Competitive Equilibria in Two Sided Matching Markets with Non-transferable Utilities *

Saeed Alaei †, Kamal Jain‡, Azarakhsh Malekian§

Abstract

We consider two sided matching markets consisting of agents with non-transferable utilities; agents from the opposite sides form matching pairs (e.g., buyers-sellers) and negotiate the terms of their math which may include a monetary transfer. Competitive equilibria are the elements of the core of this game.

We present the first combinatorial characterization of competitive equilibria that relates the utility of each agent at equilibrium to the equilibrium utilities of other agents in a strictly smaller market excluding that agent; thus automatically providing a constructive proof of existence of competitive equilibria in such markets.

Our characterization also yields a group strategyproof mechanism for allocating indivisible goods to unit demand buyers with non-quasilinear utilities that highly resembles the Vickrey–Clarke-Groves (VCG) mechanism. As a direct application of this, we present a group strategyproof welfare maximizing mechanism for Ad-Auctions without requiring the usual assumption that search engine and advertisers have consistent estimates of the clickthrough rates.

*Id: ce.tex 312 2012-11-13 21:32:38Z saeed

†Dept. of Computer Science, Cornell University, Ithaca, NY 14850. saeed@cs.cornell.edu. Part of this work was done when the author was visiting Microsoft Research, Redmond. This work was partially supported by the NSF grant CCF-0728839

‡eBay Research, 2065 Hamilton Ave, San Jose, 95125 kamaljain@gmail.com.

§Laboratory of Information and Decision Systems, MIT, Cambridge, MA 02142. malekian@mit.edu.

1 Introduction

In markets with a common currency it is typical to assume transferable utilities, i.e., that an agent can transfer some of its utility to another agent at no loss; however this assumption rarely holds in practice as wealthy agents often derive less marginal utility from a fixed increase to their wealth compared to poor agents. The situation is more pronounced when monetary transfers are significant compared to agents wealth. We consider non-transferable utilities in the context of matching markets.

In this paper we consider a two-sided matching market with non-transferable utilities. Agents are divided in two sides (e.g., sellers-buyers, men-women, etc). Each agent can partner with at most one other agent from the opposite side, and they mutually decide about the terms of their partnership which may involve a monetary transfer. We are interested in the "*stable* outcomes of this market, i.e., outcomes in which no coalition of agents could rearrange in such a way that would strictly benefit every member of the coalition. Any such stable outcome corresponds to a *competitive equilibrium*.

A competitive equilibrium is a feasible matching of agents (including the terms of matching and monetary transfers) such that (i) there are no agents who could form a matching pair in such a way that would benefit both of them better than their current state, and (ii) there is no matched agent who would prefer to be unmatched.

Throughout this paper we consider markets in canonical form consisting of buyers on one side and goods (or sellers) on the opposite side such that (i) utility of each buyer depends on the choice of good received and the payment charged (utilities are decreasing in payment but not necessarily quasilinear), and (ii) utility of each seller is equal to the payment received. A competitive equilibrium then corresponds to an assignment of prices to goods together with envy–free matching of goods to buyers, i.e., such that any buyer receives the good that she prefers the most at the assigned prices. In particular, any unassigned good should have a price of zero. Any market can be represented in the canonical form by treating the utilities of agents of one side of the market as prices of goods (i.e., each agent represented by a good), and by considering agents of the opposite side as buyers; the utility of a buyer for a good at a price x then corresponds to the highest utility the agent represented by the buyer could get if matched with the agent represented by the good under any set of terms or monetary transfer that would give the agent represented by the good a utility of at least x.

Competitive equilibria and their various properties have been studied extensively in the past, e.g., Demange and Gale (1985), Gale (1984), Quinzii (1984). It has been shown that a competitive equilibrium always exists, however previous proofs are non-constructive and provide little insight into equilibrium structure. There has been independent recent work on algorithmic and/or graph theoretic characterizations of competitive equilibria, e.g., Caplin and Leahy (2010b,a) and Mani et al. (2010a,b); which have been focused on obtaining iterative processes for bottom-up construction of competitive equilibrium by growing an indifference graph while making gradual adjustments to prices. The current paper provides a simple and natural characterization of competitive equilibria based on a top-down construction, also yielding a simple recursive algorithm for computing them; it also provides deep insight into equilibrium structure and how it reacts to changes in supply/demand.

A competitive equilibrium can be concisely represented by its price vector. A price vector induces a utility for each buyer which is equal to the highest utility the buyer could get from any of the goods at the specified prices. At a competitive equilibrium, utilities of buyers are exactly equal to their induced utilities from the equilibrium prices. Similarly, a utility vector induces a price on each good, which is the highest price any buyer would pay for that good given the specified utility vector. Gale (1984) showed that the set of equilibrium price vectors form a complete lattice (similarly the set of equilibrium utility vectors). In particular, there is a competitive equilibrium with the highest prices (and lowest utilities), and there is one with the lowest prices (and highest utilities).

In this paper we present a simple combinatorial characterization of competitive equilibria in two-sided matching markets with non-transferable utilities. Our characterization is inductive, i.e., it relates the equilibrium prices/utilities to those of strictly smaller markets, thereby providing a constructive proof of existence of competitive equilibria, and yielding a purely combinatorial approach for computing the equilibria. We summarize our results below:

- The utility of a buyer at the lowest priced competitive equilibrium is equal to her induced utility from the prices of the highest priced competitive equilibrium of the market without her. Similarly the price of a good at the highest priced competitive equilibrium is equal to its induced price from the utilities of the lowest prices competitive equilibrium of the market without that good. Combining the previous two characterization implies that the price a buyer pays for a good she receives at the lowest priced competitive equilibrium is equal to the highest price the rest of the market would be willing to pay for that good, i.e., it is the price at which the rest of the market becomes indifferent between buying and not buying the good. The previous characterization also yields a mechanism for obtaining the lowest priced competitive equilibrium which highly resembles the Vickrey–Clarke-Groves (VCG) mechanism, albeit for non-transferable utilities.

- There is a continuum of competitive equilibria between the lowest priced competitive equilibrium and the highest priced competitive equilibrium. In particular, given any lower bound on the price vector or utility vector, one could obtain a competitive equilibrium satisfying the bounds, if any, using a similar combinatorial characterization as above.

- At the lowest priced competitive equilibrium, every set of goods with strictly positive prices is weakly over demanded, i.e., there is a buyer who is not matched to any of the goods in that set, but is weakly interested in at least one of the goods in that set. Similarly, at the highest priced competitive equilibrium, in every set of buyers with strictly positive utilities there is a buyer who is weakly interested in a good that is not among the goods matched to that set of buyers. The next two properties can be derived from the previous two properties.

- At the highest priced competitive equilibrium, the market is indifferent between keeping or losing any single good. In other words, if any single good is removed from the market, there is matching of the remaining goods to buyers at the current prices such that the market remain at equilibrium with the same prices/utilities as before. Similarly, at the lowest priced competitive equilibrium, the market is indifferent between keeping or losing any single buyer. In other words, if any single buyer is removed from the market, there is matching of the goods to remaining buyers at the current prices such that the market remain at equilibrium with the same prices/utilities as before.

Roadmap We start by defining our model and assumptions in section 2. In section 3, we provide a brief overview of competitive equilibria and their structure in the case of transferable utilities, we then identify the main difficulties of applying existing techniques to non-transferrable utilities. We then present our main results in section 4 . Finally, in section 5 we present an immediate application of our results to Ad-Auctions. The next section provides a quick review of the related literature.

