

Walrasian Equilibria

Thomas Kesselheim

Last Update: June 14, 2025

We have seen many different approaches to mechanism design so far, all talking about auctions in some form. The most common form of a mechanism, however, is very different: Whenever we go shopping, we are not asked to bid for items. Instead, each of them has a price tag. We may either buy the item at this price or leave it. Therefore, we now turn to the question how well such prices can coordinate markets. Today, we will start with some classic economic theory about this.

1 Setting and Definition

We consider the standard setting of combinatorial auctions. There are n bidders N and m items M . Feasible allocations are vectors $S = (S_1, \dots, S_n)$, $S_i \subseteq M$ for all $i \in N$, and $S_i \cap S_{i'} = \emptyset$ for $i \neq i'$. Each bidder has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$. We consider *full information*. That is, the valuation functions are fixed and known.

A *Walrasian equilibrium* is an equilibrium in the sense that it is a stable state. In contrast to the equilibrium concepts that we got to know so far, it does not talk about players' strategies but rather about prices making an allocation stable.

Definition 17.1. A pair of a price vector $q \in \mathbb{R}_{\geq 0}^m$ and an allocation $S = (S_1, \dots, S_n)$ is a Walrasian Equilibrium if

(a) Each bidder i gets a bundle that maximizes utility:

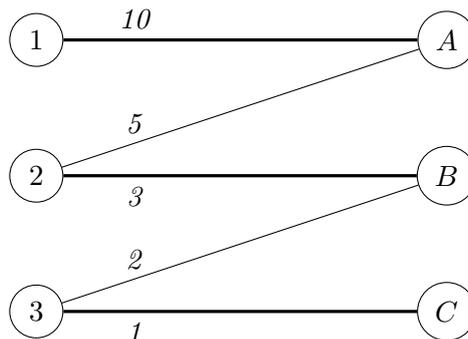
$$v_i(S_i) - \sum_{j \in S_i} q_j \geq v_i(S'_i) - \sum_{j \in S'_i} q_j \quad \text{for all } S'_i \subseteq M .$$

(b) If an item j is unallocated, i.e., $j \notin \bigcup_{i \in N} S_i$, then $q_j = 0$.

Example 17.2. If there is only a single item and $v_1 \geq v_2 \geq \dots \geq v_n$, then the Walrasian equilibria are exactly the prices $q_1 \in [v_2, v_1]$ paired with the allocation that assigns the item to bidder 1.

Example 17.3. We now consider multiple items with unit-demand valuations. That is, a bidder's valuation is of the form $v_i(S) = \max_{j \in S} v_{i,j}$ for $v_{i,j} \geq 0$. Assigning the items is just the same as finding a matching in a complete bipartite graphs whose vertices are $N \cup M$. The edge between $i \in N$ and $j \in M$ has weight $v_{i,j}$.

We consider an example with three bidders 1, 2, 3 and three items A, B, C. We only draw edges of positive value.



The allocation is given by the thick edges. One choice for q would be $q_A = 3, q_B = 1, q_C = 0$. These are the prices that come out of the VCG payments, a connection that we will see later. But it is not the only feasible choice for q . An alternative would be $q_A = 10, q_B = 3, q_C = 1$.

Example 17.4. There are profiles of valuation functions for which no Walrasian equilibrium exists. Consider the example of three single-minded bidders. Bidder 1 wants items 1 and 2, bidder 2 wants items 1 and 3, bidder 3 wants items 2 and 3. Each of them has a value of 1 for getting both items and 0 otherwise.

We now have to distinguish multiple cases how the allocation $S = (S_1, S_2, S_3)$ is chosen. For no such allocation there is a price vector q that fulfills both conditions. Consider, for example, $S_1 = \{1, 2\}, S_2 = S_3 = \emptyset$. Then $q_3 = 0$ because it is not allocated. This means that $q_1 \geq 1$ because otherwise bidder 2 would not be happy. Analogously, $q_2 \geq 1$. This, however, means that $v_1(S_1) - \sum_{j \in S_1} q_j \leq 1 - 2 = -1 < 0 = v_1(\emptyset) - \sum_{j \in \emptyset} q_j$. This is a contradiction to condition (a).

2 First Welfare Theorem

Our first theorem is a very famous one: It tells us that the allocation of any Walrasian equilibrium maximizes social welfare. This has often been interpreted as “markets are efficient”. Undoubtedly, this is a little questionable. One of many reasons is that Walrasian equilibria do not always exist.

