Lecture 16: Competitive Pricing in Networks

Alexander Wolitzky

MIT

6.207/14.15: Networks, Spring 2022

Networked Markets

We continue our study of how network phenomena affect economic interactions, now from the perspective of **markets**.

Traditional economic theory assumes markets consist of a large number of participants who can all trade freely with each other.

Many markets do work this way: financial markets, commodity markets, markets for relatively non-specialized labor.

But in many other markets, only some agents can trade with each other, due to various factors.

- Geography
- Technological capability
- Limited information

Even when in principle all agents can trade with each other, many markets have a small number of participants, and buyers can have very heterogeneous valuations for different sellers' goods.

- Housing market.
- Labor markets for unique positions (CEO's, NFL quarterbacks)

Networked Markets (cntd.)

Markets where only some agents can (or want to) trade with others are more like networks than traditional economic markets. (**Note:** networks of "who can/wants to trade with whom" is a completely different issue than the "network effects"/externalities in markets that we studied last week.)

This week, we'll analyze several features of these "networked markets."

- What are **competitive** or **market-clearing** prices in a network of buyers and sellers, and what are the properties of the resulting **competitive equilibrium**? (Today)
- If we view prices as being strategically determined by individual agents rather than competitive market forces, how does the network structure determine who has more **bargaining power**? (Wednesday)
- How is bargaining power and the division of economic value determined in **intermediated markets**, where goods must pass through a network of intermediaries on their way from producers to final consumers? (Also Wednesday)

Matching Markets

Today, we study competitive pricing in a network of buyers (I) and sellers (J), with $|I| = |J|$ (for simplicity).

Make the simplifying assumption that each buyer wants to consume (at most) **one good**, and each seller has **one good** to sell.

- **Examples:** housing markets (only want one house), labor markets for unique positions (only want one CEO or quarterback).

Goods are **heterogeneous**: buyer i values seller j's good at $v_{ij} \geq 0$ (and obtaining a second good has no value).

Such a "unit demand/unit supply" environment is called a **matching market** (or an **assignment game**). Why?

- Consider the bipartite graph of buyers and sellers.
- An outcome of trade is a **matching** (which buyer obtains each seller's good), together with a **price** paid for each good.

Matching Markets (cntd.)

If buyer i buys seller j's good at price p_j,

$$\begin{aligned} i\text{'s payoff} &= v_{ij} - p_j \\ j\text{'s payoff} &= p_j \end{aligned}$$

Questions:

1. Ignoring prices for a moment, what can we say about the optimal (total value-maximizing) way to assign goods to buyers?
2. Does there always exist a vector of prices $p = (p_j)_{j \in J}$ that supports the optimal assignment as a "competitive equilibrium" (meaning that each buyer i prefers his assigned good j at price p_j to any other good j' at price $p_{j'}$)? Conversely, are the goods assigned optimally in *every* competitive equilibrium?

3. How can we find the competitive equilibrium prices?

Aside: Matching Markets and Market Design

The assignment game was introduced by Koopmans and Beckman (1957) and Shapley and Shubik (1972).

Along with related allocation problems, it is a core model in the field of **market design**, where economists help design the rules of markets or mechanisms (e.g., matching mechanisms or auctions) to try to improve market functioning.

- Compare to 2-sided matching models without money (Gale-Shapley 1962), assignment problems without money (Scarf-Shapley 1973).
- We'll learn more about auctions later in the course. For much more on market design (both theory and real-world applications), see 14.19.

Competitive Equilibrium: Definition

Formally, a **competitive equilibrium** is a bijection $M : I \rightarrow J$ (which specifies the good $j = M(i)$ consumed by each buyer i) and a price vector $p = (p_j)_{j \in J}$ such that, for each buyer i, if $j = M(i)$ then

$$\begin{aligned} v_{ij} - p_j &\geq v_{ij'} - p_{j'} \text{ for all } j' \in J, \text{ and} \\ v_{ij} - p_j &\geq 0. \end{aligned}$$

In economic terminology,

- M is the **equilibrium assignment** (or **allocation**).
- p is the **market-clearing price vector.**
- The condition that M is a bijection (so each good is assigned to a single buyer) is the **market-clearing condition.**
- The condition that each buyer prefers her assigned good to any other good at the market-clearing price vector p is the **individual optimization condition**.

Questions about Competitive Equilibrium

Two central questions about competitive equilibrium:

1. Is it **efficient**—does it assign goods to buyers who value them the most?
2. Does it always **exist**—is there actually an assignment and prices that satisfy the definition?

Efficiency

Forgetting about prices and individual optimization for the moment, we can define the **total value** generated by any bijection $M : I \rightarrow J$ as

$$V(M) = \sum_{i \in I} v_{iM(i)}.$$

Since there are finitely many ways to assign the goods to the buyers, we can define the **optimal total value**

$$V^* = \max_M V(M).$$

By definition, V^* is the maximum value that can be created by assigning goods to buyers in any way. It says nothing about how this assignment can be implemented (e.g. whether there's any way to get the individual buyers and sellers to go along with it).

