Profit Maximization for Online Advertising Demand-Side Platforms

Paul Grigas, UC Berkeley

Joint with Alfonso Lobos, Zheng Wen, Kuang-chih Lee

2017 AdKDD & TargetAd Workshop at KDD 2017, Halifax, Canada

Outline

- Problem motivation and perspectives
- Optimization model preliminaries, assumptions, and properties
- Solution approach based on Lagrangian duality
- Synthetic computational results
- Conclusions and ongoing work

+

Problem Motivation

Problem Motivation

- What is a demand side platform? (DSP)
- DSPs *manage* the campaigns of many different advertisers and play a crucial role *connecting* them with publishers

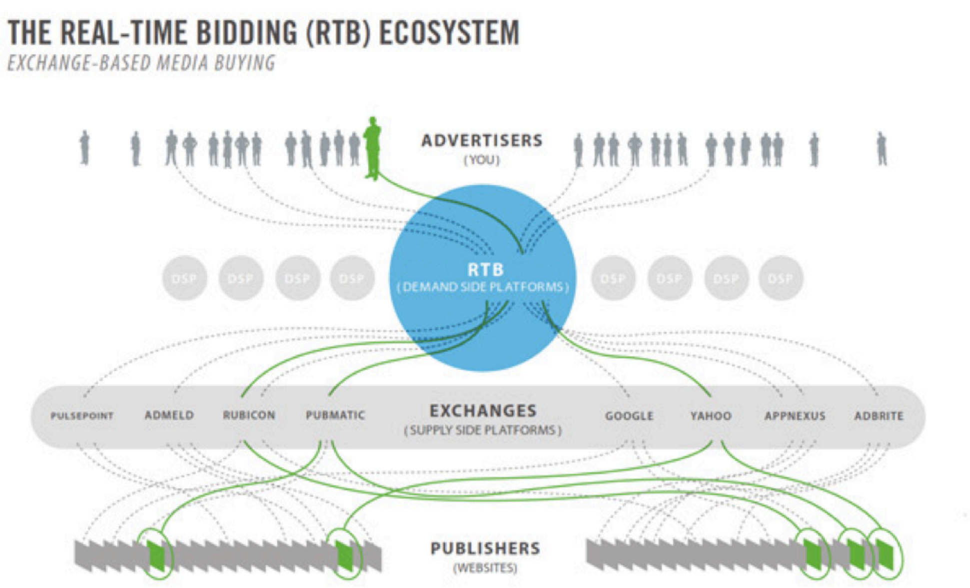

+ Problem Motivation cont.

- DSPs are faced with the challenge of managing advertisers' campaigns by interacting with ad exchanges in a real time bidding environment
  - Effective management requires forecasting the landscape of ad exchanges
- We focus on campaign management, particularly how to balance:
  - Meeting advertisers' goals and constraints
  - Profitability for the DSP
- DSPs may receive as many as a million ad requests per second and need to make decisions in real time
  - Thus simple greedy heuristics are often employed

Problem Formulation (in words)

- DSP profit maximization
- CPC/CPA pricing model
- Decision variables:
  - When a new impression arrives, who (among all the campaigns for the DSP) do we bid on behalf of and how much should we bid?
- Objective: maximize profit
- Constraints:
  - Campaign budget/pacing constraints
  - Targeting constraints
  - Supply (impression) availability constraints

Perspective and Contributions

- We develop a mathematical optimization formulation that:
  - Carefully models stochasticity in the real-time bidding process
  - Jointly optimizes over allocation strategies and bid prices
  - Accounts for limited supply of impression type inventory
- Our approach has several important features:
  - Scalability to the large-scale size of the problem
  - We address the stochastic nature of the problem
  - We account for the dynamic nature of the problem via model predictive control

"DSP Analytics Pipeline"

- A crucial input to our methodology is accurate forecasting of the value of an incoming impression, and how this value varies across different campaigns (e.g., CTR prediction)

