

McMaster University

Advanced Optimization Laboratory



Title:

Chance Constrained Optimization for
Targeted Internet Advertising

Authors:

Antoine Deza, Kai Huang, and Michael R. Metel

AdvOL-Report No. 2013/3

December 2013, Hamilton, Ontario, Canada

Chance Constrained Optimization for Targeted Internet Advertising

Antoine Deza^a, Kai Huang^b, Michael R. Metel^c

^a*Advanced Optimization Laboratory, Department of Computing and Software, McMaster University, Hamilton, Ontario, Canada*

^b*DeGroot School of Business, McMaster University, Hamilton, Ontario, Canada*

^c*School of Computational Science and Engineering, McMaster University, Hamilton, Ontario, Canada*

Abstract

We introduce a chance constrained optimization model for the fulfillment of guaranteed display Internet advertising campaigns. The proposed formulation for the allocation of display inventory takes into account the uncertainty of the supply of Internet viewers. We discuss and present theoretical and computational features of the model via Monte Carlo sampling and convex approximations. Theoretical upper and lower bounds are presented along with a numerical substantiation.

Keywords: Internet advertising, chance constrained programming, sample approximation, convex programming

1. Introduction

Internet advertising has witnessed growth of 15% in 2012, reaching \$36.6 billion in the United States [11]. This field is markedly different from traditional media used by advertisers such as radio, television and newspaper. Information such as a user's profile, data input and past Internet activity allow marketers to display their advertisements to targeted audiences, resulting in an efficient use of their advertising budget and an improved experience for users.

Email addresses: `deza@mcmaster.ca` (Antoine Deza), `khuang@mcmaster.ca` (Kai Huang), `metelm@mcmaster.ca` (Michael R. Metel)

Our work is interested in the planning of guaranteed display Internet advertising by an ad network, which acts as an intermediary between website publishers and advertisers. Advertisers purchase an advertising campaign from the ad network consisting of a guaranteed campaign goal, which is the number of ads to be displayed, and a set of viewer types, which describe who to show the campaign’s ads to. Guaranteed display advertising campaigns are typically for brand awareness where the industry practice is for ad networks to maximize representativeness, which is accomplished by displaying ads of each campaign as proportionally as possible to all targeted viewer types, see Yang et al. [20].

Quadratic optimization programs for this problem have recently been developed by Turner [18] and Yang et al. [20]. In particular, Turner showed that performance metrics are maximized using a proposed allocation methodology assuming the viewer supply follows a certain distribution. Our work addresses the uncertainty in viewer supply using a chance constrained framework. Bharadwaj et al. [4] presented an extension to [20] tangential to our research, using a two-stage stochastic program with recourse, with the second stage selling or purchasing ads on the spot market if the realized supply is greater or less than expected. An alternative objective to spread ads across campaigns is to maximize entropy, see Tomlin [17]. We pursue the quadratic objective value approach motivated in part by the availability of advanced and efficient solvers.

We introduce the model in Section 2 and formulate the joint chance constrained optimization program to solve the ad network’s problem. Section 3 discusses how lower and upper bounds can be found through sample approximations. In Section 4, a convex approximation program is presented which can be used to find lower and upper bounds under different Internet viewer distribution assumptions. We conclude with Section 5, which discusses the results of a computational substantiation of the introduced bounds. A nomenclature can be found at the end of the text.

2. Chance Constrained Optimization Model

2.1. Definitions and Notation

An online ad network is an aggregator of display ad slots, which it sells to advertisers in partnership with website publishers. For each guaranteed

display advertising campaign, the ad network displays ads to a targeted set of viewers that fit certain criteria, such as by demographic or interest. Advertisers are able to choose their targeted set of viewers from the set V of viewer types, which partitions the publishers' viewers by a predefined set of attributes. Namely, the supply of viewers is modeled as a $|V|$ -dimensional random variable with mean vector μ and covariance matrix Σ . Let \mathcal{S}_v denote the supply of incoming ad slots across all websites in the ad network loaded by individuals of viewer type $v \in V$, with μ_v and σ_v being its mean and standard deviation respectively. Let K denote the set of advertising campaigns. For a campaign $k \in K$, the campaign goal g_k is the number of ads to be displayed to viewers, which we assume is given. For research concerning optimal campaign goal sizes, see [1]. The subset of viewer types $V_k \subseteq V$ are the viewer types targeted by advertiser k . The subset of campaigns $K_v \subseteq K$ are the campaigns which target viewer type v .

This problem can be viewed as a stochastic transportation problem with each viewer type as a source with random supply and each advertising campaign as a sink with known demand. Each time a user loads a website affiliated with the ad network, a decision must be made as to which advertisement to display. This paper focuses on the high level planning stage at the beginning of each optimization time period, determining what proportion of ads from each viewer type to allocate to each applicable campaign. The decision variables of the ad network are p_{vk} , the proportion of each viewer type v 's supply allocated to each campaign $k \in K_v$. Another means of planning, especially when dealing with campaigns over short time periods, is by allocating ads to one minute time slots, whereby all visitors during each time period are shown the same ads, see [8].

