DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Expressive Banner Ad Auctions and Model-Based Online Optimization for Clearing

(Article begins on next page)

Expressive Banner Ad Auctions and Model-Based Online Optimization for Clearing

Craig Boutilier
Dept. of Computer Science
University of Toronto
Toronto, ON, CANADA
cebly@cs.toronto.edu

David C. Parkes
SEAS
Harvard University
Cambridge, MA, USA
parkes@eecs.harvard.edu

Tuomas Sandholm
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA
sandholm@cs.cmu.edu

William E. Walsh
CombineNet, Inc.
Fifteen 27th St.
Pittsburgh, PA, USA
wwalsh@combinenet.com

Abstract

We present the design of a banner advertising auction which is considerably more *expressive* than current designs. We describe a general model of expressive ad contracts/bidding and an allocation model that can be executed in real time through the assignment of fractions of *relevant ad channels* to specific advertiser contracts. The uncertainty in channel supply and demand is addressed by the formulation of a stochastic combinatorial optimization problem for channel allocation that is rerun periodically. We solve this in two different ways: fast deterministic optimization with respect to expectations; and a novel online sample-based stochastic optimization method—that can be applied to continuous decision spaces—which exploits the deterministic optimization as a black box. Experiments demonstrate the importance of expressive bidding and the value of stochastic optimization.[1]

Introduction

The prevalence and variety of online advertising in recent years has led to the development of an array of services for both advertisers and purveyors of online media. Because matching an advertiser's needs (demand) with a content provider's properties (e.g., locations on displayed web pages) is a complex enterprise, often automated matching is used to match ad channels with advertisers. One famous example is the dispatch of (typically textual) ads in response to keyword-based web searches, such as those on *Google*, *Yahoo!*, and *MSN*. In those settings, auctions are used to match the supply and demand (see, e.g., [7, 22]). Internet auctions of traditional advertising (TV, radio, print) are also emerging (e.g., via companies like *Google* and *Spot Runner*). Auctions and exchanges for banner ads have also been established—e.g., *Right Media* (now part of *Yahoo!*) and *DoubleClick* (now part of *Google*)—although many banner ad bulk contracts are still manually negotiated.

There has been considerable research on developing auction mechanisms for allocating ad channels, with a focus on issues like auction design [7, 14, 16, 10], charging schemes (e.g., per impression or per *click-through (CT)*) [13, 7, 22], bidder strategies [5, 9, 18], and so on. However, attention has focused almost exclusively on improving single-period expressiveness, still with per-impression or per-CT prices. As has been well-documented in other auction domains, requiring bidders and bid takers to shoehorn their preferences into the impoverished language of per-item bids is usually unnecessarily and undesirably restrictive. Significant increases in efficiency and revenue have been reported from auction designs that enable the participants to express their preferences in richer ways (e.g., [3, 19]).

In this paper, we explore the use of *expressive bidding* for online banner ad auctions.[2] In many domains, the value of a *set* of ads may not be an additive function of value of its individual elements. For instance, in an advertising campaign, *campaign-level* expressiveness is important. Advertisers may value particular *sequences* of ads, rather than individual ads per se. Allocative efficiency and revenue maximization in such an environment demand that we allow bidders to express bids (propose *contracts*) on complex allocations, and that bid takers optimize over *sequences* of allocations to best match bidder preferences, in a way that cannot be accommodated using per-item bidding.

The key technical challenge for expressive ad auctions is optimization: determining the optimal allocation of ad channels to very large numbers of complex bids in real-time. This is further complicated by the stochastic nature of the domain—both *supply* (number of impressions or CTs) and *demand* (future bids) are uncertain—which necessitates online allocation. To address these issues, we model the problem as a Markov decision process (MDP), whose solution is approximated in several ways. First we perform optimization only periodically. Following the general *optimize-and-dispatch* framework of Parkes and Sandholm [16], our optimization generates an on-line *dispatch policy* that assigns ad channels to advertisers in real-time. Our dispatch policies use the fractional assignment of (dynamically defined) channels to specific contracts. To approximate the optimization itself, we consider two approaches. The first is deterministic optimization using expectations of all random variables and exploiting powerful mixed integer programming (MIP) algorithms for expressive market clearing [19].

[2]For ease of presentation, we discuss banner ads, but the general principles and specific techniques we propose can be applied to other forms of online advertising (keyword search auctions, electronic auctions of TV and radio ads, etc.) as well.

We propose a second, sample-based approach derived from van Hentenryck and Bent's [12] online model for stochastic optimization—but with novel adaptations to a continuous decision space. This approach is able to leverage the MIP framework, applying it to multiple possible future scenarios in order to form a dispatch policy. In both cases, periodic reoptimization is used to overcome the approximate nature of the methods. We provide experiments to evaluate the benefits of expressive bidding for ad auctions over various per-item strategies, and the value and efficacy of our stochastic optimization techniques.

Expressive ad markets

We consider an ad network (or *seller*) that is charged with serving banner ads over a number of sites. Given a particular page view, the seller can display ads in various locations on that page. We assume a fixed set of locations L corresponding to particular page-location pairs. Constraints on location allocations can be imposed (e.g., to prevent allocation of overlapping locations). Each location has static and transient properties of potential interest. Static properties include the page identity (e.g., NY Times main page), page category (major news site), expected demographics, identity and size of the location (top banner, wide skyscraper, micro bar, etc.) and so on. Transient location (or allocation) features include time of day, page content, the presence of a competitor's ad on the same page, etc.

We assume time is divided into a discrete set of decision periods of suitable duration (e.g., several minutes, an hour, or even a day).[3] Ads are allocated to locations—possibly fractionally, so that the multiple impressions of each location are allocated across multiple ads—over entire periods.

