

Posted Prices Exchange for Display Advertising Contracts

Yagil Engel

Advanced Technology Lab, MSR
Herzliya, Israel
yagile@microsoft.com

Moshe Tennenholtz

Microsoft Research
Herzliya, Israel
moshet@microsoft.com

Abstract

We propose a new market design for display advertising contracts, based on posted prices. Our model and algorithmic framework address several major challenges: (i) the space of possible impression types is exponential in the number of attributes, which is typically large, therefore a complete price space cannot be maintained; (ii) advertisers are usually unable or reluctant to provide extensive demand (willingness-to-pay) functions, (iii) the levels of detail with which supply and demand are specified are often not identical.

Introduction

Display advertising is a complex environment. Many of the advertising deals are still made by so called reserved contracts, in which an advertiser spends its budget on some predefined types of inventory. Each type of inventory has its price quote, where this price may be open for some negotiation between the advertiser and a sales person. The negotiation may typically exchange some quantity of a certain type of good (e.g. users' segments) for some quantity of another type of good, exploiting flexibility and ability for substitution. On the other hand auction markets (or *ad exchanges*), for the so called non-reserved markets, where real-time bidding for particular ad impression is conducted, are flourishing. Ad exchanges have been recently studied (Muthukrishnan 2009; Emek et al. 2012), but it is well accepted that while automated and computationally tractable they lack the expressive power of the classical reserved markets. Indeed, the goal of having an efficient reserved-like market is a highly desired one. In service of this goal, researchers in AI and electronic commerce have advocated the use of expressive auction mechanisms for whole display advertising campaigns (Lahaie, Parkes, and Pennock 2008; Boutilier et al. 2008; Walsh et al. 2010). However, despite their appeal and potential, these proposals push us in somewhat different direction to the one currently used in practice: full demand function should be specified, and the mechanism of allocation does not appeal to posted prices. Moreover, the mechanisms suggested do not fit the recent trend of having many publishers competing for serving an ad opportunity, which becomes more and more popular in the context of *Ads in Apps*.

We suggest the use of a posted-prices mechanism, and the concept of market equilibrium, to deal with the above challenges. We observe that the situation where many advertisers have their own (expressive) campaigns, and many publishers (e.g. apps owners) provide ad impressions for sale (where an ad impression is associated with a user in a particular context), can in principle be modelled in a market equilibrium setting. However, since the number of possible impressions (i.e. possible goods in the system) is huge, and typically exponential in the number of attributes defining an impression, and since advertisers can not be expected to provide full demand functions, and because advertisers and suppliers may each use some (other) subset of the attributes to describe impressions, this theoretically appealing setting should be carefully re-visited.

Our main idea is to associate the goods in the system with the demand statements that have explicit and significant demand from advertisers along a purchase period. This may create overlapping goods, and inconsistency between the description of goods by advertisers to the ones used by the sellers. However, we show how one can deal with these issues in an effective manner, resulting in an efficient procedure for computing market equilibrium. As part of this, lack of information about the demand function is dealt with using a simple extraction procedure based on very minimal input by the advertisers, assuming *constant elasticity of substitution (CES)* of the demand function. Interestingly, this constraint is much less restrictive than the one required in previous work (Lahaie, Parkes, and Pennock 2008).

Other literature focus on pricing in the context of the one-to-one contracting which is in practice today. Radovanovic and Heavlin (2012) consider the problem from the publishers' side, and propose a scheme to find revenue maximizing prices based on empirical evaluation of demand. Another approach for computing prices for reserved contracts was introduced by Bharadwaj et al. (2010); they base pricing on an assessment of the value of an impression, which is based on the history of prices resulting from negotiation between the publisher and the advertiser. Our work, in contrast, is pursuing a global market with posted prices, as alternative to today's bilateral negotiation practice. The benefits of such a market are transparency, automation, low entry cost, and higher allocational efficiency; all this while maintaining the appeal of dynamic and on demand trading.

To establish the theoretical properties of our solution we introduce a novel market model, in which the utility functions of players are not defined directly on the market goods, but rather on items that relate to the market goods via a binary *satisfaction* relation. Beyond introducing the model, our work includes the following contributions: (1) showing how elicitation of CES demand can be performed efficiently from just two observations, (2) introducing a convenient and compact decision-tree based representation of future inventory, (3) proving that approximate equilibrium prices can be computed in polynomial time for our proposed market, despite the complexity of the model.

Preliminaries

Basic Notation

Display advertising impressions are characterized by a vector of *attributes*, denoted α . As examples, we use the attributes *State* (S) (e.g., MI , OH , CA), *Income* (I) (roughly discretized to *high* (H) or *low* (L)), *Gender* (G) (with values F or M), and whether the person is a *Cyclist* (C) or not.

The cartesian product of the attributes domains is denoted $\mathcal{A} = \times \theta_i$. Each element of \mathcal{A} is an instantiation of all the attributes in α . The number of attributes can be in the thousands; therefore, the set \mathcal{A} is typically extremely large.