2.2. Chance Constrained Optimization Program

We introduce the optimization program to find the proportion allocations, p_{vk} , for all viewer types and targeting campaigns, with an explanation following.

$$\begin{aligned}
& \min \sum_{k \in K} \frac{w_k}{|V_k|} \sum_{v \in V_k} (p_{vk} - q_k)^2 && \text{(CC)} \\
& \text{s.t.} \quad \sum_{k \in K_v} p_{vk} \leq 1 && \forall v \in V \\
& \quad \mathbb{P} \left(\sum_{v \in V_k} \mathcal{S}_v p_{vk} \geq g_k, \forall k \in K \right) \geq 1 - \alpha \\
& \quad q_k = \frac{1}{|V_k|} \sum_{v \in V_k} p_{vk} && \forall k \in K \\
& \quad p_{vk} \geq 0 && \forall k \in K, v \in V_k
\end{aligned}$$

The first constraint ensures that no more than 100% of a viewer type's supply is allocated. The fourth constraint ensures that proportions are non-negative. The second constraint captures the idea of guaranteed campaign fulfillment, which is interpreted as fulfillment with high probability. In particular, the second constraint ensures that all campaigns are fulfilled with a probability of at least $1 - \alpha$, where $\alpha < 0.5$ is the un-fulfillment tolerance.

Chance constrained programming has been used in many different fields such as finance [12] and water resource management [19]. We model campaign fulfillment using a chance constraint for two reasons. The first is that the success of an advertising campaign is unlikely to change dramatically if g_k ads or $(1 - \epsilon)g_k$ ads are displayed for some small percentage ϵ , whereas strictly requiring the former may significantly limit the number of advertising campaigns the ad network can accept. With the parameter α , the ad network is able to balance advertiser satisfaction with the total number of advertising campaigns executed. The second, more fundamental reason is that robust solutions are unlikely to exist without making strong assumptions on the underlying distribution of Internet viewers. $\mathbb{P}(\cup_{k \in K} \{\sum_{v \in V_k} \mathcal{S}_v < g_k\}) = 0$ is a necessary condition for the existence of a robust solution. For distribution assumptions of viewer type supply where this condition does not hold, e.g., normal, Poisson, log-normal, there exists a minimal $\hat{\alpha} > 0$ such that $\alpha \geq \hat{\alpha}$ for (CC) to be a feasible program.

The objective of the ad network is to maximize representativeness, which is achieved by allocating each campaign k 's ads across all $v \in V_k$ propor-

tionally to the supply. Objectives of the following general form have been proposed for guaranteed advertising campaigns, see [18, 20],

$$\min \sum_{k \in K} \sum_{v \in V_k} w_{vk} \left(p_{vk} - \frac{g_k}{\bar{\mu}_k} \right)^2$$

where the w_{vk} 's are weights, $\bar{\mu}_k = \sum_{v \in V_k} \mu_v$ is the total expected supply from the viewer types targeted by campaign k , and $\frac{g_k}{\bar{\mu}_k}$ is the target proportion. The objective maximizes weighted representativeness of campaigns, assuming the ad network is constrained to fulfill campaigns in expectation. Given the chance constraint, an ideal feasible allocation is unknown a priori. To construct a convex objective function reflecting the representativeness, we propose to minimize the variance of each campaign's allocation proportions. The objective is then

$$\sum_{k \in K} \frac{w_k}{|V_k|} \sum_{v \in V_k} (p_{vk} - q_k)^2$$

where q_k is the mean of the proportions allocated to campaign k from viewer types in V_k , enforced in the third constraint, and the weights w_k represent the campaign's priority to the ad network. For example, assume campaign k targets 5 viewer types, and a feasible solution to (CC) includes the vector of proportions allocated to campaign k , $p_k = [0.2, 0.3, 0.1, 0.4, 0]$. The objective then attempts to set $p_k = [\phi, \phi, \phi, \phi, \phi]$ for some unknown ϕ , which would achieve perfect representativeness for campaign k . We decided to minimize variance, but other measures of dispersion could also have been used.

Joint chance constraint programs are in general difficult to solve due to their non-convexity and the numerical integration required to calculate $\mathbb{P}(\sum_{v \in V_k} \mathcal{S}_v p_{vk} \geq g_k, \forall k \in K)$, see Pagnoncelli et al. [16]. Sample Approximations provide theoretically well founded solution approaches where Monte Carlo sampling is used to generate approximate mixed integer programs, see [5]. In Section 3, we present results enabling the construction of sample approximations which, when solved, achieve lower and upper bounds with high probability for (CC).

