TTIC 31260 - Algorithmic Game Theory (Winter 2026)

Lecturer: Avrim Blum Lecture 15: March 4, 2026

Market Equilibrium (Walrasian Equilibrium)

1 Overview

Today we will talk about market equilibrium, a mathematical formalization of Adam Smith's "invisible hand" of the market, the idea that the market, through the self-interested actions of agents involved, will allocate resources efficiently, i.e., to maximize social welfare.

2 The setting

We consider n agents and m items (goods) for sale. Each agent i has some valuation function v_i over bundles, and has quasi-linear utilities. So, this is the same setup as when we were looking at combinatorial auctions.

Definition 1 *Consider some pricing $p_1, \ldots, p_m$ on items. We define the demand set D_i for bidder i to be $\arg\max_S v_i(S) - p(S)$ where $p(S) = \sum_{j \in S} p_j$.*

Definition 2 *A Walrasian equilibrium (also called a market equilibrium) is a set of prices $p_1, \ldots, p_m$ and an allocation $S_1, \ldots, S_n$ (set S_i allocated to bidder i and all sets are disjoint) such that S_i is a demand set for buyer i. Furthermore, any unallocated item has price zero.*

Note that if there are no ties in defining the demand sets (each buyer has a unique prefered set) then a Walrasian equilibrium means that all the buyers can come in at the same time and buy what they want, and there will be no contention. Moreover, even if there are ties, we can still assign sets to buyers so that they each are getting a favorite set at these prices. Also, if we imagine a bipartite graph with buyers on one side and goods on the other, and each good being sold by a different seller, there's no incentive to lower prices: the only items not being sold *can't* be sold: they already have price 0.

3 The First Welfare Theorem

Theorem 1 *If a Walrasian Equilibrium exists, then the allocation maximizes social welfare.*

Proof: Here is an LP-relaxation of the social welfare maximization problem. It will have an exponential number of variables, so just think of this as a tool for our analysis. Let x_{iS} indicate whether

S is the set that we allocate to buyer i. Normally this would be $\{0,1\}$ but we will relax this to $[0,1]$. So, instead of saying that we give buyer i exactly one set we will say that for all i we have $\sum_S x_{iS} = 1$. Also, instead of saying that the sets given to different buyers are disjoint, we will say that no item is given in total more than the one unit we have of it (thinking of the item as if it were a divisible good): for all j we have $\sum_i \sum_{\{S:j\in S\}} x_{iS} \leq 1$.

$$\begin{aligned}
&\max \sum_i \sum_S x_{iS}\, v_i(S), \\
\text{s.t.} \quad &\sum_i \sum_{\{S:j\in S\}} x_{iS} \leq 1 && \forall \text{ items } j \text{ (j goes to at most 1 buyer)} \\
&\sum_S x_{iS} = 1 && \forall \text{ buyers } i \text{ (buyer i gets exactly 1 set)} \\
&x_{iS} \geq 0 && \forall i, S
\end{aligned}$$

Let $S_1^*, \ldots, S_n^*$ be the allocation at Walrasian equilibrium and let $\{x_{iS}^*\}$ be the optimal LP solution. Then for all sets S and buyers i we have:

$$v_i(S_i^*) - p(S_i^*) \geq v_i(S) - p(S) \tag{1}$$

$$v_i(S_i^*) - p(S_i^*) \geq \sum_S x_{iS}^* \left(v_i(S) - p(S)\right), \tag{2}$$

where the first inequality comes from the definition of a Walrasian equilibrium and the second inequality follows from the fact that $\sum_S x_{iS}^* = 1$. Summing over all buyers,

$$\sum_i \left(v_i(S_i^*) - p(S_i^*)\right) \geq \sum_i \sum_S x_{iS}^* \left(v_i(S) - p(S)\right)). \tag{3}$$

$$\text{(social welfare at equilibrium)} - \sum_j p_j \geq \text{(optimal social welfare)} - \sum_j p_j. \tag{4}$$

In going from (3) to (4) we are using the fact that $S_1^*, \ldots, S_n^*$ is a Walrasian equilibrium, so that on the left-hand-side, summing the prices of all the sets allocated on is the same as summing the prices of all the items, and on the right-hand-side we are using the fact that each item in our LP relaxation is allocated in total at most one unit worth, so the sum of prices of the fractional allocation is at most the sum of the prices of the items. ■

4 Existence of Walrasian Equilibria for Gross Substitutes Valuations

As we saw in a previous class, Walrasian equilibria might not exist for general monotone valuation functions. Here will show they exist (and give an algorithm for finding them) in the case of buyers with valuations satisfying the *gross substitutes* property, which includes additive and unit-demand buyers, and is a special case of submodular valuations. Perhaps even more importantly, the algorithm will correspond to a natural market dynamic.