- Note that the sellers don't show up in the definition of V^*, since they don't care how the goods are assigned.

Competitive Equilibrium: Efficiency

Even though V^* is defined without any reference to how the optimal assignment can be implemented, it turns out that a competitive equilibrium (if it exists) always attains value V^*.

Theorem

If (M, p) is a competitive equilibrium, then $V(M) = V^*$.

The result that competitive equilibria are efficient is a fundamental result in economics, generally known as the **first welfare theorem**.

- **Remark for economics students:** The theorem on this slide is a bit different from the classical first welfare theorem, in that it assumes a particular kind of preferences (unit demand and linear utility in money) but establishes utilitarian efficiency rather than only Pareto efficiency.

Competitive Equilibrium: Efficiency (cntd.)

Theorem

If (M, p) is a competitive equilibrium, then $V(M) = V^*$.

Proof:

- Since each buyer chooses optimally given prices p, we have

$$M(i) \in \operatorname*{argmax}_j v_{ij} - p_j \text{ for all } i \in I.$$

- Therefore, given prices p, the matching M maximizes the sum of the buyers' payoffs over all matchings:

$$M \in \operatorname*{argmax}_{M'} \sum_{i \in I} \left(v_{iM'(i)} - p_{M'(i)} \right).$$

- But the sum of all prices $\sum_{i \in I} p_{M'(i)}$ doesn't depend on the matching: it always equals $\sum_{j \in J} p_j$.
- So $M \in \operatorname{argmax}_{M'} \sum_{i \in I} v_{iM'(i)}$.
  That is, M maximizes total value.

Competitive Equilibrium: Existence

We have shown that, if a competitive equilibrium exists, it's efficient.

But does a competitive equilibrium always exist?

Not obvious: need to find a single price vector $(p_j)_{j \in J}$ that simultaneously causes all buyers to choose the "right" good.

- One price vector must coordinate purchasing choices of all buyers.

Nonetheless:

Theorem

There is always at least one competitive equilibrium.

To prove this, we will relate the economic notion of **competitive equilibrium** to the graph-theoretic notion of a **perfect matching**.

Competitive Equilibrium and Perfect Matching

A **perfect matching** in a graph G is a set of edges with no common vertices, such that each vertex is an endpoint of one edge.

- That is, a matching of nodes to neighbors such that each node is matched exactly once.

Note: in a bipartite graph between nodes $i \in I$ and $j \in J$, with $|I| = |J|$, a perfect matching is a bijection $M : I \to J$.

We can rephrase the definition of a competitive equilibrium as a price vector $p = (p_j)_{j \in J}$ together with a perfect matching in the graph where there is a link between i and j if and only if j is a **preferred-seller** for i, meaning that

$$\begin{aligned} v_{ij} - p_j &\geq v_{ij'} - p_{j'} \text{ for all } j' \in J, \text{ and} \\ v_{ij} - p_j &\geq 0. \end{aligned}$$

A Subtlety

At a market-clearing price vector, there can be **ties** in agents' preferences, so a buyer i can have more than one preferred seller, and a seller j can be preferred by more than one buyer.

- In this case, to find the equilibrium assignment M, agents need to break ties in the right way, so that we get a perfect matching in the entire preferred-seller network.

A Key Fact About Perfect Matchings

Notation: given a graph G and a set of nodes $S \subseteq N$, denote the set of all neighbors of nodes in S by

$$N(S) = \bigcup_{i \in S} N_i,$$

where N_i is the set of i's neighbors.

Our proof of existence of competitive equilibrium relies on the following important fact about perfect matchings in bipartite graphs:

Theorem (Hall's Marriage Theorem)

A bipartite graph between nodes $i \in I$ and $j \in J$, with $|I| = |J|$, has a perfect matching if and only if, for all $S \subseteq I$, we have

$$|S| \leq |N(S)|.$$

Hall's Theorem

A set $S \subseteq I$ such that $|S| > |N(S)|$ is called a **constricted set**.

- Obviously, if there is a constricted set, there can't be a perfect matching, as there's no way to match everyone in S.
- Hall's theorem asserts that the converse is also true: if no set is constricted, then there is a perfect matching.
- This should be somewhat surprising at first glance: the presence of a constricted set is an obvious reason why a perfect matching could fail to exist, but you might have thought that there could also be other, more complicated reasons why there may not be a perfect matching.
  Hall's theorem says no—the presence of a constricted set is the **only** reason why a perfect matching may not exist.

This is a fundamental theorem which we'll use extensively and the proof is instructive, so we'll prove it in the next couple slides.

Proof of Hall's Theorem

Suppose a perfect matching does not exist.
We will find a constricted set. (This proves the theorem.)

Let M be a **maximal matching**: a matching that matches as many nodes as possible.

