CS364B: Frontiers in Mechanism Design
Lecture #2: Unit-Demand Bidders and Walrasian Equilibria*

Tim Roughgarden†

January 8, 2014

1 Bidders with Unit-Demand Valuations

1.1 The Setting

Last lecture we discussed two simple auction scenarios — identical items with unit-demand bidders, and non-identical items with bidders with additive valuations. In both scenarios, we gave an ascending auction that is ex post incentive compatible (EPIC) — meaning sincere bidding by every bidder is an ex post Nash equilibrium (EPNE) with guaranteed non-negative utility — and that maximizes the welfare, assuming sincere bidding (up to a discretization error). This lecture introduces a third scenario, which generalizes the first (and is incomparable to the second).

Scenario #3:

- A set U of m non-identical items.
- Each bidder i has a private valuation v_{ij} for each item j. The number n of bidders can be more or less than m.
- Each bidder i has *unit demand*, meaning its value for a bundle $S \subseteq U$ of items is

$$v_i(S) := \max_{j \in S} v_{ij}.$$

Scenario #1 (identical items) is the special case where, for every bidder i, v_{ij} is independent of j. Here, a bidder is only interested in acquiring one item, but it has preferences amongst them. Housing markets are often given as an example of unit-demand preferences. This model was first studied by Shapley and Shubik [3].

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

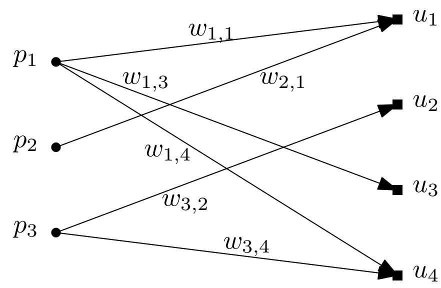

Figure 1: A graph instance of the maximum bipartite matching problem corresponding to 3 players and 4 items.

1.2 A Direct-Revelation Solution

This lecture and the next give an ascending auction for unit-demand bidders that is EPIC and that maximizes the welfare (assuming sincere bidding). We noted last lecture that a logical prerequisite to this is a DSIC direct-revelation mechanism that maximizes the welfare (assuming truthful bidding). This section provides such a direct-revelation mechanism.

As a thought experiment, suppose that all of the valuations v_{ij} are known. How could we compute a welfare-maximizing allocation of the items? The first observation is that, without loss of generality, we can assume that every bidder receives at most one item. Since every item can only be given to at most one player, welfare-maximization is exactly the *maximum-weight bipartite matching problem*. In more detail, define a weighted bipartite graph (A, B, E, w) where the vertex set A corresponds to the n bidders, the vertex set B corresponds to the m items, and for each $i \in A$ and $j \in B$ there is an edge $(i, j) \in E$ with weight v_{ij} (Figure 1). Bipartite matchings of weight W correspond to allocations of welfare W in which each bidder gets at most one item. Every maximum-weight matching yields a welfare-maximizing allocation. Recall that the bipartite matching problem can be solved in polynomial time, using linear programming or combinatorial algorithms (see e.g. [1]). We'll also give a particularly simple $(1-\epsilon)$-approximation algorithm for this problem in the next lecture.

Next we recall the *VCG mechanism*; see also Lecture #7 of CS364A. In scenario #3, this mechanism can be described as follows.

1. Collect a bid b_{ij} from each bidder i for each item j.
2. Compute an allocation corresponding to a maximum-weight bipartite matching, using the b_{ij}'s as edge weights.
3. Charge each bidder its externality — the welfare loss to the others caused by its presence.

Informally, the VCG payments align bidders' individual objectives with the welfare, resulting in a DSIC mechanism. Each bidder i's payment can be computed as the difference between the optimal solution values of two different maximum-weight matching problems (with and without i); we'll review the exact formula later, when we need it. The VCG mechanism can be defined very generally, and is always DSIC and welfare-maximizing; what's special about scenario #3 is that the VCG mechanism can be implemented in polynomial time.

2 High-Level Plan

Our goal in this lecture and the next is to develop an analog of the English auction for scenario #3. We're striving for the following properties.