Theorem 17.5. If (q, S) is a Walrasian equilibrium, then S maximizes social welfare.

Proof. Let $S^* = (S_1^*, \dots, S_n^*)$ be an allocation that maximizes social welfare. Then for each bidder i we have

$$v_i(S_i) - \sum_{j \in S_i} q_j \geq v_i(S_i^*) - \sum_{j \in S_i^*} q_j .$$

Summing this inequality over all bidder i yields

$$\sum_{i \in N} v_i(S_i) - \sum_{i \in N} \sum_{j \in S_i} q_j \geq \sum_{i \in N} v_i(S_i^*) - \sum_{i \in N} \sum_{j \in S_i^*} q_j .$$

Observe that $\sum_{i \in N} \sum_{j \in S_i} q_j = \sum_{j \in M} q_j$ because each item is allocated at most once in S and items that are not allocated in S have a zero price by property (b). Furthermore $\sum_{i \in N} \sum_{j \in S_i^*} q_j \leq \sum_{j \in M} q_j$ because also in S^* each item is allocated at most once. Unallocated items may have a non-zero price but it cannot be negative. This directly implies

$$\sum_{i \in N} v_i(S_i) \geq \sum_{i \in N} v_i(S_i^*) ,$$

which means that S also maximizes social welfare. □

If you are familiar with linear programming and duality, this argument might look familiar. Indeed, it is nothing but weak LP duality: The price vector q is a feasible solution to the dual LP that certifies optimality of S .

3 Unit-Demand VCG Outcome as Walrasian Equilibrium

As our second main result, we will now see an interesting connection between Walrasian equilibria and the VCG mechanism if bidders have unit-demand valuations.

Let S denote a social-welfare maximizing allocation. Without loss of generality, in S each bidder gets at most one item because multiple items do not increase the value. Let S^{-i} denote the same if bidder i is excluded. Recall that on truthful bids the VCG mechanism defines the payment of bidder i as $p_i(v) = \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) - \sum_{i' \neq i} v_{i'}(S_{i'})$.

We use this to define item prices. If item j is unallocated in S , set $q_j = 0$. If item j is assigned to bidder i , set its price to bidder i 's VCG payment. That is,

$$q_j = p_i(v) = \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) - \sum_{i' \neq i} v_{i'}(S_{i'})$$

Theorem 17.6. *The price vector q defined by the VCG mechanism for unit-demand valuations combined with any social-welfare maximizing allocation is a Walrasian equilibrium.*

So, this means that in particular a Walrasian equilibrium always exists if valuations are unit-demand.

Proof. To simplify notation, we assume that not only each bidder gets at most one item but also that each item is allocated S . This is without loss of generality because we can add bidders of zero value without changing the VCG outcome. In particular, items that were previously unallocated now are allocated to a zero bidder, who does not have to pay anything. Therefore, Property (b) in Definition 17.1 follows.

It remains to prove Property (a). Let buyer i receive item j in S . We have to show that for all $\ell \in M$

$$v_{i,j} - q_j \geq v_{i,\ell} - q_\ell .$$

Let $k \in N$ be the buyer who receives ℓ in S . Then the two prices are defined as

$$q_j = \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) - \sum_{i' \neq i} v_{i'}(S_{i'}) \quad \text{and} \quad q_\ell = \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) - \sum_{i' \neq k} v_{i'}(S_{i'}) .$$

Below, we will show that

$$v_{k,\ell} + \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) \geq v_{i,\ell} + \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) .$$

This then gives us

$$\begin{aligned} v_{i,j} - q_j &= v_{i,j} - \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) + \sum_{i' \neq i} v_{i'}(S_{i'}) \\ &= \sum_{i' \neq k} v_{i'}(S_{i'}) + v_{k,\ell} - \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) \\ &\geq \sum_{i' \neq k} v_{i'}(S_{i'}) + v_{i,\ell} - \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) \\ &= \sum_{i' \neq k} v_{i'}(S_{i'}) + v_{i,\ell} - \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) \\ &= v_{i,\ell} - q_\ell \end{aligned}$$

This shows property (a) in Definition 17.1. □

Lemma 17.7. *If buyer k receives item ℓ in an optimal allocation S , then for all buyers i , we have*

$$v_{k,\ell} + \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) \geq v_{i,\ell} + \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) .$$

Before proving the lemma, let us first understand what this statement means. On both sides of the inequality, we are assigning item ℓ to a buyer and then we consider the optimal allocation without this buyer. Note that in this allocation, item ℓ can be assigned again. So, we are considering a world in which we can assign two copies of item ℓ . The lemma says that social welfare is maximized if we still give one of these two copies to buyer k .