3. Sample Approximations

3.1. SA Program

The following program, (SA), is a finite approximation to (CC). For $i = 1, \dots, N$, the \mathcal{S}_v^i 's are independently sampled supply scenarios. The binary

variable $x_i = 1$ enforces the fulfillment of all campaign goals in scenario i . The second, third and sixth constraints require that all campaign goals are satisfied in at least $\lceil(1-\xi)N\rceil$ scenarios, which approximates the joint chance constraint $\mathbb{P}(\sum_{v \in V_k} \mathcal{S}_v p_{vk} \geq g_k, \forall k \in K) \geq 1 - \xi$.

$$\begin{aligned}
& \min \sum_{k \in K} \frac{w_k}{|V_k|} \sum_{v \in V_k} (p_{vk} - q_k)^2 && \text{(SA)} \\
& \text{s.t.} \quad \sum_{k \in K_v} p_{vk} && \leq 1 \quad \forall v \in V \\
& \quad \sum_{v \in V_k} \mathcal{S}_v^i p_{vk} && \geq x_i g_k \quad \forall k \in K, i = 1, \dots, N \\
& \quad \sum_{i=1}^N x_i && \geq \lceil(1-\xi)N\rceil \\
& \quad q_k && = \frac{1}{|V_k|} \sum_{v \in V_k} p_{vk} \quad \forall k \in K \\
& \quad p_{vk} && \geq 0 \quad \forall k \in K, v \in V_k \\
& \quad x_i && \in \{0, 1\} \quad \forall i = 1, \dots, N
\end{aligned}$$

We can obtain lower and upper bounds with high probability by solving (SA) with an appropriate choice for N and ξ .

3.2. SA Lower Bound

Assume (CC) is a feasible program with optimal objective value $z^{(CC)*}$ and optimal solution $p_{vk}^{(CC)*}$. Property 1 determines the probability of $p_{vk}^{(CC)*}$ being feasible in (SA), implying the optimal objective value of (SA), $z^{(SA)*} \leq z^{(CC)*}$. When (CC) is not feasible, $z^{(SA)*} \leq z^{(CC)*}$, using the convention that $z^{(CC)*} = \infty$.

Property 1 (Luedtke & Ahmed [14]). $\mathbb{P}(z^{(SA)*} \leq z^{(CC)*}) \geq \sum_{i=0}^{\lceil \xi N \rceil} \binom{N}{i} \alpha^i (1-\alpha)^{N-i}$.

3.3. SA Upper Bound

Property 2 requires that the objective is convex, the deterministic feasible region is convex and closed, and that the chance constraint mapping is closed and convex. Let (RSA) be the robust version of (SA) with $\xi = 0$. This implies all $x_i = 1$ converting (SA) into a convex quadratic program. Property 2 gives the probability that $z^{(RSA)*} \geq z^{(CC)*}$.

Property 2 (Calafiore & Campi [7]).

$\mathbb{P}(z^{(RSA)*} \geq z^{(CC)*}) \geq 1 - \binom{N}{|V_K|} (1 - \alpha)^{N - |V_K|}$, where $|V_K| = \sum_{k=1}^{|K|} |V_k|$ is the number of decision variables.

3.4. A Branching Scheme for the Branch-and-Bound Algorithm

In this subsection we discuss an aspect of the algorithm used to find sample approximation lower bounds, which enabled us to solve larger scale problems. Assuming we are in the midst of solving (SA), we must solve the following program, (SA_m) , at node m of the Branch-and-Bound algorithm.

$$\begin{aligned}
 & \min \sum_{k \in K} \frac{w_k}{|V_k|} \sum_{v \in V_k} (p_{vk} - q_k)^2 && (SA_m) \\
 \text{s.t.} \quad & \sum_{k \in K_v} p_{vk} && \leq 1 && \forall v \in V \\
 & \sum_{v \in V_k} \mathcal{S}_v^i p_{vk} && \geq x_i g_k && \forall k \in K, i = 1, \dots, N \\
 & \sum_{i=1}^N x_i && = \lceil (1 - \xi)N \rceil \\
 & x^T \text{diag}(X_m^1) && = \mathbf{1}^T \text{diag}(X_m^1) \\
 & x^T \text{diag}(X_m^0) && = \mathbf{0} \\
 & q_k && = \frac{1}{|V_k|} \sum_{v \in V_k} p_{vk} && \forall k \in K \\
 & p_{vk} && \geq 0 && \forall k \in K, v \in V_k \\
 & x_i && \in [0, 1] && \forall i = 1, \dots, N
 \end{aligned}$$

where X_m^1 and X_m^0 are binary vectors of length N which indicate the x_i set to one and zero at node m of the branching tree. We use an equality in the constraint $\sum_{i=1}^N x_i = \lceil (1 - \xi)N \rceil$, as for any integral optimal solution with $\sum_{i=1}^N x_i^* = \lceil (1 - \xi)N \rceil + r$ for some $r \in \mathbb{Z}_{>0}$, any r x_i^* 's equal to 1 can be set to 0 with $p^{(SA_m)*}$ remaining as the optimal solution.