1. **(Incentive guarantee.)** EPIC. Again, this means that sincere bidding is an ex post Nash equilibrium that guarantees all bidders nonnegative utility.

2. **(Performance guarantee.)** If all bidders bid sincerely, then the outcome of the auction coincides with that of the VCG mechanism.[1]

3. **(Tractability guarantee.)** The auction should be "simple" and terminate in a reasonable amount of time.

At a very high level, we're envisioning an ascending auction that initially sets the price of every item to be 0, and gradually raises prices until "supply equals demand." This is exactly what we did in the first two scenarios. In scenario #3, the meaning of "supply" is clear — one copy of each item. For a unit-demand bidder i, the right way to define its "demand" at prices $\mathbf{q}$ is as its favorite item — the item that maximizes $v_{ij} - q_j$ over items j. (Ignore ties for now.) If $q_j > v_{ij}$ for every item j, the demand of bidder i is the empty set. Then, if two or more bidders have the same favorite item j, then j is over-demanded. If j is nobody's favorite item, then j is under-demanded. "Supply equals demand" then translates to no item being overdemanded, and a item being underdemanded only if nobody wants it (even at price 0).

It is not obvious how to implement an ascending auction so that it is guaranteed to terminate with supply equaling demand. Further, we're shooting for a seemingly more ambitious goal — we want the auction to terminate with the same allocation and payments as the VCG mechanism. VCG payments, while computable in polynomial time as the difference between the optimal matchings of two different problems, seem fairly complex. Could the simple "computational model" of ascending auctions be powerful enough to terminate with these payments?

The answer, remarkably, is "yes." We'll prove this in two steps. In today's lecture, we characterize the VCG outcome in scenario #3 as the "smallest Walrasian equilibrium." Walrasian equilibria, defined in the next section, are in some sense the natural outcome of an ascending auction that aspires to equalize supply and demand. This characterization is important because it is far easier to envision how an ascending auction would compute

[1] We've strengthened the goal since last lecture, where we asked only for a VCG (i.e., welfare-maximizing) allocation. Here, we're also asking that the auction recreate the VCG payments. We'll see later that this stronger property is necessary for an EPIC implementation.

a Walrasian equilibrium than the VCG outcome. Next lecture, we give a simple EPIC ascending auction that is guaranteed to terminate at the smallest Walrasian equilibrium, which we'll then know is the same as the VCG outcome.

3 Walrasian Equilibria

3.1 Definitions and Examples

We now formally define Walrasian equilibria[2] in the context of scenario #3; we'll give more general definitions later.

Definition 3.1 For an instance of scenario #3, a *Walrasian equilibrium* is a price vector $\mathbf{q} \in \mathcal{R}^m$ defined on the item and a matching M of the bidders and items such that:

(WE1) each bidder i is matched to a favorite item

$$j \in \text{argmax}\{v_{ij} - q_j\}_{j \in U \cup \{\emptyset\}}; \tag{1}$$

(WE2) an item $j \in U$ is unsold only if $q(j) = 0$.

This definition does not concern incentives per se, and hence may seem like a detour from our quest for incentive-compatible auctions. The connection is that, with sincere bidding, natural ascending auctions terminate at a WE. For this reason, understanding WE are crucial for understanding what can and cannot be accomplished with ascending auctions.

We first make some technical comments about Definition 3.1, then provide an interpretation, and then give two examples. First, we allow the number n of bidders to be more or less than the number m of items. Thus, some bidders might receive no item and some items might remain unsold. For a matching M and a bidder i, we use $M(i)$ to denote the item to which i is matched; if i receives no item, we interpret $M(i)$ as $\emptyset$. Second, for visual clarity, we sometimes use the notation $v_i(M(i))$ and $q(M(j))$ instead of $v_{iM(i)}$ and $q_{M(i)}$. Third, when $M(i) = \emptyset$, we interpret $v_{iM(i)}$ and $q_{M(i)}$ as 0. Thus, a consequence of condition (WE1) is that a bidder i is matched to a item j only if $q_j \leq v_{ij}$. Fourth, the condition (WE1) by itself is uninteresting — setting all q_j's sufficiently large and M to the empty allocation satisfies this condition. Condition (WE2) asserts that the market clears on the supply side as well.

