Competitive Equilibrium in Unit-Demand Markets

CSCI 1440/2440

2025-11-12

We define competitive equilibrium, and present a method of computing these equilibria in unit demand markets. This method is based on the theory of strong duality in linear programming (LP). Specifically, the primal of the LP encodes the equilibrium allocation, while the dual encodes the equilibrium prices.

1 Competitive Equilibrium

En route to designing an EPIC auction for unit-demand markets, we take a detour leaving game-theoretic incentives behind, and instead shift our attention to economic equilibrium, specifically **competitive equilibrium**, which is also sometimes called **Walrasian equilibrium**.

Consider a market $\mathcal{M} = (n, m, (\boldsymbol{v}_i)_{i \in [n]})$ comprising n agents with quasi-linear utilities and m goods, in which each agent's valuation function is **unit demand**: i.e., agent i's value for bundle $\boldsymbol{x} \in \{0,1\}^m$ is given by:

$$\boldsymbol{v}_i(\boldsymbol{x}) = \max_{j \in [m] : \boldsymbol{x}_j = 1} v_{ij}.$$

An **allocation** $\boldsymbol{X} \in \{0,1\}^{n \times m}$ is a mapping from goods to agents, represented as a matrix s.t. $x_{ij} \in \{0,1\}$ denotes the quantity of good $j \in [m]$ allocated to agent $i \in [n]$. Goods are assigned (anonymous) **prices** $\boldsymbol{p} \in \mathbb{R}_+^m$.

A pair comprising an allocation and prices $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ is said to be a **competitive** (or **Walrasian**) **equilibrium** of such a market $\mathcal{M}$ iff

Feasibility (WE0): An allocation is **feasible** iff no good is overallocated: i.e., each good is allocated to at most one agent. Mathematically,

$$\sum_{i \in [n]} \boldsymbol{x}_i^* \leq \mathbf{1}_m$$

In more detail, for all $j \in [m]$, $\sum_{i \in [n]} x_{ij}^* \leq 1$.

Agent stability (WE1): Given prices $\boldsymbol{p}^*$, all agents $i \in [n]$ maximize their utility at allocation $\boldsymbol{X}^*$, i.e., for all $i \in [n], \boldsymbol{x}_i^* \in \arg\max_{\boldsymbol{x} \in \{0,1\}^m} u_i(\boldsymbol{x})$, where

$$u_i(\boldsymbol{x}) = \max_{j \in [m] : \boldsymbol{x}_j = 1} v_{ij} - p_j$$

Market clearance (WE2): The market clears when Walras' law holds.

Walras' law is a way of equating supply and demand. It requires that the total monetary value of demand equal the total monetary value of supply. That is, the monetary value of excess demand (and excess supply, as well) must be zero. Another (equivalent) interpretation of Walras' law is: if a good is not allocated, then it is necessarily priced at zero.

Mathematically,

$$\boldsymbol{p}^* \cdot \left(\mathbf{1}_m - \sum_{i \in [n]} \boldsymbol{x}_i^* \right) = 0$$

In more detail, for all $j \in [m]$, $p_j^* \cdot \left(1 - \sum_{i \in [n]} x_{ij}^*\right) = 0$.

Importantly, competitive equilibria have been shown to exist in broad classes of markets. Moreover, by the first welfare theorem of economics, competitive equilibria are necessarily welfare maximizing. The second welfare theorem is something like a converse to the first welfare theorem: it provide conditions under which prices can be attached to a welfare-maximizing allocation so that a competitive equilibrium ensues. This latter theorem gives rise to a method for searching for competitive equilibria called Negishi's method: first, find a welfare-maximizing allocation $\boldsymbol{X}^*$; second, try to find prices $\boldsymbol{p}^*$ that support that allocation, in the sense that together $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ comprise a competitive equilibrium.

Let's look at a couple of examples of unit-demand markets, and try to identify their competitive equilibria via Negishi's method. Recall that the first step is to find a welfare-maximizing allocation. In unit-demand markets, solving for a welfare-maximizing allocation can be posed as a maximum weight matching problem on a bipartite graph, i.e., a problem of assigning goods to agents so as to maximize welfare.

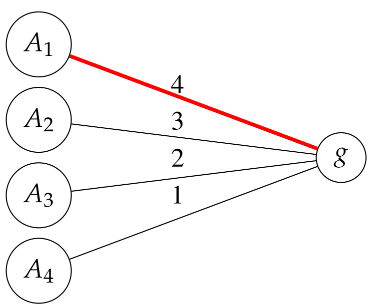

Figure 1: An example of a unit-demand market with four agents and one good.