The supply of locations is uncertain, dictated by a sequence of page hits to the sites in question, each hit "creating" a specific set of location realizations in the current period. The seller has a predictive distribution over page hits. If CTs are of interest, we also assume a model of CT probability (conditioned on location and ad features).

The seller receives bids of various forms from potential advertisers (*bidders*) that indicate their willingness to pay for specific allocation *schedules*, perhaps coupled with budget constraints. Bidders may be interested in CTs, impressions, or other actions induced by the display of the ad. As page hits occur, the seller must assign ads to the realized locations, ideally in such a way as to maximize *expected* revenue or some other objective over a horizon of interest.

Standard banner and keyword auctions allow bidders to express a cost per impression (CPI) or cost per clickthrough (CPC) together with a budget constraint over a particular period of time (e.g., hour or day). While certain forms of "local" expressiveness are provided to enable good matches to be made between an ad and instantaneous supply, little beyond budget constraints (e.g., [2]) is provided to allow for *sequential* or *campaign-level expressiveness* (but see [1, 8] for mild expressiveness extensions in keyword auctions).

A *(long-term) contract* expresses an advertiser's entire preferences (or willingness to pay) for a *set* or *sequence* of location allocations rather than individual allocations.[4] A variety of forms of sequential, campaign-level expressiveness are quite natural in banner ad auctions [16]. Examples of complex preferences that our model allows include:

- Minimum targets: a minimal target level in a specific period is desired (e.g., 100K impressions in a week) and payment occurs only if that target is reached. Or the offer may provide a small CPI for any number of impressions less than 100K, but a significant lump-sum bonus if the 100K target is met.
- Willingness to pay may be a function of multiple target levels (low saturation at 30K impressions may be of some value, high saturation at 100K of significantly greater value).
- Temporal sequencing: e.g., (a minimum) 20K impressions on the same page for each in a specific sequence of time periods (e.g., 11PM-1AM for each of the next 14 days).
- Substitution among properties: e.g., the same price for a time-limited campaign on either (but not both) of the NY Times or CNN; or offer a slightly higher price for the NY Times campaign. Note that substitution issues can benefit from our approach even in markets that have no temporal considerations.

The forms of local expressiveness (i.e., features of specific impressions or CTs) that can be handled in current auctions can also be incorporated into the conditions of a campaign-level contract. Thus context and other transient features can be incorporated into our *set preferences* (e.g., bid for 100K NY Times front page impressions this week; but offer a bonus if at least 20K of these hits include an article on health care). Additionally, expressive auctions can allow bidders to specify preferences directly in terms of their *target audience* (e.g., via demographic attributes), rather than only indirectly via ad location properties.

The following example illustrates the value of sequential expressiveness coupled with optimization.[5] There are two sites A and B. Bidder b_1 bids \$1 per thousand impressions on A and \$0.50 on B, with a budget of \$50K. Bidder b_2 bids \$0.50 per thousand impressions on A, with a budget of \$20K. Suppose supply on A is 5 times that of B for the first 50K units, but is then exhausted (only B has supply from then on). In a non-expressive auction, b_1 will win all of A's and B's supply until its budget is exhausted. Specifically, bidder 1 would win $(500/11)K$ impressions $((1)x + 0.5x/5 = 50K)$. At this point b_2 wins the remaining $(50/11)K$ impressions on A. Total revenue is $50 + (0.5)(50/11) \approx \$52.3K$. An optimal expressive auction would collect revenue of \$70K by selling 40K units of A to b_2, and 10K units of A plus 80K units of B to b_1.

Preliminaries: The optimization problem

Our optimization takes as input a set of long-term contracts that have been submitted to the auction. They specify all the

---

[3]These periods need not be of the same duration; the start of a new period may even be triggered dynamically by the occurrence of an event, such as an advertiser reaching its budget limit.

[4]Although a contract need not insist on a guarantee of a certain number of impressions/CTs, such guarantees can be handled by including penalties on targets not achieved.

[5]The example is simplistic since we do not provide equilibrium analysis of either auction. Nevertheless, it illustrates the advantage of global optimization over myopic bidding in non-expressive auctions—even when there is no uncertainty.

offers (including bids, constraints, bonuses, etc.) the bidders have made. We model the *spot market* for new bids—traditional bids without sequential expressiveness—and assume a spot market *demand distribution* P^D over location-price pairs. Spot demand can easily be treated as a standing contract containing only inexpressive bids. The seller has a predictive distribution over page hits, inducing a *supply distribution* P^S over locations for each period.

Suppose we have a set B of long-term contracts, with maximum horizon T (i.e., the final state of all contracts is determined by period T). For any $j \in B$, let $A_j^{1\cdots T}$ be a random variable denoting the set of locations assigned to contract j, and $R(j, A_j^{1\cdots T})$ the revenue generated by contract j given this realization of locations. Finally, let π denote the seller's *policy*, which assigns realized locations to contracts $j \in B$ in a history-dependent fashion. Our objective is to find a policy that maximizes expected revenue:

$$\mathbf{E}\left[\sum_{j \in B} R(j, A_j^{1\cdots T})|\pi\right], \qquad (1)$$

where the expectation is taken w.r.t. the distribution over supply (page hits) and demand (future bids).

The decision problem facing the bid taker can be modeled as a fully observable, finite-horizon Markov decision process (MDP) [6]. Ideally, a policy should take into account contract states after each "event" (e.g., page view) and determine the optimal allocation of locations to maximize expected future reward. Of course, the size of the state space in such an MDP renders its optimal solution infeasible. Even online approximations cannot be re-run at the time scale of individual events. Thus we use coarser-grained decision periods (e.g., at the level of minutes or hours). At each period, the seller assigns a fraction of a specific location (or channel as defined below) to each contract. This is the *optimize-and-dispatch* approach [16]. Given this, we can can define an MDP using the following components.