After solving (SA_m) , assume that $p^{(SA_m)*}$ is not feasible in (SA) and the optimal objective value, $z^{(SA_m)*}$ is less than the current upper bound. Thus, we want to branch on one of the x_i for $i \in \{l : X_m^1(l) = 0, X_m^0(l) = 0\}$. We use the following heuristic which finds the scenario j with the constraint which is the farthest from being satisfied on a percentage basis:

$j = \underset{i: X_m^1(i)=0, X_m^0(i)=0}{\operatorname{argmin}} \min_{k=1, \dots, |K|} \frac{\sum_{v \in V_k} \mathcal{S}_v^i p_{vk}^{(\text{SAM})^*}}{g_k}$. For the path with $x_j = 1$, the branching tree can be effectively pruned as enforcing scenario j will likely enforce other scenarios, and the path with $x_j = 0$ will lead to near optimal solutions as x_j is a promising candidate for one of the $\lfloor \xi N \rfloor$ scenarios to discard.

4. Convex Approximations

In this section we present convex constraints which can replace the joint chance constraint in (CC) to achieve bounding convex programs efficiently solvable using an interior point method. Let \mathcal{S}_k be the vector of the viewer types' supply which campaign k targets, with μ_k being the $|V_k|$ -dimensional mean vector and Σ_k being the $|V_k| \times |V_k|$ covariance matrix. In addition, let p_k be the $|V_k|$ -dimensional vector of proportions allocated to campaign k from viewer types in V_k .

4.1. Distribution-free Bounds

This subsection assumes that we only have estimates for the first two moments with no knowledge of the underlying distribution. General methodologies for generating convex relaxations and restrictions have been developed by Ahmed [2] and Nemirovski & Shapiro [15], respectively. We present lower and upper bounds based on classic probability inequalities.

Property 3 (Distribution-free Lower Bound). *Any feasible solution of (CC) satisfies the constraints $p_k^T \mu_k \geq (1 - \alpha)g_k$ for $k \in K$.*

PROOF. Assume there exists a $k' \in K$ with $p_{k'}^T \mu_{k'} < (1 - \alpha)g_{k'}$, then $\mathbb{P}(p_k^T \mathcal{S}_k \geq g_k \forall k \in K) \leq \mathbb{P}(p_{k'}^T \mathcal{S}_{k'} \geq g_{k'}) \leq \frac{p_{k'}^T \mu_{k'}}{g_{k'}} < 1 - \alpha$, where the second inequality follows from Markov's inequality. \square

Property 4 (Distribution-free Upper Bound). *The constraints*

$$g_k - p_k^T \mu_k + \sqrt{\frac{1 - \alpha_k}{\alpha_k}} \sqrt{p_k^T \Sigma_k p_k} \leq 0 \quad \forall k \in K$$

where $\sum_{k \in K} \alpha_k = \alpha$, $\alpha_k > 0$, form a conservative approximation of $\mathbb{P}(p_k^T \mathcal{S}_k \geq g_k \forall k \in K) \geq 1 - \alpha$.

PROOF. Following the reasoning in [15, Sec. 2], assume $\mathbb{P}(p_k^T \mathcal{S}_k < g_k) \leq \alpha_k \forall k \in K$, then $\mathbb{P}(\cup_{k \in K} \{p_k^T \mathcal{S}_k < g_k\}) \leq \sum_{k \in K} \mathbb{P}(p_k^T \mathcal{S}_k < g_k) \leq \alpha$, implying $\mathbb{P}(p_k^T \mathcal{S}_k \geq g_k \forall k \in K) \geq 1 - \alpha$. To show $\mathbb{P}(p_k^T \mathcal{S}_k < g_k) \leq \alpha_k$, we use the one-sided Chebyshev inequality, $\mathbb{P}(Y \leq \mathbb{E}(Y) - b) \leq \frac{\text{Var}(Y)}{\text{Var}(Y) + b^2}$ for a random variable Y and constant $b > 0$,

$$\begin{aligned} \mathbb{P}(p_k^T \mathcal{S}_k < g_k) &\leq \mathbb{P}(p_k^T \mathcal{S}_k \leq g_k) \\ &\leq \frac{p_k^T \Sigma_k p_k}{p_k^T \Sigma_k p_k + (p_k^T \mu_k - g_k)^2} \\ &\leq \frac{p_k^T \Sigma_k p_k}{p_k^T \Sigma_k p_k + \frac{1 - \alpha_k}{\alpha_k} p_k^T \Sigma_k p_k} \\ &= \alpha_k. \end{aligned}$$

Therefore the conclusion holds. \square

4.2. Bounds Assuming a Normal Distribution

The normal distribution has been proposed in the literature for modeling viewer type supply, see [4]. This subsection presents convex approximations under the assumption that \mathcal{S}_k follows a multivariate normal distribution, so that $p_k^T \mathcal{S}_k \sim N(p_k^T \mu_k, p_k^T \Sigma_k p_k)$. Let F_k denote the cumulative distribution function of $p_k^T \mathcal{S}_k$.