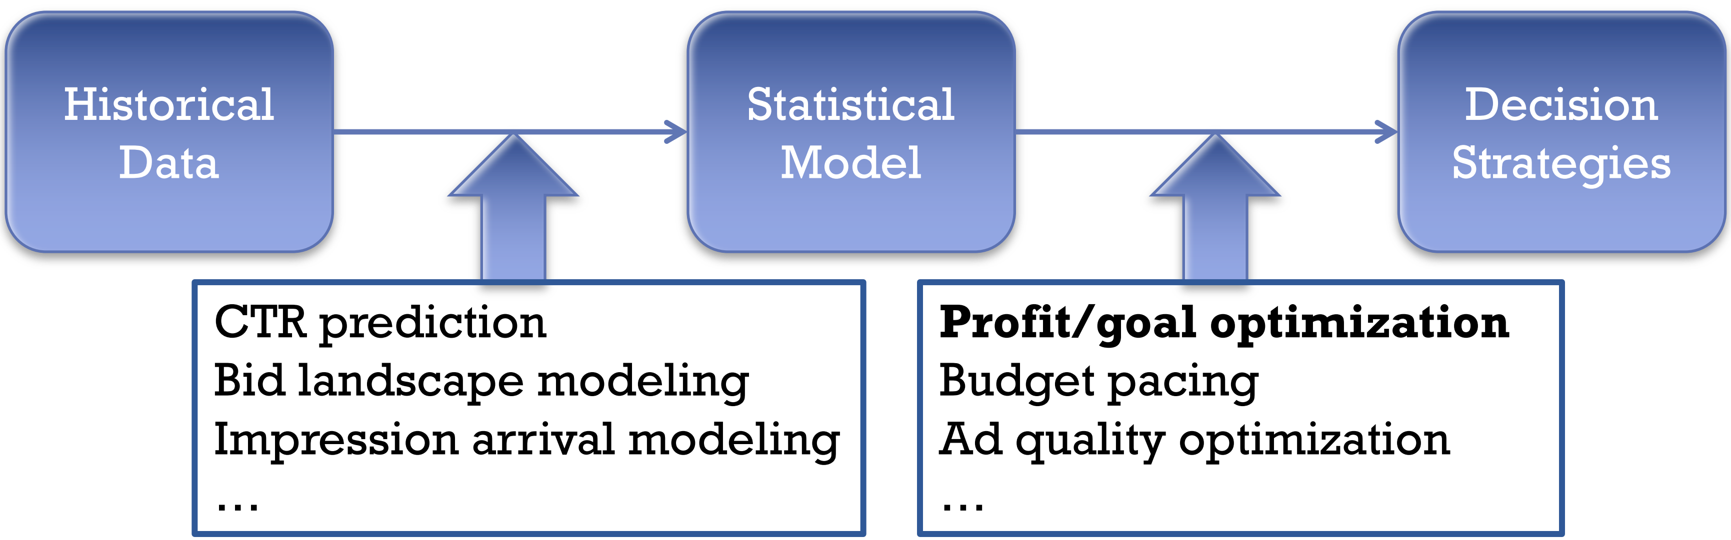

Related Literature

- Revenue Management for the Publisher
  - [Balseiro et al. 2014] and also [B. Chen et al. 2014] study how publishers should optimally trade-off guaranteed contracts with RTB
  - [Y. Chen et al. 2011] studies how a publisher should optimally allocate impressions and set up bid prices for campaigns, under an implicit "central planner" assumption
- Revenue Management for the Ad Network
  - [Ciocan and Farias 2012] provides theoretical performance guarantees for a model predictive control approach
- Profit Optimization for the Advertiser
  - [Zhang et al. 2014] studies optimal RTB bidding for an advertiser (without impression allocation)
- Others…

Model Preliminaries

Model Preliminaries

- Planning over a fixed time horizon
- $\mathcal{I}$ is the set of impression types
- $\mathcal{K}$ is the set of campaigns
- Targeting constraints are specified via a bipartite graph

