On Approximate Welfare and Revenue-Maximizing Equilibria for Size-Interchangeable Bidders

Enrique Areyan Viqueira, Amy Greenwald, and Victor Naroditskiy
Computer Science Department, Brown University, Providence, RI, United States

BROWN

(1) Motivation

We study a type of *combinatorial resource allocation* problem motivated by the **Trading Agent Competition, Ad Exchange** game [Schain and Mansour, 2013]. This game models **online ad exchanges** in which agents face the challenge of bidding for **display-ad impressions** needed to fulfill **advertisement contracts**, after which they earn the amount the **advertiser budgeted**.

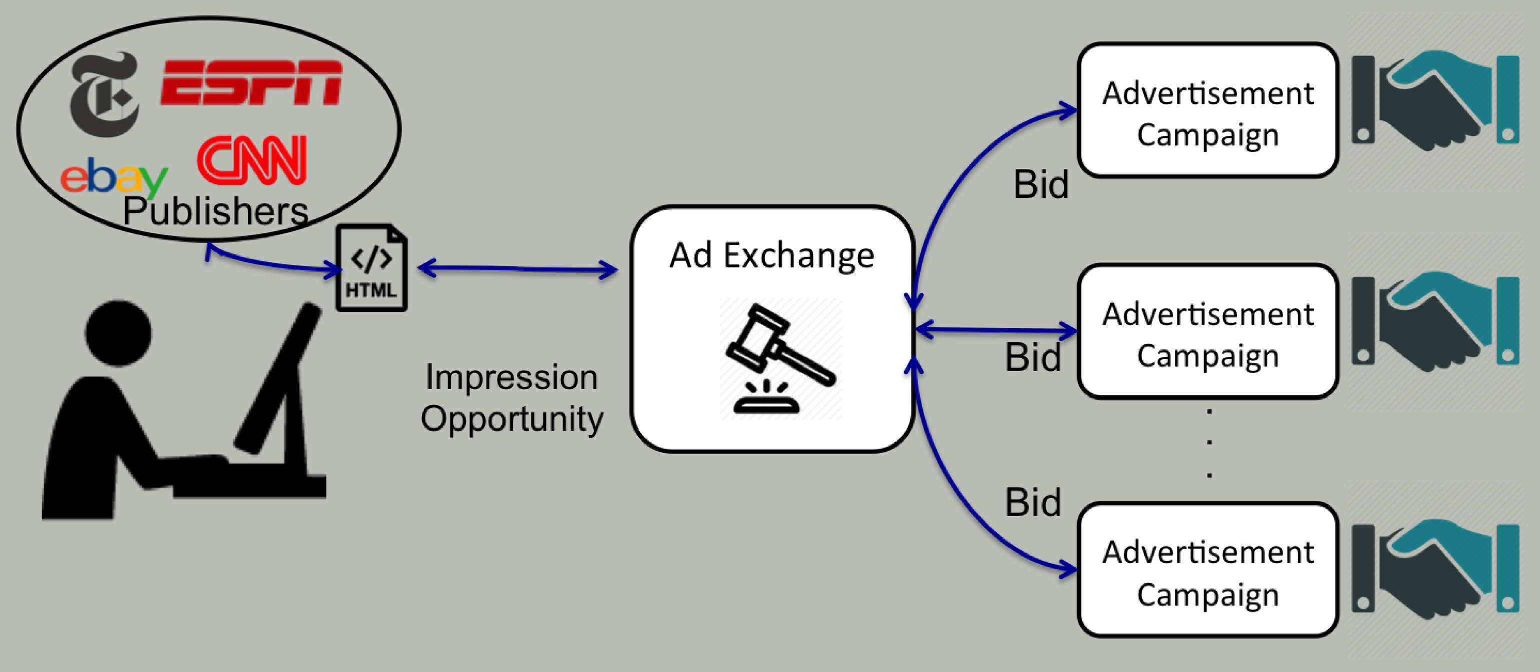

(2) Model

A **centralized combinatorial matching market** (CCMM) (or **market**, for short) is an augmented bipartite graph $M = (U, C, E, \vec{N}, \vec{I}, \vec{R})$, with a set of n types of goods U, a set of m bidders C, a set of edges E from goods to bidders indicating which goods are of interest to which bidders, a supply vector $\vec{N} = (N_1, \ldots, N_n)$, a demand vector $\vec{I} = (I_1, \ldots, I_m)$, and a reward vector $\vec{R} = (R_1, \ldots, R_m)$.
That is, there are $N_i > 0$ copies of good $i \in U$, and $I_j > 0$ *total* goods are demanded by bidder $j \in C$, where this total is a sum across all types of goods $i \in U$ such that $(i, j) \in E$. Reward $R_j > 0$ is attained by j in case its demand I_j is fulfilled.

