

Bidding for Representative Allocations for Display Advertising*

Arpita Ghosh, Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii

Yahoo! Research.

{arpita, mcafee, kpapi, sergei}@yahoo-inc.com

Abstract. Display advertising has traditionally been sold via guaranteed contracts – a guaranteed contract is a deal between a publisher and an advertiser to allocate a certain number of impressions over a certain period, for a pre-specified price per impression. However, as spot markets for display ads, such as the RightMedia Exchange, have grown in prominence, the selection of advertisements to show on a given page is increasingly being chosen based on price, using an auction. As the number of participants in the exchange grows, the price of an impressions becomes a signal of its value. This correlation between price and value means that a seller implementing the contract through bidding should offer the contract buyer a range of prices, and not just the cheapest impressions necessary to fulfill its demand.

Implementing a contract using a range of prices, is akin to creating a mutual fund of advertising impressions, and requires *randomized bidding*. We characterize what allocations can be implemented with randomized bidding, namely those where the desired share obtained at each price is a non-increasing function of price. In addition, we provide a full characterization of when a set of campaigns are compatible and how to implement them with randomized bidding strategies.

1 Introduction

Display advertising — showing graphical ads on regular web pages, as opposed to textual ads on search pages — is approximately a \$24 billion business. There are two ways in which an advertiser looking to reach a specific audience (for example, 10 million males in California in July 2009) can buy such ad placements. One is the traditional method, where the advertiser enters into an agreement, called a guaranteed contract, directly with the publishers (owners of the webpages). Here, the publisher guarantees to deliver a prespecified number (10 million) of impressions matching the targeting requirements (male, from California) of the contract in the specified time frame (July 2009). The second is to participate in a spot market for display ads, such as the RightMedia Exchange, where advertisers can buy impressions one pageview at a time: every time a user loads a page with a spot for advertising, an auction is held where advertisers can bid for

* A full version of this paper appears in [6]

the opportunity to display a graphical ad to this user. Both the guaranteed and spot markets for display advertising now thrive side-by-side. There is demand for guaranteed contracts from advertisers who want to hedge against future uncertainty of supply. For example, an advertiser who must reach a certain audience during a critical period of time (e.g. around a forthcoming product launch, such as a movie release) may not want to risk the uncertainty of a spot market; a guaranteed contract insures the publisher as well against fluctuations in demand. At the same time, a spot market allows the advertisers to bid for specific opportunities, permitting very fine grained targeting based on user tracking. Currently, RightMedia runs over nine billion auctions for display ads everyday.

How should a *publisher* decide which of her supply of impressions to allocate to her guaranteed contracts, and which to sell on the spot market? One obvious solution is to fulfill the guaranteed demand first, and then sell the remaining inventory on the spot market. However, spot market prices are often quite different for two impressions that both satisfy the targeting requirements of a guaranteed contract, since *different impressions have different value*. For example, the impressions from two users with identical demographics can have different value, based on different search behavior reflecting purchase intent for one of the users, but not the other. Since advertisers on the spot market have access to more tracking information about each user¹, the resulting bids may be quite different for these two users. Allocating impressions to guaranteed contracts first and selling the remainder on the spot market can therefore be highly suboptimal in terms of *revenue*, since two impressions that would fetch the same revenue from the guaranteed contract might fetch very different prices from the spot market².

On the other hand, simply buying the cheapest impressions on the spot market to satisfy guaranteed demand is not a good solution in terms of *fairness* to the guaranteed contracts, and leads to increasing short term revenue at the cost of long term satisfaction. As discussed above, impressions in online advertising have a *common value component* because advertisers generally have different information about a given user. This information (e.g. browsing history on an advertiser site) is typically relevant to *all* of the bidders, even though only *one* bidder may possess this information. In such settings, *price is a signal of value*—in a model of valuations incorporating both common and private values, the price converges to the true value of the item in the limit as the number of bidders goes to infinity ([8, 11], see also [7] for discussion). On average, therefore, the price on the spot market is a good indicator of the value of the impression, and delivering

¹ For example, a car dealership advertiser may observe that a particular user has been to his webpage several times in the previous week, and may be willing to bid more to show a car advertisement to induce a purchase.