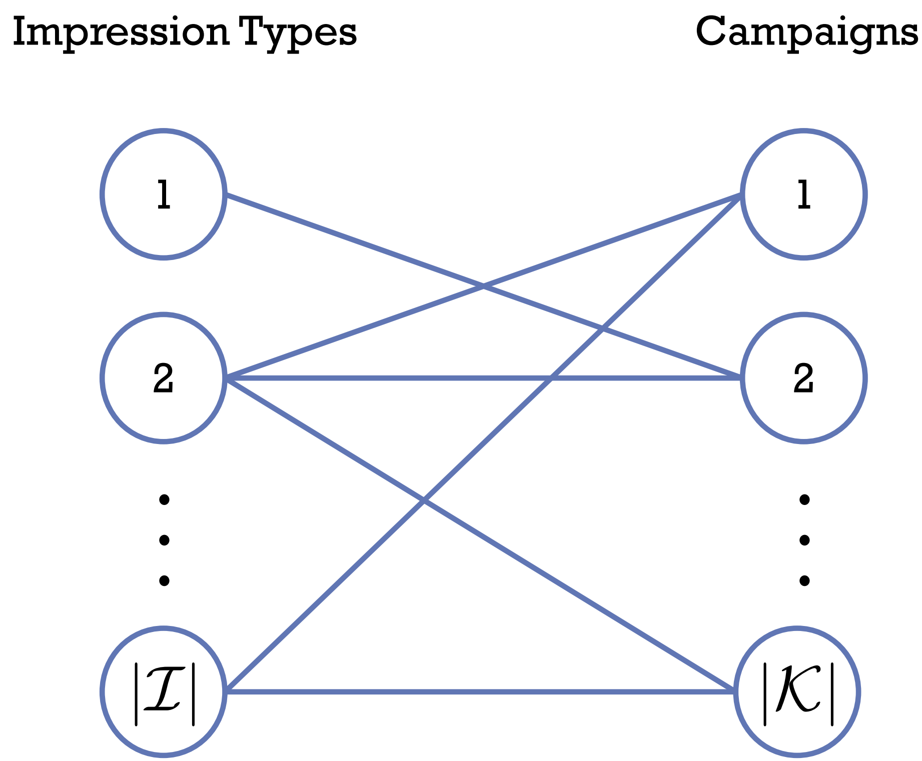

Modeling Impressions

- Impression types are defined via targeting (e.g. females, aged 25-34)
- Each arrival of impression type *i* corresponds to a real-time auction
- For each impression type, we assume that we can use a bid landscape forecasting model:
  - $\rho_i(b)$ is the probability of winning an auction for impression type *i* when entering bid b
  - $\beta_i^{\max}(b)$ is the expected second price, i.e., the expected payment if we win, as a function of the bid
- The total number of arrivals of impression type *i* is a random variable with mean s_i

Modeling Campaigns

- Each campaign has a fixed budget m_k over the time horizon
  - Budget pacing can be incorporated by controlling this input
- $\mathcal{I}_k$ is the set of impression types that campaign *k* targets
  - $\mathcal{K}_i$ is the set of campaigns targeted by impression type *i*
- $q_k > 0$ is the amount that campaign *k* is charged every time a click happens
- θ_{ik} is the predicted CTR for users of impression type *i* clicking on ads from campaign *k*
- $r_{ik} := q_k\theta_{ik}$ is the expected cost per impression (eCPI) value, which is the expected amount of revenue the DSP earns each time an ad from campaign *k* is shown to an impression *i*

+ Decision Variables and Corresponding Policy/Dynamics

- Decision variables:
  - x_{ik} is the probability of choosing campaign *k* to bid on behalf of when an arrival for impression type *i* occurs
  - b_{ik} is the corresponding bid price

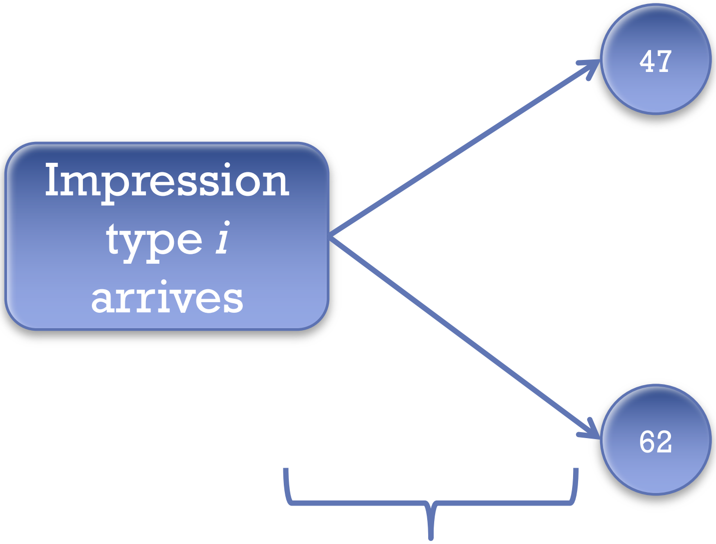

Flip coins with probabilities x_{ik} to decide which campaign to bid for

Policy Dynamics cont.