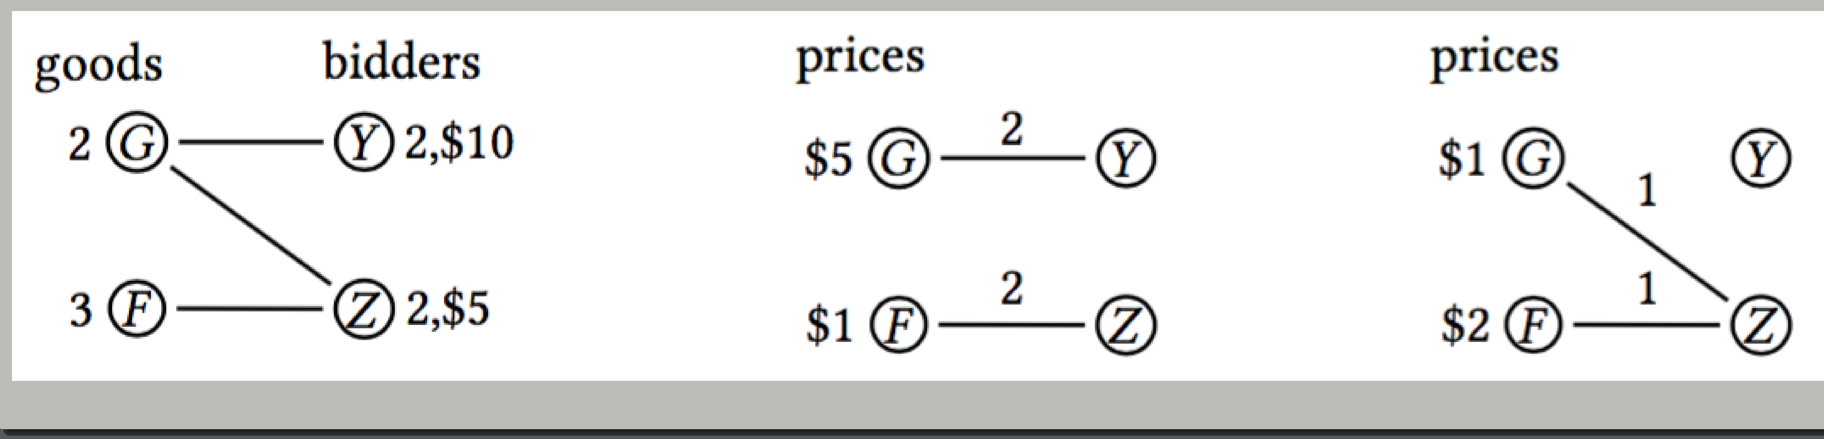

(3) Solution Concepts

A feasible outcome (X, p) is a **Walrasian Equilibrium** (WE) if:
Envy-freeness (EF): There is no bundle X_j' that any bidder j prefers to its assigned bundle X_j, i.e., for all j, $X_j \in \arg\max_{X_j' \in B(\vec{N})} \{V_j(X_j') - P_j(X_j')\}$.
Market clearance (MC): Every unallocated good is priced at zero, i.e., $\forall i \in U:$ If $\sum_{j=1}^{m} x_{ij} = 0$, then $p_i = 0$.

However, WE need not exist:

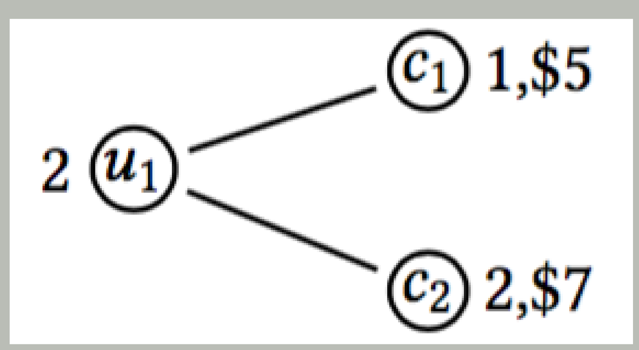

Instead we relax

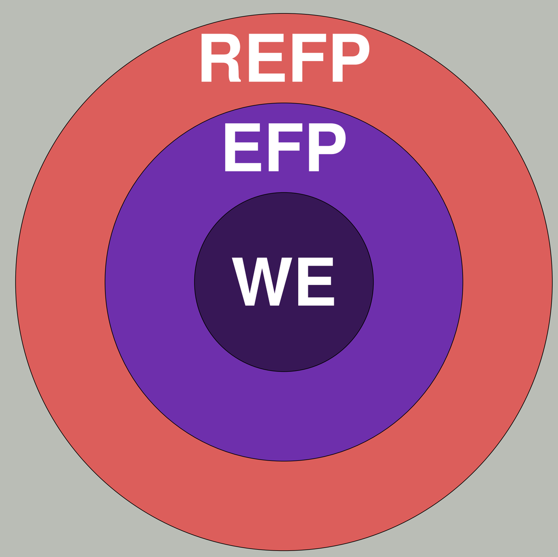

(4) Computing Restricted Envy-Free Prices

Definition

(1) A feasible outcome (X, p) is an **Envy-Free Pricing** (EFP) if EF holds.
(2) A feasible outcome (X, p) is an **Restricted-Envy-Free Pricing** (REFP) if only winners are envy-free.