Next, we define a set \mathcal{F} that includes all possible instantiations of any *subset* of α . Formally, let θ_i denote the domain of attribute α_i , and let $\hat{\theta}_i = \theta_i \cup \{\perp\}$, where \perp indicates an unspecified, or null value. We define $\mathcal{F} = \times \hat{\theta}_i$. Examples for statements in \mathcal{F} are $\langle I = H, G = F \rangle$ (high-income female) or $\langle C = \text{true}, S = MI \rangle$ (cyclist from Michigan).

\mathcal{F} is of course even larger than \mathcal{A} , but it allows us to reason about subsets of \mathcal{A} . An element $f \in \mathcal{F}$ can be considered a set that contains all the elements $a \in \mathcal{A}$ such that f and a agree on the attributes whose value in f is not \perp , that is a *satisfies* f . The satisfaction relation applies similarly between statements f and f' in \mathcal{F} , and it is defined formally as follows, where f_i indicates the i th attribute within f :

$$f \models f' \Leftrightarrow \forall i, f'_i \in \{f_i, \perp\}$$

For example, the statement $\langle I = H, G = F, S = MI \rangle$ satisfies the statement $\langle I = H, G = F \rangle$ because any high-income female from Michigan is also a high-income female.

A demand for statement f indicates that the advertiser will accept any ad opportunity that satisfies f , regardless of the values of attributes i for which $f_i = \perp$. For example, a demand statement $\langle C = \text{true}, S = MI \rangle$ can be matched with any cyclist from Michigan, regardless of her / his gender and income. Similarly, a supply statement f indicates an impression that is guaranteed to satisfy f , but can have arbitrary values for the attributes for which $f_i = \perp$.

A Posted Prices Market

In the display ad contracts market, a trade means that a publisher guarantees the delivery of a quantity of specific impression types for a specific future time window. To simplify this work, we assume all contracts within a time period are for a specific future time period (e.g., within this month, all contracts are for next month). Our approach is to base trades

on the published (posted) prices of impressions, determined according to their attributes.

Posted prices are updated periodically (e.g., each month), reflecting an equilibrium of the new supply and demand aggregated throughout the period. Each supplier provides a list of quantities in his expected inventory; the inventory list is independent of current prices. Demand, on the other hand, fluctuates more, and submitted at arbitrary times in response to published prices. As explained later, we observe an advertiser's demand request given current prices, use it to extrapolate his utility function, with which we can compute his demand at any other prices.

Our framework provides the algorithmic means for utility and inventory elicitation, and for computation of equilibrium prices. Beyond that, implementation can vary. For example, a choice should be made whether market is cleared *continuously* or *periodically*. Continuous clearing means a trade occurs whenever a demand request is submitted, according to current prices. Because demand and supply can change between periods, this can result in some unsatisfied demand or some unsatisfied supply. With periodic clearing, we recalculate and publish new prices given aggregated demand and supply at the end of each period, and only then respond to the advertisers with *take it or leave it* (TIOLI) offers, according to new prices. With accurate utility elicitation and equilibrium computation, it is likely that most of the advertisers will accept these TIOLI offers.

Price Space over Statements

Setting posted prices for each impression in \mathcal{A} is impractical due to the intractable size of \mathcal{A} . Instead, we propose to *assign prices to statements in \mathcal{F}* ; thus exploit their expressiveness. Of course, \mathcal{F} is very large as well; but we select a small subset of \mathcal{F} , denoted \mathcal{G} . A reasonable method to determine \mathcal{G} is to let advertisers express their demand naturally using statements in \mathcal{F} , and let \mathcal{G} include statements for which demand is significant. Importantly, elements in \mathcal{G} may not be mutually exclusive. For example, \mathcal{G} can include $\langle S = CA, C = \text{true} \rangle$, $\langle S = CA, C = \text{true}, G = M \rangle$, and $\langle S = CA, C = \text{true}, I = H \rangle$; demand for $\langle S = CA, C = \text{true} \rangle$ is satisfied by all three. This allows in some cases exponential reduction in the number of goods—all the combinations that do not have high demand can be left out.¹

On the other hand, a natural question is whether such price space is well defined. We address this issue from a technical standpoint in the next section; semantically, we note that statements are interpreted as *guarantees* for those attributes that appear in the statement, with no specification for the rest. An advertiser that purchases $\langle S = CA, C = \text{true} \rangle$, does not care about I , or is not willing to pay a premium for a $\langle I = H \rangle$ *guarantee*. An advertiser that purchases $\langle S = CA, C = \text{true}, I = H \rangle$ is willing to pay that premium. The choice is hence between different levels of guarantees on various attribute values. An advertiser whose utility indicates that he is willing to pay a certain price difference for the guarantee to get *high income cyclists*, as opposed to

¹Walsh et al. (2010) address a somewhat related problem: finding *mutually exclusive channels* that optimize allocation given bids.

any *cyclists*, will choose accordingly even if he has some belief on the portion of the *cyclists* (with unknown income) impressions he buys that turn out to be high income.