- Suppose that we bid b_{ik} on behalf of campaign k

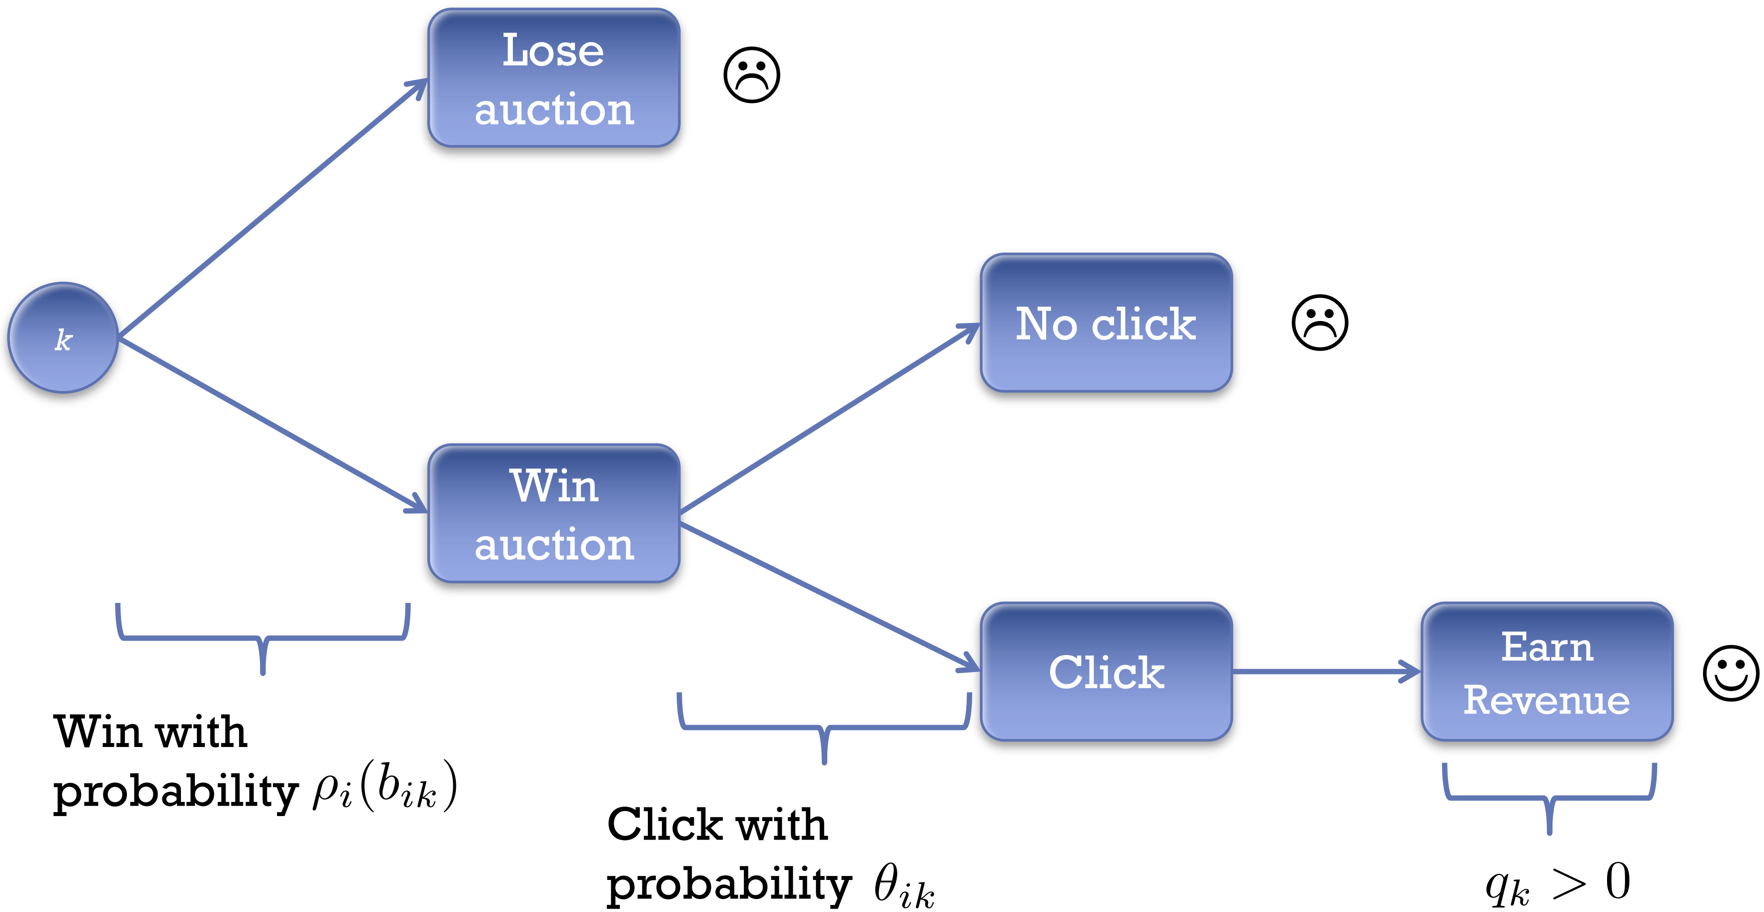

+ Optimization Formulation

- Deterministic optimization formulation, assuming all random variables take on their expected values:

$$\underset{\mathbf{x},\mathbf{b}}{\text{maximize}} \quad \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik})$$ **(Total profit)**

$$\text{subject to} \quad \textstyle\sum_{i\in\mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik}) \ \leq \ m_k \quad \forall k \in \mathcal{K}$$ **(Budget constraints)**

$$\textstyle\sum_{k\in\mathcal{K}_i} x_{ik} \ \leq \ 1 \quad \forall i \in \mathcal{I}$$ **(Supply constraints)**

$$\mathbf{x}, \mathbf{b} \ \geq \ 0 \ .$$

Properties of the Deterministic Approximation

- Due to joint optimization over allocation probabilities and bid prices, the deterministic approximation is generally non-convex
- "Difficulties" mainly arise due to the budget constraints
- Without the budget constraints, it is optimal to bid truthfully, i.e., to set $b_{ik}^* = r_{ik}$ and to greedily choose campaigns
- With budget constraints, it may be optimal for the DSP to underbid on a (relatively) less valuable impression due to the possibility of a more valuable impression arriving in the future
- For fixed bid prices, solving for the optimal allocation is a **linear optimization problem**

Solution Approach Based on Lagrangian Dual

+ Three Phase Solution Approach

- Phase 1: Solve (convex) dual problem obtained from Lagrangian relaxation of the deterministic problem