To interpret Definition 3.1, consider a *Walrasian price vector* $\mathbf{q}$, meaning a price vector that participates in some WE $(\mathbf{q}, M)$. Assume also that at prices $\mathbf{q}$ every bidder has a unique favorite item (possibly $\emptyset$). Then, $\mathbf{q}$ enables a remarkable distributed solution to a seemingly centralized optimization problem: put a price tag q_j on each item j and let each bidder independently grab its favorite item. The result will be the matching M — in particular, no item is over-demanded, and only items with price 0 are under-demanded. Bidders make no effort to coordinate, and yet their individual optimizations collective yield a feasible — and as we'll see below, optimal — allocation.

[2] These are also called "competitive equilibria."

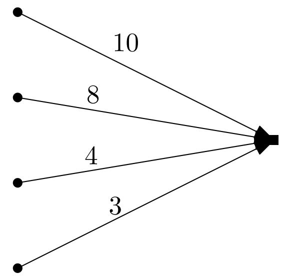

Figure 2: Single-item setting. All Walrasian price vectors are in $[8, 10]$.

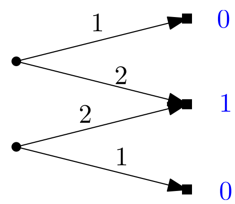

Figure 3: Multi-item setting. A Walrasian price vector is given on the right.

Example 3.2 (Single-Item Setting) Consider a single-item auction with four bidders, which corresponds to the bipartite graph in Figure 2. Edges are labeled with bidders' valuations. The Walrasian price vectors in this example are precisely the numbers in the interval $[8, 10]$. Thus, we already see that WE are not unique. To foreshadow more general results to come, observe that, when valuations are private, the DSIC Vickrey and English auctions compute the smallest of these prices. Computing the largest of these prices in some sense corresponds to a first-price auction, which of course is not DSIC.

Example 3.3 (Multi-Item Setting) Suppose there are two bidders and three items; see Figure 3. Missing edges correspond to zero valuations. We claim that setting $q(1) = q(3) = 0$ and $q(2) = 1$ yields a Walrasian price vector. To prove it, we need to exhibit a matching M such that $(\mathbf{q}, M)$ satisfies Definition 3.1. Observe that the first bidder is indifferent between the first two items, and the second bidder is indifferent between the second two items. Setting $M_1(1) = 1$ and $M_1(2)$ yields a matching such that $(\mathbf{q}, M_1)$ is a WE; the matching M_2 with $M_2(1) = 2$ and $M_2(2) = 3$ also works. Note that the matching M_3 with $M_3(1) = 1$ and $M_3(2) = 3$ does *not* work; the reason is that item 2 is unsold and has a positive price, which violates condition (WE2). Note that M_1 and M_2 are welfare-maximizing matchings, while M_3 is not.

We argued above that, without ties in bidders' demands, a Walrasian price vector obviates the need for coordination among bidders. This example shows that, when there are ties, coordination is needed to resolve them properly. Two things could go wrong if the bidders in this example make their choices independently. One, the second item might go unsold at a positive price. Two, both bidders might try to grab the second item, which does not induce a feasible allocation.

3.2 The First Welfare Theorem

In Example 3.3, only maximum-welfare matchings participate in WE. This property holds much more generally.

Proposition 3.4 (First Welfare Theorem) *In scenario #3, if* $(\mathbf{q}, M)$ *is WE, then* M *is a welfare-maximizing allocation.*

If one thinks of a WE as the natural outcome of a market, then Proposition 3.4 can be interpreted as saying "markets are efficient."[3] There are many "First Welfare Theorems," and all have this flavor. Proposition 3.4 is encouraging for goal of computing the VCG outcome (allocation and payments) using an ascending auction: as long as the auction terminates at a WE, we are guaranteed to compute the VCG (i.e., welfare-maximizing) allocation.

The proof is not difficult.