On the publisher's side, one issue is that his inventory may be specified in different terms than the market goods. In addition, the publisher may have choice between providing various levels of *guarantees*. We model this as the economic setting of *production*. A publisher that wishes to sell impressions that satisfy more than one statement in \mathcal{G} can choose between "producing guarantees" to either; naturally his choice would be such that maximizes his profit.

In the next section we formalize these semantics using a general market model, where the utility of players is defined in terms that are different from each other's and from the market goods. Afterwards, we describe our procedure to elicit consumer's CES utility function by observing current demand. In the following section we lay out existence and computational results on equilibrium in our market model. Finally, we return to our specific solution of display ad market, and propose a representation of supply inventory that facilitates the supply side of equilibrium computation.

Market Model

Consider a set of items \mathcal{F} , and a partial order \models over \mathcal{F} . Our market is a tuple $M = \langle \mathcal{G}, \{a_i\}, \{p_j\}, \{\Delta_i\}, \{u_i\}, \{\mu_i\}, \{\mathcal{I}_j\} \rangle$. $\mathcal{G} = \{f^1, \dots, f^\omega\} \subset \mathcal{F}$ is a set of goods, where $\omega = |\mathcal{G}|$. We use $\mathcal{Q} = (\mathbb{Z}^+)^{\omega}$ to denote the space of quantity vectors over \mathcal{G} . Under the interpretation of M as an ad market, a vector $Q \in \mathcal{Q}$ is an *allocation* of a set of impressions; in an allocation Q each entry $k = 1, \dots, \omega$ represents a guarantee for Q^k impressions (superscript k represents k th entry) that satisfy f^k , and therefore a total of $\sum_{k=1}^{\omega} Q^k$ impressions.

$\{a_i\} = \{a_1, \dots, a_n\}$ are the buyers (advertisers). Each buyer a_i has a *set of demand items* $\Delta_i \subset \mathcal{F}$, which we index f^1, \dots, f^{ω_i} . $\mathcal{Q}^{a_i} = (\mathbb{Z}^+)^{\omega_i}$ denotes the space of quantities over Δ_i , and $Q \in \mathcal{Q}^{a_i}$ is a quantity vector. $u_i : \mathcal{Q}^{a_i} \rightarrow \mathbb{R}^+$ is the *utility function* of a_i , that is $u_i(Q)$ designates the utility a_i obtains by receiving the number of items (impression guarantees) indicated by the quantity vector. In addition, buyer a_i has *budget* μ_i .

The relation \models connects the demanded goods in Δ_i and the market goods \mathcal{G} . Specifically, for any $f^k \in \Delta_i$ we denote,

$$\delta_{fk} = \{f' \in \mathcal{G} \mid f' \models f^k\},$$

the set of market goods that satisfy f^k . Buyer a_i does not care which of the goods in δ_{fk} satisfy his demand for f^k .

$\{p_j\} = \{p_1, \dots, p_m\}$ are the *firms* (publishers). \mathcal{I}_j denoting their inventory representation, is called an *inventory list*; it is a set of triplets $\langle f^k, Q^k, c^k \rangle$; wlog we restrict the order such that if $f^{k'} \models f^{k''}$ then $k' \leq k''$. Q^k refers to the number of items in the inventory of publisher p_j that satisfy f^k , and do not satisfy f^1, \dots, f^{k-1} . c^k is the cost to p_j associated with one unit of f^k (in display advertising, this can reflect the cost of disturbance to the user).

With slight abuse of notation, we use $\omega_j = |\mathcal{I}_j|$. Further, we map the inventory list back to market goods, analogous to the buy side; γ_{fk} refers to the set of market goods *that are satisfied* by f^k , so $\gamma_{fk} = \{f' \in \mathcal{G} \mid f^k \models f'\}$.

The goal of a market is to assign market goods to players. We use $X_i \in \mathcal{Q}$ to denote a vector of good quantities assigned (sold) to a_i , and X_i^k is the entry representing $f^k \in \mathcal{G}$. We use $Y_j \in \mathcal{Q}$ to denote a vector of quantities assigned, or sold by p_j , and Y_j^k analogous to X_i^k .

Definition 1. A *global assignment* over \mathcal{G} is a pair of matrices X of size $n \times \omega$ and Y of size $m \times \omega$ such that each row i in X is an assignment to a_i , and each row j in Y is an assignment to p_j . X, Y is *feasible* if for any $k = 1, \dots, \omega$, $\sum_{i=1}^n X_i^k = \sum_{j=1}^m Y_j^k$.

We use π to denote a price vector over \mathcal{G} , such that π^k indicates the price of f^k which is the k th item of \mathcal{G} .