1.1 Related Work

The problem we consider is a one-to-one matching with non-transferable utilities as described by Demange and Gale (1985). The problem is a generalization of the assignment game of Shapley and Shubik (1971) to non-quasilinear utilities. Shapley and Shubik studied this problem for quasilinear utilities, proved that core outcomes always exist, and showed that they form a lattice. Note that core outcomes correspond exactly to competitive equilibria in these models. The existence of competitive equilibria for non-quasilinear utilities was proved by Quinzii (1984) and by Gale (1984). Quinzii showed that the game defined by this model is a *"Balanced Game"*, where for general n-person balanced games Scarf (1967) had shown that the core is non-empty. Using a similar approach Kaneko and Yamamoto (1986) proved a similar result for a slight generalization of this problem. Gale (1984) also showed that a competitive equilibrium always exists using combinatorial topology. The proof of Gale is based on a generalization of the KKM lemma (see Knaster et al. (1929)) which is the continuous variant of the Sperner's lemma. Both of these proofs are non-constructive and only show the existence of an equilibrium. As such, they don't provide a natural characterization of equilibrium utilities [1]. Leonard (1983) showed that in one-to-one markets with quasilinear utilities, prices at the lowest competitive equilibrium equal VCG payments. The proof of Leonard is based on writing the LP for computing the welfare maximizing allocation(i.e., a maximum weight matching) and showing that every optimal assignment of the dual variables correspond exactly to the prices/utilities at a competitive equilibrium. The proof crucially depends on the quasilinearity of utilities. Demange and Gale (1985) studied various properties of these markets in the case of non-quasilinear utilities. They showed that the set of competitive equilibria is a lattice. Demange et al. (1986) later proposed an ascending auction for computing the competitive equilibrium with the lowest prices in the case of quasilinear utilities. Interestingly, their method is the same as the Hungarian method (see Kuhn (1956)) for finding a maximum weight matching. Notice that at this point there was still no combinatorial characterization of competitive equilibria for the case of non-quasilinear utilities and no constructive proof of existence except for the ascending auction of Alkan and Gale (1990) for piece-wise linear utility functions.

[1] The proof of Scarf provides an algorithm based on the pivoting algorithm of Lemke and Howson (1964), which can be combined with the construction of Quinzii to yield an algorithm that computes a competitive equilibrium in $2^{O(n!)}$ iterations running on a matrix with $O(n!)$ columns. Nevertheless, the resulting algorithm is more of an exhaustive search and does not provide any insight into the equilibrium structure.

There has been independent recent work on algorithmic and/or graph theoretic characterizations of competitive equilibria by Caplin and Leahy (2010b,a) and Mani et al. (2010a,b); which have been focused on obtaining iterative processes for bottom-up construction of competitive equilibrium by growing an indifference graph while making gradual adjustments to prices.

The are closely related works involving market with many-to-one or many-to-many matchings, either with quasilinear utilities or with ordinal preferences. For competitive equilibria in markets with many-to-one matchings and with quasilinear utilities see Mamer (1997); Gul and Stacchetti (1999); Ausubel and Milgrom (2002). For stable (core) allocations in markets with many-to-one or many-to-many matchings with ordinal preferences see Gale and Shapley (1962); Kelso and Crawford (1982); Hatfield and Milgrom (2005); Hatfield and Kominers (2010).

2 Preliminaries

Model Let $M = (I, J, \mathrm{u})$ denote a market where I is a set of unit demand buyers, J is a set of goods, and $\mathrm{u}_i^j(x)$ denote the utility of buyer $i \in I$ for receiving good $j \in I$ and making a payment of x. The utility functions are private information of the respective buyers. We assume that every $\mathrm{u}_i^j(x)$ is continuous and decreasing and there exists a large enough $m_i^j \in \mathbb{R}_+$ such that $\mathrm{u}_i^j(m_i^j) \leq 0$. Without loss of generality, we assume that $\mathrm{u}_i^j(x)$ is defined for all $x \in \mathbb{R}$ and its range also covers the whole $\mathbb{R}$ [2], and therefore $\mathrm{u}_i^j(x)$ is invertible everywhere. We will use p_i^j to denote the inverse of u_i^j, i.e., in order to give buyer i a utility of x from good j, the buyer should be charged a payment of $\mathrm{p}_i^j(x)$. Observe that $\mathrm{p}_i^j(x)$ is also a continuous and decreasing function whose domain/range is the whole $\mathbb{R}$. For the rest of this paper, we adopt the notation of using u and p to denote functions returning utilities and prices, and $\mathbf{u}$ and $\mathbf{p}$ to denote utility vector and price vector. We use subscripts to index agents and superscripts to index goods (i.e., $\mathbf{u}_i$, $\mathbf{p}^j$, etc). Negative subscripts/superscripts are used to exclude the specified index (e.g., $\mathbf{u}_{-i}$ is the same as $\mathbf{u}$ but with the i^{th} entry removed).

Definition 1 (Competitive Equilibrium (CE))**.** *A* "Competitive Equilibrium" *(henceforth abbreviated as CE) is an assignment of prices to goods together with a feasible matching of goods to buyers such that each buyer receives her most preferred good at the assigned prices, and such that every unmatched good has a price of* 0*. Formally, for a market $M = (I, J, \mathrm{u})$ we say that $W = (\mathbf{u}, \mathbf{p})$ is a CE of M with price vector $\mathbf{p}$ and utility vector $\mathbf{u}$, if and only if there exists a* "supporting matching" $\boldsymbol{\mu}$ *such that:*

- *For every buyer i and good j, if i and j are matched in $\boldsymbol{\mu}$, then $\mathbf{u}_i = \mathrm{u}_i^j(\mathbf{p}^j)$, otherwise $\mathbf{u}_i \geq \mathrm{u}_i^j(\mathbf{p}^j)$.*
- *Every buyer i that is unmatched in $\boldsymbol{\mu}$ should have a zero utility (i.e., $\mathbf{u}_i = 0$). Similarly, Every good j that is unmatched in $\boldsymbol{\mu}$ should have a zero price (i.e., $\mathbf{p}^j = 0$).*

[2] Suppose m_i^j is the smallest non-negative payment for which $\mathrm{u}_i^j(m_i^j) \leq 0$. Since at the equilibrium all prices/utilities are non-negative, we can redefine u_i^j for values of x outside of $[0, m_i^j]$ as follows without affecting the equilibrium: for $x < 0$ redefine $\mathrm{u}_i^j(x) \leftarrow \mathrm{u}_i^j(0) - x$, and for $x > m_i^j$ redefine $\mathrm{u}_i^j(x) \leftarrow \mathrm{u}_i^j(m_i^j) - (x - m_i^j)$.

- *All prices/utilities must be non-negative, i.e.,* $\mathbf{u}_i \geq 0, \mathbf{p}^j \geq 0$.

For notational convenience, we often use $\mathrm{u}(W)$, $\mathrm{p}(W)$ *to refer to* $\mathbf{u}$, $\mathbf{p}$ *respectively.* $\boldsymbol{\mu}(i)$ *will denote the good matched to buyer* i *(if unmatched,* $\boldsymbol{\mu}(i) = \emptyset$*), and* $\boldsymbol{\mu}^{-1}(j)$ *will denote the buyer matched to good* j *(if unmatched,* $\boldsymbol{\mu}^{-1}(j) = \emptyset$*).*

Note that any CE, say W, can be specified by either $\mathrm{u}(W)$ or $\mathrm{p}(W)$; given either the price vector or the utility vector, the other one can be uniquely computed by taking the induced prices/induced utilities as defined next.

Definition 2 (Induced utilities/Induced Prices)**.** *Given a market* $M = (I, J, \mathrm{u})$ *and a price vector* $\mathbf{p}$*, we define* $\mathrm{u}_i(\mathbf{p})$ *to denote the utility induced on buyer* i *by offering the goods at prices* $\mathbf{p}$*. We can formally define* u_i *as follows.*

$$\mathrm{u}_i(\mathbf{p}) = \max(0, \max_{j \in J} \mathrm{u}_i^j(\mathbf{p}^j)) \tag{2.1}$$

Similarly, given a utility vector $\mathbf{u}$*, we define* $\mathrm{p}^j(\mathbf{u})$ *to denote the price induced on good* j *which is the highest price any buyer would pay for good* j *assuming their utilities are fixed at* $\mathbf{u}$*, i.e.,*

$$\mathrm{p}^j(\mathbf{u}) = \max(0, \max_{i \in I} \mathrm{p}_i^j(\mathbf{u}_i)). \tag{2.2}$$

It is easy to see that if $W = (\mathbf{u}, \mathbf{p})$ is a CE, then $\mathbf{p}^j = \mathrm{p}^j(\mathbf{u})$ and $\mathbf{u}_i = \mathrm{u}_i(\mathbf{p})$[3]. Throughout the rest of this paper, we often specify a CE such as W by specifying either $\mathrm{p}(W)$ or $\mathrm{u}(W)$; note that once $\mathbf{p}$ and $\mathbf{u}$ are determined, it is straightforward to find a supporting matching $\boldsymbol{\mu}$[4]. Note that there could be multiple feasible supporting matchings for a given set of prices/utilities; in that case we could pick any such matching arbitrarily.

3 Transferable vs. Non-Transferable Utilities

We start by describing the structure of CEs for transferable utilities and then compare it with the more general case of non-transferable utilities. We show that the standard approaches for analyzing and/or computing CEs with transferable utilities fail to generalize to non-transferable utilities; in particular we show that ascending auctions — even when carefully modified to work for non-transferable utilities — may take arbitrarily long time to converge.