Channels. The supply of locations can be dynamically abstracted into *channels* based on the current contracts B. A channel *aggregates* locations: two locations will be part of the same channel if they are indistinguishable from the point of view of fulfilling the demands of any contract. Bids are then assigned to specific channels rather than specific locations, thus dramatically reducing the size of the decision space. Bidders do not specify channels in their contracts, only location properties of interest to them. The relevant channels are constructed automatically by developing a suitable algebra of location properties. We do not provide details for lack of space, but illustrate with a simple example: if bid b_1 makes an offer for banner ads on any page of the NY Times site (NYT), while b_2 makes an offer for any page with a medical article (Med), then three aggregate channels are created: one corresponding to a NYT and Med (i.e., a Times page with a medical article, with the potential to satisfy both bids), one to NYT without Med, and one to Med on a non-NYT site. This approach can render the space of channels exponentially smaller than the number of potential locations. We use subsumption and inconsistency based on the semantics of page properties to further reduce the number of relevant channels.

Supply (and demand) uncertainty. Precise supply, or *channel size* at any period t is not generally known in advance. For instance, we may not know the number of page hits for the NY Times Business front page between 2PM and 3PM. However, we assume that a distribution over channel size is derivable from the distributional information P^S over page hits. The general model also allows for anticipation of uncertain future demand via a distribution P^D on the spot demand, as described above. (We do not explicitly model new demand from expressive bids because there are infinitely many possible expressive bid types, but a rough model of them can be incorporated into spot demand.)

Decision space. The decision space consists of the assignment in each period of a percentage of the capacity of each channel to each contract. Decision variables are then $\{x_{ij}^t : i \leq C, j \leq B, t \leq T\}$, where x_{ij}^t denotes the percentage of channel i assigned to contract j at period t. Some channels will not be "relevant" to a particular contract (i.e., do not contribute to the satisfaction of that contract) and the corresponding x_{ij}^t are removed. With B contracts and an average of $\overline{C}$ relevant channels per contract, we have $O(B\overline{C})$ decision variables per period.

Abstraction techniques. The MDP decision space is dictated by the number of channels and time periods. If the time resolution specified by bids results in too large an action space to solve effectively (or too many decision variables for the optimization methods discussed below), we can aggregate time into larger intervals. The potential impact on optimality can be mitigated by performing the aggregation only for distant times, while maintaining finer-grained resolution in the near-term, especially since supply/demand predictions will tend to be less accurate deeper into the future.

Potentially more problematic is that certain sets of bids could cause an exponential explosion in the number of channels, even with a judicious channel construction algorithm. Here we have developed methods to further abstract channels beyond the granularity implied by the bids. For instance, in the example above, we could merge the two channels "NYT and Med" and "NYT without Med" into a single channel "NYT" for optimization. The dispatcher can correct for the loss in optimality to some degree by ensuring that bids are dispatched based on the actual features specified. That is, a bid for "NYT and Med" would be dispatched only to NY Times pages with medical articles, even if the optimizer-computed policy suggests otherwise. Variants of our optimization methods to account for such abstraction is beyond the scope of this paper.

MDP formulation. With these components in place, we can formulate the stochastic optimization problem as an MDP. For each contract $j \in B$, let S_j denote the set of *contract states*. A contract state $s_j \in S_j$ is a sufficient statistic summarizing relevant aspects of all past location allocations to j that enables the accurate prediction of contract satisfaction or revenue given any future sequence of allocations. The state space of the MDP is $S = \prod_{j \in B} S_j$.

Let X be the set of mappings

$$X = \{x : B \times C \to [0, 1] | \sum_{j \in B} x_{ij} \leq 1, \forall i \in C\} \qquad (2)$$

Here x_{ij} is the fraction of channel $i \in C$ is assigned to contract $j \in B$, and X is the decision space of the MDP (for period t). A nonstationary policy $\pi = \langle \pi^1, \ldots, \pi^T \rangle$ is a sequence of state-dependent fractional channel assignments to contracts: $\pi^t : S \rightarrow X$. The dynamics of the MDP is given by (time-dependent) transition functions P^t where $P^t(s'|x,s)$ denotes the probability of reaching state s' at the end of period t when the state entering period t is s and allocation x is used during period t. This transition function can be defined using the supply distribution P^S, CT rates, and (if appropriate) the demand distribution P^D.

Bellman equations can be used to define both the optimal value function and the optimal policy:

$$V^{T+1}(s) = R(s) \text{ and for } t \in [1,T] \text{ we have}$$
$$V^t(s) = \max_x R^t(s,x) + \sum_{s' \in S} P^t(s'|x,s)V^{t+1}(s')$$
$$\pi^t(s) = \arg\max_x R^t(s,x) \sum_{s' \in S} P^t(s'|x,s)V^{t+1}(s')$$

Here $R(s) = \sum_{j \in B} R_j(s_j)$ denotes the terminal value (at the end of period T) associating with realizing joint contract state s (reflecting any sequential or set-based revenue, e.g., bonuses). $R^t(s,x) = \sum_{j \in B} R^t(s_j, x_j)$ denotes the expected item-based revenue generated during period t.

The optimal policy will maximize expected revenue for the seller across the T-stage decision process. The key difficulty in solving this MDP is the size of the state space, consisting of the cross-product of the individual contract states as well as the high-dimensional continuous action space.