  - The main algorithm we use is subgradient descent (or some simple [e.g., stochastic] variant)
- Phase 2: Use optimal dual variables from Phase 1 to set bid prices
- Phase 3: Recover a "good" allocation strategy by solving the linear optimization problem obtained by fixing the bid prices determined from Phase 2
  - Solve using commercial LP solvers, or ADMM for large-scale problems

Useful Observations

$$\underset{\mathbf{x},\mathbf{b}}{\text{maximize}} \quad \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik})$$

$$\text{subject to} \quad \sum_{i\in\mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik}) \le m_k \quad \forall k \in \mathcal{K}$$

$$\mathbf{x} \in \mathcal{S}$$

($\mathcal{S}$ **denotes supply constraints)**

$$\mathbf{x}, \mathbf{b} \ge 0 \; .$$

- Phase 1 is based on the following (previous) observations:
  - The objective function is just the total expected profit in a second price auction
  - Without budget constraints, the optimal setting of bid prices is $b_{ij}^* = r_{ij}$, i.e. bidding truthfully
  - With budget constraints, it may be optimal to under bid – budget constraints are making the problem hard

Lagrangian Relaxation

- We put Lagrange multipliers $\lambda \in \mathbb{R}^m$ on the budget constraints and form the Lagrangian function:

$$
\begin{aligned}
L(\mathbf{x}, \mathbf{b}, \lambda) \;:=\; & \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\
& + \sum_{k\in\mathcal{K}} \lambda_k \left[m_k \;-\; \sum_{i\in\mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik}) \right]
\end{aligned}
$$