Consider a market $M = (I, J, \mathrm{u})$ with transferable utilities. Without loss of generality we may assume that payments(prices) are also measured in the same units as the utilities, therefore we may assume every utility function is of the form $\mathrm{u}_i^j(x) = v_i^j - x$, where v_i^j can be thought of as the valuation of the buyer i for good j. Then every CE of the market corresponds to a maximum weighted matching and a minimum weight covering of the bipartite

[3]The inverse is not true, i.e., given non-negative vectors $\mathbf{p}$ and $\mathbf{u}$ such that for all i and j we have $\mathbf{p}^j = \mathrm{p}^j(\mathbf{u})$ and $\mathbf{u}_i = \mathrm{u}_i(\mathbf{p})$, there may exist no feasible matching $\boldsymbol{\mu}$ such that $(\mathbf{u}, \mathbf{p})$ is a CE.

[4]If $\mathbf{p}$ and $\mathbf{u}$ correspond to a CE, then a supporting matching $\boldsymbol{\mu}$ can be computed as follows: consider the bipartite graph in which there is an edge between each buyer i and good j iff $\mathbf{u}_i = \mathrm{u}_i^j(\mathbf{p}^j)$ and find a maximum matching covering all vertices with strictly positive price/utility.

graph consisting of buyers/goods in which the edge between buyer i and good j has weight v_i^j. The proof of this claim follows from the linear programming relaxation of the maximum weight matching problem and its dual as shown below. In the following, y_i^j is the variable corresponding to the allocation of good j to buyer i.

(Primal)

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in I} \sum_{j \in J} v_i^j y_i^j \\ \text{subject to} \quad & \textstyle\sum_{j \in J} y_i^j \leq 1, \quad \forall i \in I \\ & \textstyle\sum_{i \in I} y_i^j \leq 1, \quad \forall j \in J \\ & y_i^j \geq 0 \end{aligned}$$

(Dual)

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in I} \mathbf{u}_i + \sum_{j \in J} \mathbf{p}^j \\ \text{subject to} \quad & \mathbf{u}_i + \mathbf{p}^j \geq v_i^j, \qquad \forall i \in I, \forall j \in J \\ & \mathbf{u}_i \geq 0 \\ & \mathbf{p}^j \geq 0 \end{aligned}$$

The optimal assignments of the above linear program and its dual correspond to CEs as follows. Consider an extreme point optimal assignment of the primal and any optimal assignment for the dual. The extreme points of the matching polytope are integral so every y_i^j is either 0 or 1. By strong duality both the primal and the dual have the same optimal value which happens to be equal to the optimal social welfare. By complementary slackness, if $y_i^j = 1$, then the corresponding dual constraint must be tight which implies $\mathbf{u}_i = v_i^j - \mathbf{p}^j$. Similarly, complementary slackness implies that every good with a non-zero price (similarly ever buyer with a non-zero utility) should be matched in the primal. Finally, observe that the dual constraints ensure that the utility of any buyer from her current assignment is the maximum she can get from any good under the current prices.

The previous argument also implies that every CE yields a welfare maximizing allocation. Furthermore, among all CEs (i.e., all optimal assignments of the dual program), the one with the minimum prices coincides with the outcome of the Vickrey–Clarke-Groves (VCG) mechanism.

A maximum weight matching can alternatively be computed using the *Hungarian Method* of Kuhn (1956). Interestingly, the Hungarian method is equivalent to the following ascending price auction proposed by Demange et al. (1986) for computing the CE with the lowest prices:

Definition 3 (Ascending Price Auction)**.** *Set all the prices equal to* 0*. Find a minimally over demanded subset of goods at the current prices, i.e., a subset* T *of goods such that there is a subset* S *of the buyers who strictly prefer one or more of the goods in* T *to goods outside of* T *at the current prices and such that* $|S| > |T|$*; increase the prices of goods in* T *simultaneously (at a uniform rate) until one of the buyers in* S *becomes indifferent between a good outside of* T *and her preferred good(s) in* T*; at that point, recompute the minimally over-demanded subset and repeat the process until there is no over demanded subset of goods.*

Unfortunately none of the above approaches can be effectively generalized to handle non-transferable utilities. The linear programming relaxation and the VCG mechanism heavily rely on quasi-linearity of utilities. Furthermore, the ascending auction – even if modified to work with non-transferable utilities – may take arbitrarily long time to reach equilibrium. Constructive approaches have been proposed for special classes of non-transferable utilities (e.g., Alkan and Gale (1990)), all of which essentially boil down to an ascending auction of the following form: raise the prices at some rate (potentially non-uniform) to the next point at which there is a change in the demand structure [5]; then, recompute the demand structure and repeat. For quasilinear utilities, the ascending auction stop after $O(|I| + |J|)$ iterations (a new iteration starts each time the demand sets change). For special cases of non-quasilinear utilities (for example quasilinear but with a hard budget constraints) similar ascending auctions have been applied. However such auctions may not reach an equilibrium in finite time as we illustrate with an example next.

The main problem with ascending auctions for non-transferable utilities occurs when the prices of an over demanded subset of goods, say T, is to be raised. For transferable (quasilinear) utilities, the prices of all of the goods in T are raised at the same rate without affecting the relative preferences of buyers over the goods in T; however, that is not the case for non-transferable (non-quasilinear) utilities. In the general case, one may need to raise the prices of goods in T at different and possibly variable rates and even then the preferences of buyers over the goods in T may change an unbounded number of times. We demonstrate the problem in the following example:

Example 1. *Suppose there are 3 goods and 4 buyers whose utility functions are given in the following table in which v is some constant (at least 2) and x is the payment(price):*

	good 1	*good 2*	*good 3*
buyer 1	$\mathrm{u}_1^1(x) = v + 1 - x$	$\mathrm{u}_1^2(x) = v + 1 - x$	$\mathrm{u}_1^3(x) = v + 1 - x$
buyer 2	$\mathrm{u}_2^1(x) = 0 - x$	$\mathrm{u}_2^2(x) = v + 1 - x$	$\mathrm{u}_2^3(x) = 0 - x$
buyer 3	$\mathrm{u}_3^1(x) = 0 - x$	$\mathrm{u}_3^2(x) = 0 - x$	$\mathrm{u}_3^3(x) = v + 1 - x$
buyer 4	$\mathrm{u}_4^1(x) = v - x$	$\mathrm{u}_4^2(x) = v - x - \frac{v-x}{v}\sin(v\log(v-x))$	$\mathrm{u}_4^3(x) = v - x - \frac{v-x}{v}\cos(v\log(v-x))$

All buyers have quasilinear utilities except buyer 4. Moreover, all utilities are strictly decreasing in payment (assuming $v \geq 2$). Figure 1 shows the prices of goods throughout the ascending auction. We should emphasis that in this particular example the ascending path of prices is unique. The ascending auction can only increase the prices of goods that are over demanded, i.e., demanded by at least two buyers. Furthermore, it can only raise the price of a good to the point where the demand of that good is about to drop to 1. Therefore, for every good with a positive price there should be at least a (weak) demand of 2 at any point in the auction. Observe that the demand set of buyer 1 and 4 changes an infinite number of times during the ascending auction. Specifically, the demand set of both buyer 1 and 4 include good 1 at all times. However, the demand set of buyer 1 includes good 2 and/or good 3 only at the times the price curves of those goods overlap with the price curve of good 1. Similarly, the demand set of buyer 4 includes good 2 and/or good 3 only at the times the price curves of those goods do not overlap with the price curve of good 1. Observe that the demand structure changes an infinite number of times as the price of the goods approach v. So an ascending auction does not stop in finite time.

The previous example, although contrived, illustrates a fundamental problem that arises with ascending auctions and constructive proofs that are based on them. In general, ascending auctions are quite sensitive to the structure of utility functions. One of the main

[5] Note that the prices can be jumped (discretely) to the next point at which there is a change in the demand sets.