Proof of Proposition 3.4: Let M^* be a welfare maximizing allocation. Let Q denote the sum $\sum_{j \in U} q(j)$ of item prices. For every bidder i, we can invoke property (WE1) of the WE $(\mathbf{q}, M)$ to argue that i prefers item $M(i)$ over item $M^*(i)$ at the prices $\mathbf{q}$:

$$v_i(M(i)) - q(M(i)) \geq v_i(M^*(i)) - q(M^*(i)).$$

Summing this inequality over all bidders i yields

$$\underbrace{\sum_{i=1}^{n} v_i(M(i))}_{\text{welfare of } M} - \underbrace{\sum_{i=1}^{n} q(M(i))}_{= Q \text{ by (WE2)}} \geq \underbrace{\sum_{i=1}^{n} v_i(M^*(i))}_{\text{welfare of } M^*} - \underbrace{\sum_{i=1}^{n} q(M^*(i))}_{\leq Q},$$

where we are using the fact that $\sum_{i=1}^{n} q(M(i))$ sums over all of the items that have a non-zero price. Rearranging terms shows that the welfare of M is at least that of M^*, and hence M is also welfare-maximizing. ■

The ascending auctions we describe will actually terminate with an *ϵ-Walrasian equilibrium (ϵ-WE)*, meaning that the condition (1) in (WE1) only holds up to an ϵ. The approximate version of the First Welfare Theorem asserts that if $(\mathbf{q}, M)$ is an ϵ-WE, then the welfare of M is within $\epsilon \cdot \min\{n, m\}$ of the maximum possible (see Exercises).

3.3 Characterization of VCG Payments

Thus far, we've related the allocations that participate in a WE to those computed by the VCG mechanism — both are the welfare-maximizing allocations. We next relate the VCG payments to WE price vectors. The former are uniquely defined, given the allocation, while the latter are not (Example 3.2). The best we can hope for is that there is always a "special" WE price vector that coincides with the VCG payments. This is indeed the case.

[3] Needless to say, much blood and ink has spilled over this interpretation over the past couple of centuries.

Let's now recall how VCG payments are computed. For each bidder i, we perform a counterfactual computation: had bidder i not shown up, how much welfare would have been enjoyed by the other $n-1$ bidders? Bidder i's payment is the difference between this and the actual welfare realized by the other $n-1$ bidders in the VCG outcome. That is, a bidder pays the welfare loss to others for which it is responsible.

In the context of scenario #3, for each bidder i we compare two matchings: the VCG outcome M, and the welfare-maximizing matching M^{-i} that leaves i unmatched. The VCG payment of i is then defined as

$$p_i = \sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} v_k(M(k)). \tag{2}$$

Observe that this number is always nonnegative and at most i's value $v_i(j)$ for the good j that it is assigned in M.

VCG payments are defined for bidders, not for items. In scenario #3, however, the VCG outcome naturally induces a vector $\mathbf{p}$ of item prices. If a good j is unsold, we define $p(j) = 0$; otherwise, we define it as the price paid (2) by the (unique) winner i of the good j

Theorem 3.5 (VCG Payments Lower Bound WE) *In an instance of scenario #3, let* $\mathbf{p}$ *denote the induced item price vector of the truthful-revelation VCG outcome and* $\mathbf{q}$ *a Walrasian price vector. Then,* $p(j) \leq q(j)$ *for every item* j*.*

Proof: The proof uses some of the same ideas as in the First Welfare Theorem (Proposition 3.4). Let Q denote the sum of prices $\sum_{j \in U} q(j)$ in $\mathbf{q}$. Let M denote the allocation computed by the VCG mechanism. For unsold items, $p(j) = 0$ by definition and there is nothing to prove.

Consider an item j allocated to a bidder i. With an eye toward (2), let M^{-i} denote a welfare-maximizing allocation among allocations that leave i unmatched.