- Moving the budget constraint to the objective makes the problem "easy"

- Lagrangian may be re-written as:

$$
L(\mathbf{x}, \mathbf{b}, \lambda) = \sum_{(i,k)\in\mathcal{E}} [(1-\lambda_k) r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \;+\; \sum_{k\in\mathcal{K}} \lambda_k m_k
$$

Phase 1 – Dual Problem

$$L(\mathbf{x}, \mathbf{b}, \lambda) = \sum_{(i,k)\in\mathcal{E}} [(1-\lambda_k) r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) + \sum_{k\in\mathcal{K}} \lambda_k m_k$$

- The dual function is then $L^*(\lambda) := \max_{\mathbf{x}\in\mathcal{S}, \mathbf{b}\geq 0} L(\mathbf{x}, \mathbf{b}, \lambda)$
  - The dual function is convex
  - When computing the dual function, "it is optimal to bid truthfully and allocate impressions greedily"
  - Thus the dual function and subgradients of the dual function may be computed **efficiently**
- And the dual problem is:

$$\begin{aligned} \underset{\lambda}{\text{minimize}} \quad & L^*(\lambda) \\ \text{subject to} \quad & 0 \leq \lambda_k \leq 1 \quad \text{for all } k \in \mathcal{K} \;. \end{aligned}$$

Phase 1 – Dual Problem

$$\begin{array}{ll} \underset{\lambda}{\text{minimize}} & L^*(\lambda) \\ \text{subject to} & 0 \le \lambda_k \le 1 \quad \text{for all } k \in \mathcal{K} \ . \end{array}$$

- We use projected subgradient descent to solve this problem
- Effective when the number of impressions and the number of campaigns are very large
- This yields a vector of (approximately) optimal dual variables λ^*
- Importantly, $L^*(\lambda^*)$ provides an upper bound on the optimal profit

Phase 2 – Set Bid Prices

- Recall that the Lagrangian satisfies:

$$L(\mathbf{x}, \mathbf{b}, \lambda) = \sum_{(i,k)\in\mathcal{E}} [(1-\lambda_k) r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) + \sum_{k\in\mathcal{K}} \lambda_k m_k$$

- Thus, given optimal dual variables λ^*, we interpret $(1-\lambda_k^*) r_{ik}$ as a modified valuation/bid price that accounts for the budget constraint

- In Phase 2, we set $\hat{b}_{ik} := (1-\lambda_k^*) r_{ik}$
  - This setting is not necessarily optimal but should have good performance guarantees
  - Proving good performance guarantees is ongoing

Phase 3 – Allocation Recovery

- Phase 2 gives us a setting of the bid prices as $\hat{b}_{ik} := (1 - \lambda_k^*) r_{ik}$
- Phase 3: fix these bid prices and solve the deterministic linear optimization problem to recover allocation probabilities $\hat{\mathbf{x}}$
  - We end up with an approximate solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ to the original problem
- Linear optimization solution approaches:
  - Commercial solvers – scales to moderate to large size problems
  - Distributed/parallel ADMM based on decomposition across campaigns and impression types – scales to huge size problems

Synthetic Computational Set-up

- Each impression type and each campaign line has a quality score (QS_i and QS_k) that is uniformly distributed on [0,1]
- # of campaigns targeting impression type *i* is $\mathrm{Bin}(m, \mathrm{QS}_i)$
- $\rho_i(\cdot)$ is a maximum of $\mathrm{Bin}(k, \mathrm{QS}_i)$ uniform RVs
- The CTR is given by $\theta_{ik} := \mathrm{QS}_i \cdot \mathrm{QS}_k$
- Campaigns pay $1 for each click

Synthetic Computational Set-up cont.

- We compare the policy implied by our approach to a simple greedy baseline policy
- Simulations are based on our distributional assumptions (i.e., assuming perfect forecasting)
- Recall that campaigns pay $1 for each click
- The baseline policy chooses the campaign *j* with largest value of θ_{ik} and uses θ_{ik} as the corresponding bid price

Synthetic Computational Results

- Example 1: 100 campaigns, 100 impression types; budget is constant across campaigns, forecasted supply is constant across impression types
- Our approach solved the joint (allocation, bid price) problem to within 13% of optimality
- Simulation Statistics (averaged over 500 runs):

Relative Profit (our policy/baseline)	1.257
Relative cost	0.286
Relative Revenue	0.759

	Our policy	Baseline policy
Budget utilization	0.483	0.636
Profit/Revenue	0.807	0.487

Synthetic Computational Results cont.