Expectation-based (re)optimization

One way of dealing with the complexity of solving this MDP is to ignore the uncertainty, solving a deterministic model in which uncertain channel sizes are replaced by their expectations. Let z_i^t denote the *expected* size of channel $i \in C$ at time t. Optimal allocation of channel capacity to contracts can readily be formulated as a mixed-integer program (MIP) for most natural forms of expressiveness. Specifically, we will have decision variables $x_{ij}^t \geq 0$ for each t within the horizon of the contract $j \in B$ and relevant channel $i \in C$ with the constraint $\sum_j x_{ij}^t \leq 1$ for all i,t; then $x_{ij}^t z_i^t$ denotes the quantity of channel i assigned to j in period t. We encode the objective to allow for accurate assessment of the payment of each contract j. The encoding depends on the contract language/expressiveness permitted; to give a flavor, consider a very simple example.

Suppose we have a contract $j \in B$ which pays for impressions on channels c_1 and c_2:

1. nothing for c_1 impressions if fewer than τ CTs
2. \$10,000 if at least τ CTs are achieved on c_1 by period t
3. \$8,000 if at least τ CTs are achieved on c_1 by $t' > t$.
4. \$0.50 per CT on c_1 after τ CTs have been achieved.
5. \$0.25 per impression on c_2 prior to time t''.

We encode the following as part of the objective in the MIP:
$$10000I_1 + 8000I_2 + 0.5T_1 + 0.25X_2,$$
where I_1 is an indicator variable denoting that τ CTs are achieved (in expectation) by t, I_2 denotes that τ CTs are achieved by t' (but not t), T_1 denotes how many CTs beyond τ have been achieved, and X_2 denotes the number of impressions on channel c_2.

The speed associated with moving from an MDP to a MIP is often dramatic. MIP solvers customized for auction-clearing can handle problems with tens of thousands of distinct items (multiple units of each), millions of bids, and hundreds of thousands of side constraints [19]. Our deterministic MIP formulation using expected channel size is reasonably tractable: the decision space is large, $O(B\overline{C}T)$, but manageable with suitable choice of period size and appropriate aggregation of channels.

If the distributions over expected supply (or future demand) have sufficiently high variance, then this expectation-based approach may be far from optimal, in particular, if we adhere to the expectation-based policy in the face of actual supply realizations that differ significantly from their means. Reoptimization offers a way of recovering from such deviations, and requires simply re-solving the MIP, using the updated contract states and updated supply (and demand) projections. While this does not allow one to account for risk optimally, it does allow a form of recovery from unexpected events. Reoptimization can be triggered at any time (e.g., when demand has drifted far from projection/expectation), and need not be tied to the time discretization used in the model. (In the experiments in this paper, we trigger reoptimization once every time period.)[6]

Online stochastic optimization

Another approach to solving a subclass of large scale MDPs is *sample-based online stochastic optimization* [12]. Samples are drawn from the distribution of uncertain events, and a deterministic optimization problem, or *scenario*, is constructed using each sampled realization. Each scenario is solved and the results are aggregated to construct an approximately optimal decision at the current period in the underlying MDP. The method is *online* in that the sample-based optimization is repeated after the current realization of uncertain events. That is, the approach determines the only the *next* action (in our case a fractional dispatch decision) rather than an entire policy for the MDP.

This approach can be extremely effective on problems for which good algorithms exist for the deterministic problem [12]. A critical aspect of the model is the requirement that domain uncertainty is exogenous; that is, *the distribution over future events should not be influenced by the actions taken by the decision maker*. This is, fortunately, roughly true in our domain: the assignment of channels to advertisers will have little discernible effect on the realization of future supply or demand. This *action independence* is vital as it allows valid sampling of scenarios prior to the optimization of these scenarios (i.e., action choice by the de-

---

[6]If time periods were so short that deterministic optimization could not be realized online between consecutive periods t and $t + 1$, it can be applied over multiple periods (e.g., the updated state after period t is used to compute a new deterministic policy that is put in place in period $t + q$ for some $q > 1$).

cision maker), a fact exploited to great effect below.

We adapt the REGRETS algorithm [4] to our banner ad optimization setting. In its original formulation, it assumes a set $\mathcal{X}$ of decisions are to be made at time t and that there is a generative model that can be used to sample uncertain events over horizon $[t, \ldots, T]$. The algorithm samples K scenarios, solving the deterministic optimization problem for each. Given scenario $\omega_k := GetSample(t + 1, T)$ and current state s_t, suppose $OptimalSoln(s_t, \omega_k)$ is "easily" solvable using some deterministic combinatorial optimization algorithm. In our setting, the offline problem consists of allocating known location supply to known ad demand over the planning horizon. Let x^* denote a solution to the offline problem with total value $w(x^*)$ and decision $x^*(t)$ in the current period. REGRETS works as follows:

Input: Current time t, decisions in current time $\mathcal{X}$,
$\quad$ K scenarios, s^t current state.
for $x \in \mathcal{X}$, $f(x) \leftarrow 0$
for $k = 1 \ldots K$:
$\quad \omega_k \leftarrow GetSample(t + 1, T)$
$\quad x^* \leftarrow OptimalSoln(s_t, \omega_k)$
$\quad f(x^*(t)) \leftarrow f(x^*(t)) + w(x^*)$
$\quad$ for $x \in \mathcal{X} \setminus \{x^*(t)\}$
$\quad\quad f(x) \leftarrow f(x) + w(x^*) - Regret(x^*, x, s_t, \omega_k)$
Output: Decision for time t is $\arg\max\{f(x) : x \in \mathcal{X}\}$
$\quad$ along with estimated expected value $f(x)/K$.