Theorem

Given a market M and a feasible allocation X, the following conditions are necessary and sufficient for p to be restricted envy-free.
Individual Rationality: $\forall j \in W : P_j(X_j) \leq V_j(X_j)$.
Compact Condition: $\forall i \in U, j \in C :$ *If* $x_{ij} > 0$, *then* $\forall k \in U :$ *If* $(k, j) \in E$, *and* $x_{kj} < N_k$ *then* $p_i \leq p_k$.

These conditions comprise a **polynomial-time algorithm** for finding a REFP, assuming a linear or quadratic objective.

(5) Revenue Maximizing Equilibria

Given a market, how do we find a **revenue-maximizing** REFP?
We use a simple heuristic based on reserves prices, i.e, prices below which a good cannot be allocated.
For various choices of r, corresponding to various (greedy) allocations X, we find a REFP where goods are priced at least at r, and then we output a REFP which is revenue-maximizing among all those considered.

$$r \implies X_r \implies (X_r, p)$$

(6) Results for Special cases

(1) Our algorithm, when optimizing for seller-revenue, produces *unrestricted* envy-free prices p on input (M, X), where M is a **singleton** market (where each bidder demands exactly one good) and X is a welfare-maximizing allocation.

(2) In the case of **single-minded bidders**, the problem of finding a revenue-maximizing REFP reduces to the problem of finding a welfare-maximizing allocation.

(7) Experiments and Results

We tested our algorithms on both synthetic data and real-world data. The following table contains the results for the second case.

	Welfare	Revenue	EF	EF Loss	MC	MC Loss	Time
Over Demanded							
UnlimitedSupply	0.8544	0.7580	0.0199	0.0595	0.5507	0.5507	0.0908
LP Greedy Welfare	**0.9325**	**0.9289**	**0.0803**	**0.4915**	**0.0439**	**0.1965**	**0.0330**
LP Greedy Egalitarian	0.7951	0.7600	0.0679	0.3751	0.1731	0.2125	0.0599
LP Optimal Welfare	0.9992	0.9592	0.0507	0.2970	0.1604	0.3524	1.4110
LP Optimal Egalitarian	0.8738	0.8395	0.0518	0.2708	0.3200	0.4611	1.7400
Under Demanded							
UnlimitedSupply	0.8762	0.8171	0.0089	0.0218	0.6896	0.6896	0.0204
LP Greedy Welfare	**0.9250**	**0.9231**	**0.0604**	**0.3134**	**0.0327**	**0.1904**	**0.0176**
LP Greedy Egalitarian	0.8638	0.8368	0.0466	0.2368	0.0994	0.1890	0.0241
LP Optimal Welfare	0.9919	0.9582	0.0246	0.1166	0.1515	0.4406	0.8760
LP Optimal Egalitarian	0.9049	0.8789	0.0358	0.1717	0.2180	0.5019	0.7871

Table 1. Results, *TAC-Markets*.

Our algorithms perform well across markets metrics (revenue, welfare), in reasonable time, with very few violations of the EF and MC conditions.

(8) Towards Principled Autonomous Decision Making for Markets

Although our heuristic searches only REFPs, we nevertheless obtain outcomes that are close to EFPs, even when we seed our heuristic with a welfare-maximizing allocation, rather than an egalitarian one.

Is our relaxation a good idea? Is it useful? **Yes! (we think so!)**

Our first prototype for the TAC AdX game of an automated agent bidding based only on a **REFP** was capable of accumulating a positive score against other non-principled bidding agents.

In ongoing experiments, we are testing the hypothesis that bidding based on an approximation of a **Walrasian Equilibrium** can produce a robust agent against other agents that are highly optimized for the specific rules of this game, shedding light on the functioning of the game and thus, on the functioning of real ad-exchange markets.

eareyan@brown.com

On Approximate Welfare and Revenue-Maximizing Equilibria for Size-Interchangeable Bidders

Enrique Areyan Viqueira, Amy Greenwald, and Victor Naroditskiy

Computer Science Department, Brown University, Providence, RI, United States

BROWN

eareyan@brown.com