Proof. Recall that we can understand every allocation as a bipartite matching with the buyers being one side of the graph and the items being the other.

The high-level idea of the proof will be as follows: We start with the matching S^{-i} . From this, we construct a matching S^* which is a feasible choice for S^{-k} and has a social welfare which is sufficiently high. The construction of S^* will be done via alternating paths between edges in S^{-i} and S .

First, without loss of generality, we assume that $n = m + 1$. This is without loss because we can add buyers or items without any value. Interpreting allocations as matchings, S^{-i} is even a perfect matching because it matches $n - 1$ buyers to m items. Also S is a matching in this graph, in which each item is matched.

Let's furthermore define $\tilde{S} = S^{-i} \cup \{(i, \ell)\}$. We would like to bound the sum of the edge weights in \tilde{S} because it is exactly

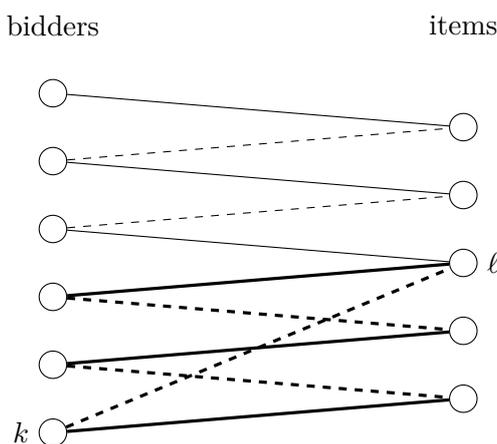
$$\sum_{i'} v_{i'}(\tilde{S}_{i'}) = v_{i,\ell} + \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) .$$

Note that \tilde{S} is “almost” a perfect matching. Only item ℓ is allocated twice, namely to i and one other buyer. Let i_2 be this other buyer who is assigned item ℓ in S^{-i} .

We will construct two sets of edges P and P' as follows. Start from i and i_2 respectively and add the edge to the item that this buyer receives in S (if any). Then add the edge to the buyer who receives this item in \tilde{S} and continue. This procedure ends when reaching a buyer who does not get any item in S . It never stops at an item because every item is allocated in \tilde{S} . If the sequence does not stop, we reach buyer k , then item ℓ , and then we go back to the buyer i or i_2 we started from and close the cycle.

These two sets of edges are on disjoint sets of items and buyers because each buyer gets only one item and each item is allocated at most once. Furthermore, there is at most one buyer who does not get any item in S and therefore one of the two sets is a cycle. Call this set C .

The construction is depicted in the following illustration, where solid edges denote the allocations in \tilde{S} and dashed edges denote the allocations in S . The thick edges are the ones that belong to C .



The key property is now that

$$\sum_{(i',j') \in C \cap S} v_{i',j'} \geq \sum_{(i',j') \in C \cap \tilde{S}} v_{i',j'}$$

because otherwise, we could improve S by replacing $C \cap S$ with $C \cap \tilde{S}$. By the way of constructing C , each buyer will still get at most one item and each item is allocated only once.

Now, construct an allocation $S^* = (\tilde{S} \setminus (C \cap \tilde{S}) \cup (C \cap S)) \setminus \{(k, \ell)\}$. Note that in S^* each buyer gets only one item, each item is allocated only once, and additionally, buyer k does not get any item.

The social welfare of S^* is given by

$$\sum_{i'} v_{i'}(S_{i'}^*) = \sum_{i'} v_{i'}(\tilde{S}_{i'}) - \sum_{(i',j') \in C \cap \tilde{S}} v_{i',j'} + \sum_{(i',j') \in C \cap S} v_{i',j'} - v_{k,\ell} \geq \sum_{i' \neq i} v_{i'}(S_{i'}^{-i}) - v_{k,\ell} + v_{i,\ell} .$$

Finally, note that S^* is a feasible allocation for S^{-k} because it does not allocate anything to k . Therefore

$$\sum_{i'} v_{i'}(S_{i'}^*) \leq \sum_{i' \neq k} v_{i'}(S_{i'}^{-k}) .$$

In combination, this shows the claim. □

Recommended Literature

- Tim Roughgarden's lecture notes <http://timroughgarden.org/w14/1/122.pdf> and lecture video <https://youtu.be/-xX1z5K5KkM>