar, Mansour, Mirrokni, Muthu, Nedev WWW 09.

Grand XO

- More general problem is to combine
  - Keyword and match type choice
  - Target ad delivery and scheduling metrics
  - Learn CTRs
  - Optimize clicks, conversions, profit, brand effectiveness, …
  - For given budget.
- Alternatively, think at higher level of abstraction of supply curve: (cost, value).
  - The knobs like max cpc bids are just implementations.
  - For each budget, Auctioneer can run BO, PO, etc.
  - Advertiser needs to just pick a point.

Grander XO

- Advertisers have to optimize across channels.
  - Across search engines.
    - YMGA problem.
  - Across search and display.
  - Across online and offline.
- Formal models will be useful.

Dynamics

Bidding Dynamics

- How should advertisers bid?
  - Vickrey-Clarke-Groves (VCG), Truthfully.
  - Reality:
    - Other auctions (eg., Generalized Second Price, or GSP) and strategies in repeated auctions.
    - Portfolio of auctions.
  - Dynamics becomes important.

GSP: Static Game

- There exists an GSP equilibrium that has prices identical to VCG. It is the cheapest envy-free equilibrium.

B. Edelman, M. Ostrovsky and M. Schwarz. AER 07.
H. Varian. IJIO 07. G. Aggarwal, A. Goel and R. Motwani. EC06.

- GSP with bidder-specific reserve prices. There exists an envy-free equilibrium, even though we don't have local envy-free property.

E. Even-Dar, J. Feldman, Y. Mansour and Muthu, WINE08.

GSP: Dynamic Game

- Balanced Bidding (BB): Target the slot which maximizes the utility, and choose bid so you don't regret getting the higher slot at bid value.
- If all bidders follow BB, there exists a unique fixed point. Then revenue is VCG equilibrium revenue.

B. Edelman, M. Ostrovsky and M. Schwarz. AER 07.

- Asynchronous, random bidders with BB converges to this fixed point with prob. 1 in poly (k^2^k, max v_i, n) steps.

M. Carey, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu and M. Schwarz. EC07.

FP, GSP Dynamics: Multiple Keywords

- Budget limited bidders with multiple keywords.
- Bidding such that the marginal return on investment is same for all keywords.
- Equlibirium analysis
  - To avoid cycling, need perturbation of bids.
  - With first price and uniform bidding, prices, utilities and revenue converge to Arrow-Debreu market equilibrium.

C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain and M. Mahdian WWW07.

Competition

- A lot of auction design really deals with competitive behavior.
- Advertisers seem to ask about individual competitors.
  - Monitor for bids, quality, brand words,
  - Who are the competitors?
    - Micro competitors.
  - Why?
    - Relative bidding
    - Malicious bidding.

Y. Zhou and R. Lukose, WSAA06.

G. Iyengar, D. Phillips and C. Stein, SMC 07.

Summary

- [Jon] The Knobs.
- [Muthu] Controling the knobs wrt bidding.
  - Optimization: BO, PO, XO, …
  - Dynamics
  - Competition
- Rest
- Acknowledgements:
  - Martin Pal
  - Vahab Mirrokni, Eyal Even-Dar, Yishay Mansour, Hal Varian, Noam Nisan.
  - Uri Nadav, Cliff Stein, Bhaskar Dasgupta, Zoya Svitkina.
  - Team