Here, $Regret(x^*, x, s_t, \omega_k)$ is an upper bound on the loss associated with taking decision x at time t (the current period) in scenario ω_k and then adopting the policy dictated by the deterministic solution x^* at future periods $t + 1, \cdots, T$, rather than executing x^* from the current period (i.e., acting optimally for ω_k). Using MDP terminology, regret bounds the (negative) *advantage* of action x relative to the optimal offline solution x^*; we can interpret regret as:

$$Regret(x^*, x, s_t, \omega_k) \geq \overline{V}_\pi^t(s_t) - \overline{Q}_x^t(s_t),$$

where $\overline{Q}_x^t(s_t)$ is the value of performing action x at time t and then "acting optimally" thereafter, and $\overline{V}_\pi^t(s_t)$ is the value of acting optimally (both relative to sampled scenario ω^k). Regret is then used in the estimate $f(x)/K$ of the Q-value of each action x in the current state.

We now consider the application of REGRETS to banner ad optimization. The algorithm is run once per period and is used to select the decision at the current period (time t). Once the decision is taken, new supply and demand information is observed and REGRETS reoptimizes for subsequent periods. A key feature of the REGRETS algorithm is the fact that, by assigning a value to each decision at time t, we can take advantage of a single optimization (for one sample) providing us with (perhaps crude) information about the expected value of $|\mathcal{X}|$ potential decisions, rather than just one. Of course, the effectiveness of this approach depends heavily on having a quality regret bound. Additionally, for REGRETS to be effective, regret computation must be much faster than full (deterministic) optimization.

The REGRETS method in [4] cannot be directly applied to ad optimization because *it requires the set of decisions $\mathcal{X}$ to be small and discrete.* A sampled scenario is a realization of channel sizes over time, and the deterministic optimization

finds the optimal allocation of capacity to contracts for that realization. Unfortunately, we must also estimate the value of all alternative stage t decisions for the scenario: in the context of our problem this is the *continuous fractional allocation of the supply of each channel to bids*, preventing the direct use of the REGRETS algorithm. Discretization of decision variables would be ineffective due to the dimensionality of the problem.

We propose a new technique for estimating the value of alternative decisions with large numbers of decision variables in continuous spaces, without enumeration of $\mathcal{X}$. We thus extend the applicability of the REGRETS algorithm. As in REGRETS we generate K scenarios (realizations of channel sizes) over the period $[t, \ldots, T]$, and solve the associated offline optimization problem for each. This gives, for each scenario $k \leq K$, a fractional allocation of each channel at each period to each contract. Denote this solution by $\dot{\mathbf{x}}_k = \langle \dot{\mathbf{x}}_k^t, \dot{\mathbf{x}}_k^{t+1}, \ldots, \dot{\mathbf{x}}_k^T \rangle$, where each $\dot{\mathbf{x}}_k^{t'}$ is a vector of allocations for period t': $\langle \dot{x}_{ij}^{t'} \rangle_{i \leq C, j \leq B}$. We cannot evaluate the continuous space of all *alternative* stage t decisions for each scenario. However, the ultimate goal is not actually to evaluate each such decision, but to find the stage t decision that has the highest (estimated) expected value over sampled scenarios. This can be accomplished by solving a single, relatively simple "scenario-aggregating" MIP *without* explicit enumeration of the decision space.

Let $\mathbf{x}^t = \langle x_{ij}^t \rangle_{i \leq C, j \leq B}$ be any stage t decision. We can estimate the Q-value of this alternative decision in scenario k by simply pinning down the allocation schedule $\dot{\mathbf{x}}_k$ at stages $[t + 1, \ldots, T]$ and replacing $\dot{\mathbf{x}}_k^t$ with this new value; note that the decisions in $[t + 1, \ldots, T]$ remain feasible in our setting because they specify a fractional allocation policy of whatever supply is realized. Indeed, this value is linear in the variables x_{ij}^t (as in the original MIP). Note that only stage t allocations are variable now; all decisions at later stages are fixed by $\dot{\mathbf{x}}_k$. Denote this value $Q_k^t(\mathbf{x}^t)$: this is equivalent to the (lower bound) estimate of Q-values for deterministic scenarios in the REGRETS algorithm (i.e., $w(x^*) - Regret(x^*, x, s_t, \omega_k)$). This provides an underestimate of the value of the new decision in scenario k.

We can now compute the stage t alternative decision with *maximum* expected "Q-value" over the K sampled scenarios by solving the following optimization problem (subject to the channel capacity constraints):

$$\max_{\mathbf{x}^t} \quad \frac{1}{K} \sum_{k \leq K} Q_k^t(\mathbf{x}^t) \tag{3}$$

This is the "optimal" decision for stage t and involves only decision variables for a single stage of the process (rather than for each stage). Thus once we have run K full optimizations for the K sampled scenarios, computing the "regret-sanctioned" optimal decision is straightforward.

There is some subtlety in dealing with budgets when solving the scenario-aggregating MIP. Let D_j be the (remaining) budget of contract $j \in B$, and let $Q_{k,j}^t(\mathbf{x}^t)$ be the portion of the Q-value ascribed to contract j under decision $\mathbf{x}^t$ in scenario k. To account for budgets, it is not appropriate to add constraints $Q_{k,j}^t(x^t) \leq D_j$ for each contract j.

For any decision $\mathbf{x}^t$, generally $Q_{k,j}^t(\mathbf{x}^t) \neq Q_{k',j}^t(\mathbf{x}^t)$ for different scenarios k and k', leading to the possibility that $Q_{k,j}^t(x^t) \leq D_j$ while $Q_{k',j}^t(x^t) > D_j$. However, any reasonable dispatch algorithm would stop serving ads to a contract once its budget limit is reached. Thus, we interpret $\mathbf{x}^t$ as specifying upper bounds on the allocation of supply to contracts; otherwise, the MIP will discard an allocation that is very good on average if the budget constraint is violated even in a single scenario. Let $\hat{Q}_{k,j}^t(x^t)$ be the "budget-independent" value obtained by j: we simply impose that $Q_{k,j}^t(x^t) = \max(\hat{Q}_{k,j}^t(x^t),\ D_j)$.