Example 1.1. Consider the market in Figure 1. The welfare-maximizing allocation $\boldsymbol{X}^*$ in this market is simply to assign the good g to agent A_1. The question then becomes, what prices support this allocation, so that together the allocation and prices form a competitive equilibrium?

WE2 is irrelevant in this market under this allocation, since all goods are allocated. The only question then is, what prices ensure agent stability?

Does a price of 2.5, for example, guarantee agent stability? No, because at that price, agents A_1 and A_2 would both include the good g in their demand sets, but this good is allocated only to A_1. Therefore, agent A_2's value for the good serves as a lower bound on the price that can guarantee agent A_1's stability, while agent A_1's value for the good serves as an upper bound. Indeed all prices in the range $[3, 4]$ support the allocation $\boldsymbol{X}^*$.

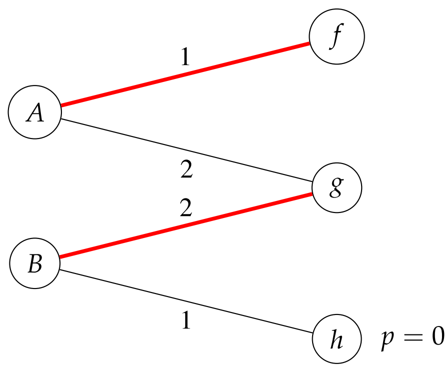

Figure 2: An example of a market with two agents and three goods.

Example 1.2. Now consider the market in Figure 2 with two agents, A and B, and three goods, f, g, and h, where the agents' valuations are indicated by the weights on the edges. That is, agent A values good f at 1 and good g at 2, while agent B values good g at 2 and good h at 2. The edges that are omitted from the graph indicate that an agent has no value whatsoever for the good (e.g., agent A has 0 value for good h).

First, we find a welfare-maximizing allocation. There are two. Agent A can be allocated f, and B can be allocated g; or agent A can be allocated g, and B can be allocated h. Let's proceed with the first of these allocations, which is colored in red in the figure. What are supporting prices? Well, since h is not allocated, it is necessarily priced at 0. But what about f and g?

Since h is priced at 0, agent B could achieve utility 1 if her allocation were h instead of g. As a result, the price of g is upper bounded by 1: her value of 2 less a price such that her utility is guaranteed to be at least 1. On the other hand, the maximum utility that agent A can achieve is 1, when the price of f is 0. This implies that the price of good g is also upper bounded by 1, because otherwise, A would prefer g to f. Therefore, good g must be priced at 1, and consequently, good f must be priced at 0. The unique set of prices that support this allocation are $(0, 1, 0)$, for goods f, g, and h, respectively.

As this problem is symmetric, these competitive equilibrium prices also support the other welfare-maximizing allocation, in which agent A is instead allocated good g and agent B is allocated good h.

2 *Constrained Optimization*

Mathematical programming is a means of formulating and solving a form of decision problem called a **constrained optimization problem**.

Common examples include **portfolio optimization**: allocate funds to stocks in a diversified manner so as to maximize profit while respecting risk limitations; **job scheduling**: complete all tasks in the minimum amount of time, subject to the constraint that only some machines can complete some jobs; and **course scheduling**: complete all course requirements within four years, subject to the constraint that only some courses satisfy the various requirements and not all courses are offered each semester—not to mention that some course time slots conflict!

Two other examples, in more detail, are:

- **Maximizing profit with budget constraints**: Imagine a factory that can produce two products, each of which accrue some revenue. Each also requires a certain amount of raw materials and a certain amount of labor, all of which comes with a cost. The decision problem is to allocate the raw materials and labor to the production of the two products so as to maximize profit: i.e., revenue minus cost.
- **Diet Problem**: The FDA recommends a certain amount of nutrient intake (proteins, fibers, vitamins, etc.) per day, but at the same time, it also recommends a maximal caloric intake (e.g., 2000). Moreover, some people have allergies. The decision problem is to recommend a diet that meets the FDA requirements (the lower bounds on nutrient intake and the upper bound on caloric intake), while ensuring that no allergic reactions ensue.

All of these problems are characterized by an **objective**, such as profit maximization, and **constraints**, such as adhering to a budget. The constraints define the space of **feasible solutions** to the problem.

When the objective function and the constraints are all linear, a constrained optimization problem is called a **linear program**. In today's lecture, we formulate the problem of solving for a competitive equilibrium in unit demand markets as a linear program.