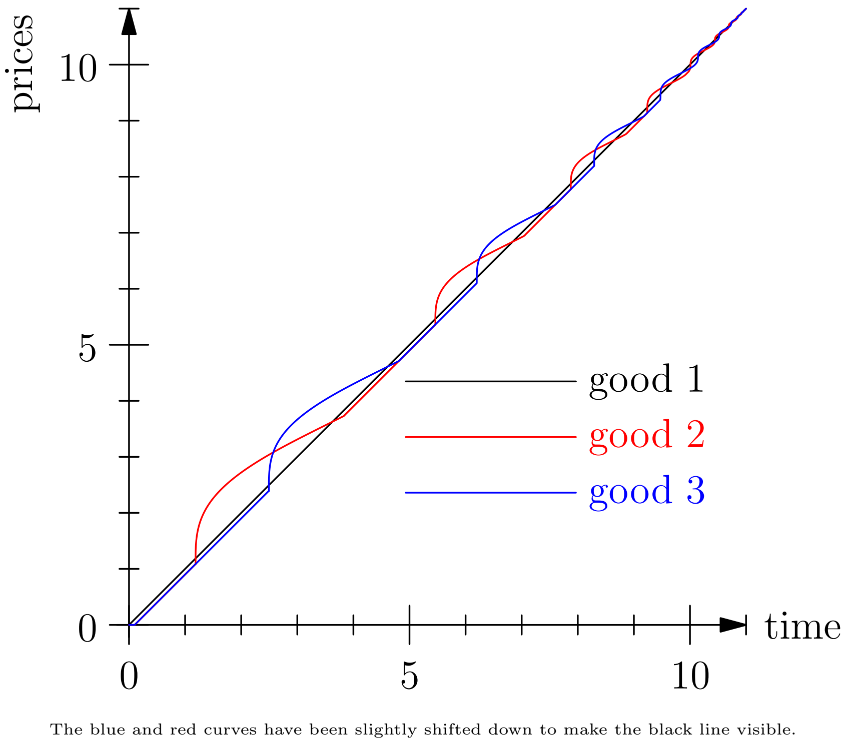

Figure 1: prices of goods in the ascending auction of Example 1 assuming that $v = 11$ and assuming the price of good 1 increases at the rate of 1 .

contributions of the current paper is a direct approach for computing the lowest CE without running an ascending auction (see Theorem 2).

Quasilinear with hard budget limit. In this case the utility function is quasilinear as long as the payment does not exceed the budget limit; and if the payment exceeds the budget, the utility function goes to negative infinity [6]. Notice that the issue outlined in the previous example does not arise with quasilinear utilities with hard budget constraints. In fact, ascending auctions with hard budget constraints converge almost as fast as ascending auction with quasilinear utilities and no budget constraints because each buyer may hit her budget limit at most $|J|$ times (once per each good) and beyond that point she never demands the same good again. For the same reason, quasilinear utilities with hard budget limits are easier to deal with than general non-quasilinear utilities.

4 Characterization of Competitive Equilibria

This section studies structural properties of CEs. The main result of this section is Theorem 2 which characterizes the equilibrium prices/utilities.

We start by showing that the set of CEs is a complete lattice. Consider a market $M = (I, J, \mathrm{u})$, and let $W = (\mathbf{u}, \mathbf{p})$ and $W' = (\mathbf{u}', \mathbf{p}')$ be two arbitrary CEs of M with supporting matchings $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ respectively. We define the min and max operators for CEs as follows.

[6]This ensures that an individually rational mechanism never charges the buyer more than her budget limit.

$$\min(W, W') = (\max(\mathbf{u}, \mathbf{u}'), \min(\mathbf{p}, \mathbf{p}')) \qquad \text{with supporting matching} \quad \boldsymbol{\mu}'' = \begin{cases} \boldsymbol{\mu}(i) & \mathbf{u}_i \geq \mathbf{u}'_i \\ \boldsymbol{\mu}'(i) & \mathbf{u}_i < \mathbf{u}'_i \end{cases}$$

$$\max(W, W') = (\min(\mathbf{u}, \mathbf{u}'), \max(\mathbf{p}, \mathbf{p}')) \qquad \text{with supporting matching} \quad \boldsymbol{\mu}''' = \begin{cases} \boldsymbol{\mu}'(i) & \mathbf{u}_i \geq \mathbf{u}'_i \\ \boldsymbol{\mu}(i) & \mathbf{u}_i < \mathbf{u}'_i \end{cases}$$

Note that we assume min and max operators applied to vectors (e.g., price vector, utility vector) return respectively the component-wise minimum and maximum of those vectors. The following theorem was originally proved by Demange and Gale (1985).

Theorem 1 (Equilibrium Lattice Demange and Gale (1985))**.** *Given a market $M = (I, J, \mathrm{u})$, if W and W' are any two CEs of M, then $W'' = \min(W, W')$ and $W''' = \max(W, W')$ are also CEs of M. Consequently the set of all CEs is a complete lattice and has a unique minimum and a unique maximum* [7]*.*

We will refer to the CE with the lowest prices as the *lowest CE* and the one with the highest prices as the *highest CE*. Observe that the lowest CE has the highest utilities and the highest CE has the lowest utilities. Throughout the rest of this paper we implicitly use the lattice structure of the set of CEs without making explicit references to Theorem 1.

Quinzii (1984) and Gale (1984) presented existential arguments (non-constructive) proving that set of CEs is always non-empty. The main contribution of this paper is an inductive characterization of equilibria that relates the prices/utilities of the highest/lowest CE of a market to those of a strictly smaller market by removing either a buyer or a good. This inductive characterization automatically yields a constructive proof of the non-emptiness of the set of CEs.

For any given market $M = (I, J, \mathrm{u})$, let $M_{-i} = (I - \{i\}, J, \mathrm{u})$ (i.e., the same market without buyer i), and let $M^{-j} = (I, J - \{j\}, \mathrm{u})$ (i.e. the same market without good j). The following theorem relates the equilibrium price/utility of any good/buyer to those of a strictly smaller market which excludes the respective good/buyer; hence yielding an inductive approach for computing equilibrium prices/utilities.

Theorem 2 (Inductive Equilibrium Characterization)**.** *Consider an arbitrary market $M = (I, J, \mathrm{u})$. Let $\underline{W}$ be the lowest CE of M, $\overline{W}$ be the highest CE of M, $\underline{W}^{-j}$ be the lowest CE of M^{-j}, and $\overline{W}_{-i}$ be the highest CE of M_{-i}. Then*

I. $\mathrm{u}_i(\underline{W}) = \mathrm{u}_i(p(\overline{W}_{-i}))$,

II. $\mathrm{p}^j(\overline{W}) = \mathrm{p}^j(u(\underline{W}^{-j}))$.

Furthermore,

III. $\mathrm{p}^j(\underline{W}) \leq \mathrm{p}^j(\overline{W}_{-i})$*, in particular, if i and j are matched in $\underline{W}$, then* $\mathrm{p}^j(\underline{W}) = \mathrm{p}^j(\overline{W}_{-i})$,

[7] That is because the set of all CEs is a closed and compact set, which can be proved by applying the BolzanoWeierstrass theorem and using the fact that the set of all possible matchings is finite.

IV. $\mathrm{u}_i(\overline{W}) \leq \mathrm{u}_i(\underline{W}^{-j})$*, in particular, if* i *and* j *are matched in* $\overline{W}$*, then* $\mathrm{u}_i(\overline{W}) = \mathrm{u}_i(\underline{W}^{-j})$*.*

Note that equilibrium prices/utility can be fully computed by recursive application of the first two equations of the above theorem(recall that a CE can specified by either of its price vector or utility vector). The result of the theorem can be interpreted as follows.

- Equation 2.I: utility of buyer i at the lowest CE of market M can be computed as follows. Remove i from the market. Compute the prices at the highest CE of the rest of the market. Offer those prices to buyer i and compute her utility from her most preferred good at those prices.
- Equation 2.II: The price of any good j at the highest CE of the market M can be computed as follows. Remove good j from the market. Compute the buyers' utilities at the lowest CE of the rest of the market. Ask each buyer to name a price for good j that would give her the same utility as what she currently gets. Take the maximum among the named prices.

Combining the two equations of Theorem 2 yields the following characterization of payments.

Corollary 1. *At the lowest CE of a market, the payment of each buyer for the good she has received is equal to the highest price the rest of the market would be willing to pay for that good given their current utilities.*

Notice the striking similarity between the above characterization and the payments in the Vickrey–Clarke-Groves (VCG) mechanism, i.e., that the payment of each buyer is equal to her externality on the rest of the buyers; however recall that VCG cannot be applied here as it crucially depends on quasi-linearity of utilities.

Next we show there is a continuum of CEs between the highest CE and the lowest CE. We also present a inductive characterization for obtaining such CEs.