² Consider the following toy example: suppose there are two opportunities, the first of which would fetch 10 cents in the spot market, whereas the second would fetch only ϵ ; both opportunities are equally suitable for the guaranteed contract which wants just one impression. Clearly, the first opportunity should be sold on the spot market, and the second should be allocated to the guaranteed contract.

cheapest impressions corresponds to delivering the lowest quality impressions to the guaranteed contract³.

A publisher with access to both sources of demand thus faces a trade-off between revenue and fairness when deciding which impressions to allocate to the guaranteed contract; this trade-off is further compounded by the fact that the publisher typically does not have access to all the information that determines the value of a particular impression. Indeed, publishers are often the least well informed participants about the value of running an ad in front of a user. For example, when a user visits a politics site, Amazon (as an advertiser) can see that the user recently searched Amazon for an ipod, and Target (as an advertiser) can see they searched target.com for coffee mugs, but the publisher only knows the user visited the politics site. Furthermore, the exact nature of this trade-off is unknown to the publisher in advance, since it depends on the spot market bids which are revealed only *after* the advertising opportunity is placed on the spot market.

The publisher as a bidder. To address the problem of unknown spot market demand (*i.e.*, the publisher would like to allocate the opportunity to a bidder on the spot market *if* the bid is “high enough”, else to a guaranteed contract), the publisher acts, in effect, as a *bidder* on behalf on the guaranteed contracts. That is, the publisher now plays two roles: that of a seller, by placing his opportunity on the spot market, and that of a bidding agent, bidding on behalf of his guaranteed contracts. If the publisher’s own bid turns out to be highest among all bids, the opportunity is won and is allocated to the guaranteed contract. Acting as a bidder allows the publisher to probe the spot market and decide whether it is more efficient to allocate the opportunity to an external bidder or to a guaranteed contract.

How should a publisher model the trade-off between fairness and revenue, and having decided on a trade-off, how should she place bids on the spot market? An ideal solution is (a) easy to implement, (b) allows for a trade-off between the quality of impressions delivered to the guaranteed contracts and short-term revenues, and (c) is robust to the exact tradeoff chosen. In this work we show precisely when such an ideal solution exists and how it can be implemented.

1.1 Our Contributions

In this paper, we provide an analytical framework to model the publisher’s problem of how to fulfill guaranteed advance contracts in a setting where there is an alternative spot market, and advertising opportunities have a common value component. We give a solution where the publisher bids on behalf of its guaranteed contracts in the spot market. The solution consists of two components: an *allocation*, specifying the fraction of impressions at each price allocated to a contract, and a *bidding strategy*, which specifies how to acquire this allocation by bidding in an auction.

³ While allocating the cheapest inventory to the guaranteed contracts is indeed revenue maximizing in the short term, in the long term the publisher runs the risk of losing the guaranteed advertisers by serving them the least valuable impressions.

The quality, or value, of an opportunity is measured by its price⁴. A perfectly representative allocation is one which consists of the same proportion of impressions at every price—*i.e.*, a mix of high-quality and low quality impressions. The trade-off between revenue and fairness is modeled using a budget, or *average target spend* constraint, for each advertiser’s allocation: the publisher’s choice of target spend reflects her trade-off between short-term revenue and quality of impressions for that advertiser (this must, of course, be large enough to ensure that the promised number of impressions satisfying the targeting constraints can be delivered.) Given a target spend⁵, a maximally representative allocation is one which minimizes the distance to the perfectly representative allocation, subject to the budget constraint. We first show how to solve for a maximally representative allocation, and then show how to implement such an allocation by purchasing opportunities in an auction, using randomized bidding strategies.