Empirical evaluation

To investigate the effectiveness of our expressive model and optimization techniques, we tested our methods on four sets of randomly generated problems. On two of these sets, we also compared our expressive methods to more "classic" auctions. The latter comparison was necessarily limited by the ability to understand how bids would be constructed for inexpressive auctions by bidders with expressive preferences. Our comparison is also complicated by the fact that equilibrium strategies are not known in expressive, dynamic first-price auctions of the kind studied here, nor in standard (non-expressive) dynamic auctions (first-price or otherwise) when bidders have non-linear valuations on sequences of allocations, or even when bidders have budget constraints.[7]

Ideally, we would generate bidder *preferences* for various campaign types, map these to suitable bids, and compare efficiency and revenue in different models. However, as discussed above, equilibrium bidding strategies are not known even for mild forms of expressiveness, so this is not feasible. Instead, we generate expressive bids/contracts directly for our expressive auctions. We compare revenue generated by both our expectation-based and stochastic optimization methods for such bids. These bids can also be viewed as surrogates for bidder preferences: thus we also use them as input to two heuristic bidding strategies we consider for traditional, non-expressive auctions. These heuristics are inspired in part by existing observations about bidding strategies in Internet advertising markets, as discussed later. This allows us to compare, subject to the appropriateness of our assumptions, the revenue properties of traditional and expressive auctions, illustrating the potential advantages of expressive bidding with respect to revenue (and, to the extent revenue reflects allocative efficiency, social welfare).

Our tests are divided along the following dimensions, each elaborated below. We first consider two different forms of expressive contracts, *Flat* and *Bonus*, reflecting different types of bidder *preferences*. We also consider allocating channel supply either using *classic* per-item auctions or our *expressive* techniques. As explained below, constructing bidding strategies for traditional auctions for bonus contracts

is ill-understood, so these are not tested on traditional auctions. For traditional bidders, we consider two different bidding strategies, *myopic optimization (MO)* and *bid-all (BA)*, which map expressive contracts or "preferences" into per-item bids in ways described later. For expressive auctions, the contracts are simply taken as given, but we compare our two different optimization techniques. The following table summarizes the major classes of experiments:

	Auction: Classic		Auction: Expressive	
	Bid: MO	Bid: BA	Opt: exp.	Opt: stoc
Pref: flat	✓	✓	✓	✓
Pref: bonus	×	×	✓	✓

Within each class, we also vary the supply distribution.

Preferences/contracts. We created four sets of ten randomly generated problems, each of which was characterized by one of two contract distributions, *flat* or *bonus*, and one of two channel supply distributions, *unimodal* or *bimodal*. All problems have 10 channels and 50 bidders. Each bidder $j \in B$ has a contract that is valid during time window $[T_j^-, T_j^+]$, with $T_j^- < T_j^+$ each drawn from $U[1, 10]$. A bidder has a positive bid on a subset of channels C_j, with $|C_j| \sim U[1, 10]$.

The *flat* contract distribution models the type of expressiveness supported in traditional ad auctions. A bidder has flat, per-unit bids on a set of channels, along with a budget over all its bids. Specifically, a bidder j has a per-impression bid $b_{j,i} \sim U[0.1, 1]$ for channel i. It also has random parameter $\alpha_j \sim U[0.1, 1.0]$, and its budget is set to $\alpha_j T_j \max_{i \in C} b_{j,i} \mu_i$, where μ_i is the mean supply of channel i in a single period and $T_j = T_j^- - T_j^+ + 1$, i.e., the number of periods for which a bid is valid. We model the spot market as a single bidder that bids 0.1 on all channels with no budget constraint. An example bidder j with a flat contract might have positive valuations on channels i and i' during time window $[3, 7]$. If the mean channel supplies are $\mu_i = 200$ and $\mu_{i'} = 100$, if j's bids are $b_{j,i} = \$0.30$ and $b_{j,i'} = \$0.70$, and if $\alpha_j = 0.6$, then j's budget is $0.6 \cdot (7 - 3 + 1) \cdot \max(0.3 \cdot 200, 0.7 \cdot 100) = \$210$.

The *bonus* contract distribution includes expressiveness not supported in traditional ad auctions. The distribution includes two types of bidders: bonus bidders and flat bidders. A bonus bidder offers a large payment if it reaches a bonus target for its bids, but offers only a small payment (below even the spot market value) if it misses the target. In contrast, a flat bidder offers higher per-unit bids, but no bonus. Uncertainty in channel supply makes it particularly challenging to maximize revenue given the bonus contract distribution. Since bonus bidders pay little if they miss their bonus targets, the optimizer must adequately account for risk when deciding what fraction to allocate to them.

The specific parameters of the bonus contract distribution are as follows. With probability 0.5, a bidder is as in the flat distribution, but with $b_{j,i} \sim U[0.5, 1]$. Otherwise, the bidder has a per-impression bid $b_{j,i} \sim U[0, 0.5]$ for $i \in C_j$, but is willing to pay an additional bonus $\hat{b}_j q_j$ if it obtains a total of q_j impressions on those channels C_j for which it has positive bids, with $\hat{b}_j \sim U[1, 5]$ and $q_j = \alpha_j T_j \sum_{i \in C_j} \mu_i$, where $\alpha_j \sim U[0.1, 1]$. The budget for a bonus bidder is

---

[7]In various special cases, equilibrium of generalized second-price pay-per-click auctions have been analyzed [22, 7, 13]; online generalizations of VCG for expressive, dynamic domains have been proposed [17]; and means for dealing with approximate policies in mechanisms exist [11, 15]. But none of these methods or analyses apply to expressiveness forms we consider here.