Demand, Supply, and Equilibrium

The definitions in this section follow consequentially from the market model presented above. First we denote,

$$\bar{\pi}^k = \min_{f^{k'} \in \delta_{fk}} \{\pi^{k'}\}. \quad (1)$$

Definition 2. The *utility maximization problem* of a buyer a_i with utility u_i and budget μ_i , given prices π , returns a demand vector $\hat{D}_i(\pi)$ and the set $\hat{\Delta}_i(\pi, k) \subseteq \mathcal{G}$ (for each $k = 1, \dots, \omega_i$), where

$$\hat{D}_i(\pi) = \arg \max_{Q \in \mathcal{Q}^{a_i}} \{u_i(Q^1, \dots, Q^{\omega_i}) \mid \sum_{k=1}^{\omega_i} Q^k \bar{\pi}^k \leq \mu_i\},$$

$$\text{and } \hat{\Delta}_i(\pi, k) = \{f^{k'} \in \delta_{fk} \mid \pi^{k'} = \bar{\pi}^k\},$$

(which is the subset of δ_{fk} whose price is minimal).

$\hat{D}_i(\pi)$ is a quantity vector, which is standard output of a utility maximization problem. In addition, we require as output a mapping $\hat{\Delta}_i(\pi, k)$ back to the market goods, so that any $\hat{D}_i(\pi)^k$ (k th entry of $\hat{D}_i(\pi)$) can be divided in any way among the market goods in $\hat{\Delta}_i(\pi, k)$. It implies a space of allocations over \mathcal{G} between which a_i is indifferent; this space is called the *demand correspondence*. In order to formally translate $\hat{\Delta}_i(\pi, k)$ into demand over \mathcal{G} , we introduce the variables $X_i^{k,k'}$, indicating a quantity of $f^{k'} \in \mathcal{G}$ that satisfies $f^k \in \Delta_i$. The double superscript is required because one market good $f^{k'}$ may satisfy multiple demand goods f^k . We use the vector notation $\vec{X}_i^k = (X_i^{k,1}, \dots, X_i^{k,|\hat{\Delta}_i(\pi,k)|})$.

Definition 3. The *demand correspondence* of a_i is

$$D_i(\pi) = \{Q \in \mathcal{Q} \mid \forall k, \exists \vec{X}_i^k, \text{ s.t. } \forall k', \sum_{k=1}^{\omega_i} X_i^{k,k'} = Q^{k'},$$

$$\text{and } \sum_{f^{k'} \in \hat{\Delta}_i(\pi,k)} X_i^{k,k'} = \hat{D}_i(\pi)^k\}$$

For the publisher, profit maximization is obtained by assigning each of its items in the inventory to the highest priced market good it satisfies, unless that price is less than his cost. First, we set $\pi^k = \max_{f^{k'} \in \gamma_{fk}} \{\pi^{k'}\}$; then define

Definition 4. The *profit maximization problem* of a publisher p_j with inventory list \mathcal{I}_j , given prices π , returns a vector $\hat{S}_j(\pi)$ and sets $\hat{\mathcal{I}}_j(\pi, k)$ such that for each $\langle f^k, Q^k, c^k \rangle \in \mathcal{I}_j$,

$$\hat{S}_j(\pi)^k = \begin{cases} Q^k, & \bar{\pi}^k \geq c^k \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and } \hat{X}_i(\pi, k) = \{f^{k'} | \pi^{k'} = \bar{\pi}^k\}.$$

Here too, the output is a space of allocations, each of which maximizes the publisher's profit at prices π .

Definition 5. The *supply correspondence* of p_j is,

$$S_j(\pi) = \{Q \in \mathcal{Q} | \forall k, \exists \bar{Y}_j^k, \text{ s.t. } \forall k', \sum_{k=1}^{\omega_j} Y_j^{k,k'} = Q^{k'} \\ \text{and } \sum_{f^{k'} \in \hat{X}_j(\pi, k)} Y_j^{k,k'} = S_j(\pi)^k\}$$

The demand and supply correspondences are never computed explicitly, but they allow us to define equilibrium.

Definition 6. A price vector π is a *Market Equilibrium* if there exists a feasible global assignment X, Y such that $X_i \in D_i(\pi)$ for $i = 1, \dots, n$ and $Y_j \in S_j(\pi)$ for $j = 1, \dots, m$.

Definition 7. A price vector π is an ϵ -*approximate equilibrium* (for $\epsilon > 0$) if there exists a market equilibrium price vector $\hat{\pi}$ such that $\forall k, (1 - \epsilon)\hat{\pi}^k \leq \pi^k \leq (1 + \epsilon)\hat{\pi}^k$.²

Demand Elicitation

Advertisers Utility Model

We introduce several assumptions on advertisers behavior.

Inelastic budget A buyer operates under specific budget (μ_i) per time period, and he always exhausts this budget. The buyers' objective under this model is to maximize their utility under the budget constraint. This is normally assumed in the world of online advertising.