Organization. We start out with the single contract case, where the publisher has just one existing guaranteed contract, in Section 2; this case is enough to illustrate the idea of maximally representative allocations and implementation via randomized bidding strategies. We move on to the more realistic case of multiple contracts in Section 3; we first prove a result about which allocations can be implemented in an auction in a decentralized fashion, and derive the corresponding decentralized bidding strategies, and comment on solution of the optimal allocation. Full details, along with experimental validations of these strategies appear in [6].

Related Work The most relevant work is the literature on designing expressive auctions and clearing algorithms for online advertising [9, 2, 10]. This literature does not address our problem for the following reason. While it is true that guaranteed contracts have coarse targeting relative to what is possible on the spot market, most advertisers with guaranteed contracts choose not to use all the expressiveness offered to them. Furthermore, the expressiveness offered does not include attributes like relevant browsing history on an advertiser site, which could increase the value of an impression to an advertiser, simply because the publisher *does not have* this information about the advertising opportunity. Even with extremely expressive auctions, one might still want to adopt a mutual fund strategy to avoid the ‘insider trading’ problem. That is, if some bidders possess good information about convertibility, others will still want to randomize their bidding strategy since bidding a constant price means always losing on some good impressions. Thus, our problem cannot be addressed by the use of more expressive auctions as in [10] — the real problem is not lack of expressivity, but lack of information.

⁴ We emphasize that the assumption being made is *not* about price being a signal of value, but rather that impressions do have a common value component – given that impressions have a common value, price reflecting value follows from the theorem of Milgrom [8]. This assumption is commonly observed in practice.

⁵ We point out that we do not address the question of how to set target spends, or the related problem of how to price guaranteed contracts to begin with. *Given* a target spend, we propose a complete solution to the publisher’s problem.

Another area of research focuses on selecting the optimal set of guaranteed contracts. In this line of work, Feige et al. [5] study the computational problem of choosing the set of guaranteed contracts to maximize revenue. A similar problem is studied by in [3, 1]. We do not address the problem of how to select the set of guaranteed contracts, but rather take them as given and address the problem of how to fulfill these contracts in the presence of competing demand from a spot market.

2 Single contract

We first consider the simplest case: there is a single advertiser who has a guaranteed contract with the publisher for delivering d impressions. There are a total of $s \geq d$ advertising opportunities which satisfy the targeting requirements of the contract. The publisher can also sell these s opportunities via auction in a spot market to external bidders. The highest bid from the external bidders comes from a *distribution* F , with density f , which we refer to as the bid landscape. That is, for every unit of supply, the highest bid from all external bidders, which we refer to as the price, is drawn i.i.d from the distribution⁶ f . We assume that the supply s and the bid landscape f are known to the publisher⁷. Recall that the publisher wants to decide how to allocate its inventory between the guaranteed contract and the external bidders in the spot market. Due to penalties as well as possible long term costs associated with underdelivering on guaranteed contracts, we assume that the publisher wants to deliver all d impressions promised to the guaranteed contract.

An *allocation* $a(p)$ is defined as follows: $a(p)/s$ is the proportion of opportunities at price p purchased on behalf of the guaranteed contract (the price is the highest (external) bid for an opportunity.) That is, of the $sf(p)dp$ impressions available at price p , an allocation $a(p)$ buys a fraction $a(p)/s$ of these $sf(p)dp$ impressions, *i.e.*, $a(p)f(p)dp$ impressions. For example, a constant bid of p^* means that for $p \leq p^*$, $a(p) = 1$ with the advertiser always winning the auction, and for $p > p^*$, $a(p) = 0$ since the advertiser would never win.

Generally, we will describe our solution in terms of the allocation $a(p)/s$, which must integrate out to the total demand d : a solution where $a(p)/s$ is larger for higher prices corresponds to a solution where the guaranteed contract is allocated more high-quality impressions. As another example, $a(p)/s = d/s$ is a perfectly representative allocation, integrating out to a total of d impressions, and allocating the same fraction of impressions at every price point.