$\hat{b}_j q_j + \alpha T_j \max_{i \in C} b_{j,i} \mu_i$. We model the spot market as in the flat contract distribution, but with bid value 0.5.

An example bidder j with a bonus contract might have positive valuations on channels i and i' during time window $[3, 7]$. Let the mean channel supply for each be $\mu_i = 200$ and $\mu_{i'} = 100$. The per-unit bids are $b_{j,i} = \$0.10$ and $b_{j,i'} = \$0.30$, and the per-unit bonus is $\hat{b}_j = 3$. If $\alpha = 0.5$, then the bonus target is $q_j = 0.5 \cdot (7 - 3 + 1)(200 + 100) = 750$. Thus the bidder will pay a bonus of $3 \cdot 750 = \$2,250$ if it gets a total of 750 page views on channels i and i'. The bidder's budget is $750 \cdot \max(0.1 \cdot 200, 0.3 \cdot 100) + 2,250 = \$24,750$.

Supply distribution. The *unimodal* supply distribution models the case when supply is relatively steady and predictable. Here, for each period, a channel i has a supply drawn from a Poisson with mean $\mu_i \sim U[10, 1000]$.

The *bimodal* distribution (crudely) models the non-parametric nature of web traffic (e.g., how it might vary given a major news event). The supply of channel i is drawn from a mixture of two Poissons. A hidden binary variable determines which Poisson distribution is active at each stage. For a given channel, the mean of one Poisson is drawn from $U[10, 100]$, the other from $U[100, 1000]$. The state of the hidden variable persists for a random number of stages (Poisson, $\mu = 2$), after which it switches value (triggering a switch to the other distribution).

Bidding strategies. For the classic (inexpressive) auctions, we run a separate auction for each channel, but an overall budget constraint is enforced in dispatch. In all cases, the pricing rule is pay-your-bid, and, for simplicity we assume that payments are per impression. Bidding strategies for *flat* contracts have been widely studied for classic auctions; we consider two possibilities here. We refer to the *bid-all (BA)* strategy as that in which a bidder simply submits all of its positive-value bids. However, as some [5, 9, 18] have observed, if the supply is known and the highest competing bids are fixed and known (which they are not, of course), a bidder should select the bids that maximize its profit at the cost of the highest competing bids, subject to its budget constraint. We incorporate this idea into a *myopic optimization (MO)* strategy as follows. A bidder computes the set of channels that would maximize the value of its bids given its budget and the prices from the last auction round, and then submits its bids at face value. If it either won a channel in the last round or did not bid on a channel whose price was lower than its bid, the bidder assumes it could win the channel at its bid level and considers the channel in its optimization. Otherwise, it ignores the channel. The bidder then optimizes its channel selection assuming that the situation is fixed for all future periods. It can reoptimize at each stage.

To see how the *MO* strategy works, assume that bidder j's bids are $b_{j,1} = \$0.50$, $b_{j,2} = \$0.20$, $b_{j,3} = \$0.70$ and $b_{j,4} = \$0.60$ for channels 1, 2, 3, and 4, respectively. Assume that, in the previous round, j submitted bids $b_{j,1}$ and $b_{j,2}$, but not bids $b_{j,3}$ and $b_{j,4}$. Currently, j is winning channel 1 but not channel 2, and the prices for channels 3 and 4 are \$0.80 and \$0.50, respectively. When determining which bids to submit in the next auction round, j will include $b_{j,1}$ in its optimization because it is winning the bid, and will include

$b_{j,4}$ because the price of channel 4 is below its bid value. The bidder will not include $b_{j,2}$ or $b_{j,3}$ in its optimization because the prices channels 2 and 3 are above its value for them. The bidder then computes which of channels 1 and 4 will maximize value, given the mean supply of the channels and its budget, and then submits the selected bids (for either or both channels) at its bid values.

We chose not to develop bidding strategies for *bonus* contracts in classic auctions because their highly non-linear nature makes good strategies much less obvious.

Set up. In a given experiment run, each of the ten instances saw 100 trials, each with a different realization of channel supply. For the *flat* contracts, we ran all four methods on each trial (i.e., each method experienced the same realized supply), while we ran only the expressive methods on the *bonus* contracts (as explained above). The continuous REGRETS algorithm used 10 sampled scenarios in each trial to determine an allocation. For each instance, we simulate the auctions and channel realizations according to the supply distribution and bids. The classic auctions are run at each time stage. For classic auctions, we dispatch the ad of the highest bidder to a channel until its budget is depleted (during a given stage). For expressive auctions, we randomly dispatch according to the specified decision fractions. Once a contract exhausts its budget, we stop dispatching to it.

Experimental results. Table 1 compares the average *ex post* realized revenue from the bids for the two bidding strategies in traditional auctions and the two optimization methods for expressive auctions, considering *flat* contracts and for the two different models of supply. The MO strategy gives rise to greater revenue than the BA strategy, but it still realizes only $\sim 70\%$ of the revenue obtained by the expressive auctions. The stochastic and expectation-based approach perform (statistically) the same on these problems. Assuming that bids in the expressive auctions do indeed provide adequate surrogates for bidder preferences (enabling in turn a comparison with classic auctions populated with bidders with heuristic strategies), these auctions have revenue properties that are superior to traditional auctions, irrespective of whether the MO or BA bidding strategy is used by bidders. Furthermore, to the extent that increased revenue reflects improved allocative efficiency, then this advantage would also be expected to extend to efficiency.