Gross Substitutes u_i exhibits gross-substitutes if a price increase of one item does not cause the buyer to demand more of another item whose price did not change. This prevents buyers from expressing complementarities between goods, which is reasonable in the advertising market.

Constant elasticity of substitution (CES). Mathematically, this means that for any buyer there exists a vector β ($\sum \beta^k = 1$) and scalar $\rho \in \{(-\infty, 0), (0, 1]\}$ such that

$$u_i(Q) = \left(\sum_{k=1}^{\omega_i} \beta^k (Q^k)^\rho \right)^{\frac{1}{\rho}} \quad (2)$$

where β^k are *weights* that the advertiser assigns to impressions. ρ is an elasticity parameter: $\sigma = \frac{1}{1-\rho}$ is called *the elasticity of substitution* (for $\rho = 1$, we have elasticity $\sigma = \infty$). CES preferences are gross-substitutes iff $\rho > 0$.

The elasticity of substitution is a measurement of how price changes affect transferring of budget between goods. The CES assumption means that this measurement is constant over all the pairs of goods at all prices. This is of course a serious limitation, but we argue it is reasonable for the world of display advertising. In particular, previous literature often takes the assumption of linear utility, which is the case in CES when $\rho = 1$; this implies that the advertiser is completely oblivious about the mixture of impressions as

²This notion of approximation can be translated (by proof of Thm. 2) to the one used by Codenotti et al. (2005) in Lemma 13.

long as he maximizes the sum of the independent value per impression. On the other extreme, another popular utility form is *Cobb-Douglas* (obtained when $\rho \rightarrow 0$); it means that goods cannot be substituted at all – if the price of one rises, the consumer just consumes less of that good (elasticity is 1). Each of the two examples is too limiting; advertisers do care about balancing their exposure to different markets, but on the other hand perform substitutions when price differences make it worthwhile. The CES form, in contrast, covers the whole spectrum of elasticity between those two extremes.

Finally, as most natural with posted prices, we assume buyers (and suppliers) are **price takers**, that is they purchase (or sell) the bundle of goods that is optimal to them, taking prices as given. When the number of traders is large (as we expect here), traders cannot gain much by behaving otherwise (Al-Najjar and Smorodinsky 2007).

Elicitation of Utility

Interestingly, CES utility elicitation over \mathcal{G} can be effectively performed by simply observing two purchase decisions of the buyer. Note that we can infer nothing, and need to infer nothing, regarding preferences that do not affect demand over \mathcal{G} ; the meaning of u_i as elicited below is hence that it results in the same demand correspondence as the “true” u_i .

Theorem 1. Assume: (i) a_i is maximizing utility consistently according to CES utility function u_i and budget μ_i , (ii) if $k \neq k'$ and $f^k \models f^{k'}$, then $\pi^k > \pi^{k'}$.³ Then u_i can be determined by observing bundles Q and Q_0 selected at two price points π and π_0 (respectively), which differ on at least one index k for which $Q^k > 0$.

Proof. The observed choices must first be translated to prices and quantities over Δ_i . We obtain $\bar{\pi}$ and $\bar{\pi}_0$ according to (1). We obtain \bar{Q} over Δ_i by defining $\bar{Q}^k = \sum_{k' \in \nu_k} Q^{k'}$, where ν_k is obtained as follows: iterate over k' such that $Q^{k'} > 0$. By (i), there exists $f^{k''} \in \Delta_i$ such that $f^{k'} \models f^{k''}$. If there exists an additional $f^{k*} \in \Delta_i$ such that $f^{k'} \models f^{k*}$, then if f^{k*} or $f^{k''}$ are also in \mathcal{G} , WLOG f^{k*} , we denote its index in \mathcal{G} as k^* as well, and by (ii) $\pi^{k'} > \pi^{k*}$, hence $f^{k'}$ is not purchased to satisfy demand for f^{k*} . If f^{k*} and $f^{k''}$ are not in \mathcal{G} , we do not need to distinguish whether $f^{k'}$ is purchased to satisfy demand for $f^{k''}$ or f^{k*} . In both cases, we add k' to $\nu_{k''}$. We obtain \bar{Q}_0 from Q_0 similarly.

CES utility u_i takes the form of Eq. (2). Taking first-order condition, it is easy to derive for each pair β^k, β^j in (2):

$$\frac{\beta^k (Q^k)^{\rho-1}}{\beta^j (Q^j)^{\rho-1}} = \frac{\pi^k}{\pi^j} \quad (3)$$

We instantiate (3) with \bar{Q} and $\bar{\pi}$ for each pair $k, k+1$ for $k = 1, \dots, \omega_i - 1$. By the assumptions, there exists k' such that $\pi^{k'} \neq \pi_0^{k'}$ and there exists k'' such that $k' \in \nu_{k''}$. We obtain another equation of the form (3) for $\bar{\pi}_0^{k''}, \bar{\pi}_0^{k''+1}, \bar{Q}_0^{k''}, \bar{Q}_0^{k''+1}$. We now have two equations (3) for the pair

³An equilibrium guarantees that $f^k \models f^{k'} \Rightarrow \pi^k \geq \pi^{k'}$; usually the inequality will be strict otherwise suppliers will prefer not to guarantee f^k (see also discussion on *Supply Tree* later).