3 *Linear Programming Duality*

A linear program is characterized by an objective function and constraints. Here is (something close to) the standard form of a linear program:

$$\max_{\boldsymbol{x} \in \mathbb{R}^n} \quad \boldsymbol{c} \cdot \boldsymbol{x} \tag{1}$$

$$\text{subject to} \quad A\boldsymbol{x} \leq \boldsymbol{b} \tag{2}$$

$$\boldsymbol{x} \geq \boldsymbol{0} \tag{3}$$

where

- $\boldsymbol{x}$ is the vector of (primal) decision variables
- $\boldsymbol{c}$ is the vector of coefficients for the objective function
- the inequalities $A\boldsymbol{x} \leq \boldsymbol{b}$ represent the constraints on the variables

Linear programming is a powerful solution technique because it is capable of modeling many many problems, and it is solvable in polynomial time.

For every so-called **primal** linear program of the above form, there is a corresponding **dual** linear program of the following form:

$$\min_{\boldsymbol{y}\in\mathbb{R}^m} \quad \boldsymbol{b}\cdot\boldsymbol{y} \tag{4}$$

$$\text{subject to} \quad A^T\boldsymbol{y} \geq \boldsymbol{c} \tag{5}$$

$$\boldsymbol{y} \geq \boldsymbol{0} \tag{6}$$

where

- $\boldsymbol{y}$ is the vector of decision variables
- $\boldsymbol{b}$ is the vector of coefficients for the objective function
- the inequalities $A^T\boldsymbol{y} \leq \boldsymbol{c}$ represent the constraints on the dual variables

The dual is constructed by associating a dual variable with every constraint in the primal and a dual constraint with every variable in the primal.

N.B. The dual of the dual is again the primal.

One straightforward observation about the primal and the dual is called **weak duality**: For all feasible solutions $\boldsymbol{x}, \boldsymbol{y}$, to the primal and the dual, respectively, $\boldsymbol{c}\cdot\boldsymbol{x} \leq \boldsymbol{b}\cdot\boldsymbol{y}$.

Proof.

$$A\boldsymbol{x} \leq \boldsymbol{b} \Rightarrow \boldsymbol{y}^T(A\boldsymbol{x}) \leq \boldsymbol{y}^T\boldsymbol{b} = \boldsymbol{b}\cdot\boldsymbol{y} \tag{7}$$

$$A^T\boldsymbol{y} \geq \boldsymbol{c} \Rightarrow \boldsymbol{x}^T(A^T\boldsymbol{y}) \geq \boldsymbol{x}^T\boldsymbol{c} = \boldsymbol{c}\cdot\boldsymbol{x} \tag{8}$$

But now, since $\boldsymbol{y}^T(A\boldsymbol{x})$ is a vector, it holds that $\boldsymbol{y}^T(A\boldsymbol{x}) = (\boldsymbol{y}^T(A\boldsymbol{x}))^T = (A\boldsymbol{x})^T\boldsymbol{y} = \boldsymbol{x}^T(A^T\boldsymbol{y})$. Therefore, $\boldsymbol{c}\cdot\boldsymbol{x} \leq \boldsymbol{x}^T(A^T\boldsymbol{y}) = \boldsymbol{y}^T(A\boldsymbol{x}) \leq \boldsymbol{b}\cdot\boldsymbol{y}$.

In other words, the value of the primal is upper bounded by the value of the dual. In fact, a much stronger[1] property holds, called strong duality. **Strong duality** states that the value of the primal and the dual are in fact equal! This result is both striking and foundational. It is usually attributed to John von Neumann, who established the result in his study of zero-sum games, but Leonid Kantorovich and George Dantzig both also played a key role in the development of the theory of linear duality.

[1] and harder to prove

Together, strong duality and our weak duality proof imply the following:

$$\boldsymbol{c}\cdot\boldsymbol{x}^* = \boldsymbol{y}^*\cdot(A\boldsymbol{x}^*) = \boldsymbol{b}\cdot\boldsymbol{y}^*$$

Rearranging this equation yields an optimality condition known as **complementary slackness**:

- $(\boldsymbol{c} - A^T\boldsymbol{y}^*)\cdot\boldsymbol{x}^* = \boldsymbol{0}$
- $\boldsymbol{y}^*\cdot(\boldsymbol{b} - A\boldsymbol{x}^*) = \boldsymbol{0}$

In more detail,

- $x_i^*(A^T y^* - c)_i = 0$, for all $i \in [n]$
- $y_j^*(b - Ax^*)_j = 0$, for all $j \in [m]$

In words, complementary slackness states that, for each primal constraint, if the optimal value of the dual variable is positive, then the constraint is tight (i.e., $(Ax^*)_j = b_j$) Alternatively, if the constraint is not tight, then the optimal value of the dual variable is necessarily 0. And likewise for each dual constraint and primal variable.