Definition 3 (Gross substitutes) *A valuation function satisfies the gross substitutes property if raising the prices on some items cannot cause items whose prices were* not *raised to drop out of the demand set.*

For example, if the price of beef goes up, then people will buy more chicken, and if the price of one kind of ice-cream goes up, people will buy more of other kinds of ice-cream. That's because different meats and different kinds of ice-cream are "gross substitutes" for each other. On the other hand, if the price of hamburger goes up, then people may buy fewer hamburger buns, because they go together (are "complements").

To make things easier, we will look at a discretized process where prices go up in units of ϵ. We will show this produces an ϵ-Walrasian equilibrium, which is just like a Walrasian equilibrium except that each set S_i may not be exactly a demand set at the current prices $p_1, ..., p_m$. Instead, each S_i is a demand set at prices in which each $j \in S_i$ is priced at p_j and each $j \notin S_i$ is priced at $p_j + \epsilon$. Think of it as: people might not be perfectly happy with the set S_i that they bought, but they wouldn't go through the hassle of trading it in if there is some cost for switching. We can make ϵ as small as we like though. Here is the discretized process, which is essentially a joint ascending auction among all the items:

Increasing Prices Market Process (ϵ-Discretized):

- Initialize all prices p_j to 0. Initialize each S_i to the empty set. Think of p_j as the highest bid so far for item j, and the seller is looking for bids at $p_j + \epsilon$.
- Define demand set D_i to be the preferred bundle for buyer i at prices in which each $j \in S_i$ is priced at p_j and each $j \notin S_i$ is priced at $p_j + \epsilon$. Break ties in favor of D_i that contain S_i (and break ties among such D_i arbitrarily).
- While there exists an agent i such that $D_i \neq S_i$ do:
  1. For every item $j \in D_i \setminus S_i$, let $p_j \leftarrow p_j + \epsilon$.
  2. Set $S_i \leftarrow D_i$.
  3. For every agent $k \neq i$, set $S_k \leftarrow S_k \setminus D_i$.
- Finally, output the allocation $S_1, ..., S_n$.

We can think of this as a joint ascending auction, where each item has a current price, and the seller is asking whether anyone is willing to pay ϵ more than that price.

To analyze this process, we first prove the following:

Claim 2 *For agents with valuation functions satisfying gross substitutes, this process maintains the property that for every agent i we have $S_i \subseteq D_i$.*

Proof: The property is certainly true at the start where each $S_i = \emptyset$. Now, consider what happens in an update step. When we update some set S_i, we set it to equal D_i according to the definition of D_i, so that set continues to satisfy the property. But what about agents $k \neq i$? For those agents, we made two updates: first, we may have removed some items from S_k (which maintains the desired property) and second, we increased prices for some items that are no longer in S_k. For

an agent with a general valuation function, this second step could potentially completely change D_k; however, by definition of gross-substitutes valuations, raising prices of items that are not in S_k cannot remove items that are in S_k from the demand set. So, the property still holds for those agents as well. ■

Notice that Claim 2 implies that once an item is allocated to some agent, it never gets dropped. All it can do is move around among agents.

Theorem 3 *This process halts, and ends with an ϵ-Walrasian equilibrium for agents with gross substitutes valuations.*

Proof: The process halts because each step increases the price of at least one item (by Claim 2, the fact that $D_i \neq S_i$ means that there is some item in $D_i \setminus S_i$) and items are never dropped so no item can become so expensive that nobody wants it. When the process halts, it's an ϵ-Walrasian equilibrium because the fact that items never become unallocated means that any unallocated items must have price 0, and by design of the process we have all $D_i = S_i$. ■

We can also give an analog of Theorem 1 for ϵ-Walrasian equilibria.

Theorem 4 *The social welfare of an ϵ-Walrasian equilibrium is within $n\epsilon$ of the maximum social welfare.*

Proof: If we look at the proof of Theorem 1, equation (1) is satisfied if on the right-hand-side we replace $p(S)$ with $p_{+\epsilon}(S)$, defined as adding ϵ to the price of everything in S. This then carries over to equations (2) and (3). For (4), the left-hand-side stays the same, but the right-hand-side now can be smaller by up to $n\epsilon$ since the prices of each item are ϵ higher. ■

One natural question is: is this process incentive compatible? Unfortunately, unlike the case of an ascending auction of 1 item, the answer is no. Here is an example from the book:

	v(a)	v(b)	v(ab)
Alice	4	4	4
Bob	5	5	10

Notice that Bob would be better off leaving Alice alone and just bidding on item b.