Not every allocation can be purchased by bidding in an auction, because of the inherent asymmetry in bidding— a bid b allows every price below b and rules out every price above; however, there is no way to rule out prices *below* a certain value. That is, we can choose to exclude high prices, but not low prices. Before describing our solution, we state what kinds of allocations $a(p)/s$ can be purchased by bidding in an auction.

⁶ Specifically, we do not consider adversarial bid sequences; we also do not model the effect of the publisher’s own bids on others’ bids.

⁷ Publishers usually have access to data necessary to form estimates of these quantities.

Proposition 1. *A right-continuous allocation $a(p)/s$ can be implemented (in expectation) by bidding in an auction if and only if $a(p_1) \geq a(p_2)$ for $p_1 \leq p_2$.*

Proof. Given a right-continuous non-increasing allocation $\frac{a(p)}{s}$ (that lies between 0 and 1), define $H(p) := 1 - \frac{a(p)}{s}$. Let $p^* := \inf \{p : a(p) < s\}$. Then, H is monotone non-decreasing and is right-continuous. Further, $H(p^*) = 0$ and $H(\infty) = 1$. Thus, H is a cumulative distribution function. We place bids drawn from H (the probability of a strictly positive bid being $a(0)/s$). Then the expected number of impressions won at price p is then exactly $a(p)/s$. Conversely, given that bids for the contract are drawn at random from a distribution H , the fraction of supply at price p that is won by the contract is simply $1 - H(p)$, the probability of its bid exceeding p . Since H is non-decreasing, the allocation (as a fraction of available supply at price p) must be non-increasing in p .

Note that the distribution H used to implement the allocation is a different object from the bid landscape f against which the requisite allocation must be acquired— in fact, it is completely independent of f , and is specified only by the allocation $a(p)/s$. That is, *given an allocation*, the bidding strategy that implements the allocation in an auction is independent of the bid landscape f from which the competing bid is drawn.

2.1 Maximally representative allocations

Ideally the advertiser with the guaranteed contract would like the same proportion of impressions at every price p , *i.e.*, $a(p)/s = d/s$ for all p . (We ignore the possibility that the advertiser would like a higher fraction of higher-priced impressions, since these cannot be implemented according to Proposition 1 above.) However, the publisher faces a trade-off between delivering high-quality impressions to the guaranteed contract and allocating them to bidders who value them highly on the spot market. We model this by introducing an average unit target spend t , which is the average price of impressions allocated to the contract. A smaller (bigger) t delivers more (less) cheap impressions. As we mentioned before, t is part of the input problem, and may depend, for instance, on the price paid by the advertiser for the contract.

Given a target spend, the *maximally representative* allocation is an allocation $a(p)/s$ that is ‘closest’ (according to some distance measure) to the ideal allocation d/s , while respecting the target spend constraint. That is, it is the solution to the following optimization problem:

$$\begin{aligned} \inf_{a(\cdot)} \int_p \mathbf{u} \left(\frac{a(p)}{s}, \frac{d}{s} \right) f(p) dp \\ \text{s.t. } \int_p a(p) f(p) dp = d \\ \int_p p a(p) f(p) dp \leq td \\ 0 \leq \frac{a(p)}{s} \leq 1. \end{aligned} \tag{1}$$

The objective, \mathbf{u} , is a measure of the deviation of the proposed fraction, $a(p)/s$, from the perfectly representative fraction, d/s . In what follows, we will

consider the L_2 measure

$$\mathbf{u} \left(\frac{a(p)}{s}, \frac{d}{s} \right) = \frac{s}{2} \left(\frac{a(p)}{s} - \frac{d}{s} \right)^2,$$

in [6] we also consider the Kullback-Leibler (KL) divergence. Why the choice of KL and L_2 for “closeness”? Only Bregman divergences lead to a selection that is consistent, continuous, local, and transitive [4]. Further, in R^n only least squares is scale- and translation- invariant, and for probability distributions only KL divergence is *statistical* [4].