Definition 4 (($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)-Bounded CE)**.** *Given a market* $M = (I, J, \mathrm{u})$ *and a lower bound price vector* $\underline{\mathbf{p}} \in \mathbb{R}_+^J$ *and a lower bound utility vector* $\underline{\mathbf{u}} \in \mathbb{R}_+^I$*, we say that* W *is a* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)-bounded CE *of* M *iff* W *is a CE of* M *and* $\mathrm{p}(W) \geq \underline{\mathbf{p}}$ *and* $\mathrm{u}(W) \geq \underline{\mathbf{u}}$*.*

Note that for a given $\underline{\mathbf{u}}$ *and* $\underline{\mathbf{p}}$*, the* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CEs of* M *form a complete sublattice of all the CEs of* M*. In particular, there is a lowest and a highest* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CE of* M*. Notice that for arbitrary* $\underline{\mathbf{u}}$ *and* $\underline{\mathbf{p}}$*, a* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CE may not necessarily exist.*

Theorem 3 (Continuity)**.** *Assume a market* $M = (I, J, \mathrm{u})$ *with equal number of buyers and goods (i.e.,* $|I| = |J|$*) and lower bounds* $\underline{\mathbf{p}} \geq \mathbf{0}$ *and* $\underline{\mathbf{u}} \geq \mathbf{0}$ *on the prices/utilities. If a* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CE exists, then at the lowest* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CE there exists at least one good* $j^* \in J$ *whose price is exactly equal to its lower bound (i.e. equal to* $\underline{\mathbf{p}}^{j^*}$*), and at the highest* ($\underline{\mathbf{u}}, \underline{\mathbf{p}}$)*-bounded CE there exists at least one buyer* $i^* \in I$ *whose utility is exactly equal to her lower bound (i.e. equal to* $\underline{\mathbf{u}}_{i^*}$*).*

The proof of the above theorem is based on defining a new market $M' = (I, J, \mathrm{u}')$ with transformed utility functions $\mathrm{u}'^j_i(x) = \mathrm{u}^j_i(x + \underline{\mathbf{p}}^j) - \underline{\mathbf{u}}_i$. There is a one to one correspondence between CEs of M and M' [8]. Given that there are equal number of buyers and goods, applying Theorem 2 and removing a buyer (or a good) from M' implies that there will be an unmatched good (or unmatched buyer) in the remaining market which will have a zero price (or zero utility) which then implies the claim in the original market.

The following corollary immediately follows the previous theorem by setting the lower bounds to zero.

Corollary 2. *Given a market $M = (I, J, \mathrm{u})$ with $|I| = |J|$:*

- *At the lowest CE, there is at least one good that has a price of* 0.
- *At the highest CE there is at least one buyer that has a utility of* 0.

As another immediate corollary of Theorem 3, a continuum of CEs can be obtained as follows. Define $\underline{\mathbf{p}}(t) = (1-t)\mathrm{p}(\underline{W}) + t\mathrm{p}(\overline{W})$. Now applying Theorem 3 and taking the lowest $(\mathbf{0}, \underline{\mathbf{p}}(t))$-Bounded CE for $t \in [0, 1]$ yields a continuum of equilibria between the lowest CE and the highest CE of M.

Corollary 3. *Given a market $M = (I, J, \mathrm{u})$ with $|I| = |J|$, there is a continuum of equilibria between $\underline{W}$ and $\overline{W}$.*

Throughout the rest of this section we present several theorems which capture important properties of lowest/highest CEs. The next two theorems capture the demand structure in the lowest/highest CEs; they also play a key role in the proof of Theorem 2.

Definition 5 (Demand Sets)**.** *Consider a market $M = (I, J, \mathrm{u})$ and let $W = (\mathbf{u}, \mathbf{p})$ be a CE of M. For each subset $S \subseteq I$ of buyers, we define the demand set of S in W as*

$$D_S(W) = \left\{ j \middle| \exists i \in S : j \in \mathrm{argmax}_{j' \in J}\, \mathrm{u}^{j'}_i(\mathbf{p}^{j'}) \right\}.$$

Similarly, for each subset $T \subseteq J$ of goods, we define the demand set of T in W as

$$D^T(W) = \left\{ i \middle| \exists j \in T : i \in \mathrm{argmax}_{i' \in I}\, \mathrm{p}^j_{i'}(\mathbf{u}_{i'}) \right\}.$$

Theorem 4 (Tightness)**.** *Consider a market $M = (I, J, \mathrm{u})$, and let W be a CE of M.*

- *W is the lowest CE of M iff $|D^T(W)| \geq |T| + 1$ for every subset $T \subseteq J$ of goods with strictly positive prices (i.e., at least $|T| + 1$ buyers are (weakly) interested in T).*
- *W is the highest CE of M iff $|D_S(W)| \geq |S| + 1$ for every subset $S \subseteq I$ of buyers with strictly positive utilities (i.e., buyers in S are (weakly) interested in at least $|S| + 1$ goods).*

[8] This statement is not true if M does not have a $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE.

To get a better intuition of Theorem 4 suppose W is the lowest CE of a market M and suppose there is a subset T of goods with strictly positive prices for which the statement of the theorem fails to hold (i.e. $|D^T(W)| < |T| + 1$); then, we could conceptually decrease the prices of goods in T down to the point where either a buyer out of $D^T(W)$ becomes indifferent between her current allocation and some good in T; or one of the goods in T hit the price of 0. But then, we get a CE lower than W which contradicts W being the lowest CE. Despite the easy intuition the formal proof of this statement is more involved.

The following is a direct consequence of Theorem 4.

Theorem 5. *Consider a market $M = (I, J, \mathrm{u})$, and let $W = (\mathbf{u}, \mathbf{p})$ be a CE of M.*

- *If and only if W is the highest CE of M, then $W^{-j} = (\mathbf{u}, \mathbf{p}^{-j})$ is also a CE for the market M^{-j} for every $j \in J$.*
- *If and only if W is the lowest CE of M, then $W_{-i} = (\mathbf{u}_{-i}, \mathbf{p})$ is also a CE for the market M_{-i} for every $i \in I$.*

Intuitively, the above theorem says that at the highest CE, after removing any single good from the market, the assignments can be modified to get a CE for the rest of the market with the same prices/same utility as before; similarly, at the lowest CE, after removing any single buyer i from the market, the assignments can be modified to get a CE for the rest of the market with the same prices/same utilities.

Proof. We only prove the first statement since the proof of the second one is similar (completely symmetric). Consider a supporting matching $\boldsymbol{\mu}$ for W. If either j is unmatched or the utility of buyer who is matched to j is 0 we are done. Otherwise, we run the following process during which we maintain a subset S of buyers with strictly positive utilities such that from each buyer $i \in S$ there is an alternating sequence of buyers/goods of the form $j'_1, i'_1, j'_2, i'_2, \ldots, j'_k, i'_k$ (of some length k) where $j'_1 = j$ and $i'_k = i$ and such that i'_r is indifferent between j'_r and j'_{r+1} and $\boldsymbol{\mu}(i'_r) = j'_r$, for every $r \in [k]$. We initialize S to be the singleton containing the buyer that is matched to j. Since all buyers in S have strictly positive utilities, they must be weakly interested in at least $|S| + 1$ goods. So there exists a buyer $i^* \in S$ who is weakly interested in a good j^* that is not currently matched to any of the buyers in S; if j^* is itself matched to another buyer with positive utility, we add that buyer to S and repeat; otherwise we assign j^* to i^* and switch the assignments along the alternating path from i^* to j which yields a supporting matching for W^{-j}. Note that the above process always finds such a an alternating path in at most $|I| - 1$ iterations. The "only if" direction is trivial by applying Theorem 4. □

Next, we show that the lowest CE is group strategyproof for buyers, i.e., that no group of buyers can collude in such a way that they all get strictly higher utilities (assuming no side payments).