Approach	*Unimodal* supply	*Bimodal* supply
Bid-all	$25,687 \pm 436$	$14,004 \pm 141$
Myopic	$30,256 \pm 437$	$15,890 \pm 175$
Expectation	$42,365 \pm 581$	$22,385 \pm 227$
Stochastic	$42,237 \pm 581$	$22,774 \pm 238$

Table 1: *Classic vs. expressive auctions on* flat *contracts. Average values shown with 95% confidence intervals.*

Table 2 compares the revenue for the expectation-based and our stochastic optimization algorithms, on the two sets of problems with *bonus* contracts. There is a pronounced advantage to using stochastic optimization when there are bonuses. The expectation-based algorithm obtains 67.1%–85.9% of the revenue that stochastic optimization yields. Note that the strong performance of the stochastic algorithm was achieved with few sample scenarios, requiring about 11

times (10 scenarios plus one aggregation optimization) as much computation as the expectation-based algorithm.

Approach	*Unimodal* supply	*Bimodal* supply
Expectation	$100,266 \pm 3,355$	$55,901 \pm 1,887$
Stochastic	$149,423 \pm 3,204$	$65,065 \pm 2,356$

Table 2: *Expectation-based vs. stochastic optimization on* bonus *contracts. Average values shown with 95% confidence intervals.*

Although performed without an equilibrium analysis, our results suggest that sequential optimization on expressive contracts can offer tremendous advantages for revenue, and we conjecture efficiency as well. These advantages are apparent *even for traditional forms of expressiveness*. Furthermore, stochastic optimization can provide a significant benefit over expectation-based optimization when additional, highly non-linear expressiveness forms are introduced.

Conclusions and future work

Within the *optimize-and-dispatch* framework [16], we presented a concrete design of a banner ad auction (applicable also to search keywords and ads for TV, radio, and newspapers), which, to the best of our knowledge, is more *expressive* than current designs. We described a general model of expressive ad contracts (or expressive bidding) and an allocation model that can be executed in real time through the assignment of a fraction of *relevant ad channels* to specific advertiser contracts. As a first practical approach to addressing the allocation problem, we presented combinatorial optimization based on expectations, accompanied by re-optimization if supply/demand differs significantly from projection. As a more refined alternative, we formulated the problem as a Markov decision process. Its solution is generally intractable (even offline). To address that, we proposed the use of sample-based online stochastic optimization techniques [12] to render the optimization problem tractable enough to admit online allocation. Our approach required the modification of these techniques to allow optimization in high-dimensional continuous action spaces.

Our experiments showed the superiority of our approaches over current non-expressive auction designs. They also showed that a sample-based approach can further improve the auction over a simple, expectations-based approach. Future work includes interviewing advertisers in different markets to determine which expressive bidding forms are (the most) important. Then, experiments are planned to determine the scalability of our expectations-based approach and our sample-based approach (under varying granularities of time discretization). Other directions include the investigation of automated abstraction techniques to minimize the number of MIP variables corresponding to channels and time periods, and examining the incentives and equilibrium properties of auctions with such substantial campaign-level expressiveness.

References

[1] Z. Abrams, A. Gosh, and E. Vee. Cost of conciseness in sponsored search auctions. In *3rd International Workshop on Internet and Network Economics (WINE'07)*, 2007.

[2] Z. Abrams, O. Mendelevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. *ACM Conf. on Elec. Comm.*, 2007.

[3] M. Benisch, N. Sadeh, and T. Sandholm. Theory of expressiveness in mechanisms In *AAAI-08*, 2008.

[4] R. Bent, P. Van Hentenryck. Regrets only! Online stochastic optimization under time constraints. In *AAAI-04*, 2004.

[5] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Bid optimization in online advertisement auctions. In *2nd Workshop on Sponsored Search Auctions*, 2006.

[6] C. Boutilier, T. Dean, and S. Hanks. Decision theoretic planning: Structural assumptions and computational leverage. *J. Artificial Intelligence Research*, 11:1–94, 1999.

[7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[8] E. Even-Dar, M. Kearns, and J. Wortman. Sponsored search with contexts. In *3rd International Workshop on Internet and Network Economics (WINE'07)*, 2007.

[9] J. Feldmann, S. Muthukrishnan, P. Martin, and C. Stein. Budget optimization in search-based advertising auctions. *AMEC-07*, 2007.

[10] J. Feng, H. Bhargava, and D. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS J. Comp.*, 19(1):137–148, 2007.

[11] M. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. *AAAI-07*, 2007.

[12] P. van Hentenryck and R. Bent. *Online Stochastic Combinatorial Optimization*. MIT Press, 2006.

[13] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. *ACM Conf. on Elec. Comm.*, 2006.

[14] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized on-line matching. *FOCS-05*, 2005.

[15] D. C. Parkes and Q. Duong. An ironing-based approach to adaptive online mechanism design in single-valued domains. *AAAI-07*, 2007.

[16] D. C. Parkes and T. Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *1st Workshop on Sponsored Search Auctions*, Vancouver, BC, 2005.

[17] D. C. Parkes and S. Singh. An MDP-based approach to Online Mechanism Design. In *NIPS-03*, 2003.

[18] P. Rusmevichientong and D. P. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. *ACM Conf. on Elec. Comm.*, 2006.

[19] T. Sandholm. Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.

[20] T. Sandholm, D. C. Parkes, and C. Boutilier. System for and method of expressive sequential auctions in a dynamic environment on a network. Patent application US 2006/0224496 A1, 10/2006 (priority from 3/2005).

[21] T. Sandholm, D. C. Parkes, C. Boutilier, and W. E. Walsh. System for and method of expressive auctions of user events. Patent application US 2008/0052219 A1, 2/2008 (priority from 3/2006).

[22] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.