The first constraint in (1) is simply that we must meet the target demand d , buying $a(p)/s$ of the $sf(p)dp$ opportunities of price p . The second constraint is the *target spend* constraint: the total spend (the spend on an impression of price p is p) must not exceed td , where t is a target spend parameter (averaged per unit). As we will shortly see, the value of t strongly affects the form of the solution. Finally, the last constraint simply says that the proportion of opportunities bought at price p , $a(p)/s$, must never go negative or exceed 1.

Optimality conditions: Introduce Lagrange multipliers λ_1 and λ_2 for the first and second constraints, and $\mu_1(p), \mu_2(p)$ for the two inequalities in the last constraint. The Lagrangian is

$$\begin{aligned} L = & \int \mathbf{u} \left(\frac{a(p)}{s}, \frac{d}{s} \right) f(p) dp + \lambda_1 \left(d - \int a(p) f(p) dp \right) + \lambda_2 \left(\int pa(p) f(p) dp - td \right) \\ & + \int \mu_1(p) (-a(p)) f(p) dp + \int \mu_2(p) (a(p) - s) f(p) dp. \end{aligned}$$

By the Euler-Lagrange conditions for optimality, the optimal solution must satisfy

$$\mathbf{u}' \left(\frac{a(p)}{s}, \frac{d}{s} \right) = \lambda_1 - \lambda_2 p + \mu_1(p) - \mu_2(p),$$

where the multipliers μ satisfy $\mu_1(p), \mu_2(p) \geq 0$, and each of these can be non-zero only if the corresponding constraint is tight.

These optimality conditions, together with Proposition 1, give us the following:

Proposition 2. *The maximally representative allocation for a single contract can be implemented by bidding in an auction for any convex distance measure \mathbf{u} .*

The proof follows from the fact that \mathbf{u}' is increasing for convex \mathbf{u} .

L_2 utility In this subsection, we derive the optimal allocation when \mathbf{u} , the distance measure, is the L_2 distance, and show how to implement the optimal allocation using a randomized bidding strategy. In this case the bidding strategy turns out to be very simple: toss a coin to decide whether or not to bid, and, if bidding, draw the bid value from a uniform distribution. The coin tossing

probability and the endpoints of the uniform distribution depend on the demand and target spend values.

First we give the following result about the continuity of the optimal allocation; this will be useful in deriving the values that parameterize the optimal allocation. See [6] for the proof.

Proposition 3. *The optimal allocation $a(p)$ is continuous in p .*

Note that we do not assume *a priori* that $a(\cdot)$ is continuous; the *optimal* allocation turns out to be continuous.

The optimality conditions, when \mathbf{u} is the L_2 distance, are:

$$\frac{a(p)}{s} - \frac{d}{s} = \lambda_1 - \lambda_2 p + \mu_1(p) - \mu_2(p),$$

where the nonnegative multipliers $\mu_1(p), \mu_2(p)$ can be non-zero only if the corresponding constraints are tight.

The solution to the optimization problem (1) then takes the following form: For $0 \leq p \leq p_{\min}$, $a(p)/s = 1$; for $p_{\min} \leq p \leq p_{\max}$, $a(p)/s$ is proportional to $C - p$, *i.e.*, $a(p)/s = z(C - p)$; and for $p \geq p_{\max}$, $a(p)/s = 0$.

To find the solution, we must find p_{\min}, p_{\max}, z , and C . Since $a(p)/s$ is continuous at p_{\max} , we must have $C = p_{\max}$. By continuity at p_{\min} , if $p_{\min} > 0$ then $z(C - p_{\min}) = 1$, so that $z = \frac{1}{p_{\max} - p_{\min}}$. Thus, the optimal allocation $a(p)$ is always parametrized by two quantities, and has one of the following two forms:

1. $a(p)/s = z(p_{\max} - p)$ for $p \leq p_{\max}$ (and 0 for $p \geq p_{\max}$).