Theorem 6 (Group Strategyproofness)**.** *Given a market $M = (I, J, \mathrm{u})$, a mechanism that solicits buyer's utility functions and computes the lowest CE of M is group strategyproof. Formally, for every subset $S \subseteq I$ of buyers who collude and misreport their utility functions, if $\underline{W'}$ denotes the lowest CE with respect to the reported utility functions, then there is at least one buyer $i \in S$ whose utility at $\underline{W'}$ is no better than her utility at $\underline{W}$.*

Proof. The proof is by contradiction. Let S be the largest subset of buyers who can collude and possibly misreport their utility functions such that all of them obtain strictly higher utilities. Let $\underline{W}$ be the lowest CE of M with respect to the true utility functions and let $\underline{W}'$ be the lowest CE with respect to the reported utility functions assuming that buyers is S have colluded. Let T be the subset of the goods that are matched to S at $\underline{W}'$. Since all the buyers in S are achieving strictly higher utilities at $\underline{W}'$, they must all be matched at $\underline{W}'$ (i.e. $|T| = |S|$) and the prices of the goods in T should be strictly lower at $\underline{W}'$. That means the goods in T must have had strictly positive prices in $\underline{W}$. By applying Theorem 4, we argue that there must have been a subset S' of buyers of size at least $|T| + 1$ who were weakly interested in some of the goods in T at $\underline{W}$. Consequently all of the buyers in S' must be getting a strictly higher utilities at $\underline{W}'$ because the prices of all the goods in T are strictly lower. But S' is larger than S which contradicts our assumption that S was the largest set of buyers who could all benefit from collusion. $\square$

Finally, we preset the proof of our main theorem.

Proof of Theorem 2. We only prove (2.I) and (2.III). The proofs of (2.II) and (2.IV) are completely symmetric to the other two.

The plan of the proof is as follows. We remove an arbitrary buyer i from the market and compute the highest CE of the rest of the market. We then show that the prices at the highest CE of the market without i leads to a valid CE for the whole market (including buyer i) but with a possibly different matching. We also show that the induced utility of buyer i from these prices is the same as her utility at the lowest CE of the whole market[9]. The detail of the construction is as follows.

Choose an arbitrary buyer $i \in I$. Let M_{-i} denote the market without buyer i and let $\overline{W}_{-i}$ be the highest CE of the market M_{-i}. Note that the market M_{-i} is of size $|I| + |J| - 1$ so by inductively applying Theorem 2 to M_{-i} — which is of a strictly smaller size — we can argue that there exists a CE for M_{-i}, thus the highest CE of M_{-i} is also well-defined. Let $\mathbf{p} = \mathrm{p}(\overline{W}_{-i})$ be the prices at $\overline{W}_{-i}$. We claim that using the prices $\mathbf{p}$ for the market M leads to a valid CE, which we denote by W. In particular, all the prices/utilities at W are the same as the prices/utilities at $\overline{W}_{-i}$ and the utility of buyer i is $\mathrm{u}_i(\mathbf{p})$, however the matching of goods/buyers might be different. To obtain a supporting matching for W, we start with a supporting matching for $\overline{W}_{-i}$ and modify it as follows. If $\mathrm{u}_i(\mathbf{p}) = 0$ then we can leave buyer i unmatched and the matching does not need to be changed. Otherwise, let j be the good from which buyer i achieves her highest utility at the current prices, i.e. $j \in \mathrm{argmax}_{j' \in J}\, \mathrm{u}_i^{j'}(\mathbf{p}^{j'})$. We match j to buyer i and apply Theorem 5 to the market M_{-i} to conclude that after removing good j the matching for the rest of the market can be modified appropriately to support a CE with the same prices/utilities as before. Next, we prove each part of the theorem statement.

- Proof of (2.III) (case 1) $\boxed{\mathrm{p}^j(\underline{W}) \le \mathrm{p}^j(\overline{W}_{-i})}$: Notice that W is a CE of market M which has the same prices as the $\overline{W}_{-i}$. The prices at the lowest CE of M are no more than the prices at W so $\mathrm{p}(\underline{W}) \le \mathrm{p}(W) = \mathrm{p}(\overline{W}_{-i})$.

[9]Note that the resulting CE is not necessarily the lowest CE of M. Only the utility of buyer i is equal to his utility at the lowest CE of M.

- Proof of (2.III) (case 2) $\boxed{\mathrm{p}^j(\underline{W}) = \mathrm{p}^j(\overline{W}_{-i}) \textbf{ when } \boldsymbol{\mu}(i) = j}$: Notice that since i and j are matched, if we remove both of them the rest of $\underline{W}$ is still a valid CE for M_{-i}^{-j}. Let $\underline{W}_{-i}^{-j}$ denote the lowest CE of M_{-i}^{-j}. Note that both $\underline{W}$ and $\underline{W}_{-i}^{-j}$ are valid CEs for M_{-i}^{-j} but $\underline{W}_{-i}^{-j}$ is the lowest, so the prices of goods $J - \{j\}$ might only be lower at $\underline{W}_{-i}^{-j}$ and so the utilities of buyers $I - \{i\}$ might only be higher at $\underline{W}_{-i}^{-j}$ and so the price induced by buyers $I - \{i\}$ on good j might only be lower at $\underline{W}_{-i}^{-j}$ than the price induced by them on good j at $\underline{W}$. However, by applying Theorem 2 inductively on market M_{-i} — which is strictly smaller — and using (2.II), we get that $\mathrm{p}^j(\overline{W}_{-i})$ is exactly the induced price of buyers $I - \{i\}$ on good j at $\underline{W}_{-i}^{-j}$. Therefore, $\mathrm{p}^j(\overline{W}_{-i})$ must be less than or equal to the induced price on good j at $\underline{W}$ which is itself less that or equal to $\mathrm{p}^j(\underline{W})$. On the other hand, from the previous paragraph we have $\mathrm{p}^j(\underline{W}) \leq \mathrm{p}^j(\overline{W}_{-i})$, so the two must be equal.

- Proof of (2.I) $\boxed{\mathrm{u}_i(\underline{W}) = \mathrm{u}_i(\mathrm{p}(\overline{W}_{-i}))}$: If i is matched with j in $\underline{W}$ then $\mathrm{u}_i = \mathrm{u}_i^j(\mathrm{p}^j(\underline{W}))$ and by the previous statement $\mathrm{p}^j(\underline{W}) = \mathrm{p}^j(\overline{W}_{-i})$. Therefore $\mathrm{u}_i(\underline{W}) = \mathrm{u}_i^j(\mathrm{p}^j(\overline{W}_{-i})) = \mathrm{u}_i(\mathrm{p}(\overline{W}_{-i}))$. The last equality follows from the fact that we chose j to be the good from which buyer i obtains her highest utility at prices $\mathrm{p}(\overline{W}_{-i})$.

□

5 Application to Ad-Auctions

In this section, we present a truthful mechanism for Ad-auctions that combines pay per click (a.k.a charge per click (CPC)) advertisers and pay per impression (a.k.a charge per impression (CPM)) advertisers, potentially with non-quasilinear utility functions. In particular, our mechanism is welfare maximizing and group strategyproof regardless of whether the search engine and the advertisers have consistent beliefs about the clickthrough rates.

Model Given a set of advertisers I and a set of slots J, the utility of advertiser i for a click on her ad being displayed on slot j at a price of x is given by $\mathrm{u}_i^j(x)$ which is a continuous and decreasing function in x [10]. Note that x specifies payment per click for a CPC advertiser, or payment per impression for a CPM advertiser. We say that a CPC advertiser i has standard utility function if for all slots j: $\mathrm{u}_i^j(x) = c_i^j(v_i^j - x)$ in which v_i^j is the advertiser's value for a click on slot j and c_i^j is the advertiser's estimate about her clickthrough rate (CTR), i.e., the probability that her ad is being clicked on if displayed on slot j. We also say that a CPM advertiser i has standard utility function if for all slots j: $\mathrm{u}_i^j(x) = v_i^j - x$ in which v_i^j is the advertiser's value for a click on slot j. $\hat{c}_i^j$ will denote the search engine's estimate of the CTR of advertiser i on slot j; this could potentially be different than c_i^j (i.e., advertisers and search engine may have different beliefs). Furthermore, v_i^j and c_i^j are advertiser's private information but $\hat{c}_i^j$ is publicly available.

Consider what happens if we applied VCG payments in this environment assuming there are only CPC advertisers with standard utility function. The first of the following linear

[10]We assume that $\mathrm{u}_i^j(x)$ reaches 0 for a high enough x.

programs compute the welfare maximizing allocation (based on advertisers' reports), and its dual computes prices/utilities.