- Same example, but now campaign budget is correlated with quality score

Relative Profit (our policy/baseline)	1.576
Relative cost	0.431
Relative Revenue	0.677

	Our policy	**Baseline policy**
Budget utilization	0.542	0.801
Profit/Revenue	0.500	0.215

Synthetic Computational Results cont.

- Example 2: 100 campaigns, 10 impression types
- Here we examine the effect of varying the average budget of each campaign (results averaged over 500 runs)

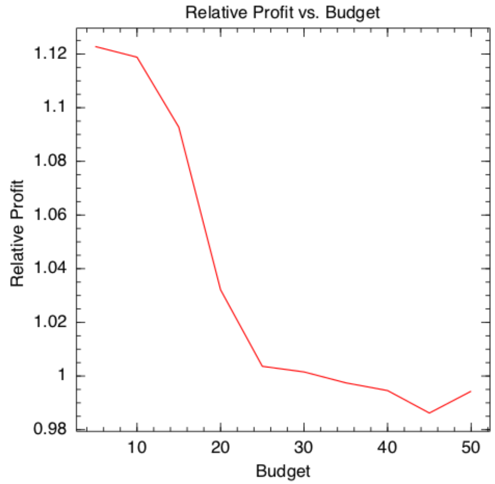

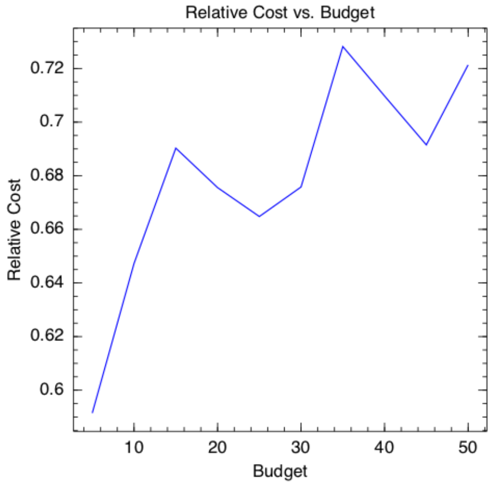

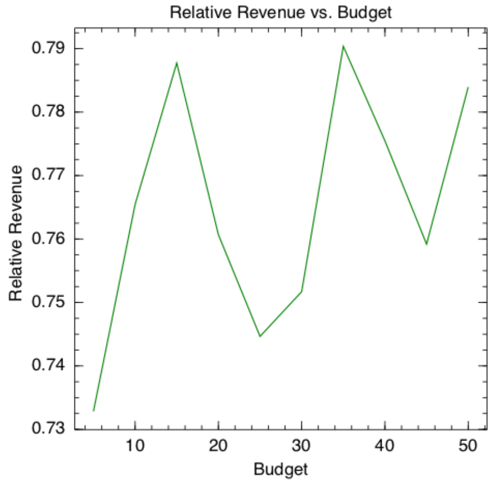

Conclusions and Ongoing Work

- Developed a mathematical optimization formulation for the management of a DSP that balances profitability with meeting advertisers' targeting goals and budget constraints
  - Our approach accounts for uncertainty in the real time bidding process
  - Jointly optimizes over allocation strategies and bid prices
- Developed a two phase solution approach to solve the non-convex joint (bid price, allocation) problem based on Lagrangian relaxation

Conclusions and Ongoing Work cont.

- Compared our policy against a baseline policy in simulations
  - Results indicate that our policy generates significantly more profit, mainly by reducing costs by avoiding overly-aggressive bidding strategies
- Ongoing work:
  - Extensions to incorporate advertisers' utility functions, model predictive control, CPM and/or oCPC pricing, robustness to uncertainty in parameter estimation, …
  - Theoretically characterize the gap between Lagrangian relaxation upper bound and recovered primal solution
  - More extensive computational evaluations based on real-world data

Thank You!