$k'', k'' + 1$; we can divide one by the other to get rid of $\beta^{k''}$ and $\beta^{k''+1}$, and solve for ρ . Together with $\sum_k \beta^k = 1$ we are left with solvable ω_i linear equations with ω_i unknowns. \square

In the practical setting of the display ad market, advertisers may not follow an exact utility function. To obtain more robustness to inconsistent behavior of advertisers, we propose a procedure (ELICITUTIL) that approximates ρ given multiple observations. First, we discretize the domain $[0, 1]$ of ρ to a reasonable size; next, we solve a linear program to obtain the most accurate estimation of β^i for fixed ρ , and then pick the ρ such that the maximal error in Eq. (3) is the smallest. Assuming the elasticity of an advertiser does not change often, and even if his weight parameters do change, we potentially improve the estimate of ρ after each new observation, and then extrapolate u_i from the most recent observation, as done in the last step of the proof of Theorem 1.

procedure ELICITUTIL

input: Price, demand pairs: $(\pi_1, Q_1) \dots, (\pi_r, Q_r)$, over Δ_i items Δ_i , discrete domain for ρ

output: Utility u_i

foreach value ρ_j of ρ

solve LP:

$\min \epsilon_j$, s.t.

$$\frac{\pi_i^k}{\pi_i^{k+1}} - \epsilon_j \leq \frac{\beta^k (Q_i^k)^{\rho-1}}{\beta^{k+1} (Q_i^{k+1})^{\rho-1}} \leq \frac{\pi_i^k}{\pi_i^{k+1}} + \epsilon_j$$

Let $j = \arg \min_j \{\epsilon_j\}$

Obtain u_i given $\rho = \rho_j$ and (π_r, Q_r) , using (3)

Equilibrium: Existence and Computation

Theorem 2. *Market equilibrium prices always exist for M .*

Sketch. We show that our model is a special case of the full-fledged production model of Arrow and Debreu (1954). In particular, the CES function u_i (which is convex and non-satiable) is translated to a convex and non-satiable utility function over \mathcal{G} ; and \mathcal{I}_j (which implies a production set with no inputs and a linear profit function) is translated to a closed and bounded production set over \mathcal{G} . Next, we transform M to a market with endowments and shares, rather than monetary budgets. We can assume continuity because the quantity of each good is in the millions. \square

The main result of this section is that approximate equilibrium prices can be computed in polynomial time. We begin by showing that the players local optimizations are tractable.

Lemma 3. *The utility maximization problem of a_i with utility u_i can be solved in polynomial time, and has a unique solution demand vector $\hat{D}_i(\pi)$.*

Proof. Profit maximization needs to solve

$$\hat{D}_i(\pi) = \arg \max_{Q \in \mathcal{Q}^i} \{u_i(Q^1, \dots, Q^{\omega_i}) \mid \sum_{k=1}^{\omega_i} Q^k \bar{\pi}^k \leq \mu_i\}$$

We instantiate (3) again; now the β^k , the π^k , and ρ are known. The unknown variables are the Q^k . To solve it, we start with assigning an arbitrary Q^1 , instantiate the rest of Q^k by applying (3) sequentially, and then normalize Q^1, \dots, Q^{ω_i} to just match the budget constraint. This is

also showing that the solution is unique. From here adding $\hat{\Delta}_i(\pi, k) = \{f^{k'} \mid \pi^{k'} = \bar{\pi}^k\}$ is trivial. \square

The next Lemma is proved by Procedure PROFITMAX.

Lemma 4. *The profit maximization problem of p_j with inventory \mathcal{I}_j can be solved in polynomial time.*

Theorem 5. *ϵ -approximate equilibrium prices for M can be computed in time polynomial in the input and in $\log \frac{1}{\epsilon}$.*

Codenotti et al. (2005) prove a similar result for the general Arrow-Debreu production model. To use their result we need: (1) the transformation of M to Arrow-Debreu in the proof of Theorem 2, (2) an *excess demand oracle*, which is developed next. (3) Excess demand must fulfill the gross-substitutes condition. First, we define the following LP:

Linear Program EXCESSDEMAND

$\min \epsilon$

$$\forall k, -\epsilon \leq \sum_i X_i^k - \sum_j Y_j^k \leq \epsilon$$

$$\forall i, k, \sum_{f^{k'} \in \hat{\Delta}_i(\pi, k)} X_i^{k, k'} = \hat{D}_i(\pi)^k, \sum_{k=1}^{\omega} X_i^{k, k'} = X_i^{k'}$$

$$\forall j, k, \sum_{f^{k'} \in \hat{\mathcal{I}}_j(\pi, k)} Y_j^{k, k'} = \hat{S}_j(\pi)^k, \sum_{k=1}^{\omega} Y_j^{k, k'} = Y_j^{k'}$$

The second and third constraints are obtained directly from Definitions 3 and 5, and ensure that X and Y reflect demand and supply (respectively); therefore, the vector $Z(\pi) = \sum_i X_i - \sum_j Y_j$ is the *excess demand vector*. $Z(\pi)$ is also minimal excess demand in the sense that its largest coordinate is minimized by EXCESSDEMAND. This serves as *excess demand oracle* for the market M .