4.2.1. Normal Lower Bound

Requiring each campaign's probability of fulfillment to be at least $1 - \alpha$ is necessary for feasibility in (CC), resulting in a convex relaxation. The chance constraint for each campaign is equivalent to a second-order cone constraint [6],

$$\begin{aligned} \mathbb{P}(p_k^T \mathcal{S}_k \geq g_k) &= 1 - F_k(g_k) \geq 1 - \alpha \\ F_k(g_k) &\leq \alpha \\ g_k &\leq F_k^{-1}(\alpha) \\ g_k &\leq p_k^T \mu_k + n_\alpha \sqrt{p_k^T \Sigma_k p_k}, \end{aligned}$$

where n_α is the α percentile of a standard normal random variable.

4.2.2. Normal Upper Bound

As in Property 4, an upper bound can be found by requiring $\mathbb{P}(p_k^T \mathcal{S}_k \leq g_k) \leq \alpha_k$ for all $k \in K$. Assuming the viewer supply follows a normal distribution, we can then use the constraints found in Subsection 4.2.1 with α replaced by α_k .

4.3. Summary

The above convex approximations can be obtained by solving the following convex approximation program (CA) with the proper choice of parameters u_k and h_k as summarized in Table 1.

$$\begin{aligned}
 & \min \sum_{k \in K} \frac{w_k}{|V_k|} \sum_{v \in V_k} (p_{vk} - q_k)^2 && \text{(CA)} \\
 & \text{s.t.} \quad \sum_{k \in K_v} p_{vk} && \leq 1 \quad \forall v \in V \\
 & \quad p_k^T \mu_k - u_k \sqrt{p_k^T \Sigma_k p_k} && \geq h_k \quad \forall k \in K \\
 & \quad q_k && = \frac{1}{|V_k|} \sum_{v \in V_k} p_{vk} \quad \forall k \in K \\
 & \quad p_{vk} && \geq 0 \quad \forall k \in K, v \in V_k
 \end{aligned}$$

Bound	\mathbf{u}_k	\mathbf{h}_k
Distribution-free Lower Bound	0	$(1 - \alpha)g_k$
Distribution-free Upper Bound	$\sqrt{\frac{1 - \alpha_k}{\alpha_k}}$	g_k
Normal Lower Bound	$-n_\alpha$	g_k
Normal Upper Bound	$-n_{\alpha_k}$	g_k

Table 1: Parameters in (CA)

(CA) was solved using a primal-dual interior point algorithm. The algorithm used generalized logarithm barriers to solve for points on the central path, see [6, Ch. 11.6-11.8]. To form the modified KKT conditions, the Jordan algebra for second-order cones is used to express the complementary slackness conditions of the second-order cone constraints, see [3]. The system of

equations to solve for the Newton steps was simplified so that only a system involving the step of p , Δp , was required to be solved, with closed form expressions for the remaining dual variable steps in terms of their current value, p , and Δp . In order to maintain stability, a universal step size was found such that all variables remained feasible. The central path parameter t is updated to equal a multiple of the reciprocal of the maximum error of the modified KKT conditions involving t . The algorithm quits when the maximum error of the modified KKT conditions is less than or equal to a small ϵ times the current objective function value. In order to find an initial feasible solution, we begin with an algebraic expression which spreads proportions relative to the campaign's goal size to expected targeted supply, which satisfies the first and fourth constraints of (CA). The Big M method is then used to find an initial solution feasible in the second set of constraints.

4.4. Setting the α_k 's

We now present an iterative method to calculate upper bounds. When finding a distribution-free upper bound, (CA) is first solved with the α_k 's set equal to $\frac{\alpha}{|K|}$, as proposed in [15, Sec. 2]. Letting p_k^* be the optimal solution, with optimal objective $z^{(CA)*}$, the approximating constraint of Property 4 can be rearranged as $\alpha_k \geq \frac{p_k^{*T} \Sigma_k p_k^*}{p_k^{*T} \Sigma_k p_k^* + (g_k - p_k^{*T} \mu_k)^2} = \hat{\alpha}_k$. For any k for which this constraint is not tight, we can set $\alpha_k = \hat{\alpha}_k$. As these tighter constraints are valid for (CA), resolving the optimization problem with the tighter constraints, (TCA), will result in an objective value $z^{(TCA)*} = z^{(CA)*}$. Assuming there was at least one constraint in (CA) which had slack, $\sum_k \alpha_k < \alpha$ in (TCA). The total slack $s = \alpha - \sum_k \alpha_k$ can be added evenly to all α_k 's of the originally tight constraints in (CA). Solving this relaxation of (TCA), (RTCA), will result in an objective value $z^{(RTCA)*} \leq z^{(TCA)*}$. This process of redistributing slack among the α_k 's is iterated until the improvement in the objective values becomes sufficiently small. The same process is used for the normal upper bound, where for all approximating constraints with slack, α_k is updated to equal $F_k(g_k)$. The algorithm for solving the distribution-free upper bound is presented below, with the necessary changes to solve for the normal upper bound in the comments.