When the solution is parametrized by z, p_{\max} , these values must satisfy

$$s \int_0^{p_{\max}} z(p_{\max} - p)f(p)dp = d \quad (2)$$

$$s \int_0^{p_{\max}} zp(p_{\max} - p)f(p)dp = td \quad (3)$$

Dividing (2) by (3) eliminates z to give an equation which is monotone in the variable p_{\max} , which can be solved, for instance, using binary search.

2. $a(p)/s = 1$ for $p \leq p_{\min}$, and $a(p)/s = \frac{p_{\max} - p}{p_{\max} - p_{\min}}$ for $p \leq p_{\max}$ (and 0 thenceforth).

When the solution is parametrized by p_{\min}, p_{\max} , these values must satisfy

$$sF(p_{\min}) + \int_{p_{\min}}^{p_{\max}} s \frac{(p_{\max} - p)}{p_{\max} - p_{\min}} f(p)dp = d \quad (4)$$

$$\int_0^{p_{\min}} spf(p)dp + \int_{p_{\min}}^{p_{\max}} sp \frac{(p_{\max} - p)}{p_{\max} - p_{\min}} f(p)dp = td. \quad (5)$$

Note that the optimal allocation can be represented more compactly as

$$\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}. \quad (6)$$

Effect of varying target spend: Varying the value of the target spend, t , while keeping the demand d fixed, leads to a tradeoff between representativeness and revenue from selling opportunities on the spot market, in the following way. The minimum possible target spend, while meeting the target demand (in expectation) is achieved by a solution where $p_{\min} = p_{\max}$ and $a(p)/s = 1$ for p less equal this value, and 0 for greater. The value of p_{\min} is chosen so that

$$\int_0^{p_{\min}} sf(p)dp = d \Rightarrow p_{\min} = F^{-1}\left(\frac{d}{s}\right).$$

This solution simply bids a flat value p_{\min} , and corresponds to giving the cheapest possible inventory to the advertiser, subject to meeting the demand constraint. This gives the minimum possible total spend for this value of demand, of

$$\underline{t}d = \int_0^{p_{\min}} spf(p)dp = sF(p_{\min})E[p|p \leq p_{\min}] = dE[p|p \leq p_{\min}]$$

(Note that the maximum possible total spend that is maximally representative while not overdelivering is $R = \int pf(p)dp = dE[p] = d\bar{p}$.)

As the value of t increases above \underline{t} , p_{\min} decreases and p_{\max} increases, until we reach $p_{\min} = 0$, at which point we move into the regime of the other optimal form, with $z = 1$. As t is increased further, z decreases from 1, and p_{\max} increases, until at the other extreme when the spend constraint is essentially removed, the solution is $\frac{a(p)}{s} = \frac{d}{s}$ for all p ; *i.e.*, a perfectly representative allocation across price. Thus the value of t provides a dial by which to move from the “cheapest” allocation to the perfectly representative allocation.

2.2 Randomized bidding strategies

The quantity $a(p)/s$ is an optimal *allocation*, *i.e.*, a recommendation to the publisher as to how much inventory to allocate to a guaranteed contract at every price p . However, recall that the publisher needs to *acquire* this inventory on behalf of the guaranteed contract by bidding in the spot market. The following theorem shows how to do this when \mathbf{u} is the L2 distance.

Theorem 1. *The optimal allocation for the L_2 distance measure can be implemented (in expectation) in an auction by the following random strategy: toss a coin to decide whether or not to bid, and if bidding, draw the bid from a uniform distribution.*

Proof. From (6) that the optimal allocation can be represented as

$$\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}.$$

By Proposition 1, an allocation $\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}$ can be implemented by bidding in an auction using the following randomized bidding strategy: with probability $\min\{zp_{\max}, 1\}$, place a bid drawn uniformly at random from the range $[\max\{p_{\max} - \frac{1}{z}, 0\}, p_{\max}]$.