Lemma 6. *If all u_i expresses gross-substitute preferences, then $Z(\pi)$ fulfills the gross-substitutes condition.*

Sketch. We exploits the particular structure of utility and production. If the price of $f^k \in \mathcal{G}$ increases, the demand on another good $f^{k'} \in \mathcal{G}$ does not decrease, due to gross substitutability of the utility, which translates to the demand correspondence. Furthermore, the *supply* of $f^{k'}$ does not *increase*, because the price increase of f^k could only shift supply towards f^k and therefore only away from $f^{k'}$. \square

Proof of Theorem 5. By Lemma 3 and 4 we compute the input to EXCESSDEMAND in polynomial time, and hence we can compute excess demand polynomially. By Lemma 6 the excess demand exhibits gross-substitutability. With that, we can use an algorithm given by Codenotti et al. (2005) for the computation of production market equilibrium. We note that equilibrium allocations may be fractional but this is not a concern for the same argument as in Theorem 2. \square

procedure PROFITMAX

input: inventory list \mathcal{I}_j , prices π

output: sets of market goods $\hat{\mathcal{I}}_j(\pi, k)$, supply $\hat{S}_j(\pi)$

Initialize $S_j(\pi) = 0$

Sort \mathcal{G} , such that $f^k \in \mathcal{G}$ are in descending order of π^k

foreach $k = 1, \dots, \omega_j$

foreach $f^{k'} \in \mathcal{G}$ by sorted order

if $f^k \models f^{k'}$ and $\pi^{k'} \geq c^k$ /* we found $\bar{\pi}^k = \pi^{k'}$ */

then $\hat{S}_j(\pi)^k = Q^k$

$\hat{\mathcal{I}}_j(\pi, k) = \{f^{k''} \in \mathcal{G} \mid f^k \models f^{k''}, \pi^{k''} = \pi^{k'}\}$, **next** k

algorithm TATONNEMENT

input: Utilities u_i Inventories \mathcal{I}_j , and δ

output: δ -Approximate equilibrium prices π

Initialize $\pi_0 = \vec{0}$, $t = 0$

loop

call excess demand oracle, get $Z(\pi)$, ϵ

if $\epsilon \leq 0$, return π_t .

else let $k = \max_k Z(\pi)^k$

 let $\pi_{t+1}^k = \pi_t^k + \delta$, $t = t + 1$

Tatonnement-Like Algorithm

Theorem 5 relies on an ellipsoid algorithm, for which practical performance is in some cases weak. As an alternative, *tatonnement*-like algorithms for the *exchange* model are known to converge under GS demand (Cheng and Wellman 1998; Codenotti, McCune, and Varadarajan 2005). In fact, using Lemma 6 we can show that a discrete, price-increasing tatonnement also converges under our model. The algorithm

TATONNEMENT is polynomial in $\frac{\max_k \hat{\pi}^k}{\delta}$ (where $\hat{\pi}$ is equilibrium price), and hence not strongly polynomial, but very simple and likely to perform well in practice because price range for display ads is small.

Theorem 7. *Algorithm TATONNEMENT terminates at price vector π , which is ϵ -approximate equilibrium.*⁴

Supply Representation

Unlike utility elicitation, the “elicitation” problem on the supply side is more specific to the display ad domain: how can a publisher translate log data into a compact and computationally convenient representation. One interesting and known approach is with *Bayesian Networks*, that provide a succinct and accurate representation of a probability distribution over the future inventory (Chickering et al. 2010). The downside is that there is no obvious solution to the problem of *inventory reduction*: how to update the Bayes-Net probabilities after some quantity of specific impression types have been deducted from the inventory. This is required for implementing profit maximization; in order to maximize the sum of profits, quantities of inventory impressions should be assigned to market goods in descending order of prices.

We propose an alternative tree representation for supply inventory, named *supply tree*.⁵ At each node of a supply tree we branch over one attribute α_i , according to the domain $\hat{\theta}_i$ (that is, all the values of α_i , along with the null, or *unknown* value \perp , which is always the rightmost branch). At each leaf we specify a pair of values: the number of impressions that satisfy the statement implied by the path from the root to that leaf, and the cost associated with it. The supply tree can serve as a graphical representation of an inventory list \mathcal{I}_j , as it can be obtained from the tree by performing leftmost-first DFS search: At each leaf k we reach we define f^k as the conjunction of the specified attribute values on the path from the root to node k (ignoring the unknowns) and add the

triplet $\langle f^k, Q^k, c^k \rangle$ to \mathcal{I}_j , where (Q^k, c^k) is the pair specified at the leaf k .