(Primal)

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in I} \sum_{j \in J} c_i^j v_i^j y_i^j \\ \text{subject to} \quad & \textstyle\sum_{j \in J} y_i^j \leq 1, \quad \forall i \in I, \\ & \textstyle\sum_{i \in I} y_i^j \leq 1, \quad \forall j \in J, \\ & y_i^j \geq 0 \end{aligned}$$

(Dual)

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in I} \mathbf{u}_i + \sum_{j \in J} \mathbf{p}^j \\ \text{subject to} \quad & \mathbf{u}_i + \mathbf{p}^j \geq c_i^j v_i^j, \quad \forall i \in I, \forall j \in J \\ & \mathbf{u}_i \geq 0 \\ & \mathbf{p}^j \geq 0 \end{aligned}$$

The set of solutions to the dual program would be the set of CEs of the market and the one with the lowest prices would correspond to the lowest CE which would also coincide with the VCG payments/utilities. However, the problem is that payments must be charged *per click* while $\mathbf{p}^j$ represents the expected payment *per impression*. So, per click payments would be given by dividing $\mathbf{p}^j$ by the probability of a click which is the CTR. Observe that dividing by c_i^j breaks the strategyproofness guarantee because c_i^j is reported by the advertiser and reporting a higher c_i^j would give the advertiser a higher chance of winning a better slot while at the same time it would lower the payment. Dividing by $\hat{c}_i^j$ also breaks the strategyproofness when $\hat{c}_i^j$ and c_i^j are different. Therefore, the straightforward application of VCG payments fails.

Consider the Ad-Auction problem as a two sided matching market with ads on one side and slots on the opposite side. For a CPC advertiser i, allocating slot j and charging a payment of x per click, yields an expected utility of $\mathrm{u}_i^j(x)$ for the advertiser and generates an expected revenue of $x/\hat{c}_i^j$ per impression for the search engine (from the perspective of the search engine based on its own estimated CTRs). For a CPM advertiser i, allocating slot j and charging a payment of x per impression, yields an expected utility of $\mathrm{u}_i^j(x)$ for the advertiser and generates a revenue of x per impression for the search engine. Consider a CE in this market (treat each slot as an independent agent even though all slots are owned by the search engine). Our proposed mechanism solicits advertisers' utility functions and outputs as its outcome the CE that has the highest advertiser utilities.

Definition 6. *(Ad-Auction Competitive Equilibrium) The mechanism solicits advertisers' utility functions* u_i^j *and defines* $M' = (I, J, \mathrm{u}')$*, where for a CPC advertiser the utility function is given by* $\mathrm{u'}_i^j(x) = \mathrm{u}_i^j(x/\hat{c}_i^j)$*, and for a CPM advertiser the utility function is given by* $\mathrm{u'}_i^j(x) = \mathrm{u}_i^j(x)$*. Compute the lowest CE of* M'*, call it* $\underline{W} = (\mathbf{u}, \mathbf{p})$*, and let* $\boldsymbol{\mu}$ *be a supporting matching. Allocate to advertiser i the slot* $\boldsymbol{\mu}(i)$*. If i is a CPC advertiser, charge her* $\mathbf{p}^{\boldsymbol{\mu}(i)}/\hat{c}_i^j$ *if there is a click on the slot; otherwise i is a CPM advertiser and should be charged* $\mathbf{p}^{\boldsymbol{\mu}(i)}$ *for an impression.*

Recall that M' is the reduced form of the original market which is obtained by redefining the utility functions of one side in terms of the utilities of the other side. Observe that $\underline{W}$ satisfies the following inequality for each CPC advertiser i:

$$\mathbf{u}_i = \max_{j \in J} \mathrm{u}_i^j(\mathbf{p}^j/\hat{c}_i^j) \qquad \forall j \in J.$$

it also satisfies the following inequality for each CPM advertiser i:

$$\mathbf{u}_i = \max_{j \in J} \mathrm{u}_i^j(\mathbf{p}^j) \qquad \forall j \in J.$$

$\mathbf{p}^j$ can be interpreted as the *virtual price* of slot j, where as $\mathbf{p}^j/\hat{c}_i^j$ can be interpreted as the weighted price of slot j for a CPC advertiser i. Observe that the virtual price $\mathbf{p}^j$ represents the expected per impression payment required for slot j computed from the perspective of the search engine. Notice that the above inequalities ensure that each advertiser receives their preferred slot according to their weighted prices.

The above mechanism can also be conceptually reinterpreted as an ascending auction in which all virtual prices $\mathbf{p}_i^j$ start from 0 and are gradually increased as long as there is excess demand; and such that at any time during the auction a CPC advertiser i observes a price of $\mathbf{p}^j/\hat{c}_i^j$ for slot j whereas a CPM advertiser i observes a price of $\mathbf{p}^j$ for slot j.

Next theorem summarizes the important properties of the above mechanism.

Theorem 7. *Mechanism of Definition 6 is group strategyproof and also maximizes the social welfare in the following sense. Let S be a group of advertisers with standard utility functions who also agree with the search engine on the CTRs and let S' be the rest of the advertisers. Let s denote the search engine. Welfare of $\{s\} \cup S$ is always maximized. In particular, if all advertisers have standard utility functions and agree with the search engine on the CTRs, then the outcome of this mechanism coincides exactly with the VCG outcome.*

Notice that mechanism 6 is group strategyproof regardless of whether the search engine and advertisers have the same estimates about the clickthrough rates.

6 Acknowledgement

We thank Rakesh Vohra, Lawrence Ausubel, John Hatfield, Scott Kominers, Jason Hartline, Nicole Immorlica, Sebastien Lahaie, David Easley, Larry Blume and Martin Pal for helpful discussions and/or comments.

References

Alkan, A. and Gale, D. (1990). The core of the matching game. *Games and Economic Behavior*, 2(3):203–212.

Ausubel, L. M. and Milgrom, P. (2002). Ascending auctions with package bidding. Working Papers 02004, Stanford University, Department of Economics.

Caplin, A. and Leahy, J. V. (2010a). Comparative statics in markets for indivisible goods. Working paper, National Bureau of Economic Research.

Caplin, A. and Leahy, J. V. (2010b). A graph theoretic approach to markets for indivisible goods. Nber working papers.

Demange, G. and Gale, D. (1985). The strategy structure of two-sided matching markets. *Econometrica*, 53(4):873–88.

Demange, G., Gale, D., and Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, 94(4):863–72.

Gale, D. (1984). Equilibrium in a discrete exchange economy with money. *International Journal of Game Theory*, 13(1):61–64.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. volume 69, pages 9–15.

Gul, F. and Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124.

Hatfield, J. W. and Kominers, S. D. (2010). Matching in networks with bilateral contracts: extended abstract. In *Proceedings of the 11th ACM conference on Electronic commerce*, EC '10.

Hatfield, J. W. and Milgrom, P. R. (2005). Matching with contracts. *American Economic Review*, 95(4):913–935.

Kaneko, M. and Yamamoto, Y. (1986). The existence and computation of competitive equilibria in markets with an indivisible commodity. *Journal of Economic Theory*, 38(1):118–136.

Kelso, Alexander S, J. and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504.

Knaster, B., Kuratowski, C., and Mazurkiewicz, S. (1929). Ein beweis des fixpunktsatzes fr n-dimensionale simplexe. *Fund. Math*, 14:132–137.

Kuhn, H. W. (1956). Variants of the hungarian method for assignment problems. page 253258.

Lemke, C. E. and Howson, J. T., J. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423.

Leonard, H. B. (1983). Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–79.

Mamer, J.and Bikhchandani, S. (1997). *Journal of Economic Theory*, 74:385–413.

Mani, A., Ozdaglar, A. E., and Pentland, A. (2010a). Existence of stable exclusive bilateral exchanges in networks. *CoRR*.

Mani, A., Ozdaglar, A. E., and Pentland, A. (2010b). Generalized stable matching in bipartite networks. *CoRR*.

Quinzii, M. (1984). Core and competitive equilibria with indivisibilities. *International Journal of Game Theory*, 13:41–60.

Scarf, H. E. (1967). The core of an n person game. *Econometrica*, 35(1).

Shapley, L. and Shubik, M. (1971). The assignment game i: The core. 1(1):111130.

A Other Results and Omitted Proofs

This section presents the proofs of the theorems which were omitted from the previous sections along with a few lemmas which are used in those proofs.

Lemma 1 (Entanglement). *Consider a market $M = (I, J, \mathrm{u})$. If there exists a CE of M at which buyer i is matched with good j, then at any other CE of M the price of good j is higher if and only if the utility of buyer i is lower and vice versa. Note that this statement is true regardless of wether buyer i and good j are actually matched to each other in other CEs.*

Lemma 2 (Conservation of Matching). *Given a market $M = (I, J, \mathrm{u})$, for any $i \in I$, if there exists a CE of M at which buyer i has a strictly positive utility, then buyer i is never unmatched in any CE of M. Similarly, for any $j \in J$, if there exists a CE of M at which good j has a strictly positive price, then good j is never unmatched at any CE of M.*

Proof of Lemma 1 and Lemma 2. Let W be a CE of M at which the hypothesis of the lemma holds, and let W' be any other CE of M. Partition the buyers to three groups S, S', S'', such that buyers in S have higher utilities at W, buyers in S' have higher utilities at W', and buyers in S'' have the same utilities at W and W'. Similarly, partitions the goods to T, T', T'', such that goods in T have higher prices at W, goods in T' have higher prices and W', and goods in T'' have the same prices at both W and W'. The following statements are easy to derive using the definition of CE and using the fact that both W and W' are CEs:

- At W, all buyers in S must be matched to goods in T' so $|S| \leq |T'|$.
- At W', all goods in T' must be matched to buyers in S so $|T'| \leq |S|$.