The pair $(\mathbf{q}, M)$ is a WE.[4] For every bidder $k \neq i$, we can invoke property (WE1) of the WE $(\mathbf{q}, M)$ to argue that k prefers item $M(k)$ over item $M^{-i}(k)$ at the prices $\mathbf{q}$:

$$v_k(M(k)) - q(M(k)) \geq v_k(M^{-i}(k)) - q(M^{-i}(k)).$$

Summing this inequality over all bidders $k \neq i$ yields

$$\underbrace{\sum_{k \neq i} v_k(M(k))}_{} - \underbrace{\sum_{k \neq i}^{n} q(M(k))}_{= Q - q(j) \text{ by (WE2)}} \geq \underbrace{\sum_{k \neq i}^{n} v_k(M^{-i}(k))}_{} - \underbrace{\sum_{k \neq i}^{n} q(M^{-i}(k))}_{\leq Q},$$

[4] This requires a proof when there are multiple welfare-maximizing allocations. Our current assumptions are that M is a welfare-maximizing allocation and that there is a welfare-maximization allocation M' such that $(\mathbf{q}, M')$ is a WE. We leave as an exercise the proof that these assumptions imply that $(\mathbf{q}, M)$ is also a WE. (Hint: make use of the First Welfare Theorem.)

where we are using the fact that $\sum_{k \neq i} q(M(k))$ sums over all of the items with non-zero q-prices save for the item matched to bidder i (namely, j). Rearranging terms yields

$$q(j) \geq \sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} v_k(M(k)) = p(j),$$

where the equation follows from the definition (2) of the item prices induced by the VCG mechanism. ■

The item prices induced by the VCG mechanism are component-wise only lower than every Walrasian price vector. Theorem 3.5 is complemented by the next result, the culmination of this lecture.

Theorem 3.6 (VCG Outcome Is a WE) *In an instance of scenario #3, let M and $\mathbf{p}$ denote the allocation and induced item price vector of the truthful-revelation VCG outcome. Then, $(\mathbf{p}, M)$ is a WE.*

Theorems 3.5 and 3.6 have several remarkable corollaries, each stronger than the previous one.

1. In scenario #3, a WE is guaranteed to exist [3].
2. In scenario #3, there is a "smallest" WE — one that is component-wise smaller than any other WE [3].[5]
3. In scenario #3, the VCG outcome coincides with the smallest WE [2].

Recall that our current goal is to design an EPIC auction such that sincere bidding yields the VCG outcome. The third corollary allows us to replace "VCG outcome" with "smallest WE" in this goal. It is this version of the goal that we'll accomplish next lecture.

Our proof of Theorem 3.6 relies on the following lemma, which is interesting in its own right. It shows that, in scenario #3, VCG payments can be equally well characterized by a different set of counterfactual computations: the welfare increase that would result from duplicating a good.

Lemma 3.7 *In an instance of scenario #3, let M and $\mathbf{p}$ denote the allocation and induced item price vector of the truthful-revelation VCG outcome. For a good $j \in U$, let M^{+j} denote a welfare-maximizing allocation after adding a second copy j' of the good j (with $v_{ij} = v_{ij'}$ for every bidder i). Then*

$$p(j) = \sum_{k=1}^{n} v_k(M^{+j}) - \sum_{k=1}^{n} v_k(M). \tag{3}$$

[5]Given existence, this can also be proved directly by showing that the set of WE form a lattice. This essentially means that taking the component-wise minimum or component-wise maximum of two WE always yields another WE; see the Exercises.

Proof: The key claim is the following:

1. If j is unsold in M, then there is a welfare-maximizing allocation M^{+j} such that both copies j, j' go unsold.

2. If j is sold to bidder i in M, then there is a welfare-maximizing allocation M^{+j} such that bidder i again receives good j.

The claim should be intuitive. For the first case, if one copy of good j is useless, why should a second copy make any difference? For the second case, if i is the winner of good j when there is only one copy of it, why wouldn't it continue to win when there is less competition for it? The proof of the claim is non-trivial but it follows from standard matching or linear programming arguments. We give a sketch in the Appendix and leave further details to the reader.

Assume now that the claim holds and choose M^{+j} accordingly. In the first case, if j is unsold in M and both j, j' are unsold in M^{+j}, then the welfare of M^{+j} and M^* are the same — each a feasible solution to the optimization problem that the other one solves. Thus, both sides of (3) read 0. In the second case, j is sold to a bidder i in both M and M^{+j}. In M^{+j}, the bidders other than i induce a matching M^{-i} that allocates at most one copy of j. Moreover, M^{-i} must be the welfare-maximizing matching of this type — a better one could use to improve the allegedly optimal M^{+j}. Thus, since i is allocated good j in both M and M^{+j}, we have

$$\sum_{k=1}^{n} v_k(M^{+j}) - \sum_{k=1}^{n} v_k(M) = \sum_{k \neq i} v_k(M^{-i}) - \sum_{k \neq i} v_k(M) = p(j),$$

with the second equation following from the definition of $\mathbf{p}$ in Equation (2). ■

We now prove Theorem 3.6.