5. Computational Substantiation

In this section we compare the solutions of the sample and convex approximations. All testing was conducted on a Windows 7 Home Premium 64-bit,

Algorithm 1 Calculating the Distribution-Free Upper Bound

```
1:  $\alpha_k = \frac{\alpha}{|K|} \forall k \in K$ 
2:  $u_k = \sqrt{\frac{1-\alpha_k}{\alpha_k}} \forall k \in K \{u_k = -n_{\alpha_k} \text{ for the normal upper bound.}\}$ 
3:  $[z^*, p^*] = CA(u, g)$ 
4:  $z_{old}^* = \infty$ 
5:  $I = 0^{|K|}$  {Indicator vector with  $k^{th}$  entry set to 1 when slack found in
   constraint associated with campaign  $k$ .}
6: while  $z_{old}^* - z^* > 0$  do
7:    $s = 0$  {Stores total slack across all constraints.}
8:   for  $k=1:|K|$  do
9:     if  $\alpha_k > \hat{\alpha}_k$  then
10:       $s = s + \alpha_k - \hat{\alpha}_k$  { $\hat{\alpha}_k = F_k(g_k)$  for the normal upper bound.}
11:       $\alpha_k = \hat{\alpha}_k$ 
12:       $I_k = 1$ 
13:     end if
14:   end for
15:   if  $s > 0$  &  $\sum_{j=1}^{|K|} I_j < |K|$  then
16:     for  $k=1:|K|$  do
17:       if  $I_k = 0$  then
18:          $\alpha_k = \alpha_k + \frac{s}{|K| - \sum_{j=1}^{|K|} I_j}$ 
19:       end if
20:     end for
21:   end if
22:    $z_{old}^* = z^*$ 
23:    $u_k = \sqrt{\frac{1-\alpha_k}{\alpha_k}} \forall k \in K \{u_k = -n_{\alpha_k} \text{ for the normal upper bound.}\}$ 
24:    $[z^*, p^*] = CA(u, g)$ 
25: end while
```

Intel Core i5-2320 3GHz processor with 8 GB of RAM. All coding was done in Matlab R2012a interfaced with CPLEX 12.4 using YALMIP [13] dated 13-Feb-2013. Ten random test problems were generated. For each test problem, the number of campaigns and viewer types were chosen randomly between [5,10] and [10,20]. Campaign targeting was achieved by generating a $|K| \times |V|$ matrix of Bernoulli random variables with $p = 0.5$, with cell $(i, j) = 1$ indicating that campaign i targets viewer type j . If there was a campaign or viewer type which was not assigned at least one viewer type or campaign, then a random cell in the appropriate row or column was set to 1. A random vector of viewer type means were generated, with each mean following a uniform distribution between [1000,10000]. Given the mean, μ_v , σ_v^2 was randomly generated uniformly within $[0.25, 0.5] \times \mu_v$. A random correlation matrix was generated using the random Gram matrix approach [10]. For each campaign, $g_k = U_{[0.5,0.75]} \sum_{v \in V_k} \frac{\mu_v}{|K_v|}$, where $U_{[0.5,0.75]}$ is uniform between [0.5, 0.75]. $w_k = 1$ for all campaigns.

The sample approximation parameters for each test problem were chosen so that the optimal solution is between the bounds with a probability of at least 99%. We tested all bounds sampling the viewer supply from a normal distribution. $\alpha = 0.1$ for the first five test problems and $\alpha = 0.05$ for the remaining five.

Let the probability of fulfillment (PF) equal $\mathbb{P}(\sum_{v \in V_k} \mathcal{S}_v p_{vk} \geq g_k, \forall k \in K)$. This probability is estimated for all solutions by generating 100,000 supply scenarios. Indicator variables, $\mathbb{1}_{\{\sum_{v \in V_k} \mathcal{S}_v^i p_{vk} \geq g_k, \forall k \in K\}}$, for each scenario i were generated and treated as a Bernoulli sample. The 99% one-sided confidence interval of the probability of fulfillment, \hat{PF} , was then estimated, $\mathbb{P}(PF \geq \hat{PF}) = 0.99$.

Results for each test problem are displayed in the Appendix. Objective values were multiplied by 1000 for readability. For the sample approximation bounds, the lower bound objective, the 99% one-sided confidence interval of the lower bound solution's probability of fulfillment, the lower bound computation time using the branching heuristic of Subsection 3.4, the lower bound computation time solving directly with CPLEX, the upper bound objective, the 99% one-sided confidence interval of the upper bound solution's probability of fulfillment, and the upper bound computation time are presented

from left to right in Table 2 . When computing the lower bound directly with CPLEX, a time limit of $20 * T_H$ was set, after which CPLEX would quit.

For the convex approximation bounds, the lower bound objective, the 99% one-sided confidence interval of the lower bound solution’s probability of fulfillment, and the lower bound computation time comprise columns 2-4 of Tables 3 and 4. With $\alpha_k = \frac{\alpha}{|K|}$, the upper bound objective, the 99% one-sided confidence interval of the upper bound solution’s probability of fulfillment, and the computation time follow in columns 5-7. Using the algorithm of Subsection 4.4, the upper bound objective, the 99% one-sided confidence interval of the upper bound solution’s probability of fulfillment, and the computation time are displayed in columns 8-10.