From the above statement, we can conclude $|S| = |T'|$ and buyers in S and goods in T' must be matched to each other at both equilibria. Similarly:

- At W, all goods in T must be matched to buyers in S' so $|T| \leq |S'|$.
- At W', all buyers in S' must be matched to goods in T so $|S'| \leq |T|$.

So, we can conclude $|S'| = |T|$, therefore buyers in S' and goods in T must be matched to each other at both equilibria. Furthermore, we can then conclude that buyers in S'' and goods in T'' may only be matched to each other. That proves the claim of Lemma 1. To complete the proof of Lemma 2 observe that any good/buyer that has positive price/utility at W either has a positive price/utility at W' in which case it must also be matched at W', or has a 0 price/utility at W' in which case it must be in T/S and therefore it must also be matched at W'. □

Proof of Theorem 3. We only prove the first claim. The proof of the second claim is similar (completely symmetric). The plan of the proof is as follows:

First, we define a new market $M' = (I, J, \mathrm{u}')$ with transformed utility functions $\mathrm{u}'^j_i(x) = \mathrm{u}^j_i(x + \underline{\mathbf{p}}^j) - \underline{\mathbf{u}}_i$. We claim that there is a one-to-one mapping between $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CEs of the original market and the CEs of the transformed market. Formally, $\overline{W} = (\mathbf{u}, \mathbf{p}, \boldsymbol{\mu})$ is a $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M if and only if $W' = (\mathbf{u} - \underline{\mathbf{u}}, \mathbf{p} - \underline{\mathbf{p}}, \boldsymbol{\mu})$ is a CE of M'. We then show that there is CE of M' in which there is a good with a price of 0 which then means in the corresponding CE of the original market the price of that good is equal to its lower bound and therefore at the lowest $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M the price of that good must also be equal to its lower bound which proves the claim.

We now prove that there is a good with a price of 0 at the lowest CE of M'. We choose an arbitrary buyer i from M' and remove it from the market. Let $\overline{W}_{-i}$ be the highest CE of the remaining market. By the assumption of the lemma, we know $|I| = |J|$ and so in $\overline{W}_{-i}$ there are more goods than there are buyers so there must be an unmatched good which we denote by j^*. Note that the price of j^* in $\overline{W}_{-i}$ must be 0. On the other hand, by applying Theorem 2 to M' and using (2.III) we have $\mathrm{p}^j(\underline{W}) \leq \mathrm{p}^j(\overline{W}_{-i})$ for every good j. Therefore, it must be that the price of j^* in $\underline{W}$ is also 0 and that completes the proof.

There is a subtlety that we should point out about the one-to-one mapping between the CEs of the original market and those of the transformed market. It is clear that every $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M can be transformed to a CE of M'. However, for the other direction, we need to show that all goods/buyers are matched, otherwise after applying the inverse transform we may end up with an unmatched good/buyer that has a positive price/utility. To show that all buyers/goods are matched in every CE of M', we can apply Lemma 2. To apply that lemma, we only need to show that there is a CE of M' in which all goods have strictly positive prices and then by that lemma all the goods must always be matched (and so do all buyers because $|I| = |J|$). Notice that if there is no CE for M' in which all goods have strictly positive prices then either in every $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M there is a good whose price is equal to its lower bound or M has no $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE at all which either way trivially proves the claim of this lemma. □

Proof of Theorem 4. We only prove the first statement. The proof of the second statement is similar (completely symmetric).

First, we prove the "only if" direction. Assume that W is the lowest CE of M. For every subset T of goods with strictly positive prices, we prove that $|D^T(W)| \geq |T| + 1$, i.e. there are at least $|T| + 1$ buyers who are interested in some good in T. The proof is as follows. Since all the goods in T have strictly positive prices, they must all be matched. Let S be the subset of buyers that are matched to T. Notice that $S \subset D^T(W)$ and $|S| = |T|$. Therefore, to complete the proof we only need to show that there is one more buyer not in S who is also interested in a good in T. Let $\underline{\mathbf{p}}$ be the prices induced by utilities of buyers not in S, i.e. $\underline{\mathbf{p}}^j = \max_{i \in I-S} \mathrm{p}^j_i(\mathrm{u}_i(W))$. Similarly, let $\underline{\mathbf{u}}$ be the utilities induced by the prices of goods not in T, i.e. $\underline{\mathbf{u}}_i = \max_{j \in J-T} \mathrm{u}^j_i(\mathrm{p}^j(W))$. Notice that W is a $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of the market $M' = (S, T, \mathrm{u})$. Furthermore, if we replaced the part of W corresponding to S and T with any other $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M', we still get a valid CE for M, but that implies W must already be the lowest $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded CE of M' as well because otherwise we could replace the part of W corresponding to S and T with the lowest $(\underline{\mathbf{u}}, \underline{\mathbf{p}})$-bounded

CE of M' and get a lower CE for M which would contradict W being the lowest CE of M. By applying Theorem 3 to the market M', we can argue that there is a good $j^* \in T$ such that $\mathrm{p}^{j^*}(W) = \underline{\mathbf{p}}^{j^*}$. Since all the goods in T, including j^*, have strictly positive prices, $\underline{\mathbf{p}}^{j^*}$ must also be strictly positive and because of the way we defined $\underline{\mathbf{p}}$ there must be a buyer i^* not in set S such that $\underline{\mathbf{p}}^{j^*} = \mathrm{p}_{i^*}^{j^*}(\mathbf{u}_{i^*})$. That means i^* must be weakly interested in good j^* and therefore $\{i^*\} \cup S \subset D^T(W)$ which proves that $D^T(W) \geq |T| + 1$.

The proof of the "if" direction is trivial. The proof is by contradiction. Let W be a CE of M such that for every subset T of goods with strictly positive prices we have $D^T(W) \geq |T| + 1$. Let $\underline{W}$ be the lowest CE of M and assume that W and $\underline{W}$ are not the same. Let T consist of all the goods that have a higher price at W compared to $\underline{W}$. We know that there are at least $|T|+1$ buyers interested in T at W and these buyers must have higher utilities at $\underline{W}$ because the prices of the goods in T are strictly lower. Therefore, the goods assigned to these buyers at $\underline{W}$ must have lower prices and so there are at least $|T|+1$ goods that have higher prices at W compared to $\underline{W}$ which contradicts the assumption that T was the set of all the goods that had higher prices at W. □

Proof of Theorem 7. The group strategyproofness follows from Theorem 6. So we only prove the second part. Suppose the mechanism has computed a $\underline{W} = (\mathbf{u}, \mathbf{p})$ which is the lowest CE of M' as its outcome. Recall that $\mathbf{p}^j$ is the virtual price of slot j and the expected utility of advertiser i from slot j is given by $\mathrm{u'}_i^j(\mathbf{p}^j)$. So for each advertiser $i \in S$ we have $\mathrm{u'}_i^j(\mathbf{p}^j) = c_i^j(v_i^j - \mathbf{p}^j/\hat{c}_i^j)$. Furthermore, since $c_i^j = \hat{c}_i^j$, we can simplify the utility function and get $\mathrm{u}_i^j(\mathbf{p}^j) = c_i^j v_i^j - \mathbf{p}^j$. Now, consider the complete bipartite graph with advertisers and slots. Let the weight of each edge (i, j) be $c_i^j v_i^j$. Note that for each advertiser $i \in S$ we have $\mathrm{u}_i(\underline{W}) + \mathrm{p}^j(\underline{W}) \geq c_i^j v_i^j$ in which $\underline{W}$ is the outcome of the mechanism. Therefore, the total expected welfare of the coalition $\{s\} \cup S$ is at least as much as the weight of the maximum weight matching in the absence of S'. Furthermore, if S' is empty (i.e. everyone agrees on the CTRs), the mechanism computes the efficient allocation (i.e., a maximum weight matching) and the outcome is the same as the VCG outcome. □