3 Multiple contracts

We now study the more realistic case where the publisher needs to fulfill multiple guaranteed contracts with different advertisers. Specifically, suppose there are m advertisers, with demands d_j . As before, there are a total of $s \geq \sum d_j$ advertising opportunities available to the publisher.⁸ An allocation $a_j(p)/s$ is the proportion of opportunities purchased on behalf of contract j at price p . Of course, the sum of these allocations cannot exceed 1 for any p , which corresponds to acquiring all the supply at that price.

As in the single contract case, we are first interested in what allocations $a_j(p)$ are implementable by bidding in an auction. However, in addition to being implementable, we would like allocations that satisfy an additional practical requirement, explained below. Notice that the publisher, acting as a bidding agent, now needs to acquire opportunities to implement the allocations for *each of the guaranteed contracts*. When an opportunity comes along, therefore, the publisher needs to decide *which* of the contracts (if any) will receive that opportunity. There are two ways to do this: the publisher submits *one* bid on behalf of all the contracts; if this bid wins, the publisher then selects one amongst the contracts to receive the opportunity. Alternatively, the publisher can submit one bid for *each* contract; the winning bid then automatically decides which contract receives the opportunity. We refer to the former as a **centralized strategy** and the latter as a **decentralized strategy**.

There are situations where the publisher will need to choose the winning advertiser *prior to* seeing the price, that is, the highest bid from the spot market. For example, to reduce latency in placing an advertisement, the auction mechanism may require that the bids be accompanied by the advertisement (or its unique identifier). A decentralized strategy automatically fulfills this requirement, since the choice of winning contract does not depend upon knowing the price. In a centralized strategy, this requirement means that the relative fractions won at price p , $a_i(p)/a_j(p)$, are *independent of the price p* —when this happens, the choice of advertiser can be made (by choosing at random with probability proportional to a_j) without knowing the price.

As before, we will be interested in implementing optimal (*i.e.*, maximally representative) allocations. We will, therefore, concentrate on characterizing allocations which can be implemented via a decentralized strategy. In the full version of the paper [6] we show how to compute the optimal allocation in the presence of multiple contracts.

3.1 Decentralization

In this section, we examine what allocations can be implemented via a decentralized strategy. Note that it is not sufficient to simply use a distribution $H_j = 1 - \frac{a_j(p)}{a_j(0)}$ as in Proposition 1, since these contracts compete amongst each

⁸ In general, not all of these opportunities might be suitable for every contract; we do not consider this here for clarity of presentation. However the same ideas and methods can be applied in that case and the results are qualitatively similar.

other as well. Specifically, using the distribution $1 - \frac{a_j(p)}{a_j(0)}$ will lead to too few opportunities being purchased for contract j , since this distribution is designed to compete against f alone, rather than against f *as well as* the other contracts. We need to show how to choose distributions in such a way that lead to a fraction $a_j(p)/s$ of opportunities being purchased for contract j , for every $j = 1, \dots, m$.

First, we argue that a decentralized strategy with given distributions H_j will lead to allocations that are non-increasing, as in the single contract case. A decentralized implementation uses distributions H_j to bid for impressions. Then, contract j wins an impression at price p with probability

$$a_j(p) = \int_p^\infty \left(\prod_{k \neq j} H_k(x) \right) h_j(x) dx,$$

since to win, the bid for contract j must be larger than p *and* larger than the bids placed by each of the remaining $m - 1$ contracts. Since all the quantities in the integrand are nonnegative, a_j is non-increasing in p .