The average optimality gap for the sample approximation bounds was 43%, with an average computation time of 320 seconds using the heuristic, which significantly improved the computation time. The average optimality gap and computation time for the distribution-free bounds was 385% and 0.21 seconds using the algorithm. The average improvement of the upper bound using the algorithm was 16%. The average optimality gap and computation time for the normal bounds was 11% and 0.15 seconds using the algorithm. The average improvement of the upper bound using the algorithm was 4%.

6. Conclusion and Future Research

This paper presents a chance constrained optimization model for guaranteed displayed Internet advertising campaigns. A sample approximation program with a branching heuristic was developed, as well as convex approximations under normal and distribution-free viewer supply assumptions, with an iterative method for improved optimality of feasible solutions. Log-normal and Poisson distributions have also been proposed to model viewer supply, see [4, 9]. Convex approximations under these assumptions is an area of potential future research.

Acknowledgements

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada, MITACS, and by the Canada Research Chairs program. The authors would like to thank the anonymous referees for their helpful comments.

Nomenclature

α	Campaign un-fulfillment tolerance.
μ	Mean vector of viewer type supply.
μ_k	Mean vector of the viewer types' supply which campaign k targets.
μ_v	Mean supply from viewer type v .
\mathcal{S}_k	Vector of the viewer types' supply which campaign k targets.
\mathcal{S}_v	Supply of viewer type v .
Σ	Covariance matrix of viewer type supply.
Σ_k	Covariance matrix of the viewer types' supply which campaign k targets.
σ_v	Standard deviation of viewer type v 's supply.
ξ	Campaign un-fulfillment tolerance for (SA).
g_k	Campaign goal of campaign k .
K	Set of advertising campaigns.
K_v	Set of campaigns which target viewer type v .
N	Number of viewer type supply scenarios for (SA).
p_k	Vector of proportions allocated to campaign k from viewer types in V_k .
p_{vk}	Proportion of viewer type v 's supply allocated to campaign $k \in K_v$.
V	Set of viewer types.
V_k	Set of viewer types targeted by campaign k .
w_k	Campaign k 's priority weighting.

- [1] Ahmed, M. T., Kwon, C., 2014. Optimal Contract-Sizing in Online Display Advertising for Publishers with Regret Considerations. *Omega* 42 (1), 201–212.

- [2] Ahmed, S., 2014. Convex Relaxations of Chance Constrained Optimization Problems. *Optimization Letters*, to appear.
- [3] Alizadeh, F., Goldfarb, D., 2001. Second-Order Cone Programming. *Mathematical Programming* 95 (1), 3–51.
- [4] Bharadwaj, V., Saunders, M. A., Tomlin, J. A., 2010. Sensitivity of an Inventory Allocation Model for Online Advertising with Stochastic Supply. Presentation at the 24th European Conference on Operational Research.
- [5] Birge, J., Louveaux, F., 2011. *Introduction to Stochastic Programming*. Springer, New York.
- [6] Boyd, S., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press, Cambridge, U.K.
- [7] Calafiore, G. C., Campi, M. C., 2006. The Scenario Approach to Robust Control Design. *IEEE Transactions on Automatic Control* 51 (5), 742–753.
- [8] Deane, J., Agarwal, A., 2012. Scheduling Online Advertisements to Maximize Revenue Under Variable Display Frequency. *Omega* 40 (5), 562–570.
- [9] Girgin, S., Mary, J., Preux, P., Nicol, O., 2012. Managing Advertising Campaigns - An Approximate Planning Approach. *Frontiers of Computer Science* 6 (2), 209–229.
- [10] Holmes, R. B., 1991. On Random Correlation Matrices. *SIAM Journal on Matrix Analysis and Applications* 12 (2), 239–272.
- [11] Interactive Advertising Bureau, 2013. IAB Internet Advertising Revenue Report: 2012 Full Year Results.
URL <http://www.iab.net/>
- [12] Li, S. X., 1995. An Insurance and Investment Portfolio Model Using Chance Constrained Programming. *Omega* 23 (5), 577–585.
- [13] Löfberg, J., 2004. YALMIP : A Toolbox for Modeling and Optimization in MATLAB. In: *IEEE International Symposium on Computer Aided Control Systems Design 2004*. pp. 284–289.