Proof of Theorem 3.6: Let M and $\mathbf{p}$ denote the allocation and induced item price vector of the truthful-revelation VCG outcome. Every unsold item j has $p(j) = 0$ by definition, so condition (WE2) of Definition 3.1 holds.

Consider an item j that is matched to a bidder i in M; in the following argument, we also allow $j = \emptyset$ to cover bidders that are allocated nothing in M. Let ℓ be some other good of U (with $\ell \neq \emptyset$). By Lemma 3.7, p_ℓ is precisely the additional welfare created if we add a second copy ℓ' of ℓ. We bound this quantity from below as follows:

- Begin with the allocation M.

- Reassign bidder i from its old good j to the new copy ℓ' of ℓ.

- Optimally reassign all bidders other than i via a matching M^{-i}.

The second step increases the welfare of M by $v_{ij} - v_{i\ell'} = v_{ij} - v_{i\ell}$. In the third step, there is precisely one copy of each good of U available for reassigning bidders other than i via the matching M^{-i}. Thus, the third step increases the welfare by precisely

$$\sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} v_k(M(k)),$$

which is the definition of the VCG payment of bidder i and hence the induced price $p(j)$ of item j. Summarizing, we have shown

$$p(\ell) \geq v_i(j) - v_i(\ell) + p(j)$$

and hence

$$v_i(j) - p(j) \leq v_i(\ell) - p(\ell)$$

for every bidder i, $j \in U \cup \{\emptyset\}$, and $\ell \in U$. This final inequality also holds for $\ell = \emptyset$ since, in the VCG mechanism, a bidder is always charged at most its bid. This verifiers condition (WE1) and completes the proof. ■

A Claim From the Proof of Theorem 3.6

We provide a proof sketch of the following claim. Let M be a welfare-maximizing allocation in scenario #3, and an expanded environment that includes a duplicate copy j' of a good j. Then:

1. If item j is unsold in M, then there is a welfare-maximizing allocation M^{+j} in the expanded such that both copies j, j' go unsold.

2. If j is sold to bidder i in M, then there is a welfare-maximizing allocation M^{+j} such that bidder i again receives good j.

Here is an outline of a matching-based argument. Let M^* be an arbitrary maximum-welfare allocation when there are two copies of j. We think of both M and M^* as matchings in the same bipartite graph G, with two copies of j, with j' unsold in M. The symmetric difference of M and M^* is a collection of vertex-disjoint paths and cycles, which have edges alternatively from M and M^*. These paths and cycles interpolate between M and M^*, meaning that applying them (i.e., taking the symmetric difference) in any order to one of them yields the other one. Because M^* is a welfare-maximizing allocation, the value of the M^*-edges in one of these alternating paths P is at least that of the M-edges — otherwise, applying P to M^* yields a better matching, a contradiction. Similarly, because M is also welfare-maximizing among allocations that don't use both j and j', the value of the M- and M^*-edges on each of the alternating cycles must be equal. Obtain M^{+j} from M^* by applying all of the alternating cycles, along with any alternating paths that have equal value contributions from M- and M^* edges. Then, M^{+j} is a welfare-maximizing allocation, and $M \bigtriangleup M^{+j}$ consists only of alternating paths whose application to M strictly improves its welfare. We complete the proof sketch for the first case of the claim; the second case is similar. Suppose for contradiction that at least one of j, j' is sold in M^{+j}. Consider an alternating path P that, when applied to M, results in a copy of j being sold. Since neither copy of j is sold in M, the matching $M \triangle P$ is feasible in the original environment and has welfare larger than that of M, a contradiction.

References

[1] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. Wiley, 1998.

[2] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 591(3):461–479, 1983.

[3] L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1972.