Figure 1: A supply tree. At the leaves: quantities (by thousands) and costs (per thousand).

For example, the set we obtain from the tree in Figure 1 (in the form: $\langle (f^k); Q^k; c^k \rangle$) is: $\mathcal{I}_j = \{ \langle (S = MI, I = H); 100; 5 \rangle, \langle (S = MI, I = L); 150; 2 \rangle, \langle (S = MI); 200; 2 \rangle, \langle (S = OH); 300; 2 \rangle, \langle (G = F, C = t); 30; 3 \rangle, \langle (G = F); 70; 2 \rangle, \langle (G = M); 150; 1 \rangle \}$. The semantics of the branch \perp is that the value of the attribute at that node is indeed *unknown*, in contrast to the other branches. Therefore, the quantities specified in the leaves are disjoint.

The supply tree (i.e. inventory list) serves in PROFITMAX to obtain the supply set of impression types (statements) that maximize the profit to the supplier under a given set of prices; the tree should therefore differentiate (branch) over attributes only as long as this could affect profit. Fortunately, impression (statement) prices can be predicted from our current market prices, or from other (e.g., spot) markets. This brings us to a useful observation: the supply tree can be interpreted as a *decision tree*, in which the target for classification is *predicted price*. More accurately, this as a regression problem because this target class is continuous. Under this interpretation, the supply tree corresponds to a hierarchical clustering of the inventory according to the predicted profit of impression types.

Specifically, we obtain supply tree from log data as follows. The log data of a publisher is a list of past page visits, characterized by values to some of the attributes in α , that are known for that visit. Each record thus corresponds to a statement f . We label each record f with its predicted price minus its cost, and perform *regression-tree clustering* to obtain a supply tree. The information gain criterion, according to which we select an attribute to branch on at each node, is to minimize the sum of squared errors between the predicted prices of impressions in a cluster and their average. Finally, we specify the data of each leaf: the quantity is the size of the cluster, and cost as determined by the publisher.

Conclusion

We propose a market for display advertising contracts based on posted prices. The space of possible goods in this market is huge; we propose a compact price space over a selection of demand statements. To that end we design and analyze a market model in which players’ utilities are defined on different sets of items, that relate to the market goods via a relation of satisfaction. We propose algorithmic means to easily elicit advertisers’ utility functions, generate supply inventory list from data, and compute equilibrium prices.

⁴The algorithm requires slight obvious changes to PROFITMAX and EXCESSDEMAND. The proof is omitted for lack of space.

⁵Supply trees superficially resemble *bid trees* (Lahaie, Parkes, and Pennock 2008), but its role and behavior are different.

References

- Al-Najjar, N. I., and Smorodinsky, R. 2007. The efficiency of competitive mechanisms under private information. *Journal of Economic Theory* 137(1):383–403.
- Arrow, K., and Debreu, G. 1954. Existence of an equilibrium for a competitive economy. *Econometrica* 22(3):265–290.
- Bharadwaj, V. e. a. 2010. Pricing guaranteed contracts in online display advertising. In *ACM international conference on Information and knowledge management (CIKM)*.
- Boutilier, C.; Parkes, D. C.; Sandholm, T.; and Walsh, W. E. 2008. Expressive banner ad auctions and model-based on-line optimization for clearing. In *AAAI Conference on Artificial Intelligence*.
- Cheng, J. Q., and Wellman, M. P. 1998. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics* 12(1):1–24.
- Chickering, D. M.; Meek, C. A.; Charles, D. X.; and Tillman, R. E. 2010. Inventory management, US patent application 12/761,961.
- Codenotti, B.; McCune, B.; and Varadarajan, K. 2005. Market equilibrium via the excess demand function. In *ACM symposium on Theory of computing*.
- Emek, Y.; Feldman, M.; Gamzu, I.; Leme, R. P.; and Tennenholtz, M. 2012. Signaling schemes for revenue maximization. In *ACM Conference on Electronic Commerce*.
- Lahaie, S.; Parkes, D. C.; and Pennock, D. M. 2008. An expressive auction design for online display advertising. In *AAAI Conference on Artificial Intelligence*.
- Muthukrishnan, S. 2009. Ad exchanges: Research issues. In *Workshop on Internet & Network Economics (WINE)*.
- Radovanovic, A., and Heavlin, W. D. 2012. Risk-aware revenue maximization in display advertising. In *International Conference on World Wide Web*.
- Walsh, W. E.; Boutilier, C.; Sandholm, T.; Shields, R.; Nemhauser, G. L.; and Parkes, D. C. 2010. Automated channel abstraction for advertising auctions. In *AAAI Conference on Artificial Intelligence*.