- [14] Luedtke, J., Ahmed, S., 2008. A Sample Approximation Approach for Optimization with Probabilistic Constraints. *SIAM Journal on Optimization* 19 (2), 647–699.
- [15] Nemirovski, A., Shapiro, A., 2006. Convex approximations of chance constrained programs. *SIAM Journal of Optimization* 17, 969–996.
- [16] Pagnoncelli, B. K., Ahmed, S., Shapiro, A., 2009. Sample Average Approximation Method for Chance Constrained Programming: Theory and Applications. *Journal of Optimization Theory and Applications* 142 (2), 399–416.
- [17] Tomlin, J. A., 2000. An Entropy Approach to Unintrusive Targeted Advertising on the Web. In: *Proceedings of the 9th International World Wide Web Conference*. pp. 767–774.
- [18] Turner, J., 2012. The Planning of Guaranteed Targeted Display Advertising. *Operations Research* 60 (1), 18–33.
- [19] Wang, S., Huang, G. H., 2014. An Integrated Approach for Water Resources Decision Making Under Interactive and Compound Uncertainties. *Omega* 44, 32–40.
- [20] Yang, J., Vee, E., Vassilvitskii, S., Tomlin, J. A., Shanmugasundaram, J., Anastasakos, T., 2012. Inventory Allocation for Online Graphical Display Advertising. *Proceedings of the 1st International Conference on Operations Research and Enterprise Systems*.

Appendix

Sample Approximation Bounds								
#	LB				UB			
	z	$\bar{P}F$	T_H (s)	T_C (s)	z	$\bar{P}F$	T (s)	
1	0.08656	0.67382	394.12189	7885.28954	0.29815	0.99883	36.11891	
2	38.47802	0.83258	271.55118	5432.76618	43.57778	0.99931	12.98950	
3	1.39652	0.72996	118.50722	2372.77228	1.98864	0.99448	20.11220	
4	147.26848	0.78323	642.50342	5631.70639	159.05476	0.99384	12.74514	
5	1.55731	0.76499	228.41868	4570.70080	2.91692	0.99728	20.72935	
6	100.38367	0.76192	139.77481	1454.69702	117.07098	0.99687	43.10787	
7	209.21887	0.78952	375.03623	3746.96343	219.04089	0.99844	65.37173	
8	29.24341	0.83597	69.03840	1128.23085	32.06909	0.99838	73.24367	
9	0.00000	0.99995	0.45513	1.50496	0.00000	0.99995	36.52844	
10	313.32316	0.80525	578.03043	9679.22504	338.78909	0.99755	57.58939	

Table 2: Results for Sample Approximation Bounds

Distribution-free Bounds									
#	LB			$UB_{\frac{\alpha}{ K }}$			UB_{ALG}		
	z	$\bar{P}F$	T (s)	z	$\bar{P}F$	T (s)	z	$\bar{P}F$	T (s)
1	0.02433	0.10208	0.04318	1.56683	0.99995	0.13366	0.70294	0.99995	0.33721
2	35.11906	0.16493	0.00708	51.48740	0.99995	0.01638	49.67440	0.99995	0.19283
3	1.07339	0.11786	0.00652	4.90718	0.99995	0.02245	3.27246	0.99995	0.27213
4	134.91058	0.05268	0.00575	199.83478	0.99995	0.01840	194.25912	0.99995	0.17045
5	0.91310	0.17005	0.00610	9.19147	0.99995	0.01769	4.96429	0.99995	0.24479
6	88.71204	0.09186	0.00600	219.27627	0.99995	0.01562	198.94794	0.99995	0.24656
7	200.41163	0.01049	0.00738	276.35676	0.99995	0.02342	266.59733	0.99995	0.20504
8	26.67377	0.10453	0.00721	54.57600	0.99995	0.02209	52.12026	0.99995	0.24255
9	0.00000	0.99995	0.01565	0.00000	0.99995	0.01508	0.00000	0.99995	0.02982
10	290.49479	0.05637	0.00673	503.95350	0.99995	0.01890	503.95350	0.99995	0.03657

Table 3: Results for Distribution-free Bounds

Normal Bounds									
#	LB			$UB_{\frac{\alpha}{ K }}$			UB_{ALG}		
	z	$\bar{P}F$	T (s)	z	$\bar{P}F$	T (s)	z	$\bar{P}F$	T (s)
1	0.08488	0.68915	0.13207	0.16233	0.96111	0.02831	0.13225	0.90751	0.11698
2	37.80323	0.72256	0.04230	39.48617	0.93148	0.01569	39.27912	0.91504	0.04832
3	1.41510	0.70816	0.02301	1.71973	0.96048	0.01950	1.61711	0.91542	0.21837
4	144.25999	0.61781	0.01768	150.74349	0.92082	0.01747	150.16448	0.90605	0.05213
5	1.53287	0.74523	0.01855	2.11183	0.96274	0.01825	1.85374	0.90628	0.05421
6	100.78383	0.80557	0.01657	107.56587	0.96945	0.02045	106.72161	0.96050	0.06145
7	208.52184	0.73829	0.02060	213.06374	0.96324	0.01995	212.60627	0.95321	0.21879
8	29.15067	0.83168	0.02237	30.56063	0.97911	0.02198	30.08620	0.95146	0.07682
9	0.00000	0.99995	0.01484	0.00000	0.99995	0.01449	0.00000	0.99995	0.02862
10	311.30318	0.73201	0.02364	323.19042	0.95818	0.02235	323.18883	0.95787	0.24979

Table 4: Results for Normal Bounds