4 *Solving Unit-Demand Markets via a Primal and a Dual*

We now return to our study of unit-demand markets. In particular, we seek a method for computing a competitive equilibrium of a unit-demand market.

Following Negishi's method, our first goal is to solve for a welfare-maximizing allocation in a unit-demand market using linear programming.

The following mathematical program solves this problem:

$$\max_{X \in \{0,1\}^{n \times m}} \quad \sum_{i \in [n]} \sum_{j \in [m]} v_{ij} x_{ij} \tag{9}$$

$$\text{subject to} \quad \sum_{j \in [m]} x_{ij} \leq 1 \qquad \forall i \in [n] \tag{10}$$

$$\sum_{i \in [n]} x_{ij} \leq 1 \qquad \forall j \in [m] \tag{11}$$

The objective, namely to maximize welfare, can also be understood as finding a maximum-weight bipartite matching, as the two sets of constraints ensure that each bidder is matched to only one good (Constraint 10) and that each good is matched to only one bidder (Constraint 11).

The present formulation has a key difficulty, however, namely, that each $x_{ij} \in \{0, 1\}$. That is, each such variable is an **integer**. The reason this presents a problem is that *integer* linear programming is NP-hard, so it is not believed to be solvable in polynomial time.

But do not despair! We already noted that maximum weight bipartite matching *is* solvable in polynomial time. Correspondingly, the integer constraints in this particular problem can be relaxed to non-negativity constraints, so that the following is an equivalent problem formulation:

$$\max_{X \in \mathbb{R}^{n \times m}} \quad \sum_{i \in [n]} \sum_{j \in [m]} v_{ij} x_{ij} \tag{12}$$

$$\text{subject to} \quad \sum_{j \in [m]} x_{ij} \leq 1 \qquad \forall i \in [n] \tag{13}$$

$$\sum_{i \in [n]} x_{ij} \leq 1 \qquad \forall j \in [m] \tag{14}$$

$$x_{ij} \geq 0 \qquad \forall i \in [n], j \in [m] \tag{15}$$

This latter problem formulation is indeed a linear program, and as such, can be solved in polynomial time. We thus have a method of solving for a welfare-maximizing allocation in a unit-demand market in polynomial time.

We now turn our attention to the dual of this linear program. Following the formulaic manner of constructing a dual yields:

$$\min_{\boldsymbol{y}\in\mathbb{R}^n,\boldsymbol{z}\in\mathbb{R}^m} \quad \sum_{i\in[n]} 1y_i + \sum_{j\in[m]} 1z_j \tag{16}$$

$$\text{subject to} \quad y_i + z_j \geq v_{ij} \qquad \forall i \in [n], \forall j \in [m] \tag{17}$$

$$y_i, z_j \geq 0 \qquad \forall i \in [n], j \in [m] \tag{18}$$

A typical linear programming solver outputs a solution not only to the primal, but to the dual as well. We ran such a solver on the market shown in Figure 2, and obtained one of the two welfare-maximizing allocations, together with the following solution to the dual: $\boldsymbol{y} = (1,1)$ and $\boldsymbol{z} = (0,1,0)$. Do these number look familiar?

They should. The variable $\boldsymbol{z}$ corresponds to the prices in the market. What about the variable $\boldsymbol{y}$? Plugging in $\boldsymbol{p}$ for $\boldsymbol{z}$ and rearranging the key constraint yields: $y_i \geq v_{ij} - p_j$. The right-hand-side of this equation looks like quasi-linear utility. Indeed, y_i is something *like* utility. It is not quite utility, but rather, it is called **indirect utility**. We thus restate the dual as follows:

$$\min_{u\in\mathbb{R}^n,\boldsymbol{p}\in\mathbb{R}^m} \quad \sum_{i\in[n]} 1u_i + \sum_{j\in[m]} 1p_j \tag{19}$$

$$\text{subject to} \quad u_i + p_j \geq v_{ij} \qquad \forall i \in [n], \forall j \in [m] \tag{20}$$

$$u_i, p_j \geq 0 \qquad \forall i \in [n], j \in [m] \tag{21}$$

It remains to argue that the solution $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ to this primal and dual is indeed a competitive equilibrium for a unit-demand market. Feasibility is embedded into the primal program, so it is certainly satisfied. Walras' law is also satisfied; it is precisely complementary slackness. It thus remains only to argue that the agent stability condition holds. But this condition is precisely the key dual constraint: for all agents $i \in [n]$, i's utility is at least that of the utility it could achieve if it were instead allocated any other good $j \in [m]$: i.e., $u_i \geq v_{ij} - p_j$, for all $i \in [n]$ and $j \in [m]$.