Now assume that a_j are differentiable a.e. and non-increasing. Let

$$\frac{A(p)}{s} := \sum_j \frac{a_j(p)}{s}$$

be the total fraction of opportunities at price p that the publisher needs to acquire. Clearly, a_j must satisfy $A(p) \leq s$, $\forall p$. Let $p^* := \inf\{p : A(p) < s\}$. Let

$$H_j(p) := \begin{cases} e^{\int_p^\infty a'_j(x)/(s-A(x))dx} & p > p^* \\ 0 & \text{else} \end{cases} \quad (7)$$

Then, $H_j(p) \geq 0$ and is continuous. Since $a'_j(p)$ is non-increasing, $H_j(p)$ is monotone non-decreasing. Further, $H(\infty) = 1$ and $H_j(p^*) = 0$. Thus, H_j is a distribution function. We can verify that bidding according to H_j will result in the desired allocations (see [6] for details).

Thus, we have constructed distribution functions $H_j(p)$ which implement the given non-increasing (and a.e. differentiable) allocations $a_j(p)$. If any a_j is increasing at any point, the set of campaigns cannot be decentralized. The following theorem generalizes Proposition 1:

Theorem 2. *A set of allocations $a_j(p)$ can be implemented in an auction via a decentralized strategy iff each $a_j(p)$ is non-increasing in p , and $\sum_j a_j(p)/s \leq 1$.*

Having determined which allocations can be implemented by bidding in an auction in a decentralized fashion, we turn to the question of finding suitable allocations to implement. As in the single contract case, we would like to implement allocations that are maximally representative, given the spend constraints.

As we show in [6], the optimal allocation is decentralizable in two cases:

1. The target spends are such that the solutions decouple. In this case the allocation for each contract is independent of the others; we solve for the parameters of each allocation as in Section 2.1.

2. The target spends are such that, for all j, k , $\frac{a_j(p)}{a_k(p)}$ is independent of p . In this case we need to solve for the common slope and p_{\min} , and the contract specific values p_{\max}^j , which together determine the allocation. This can be done using, for instance, Newton's method.

When the target spends are such that the allocation is not decentralizable, the vector of target spends can be increased to reach a decentralizable allocation. One way is to scale up the target spends uniformly until they are large enough to admit a separable solution; this has the advantage of preserving the relative ratios of target spends. The minimum multiplier which renders the allocation decentralizable can be found numerically, using for instance binary search.

4 Conclusion

Moving guaranteed contracts into an exchange environment presents a variety of challenges for a publisher. Randomized bidding is a useful compromise between minimizing the cost and maximizing the quality of guaranteed contracts. It is akin to the mutual fund strategy common in the capital asset pricing model. We provide a readily computable solution for synchronizing an arbitrary number of guaranteed campaigns in an exchange environment. Moreover, the solution we detail appears stable with real data.

References

- [1] Moshe Babaioff, Jason Hartline, and Robert Kleinberg. Selling banner ads: Online algorithms with buyback. In *4th Workshop on Ad Auctions*, 2008.
- [2] C. Boutilier, D. Parkes, T. Sandholm, and W. Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *National Conference on Artificial Intelligence (AAAI)*, 2008.
- [3] Florin Constantin, Jon Feldman, S Muthukrishnan, and Martin Pal. Online ad slotting with cancellations. In *4th Workshop on Ad Auctions*, 2008.
- [4] I. Csiszar. Why least squares and maximum entropy? an axiomatic approach to interference for linear inverse problems. *Annals of Statistics*, 19(4):2032–2066, 1991.
- [5] Uriel Feige, Nicole Immorlica, Vahab S. Mirrokni, and Hamid Nazerzadeh. A combinatorial allocation mechanism with penalties for banner advertising. In *Proceedings of ACM WWW*, pages 169–178, 2008.
- [6] Arpita Ghosh, Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii. Bidding for representative allocations for display advertising. *CoRR*, abs/0910-0880, 2009.
- [7] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, 1987.
- [8] Paul R. Milgrom. A convergence theorem for competitive bidding with differential information. *Econometrica*, 47(3):679–688, May 1979.
- [9] David Parkes and Tuomas Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *1st Workshop on Ad Auctions*, 2005.
- [10] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. 28(3):45–58, 2007.
- [11] Robert Wilson. A bidding model of perfect competition. *Rev. Econ. Stud.*, 44(3):511–518, 1977.