Since we assumed there's no perfect matching, M is not perfect, so there's some node $i \in I$ that's not matched by M.

We'll show that the set of nodes consisting of node i, together with all nodes in the set I that can be reached from i by a certain kind of path in G—called an **alternating path**—is a constricted set.

We'll see that the set must be constricted, because otherwise we could use the alternating path to find a better matching.

Alternating and Augmenting Paths

Given a matching M and a node i that is not matched by M, an **alternating path** is a path that starts from node i, never repeats a node, and alternates between using links that *are not* in M and links that *are* in M.

- The first link is not in M, since i is not matched by M.
- Every node in an alternating path is matched by M (with another node in the path), except the initial node i and possibly the final node.

An **augmenting path** is an alternating path where the final node is not matched by M.

- Suppose we take an augmenting path and delete all the "even" links in the path and instead add all the "odd" links (leaving the rest of M unchanged).
- This creates a new matching that matches all the nodes that were matched by M, plus the initial and final nodes.
- It thus creates a new, "augmented" matching that matches more nodes.

Alternating Paths and Maximal Matchings

Go back to the case where M is a maximal matching and i is node that's not matched by M.

Claim: Every alternating path starting from node i ends at a node that's matched under M.

Proof: If not, the alternating path is an augmenting path, so switching the even and odd links would give a matching that matches more nodes than M.

We use this claim to find a constricted set.

Proof of Hall's Theorem (cntd.)

Define the set of nodes $Z \subseteq J$ to be all nodes in J that can be reached from node i by **any** alternating path.

- Similarly, define the set of nodes $W \subseteq I$ to be all nodes in I that can be reached from node i by **any** alternating path (including i itself).

We will show that W is a constricted set: that is, $|N(W)| < |W|$.

Proof of Hall's Theorem (cntd.)

Note that $N(W) \subseteq Z$.
(That is, if w is reachable from i by an alternating path and $wj \in G$, then j is reachable from i by an alternating path.)

- Start with the alternating path from i to w, and add the link from w to j to get an alternating path from i to j.

Now we show that $|Z| \leq |W| - 1$. This completes the proof, as it implies that $|N(W)| \leq |W| - 1$, so W is constricted.

- Since alternating paths can't end at unmatched nodes and every node in Z is reachable by an alternating path, every node in Z must be matched by M.
- For every node in Z, its partner (in M) is reachable by an alternating path, and hence lies in W (and in fact lies in $W \setminus \{i\}$, since i is unmatched).
  Hence, $|Z| \leq |W| - 1$.

Proof of Competitive Equilibrium Existence

Now we can use Hall's theorem to prove that there is always at least one competitive equilibrium.

The proof is constructive: we give an algorithm that constructs a price vector p for which the preferred-seller graph has a perfect matching M. (So then (M, p) is a competitive equilibrium.)

Intuitively, the algorithm is a kind of **auction**, where the relative prices of "over-demanded" goods increase over time until the market clears.

- We'll learn more about auctions later in the course.

The algorithm assumes that all buyer valuations v_{ij} are whole numbers for simplicity, but this isn't important.

Algorithm for Determining Market-Clearing Prices

1. Initialize the price vector at $p_j = 0$ for every good j.
2. Given the current price vector p, construct the preferred-seller graph, and check whether there is a perfect matching.
3. If there is, stop the algorithm—we're done.
4. If not, by Hall's theorem, there is a constricted set of buyers S, with neighbors $N(S)$.
   (Intuitively, the goods in $N(S)$ are "over-demanded"—they are demanded by a set of buyers S with $|S| > |N(S)|$.)
5. Increase the price of each good in $N(S)$ by 1.
6. If every price is greater than 0, reduce all prices by $\min_{j \in J} p_j$, so the lowest price is now equal to 0. (This will ensure prices stay below valuations, so buyer payoffs are non-negative.)
7. Go back to Step 2 of the algorithm with the new price vector.

Analyzing the Algorithm

Clearly, the algorithm can only stop at a competitive equilibrium.

- Since the values v_{ij} are non-negative and the lowest price is always equal to 0, buyers always prefer assigned goods to not buying anything .

Thus, to complete the proof that a competitive equilibrium exists, it suffices to prove that the algorithm must stop at some point (i.e., prices can't cycle forever).

Proof that the Algorithm Stops

To prove that the algorithm stops, we define a non-negative quantity that starts at a finite level and decreases by at least 1 with each iteration of the algorithm.

- Since it can only decrease finitely many times, the algorithm stops after finitely many iterations.

Proof that the Algorithm Stops

To prove that the algorithm stops, we define a non-negative quantity that starts at a finite level and decreases by at least 1 with each iteration of the algorithm.

- Since it can only decrease finitely many times, the algorithm stops after finitely many iterations.

Given the current price vector p, define the **aspiration** of buyer i as the best payoff she can attain at the current prices:

$$\max_j v_{ij} - p_j.$$

Similarly, define the **aspiration** of seller j as the best payoff he can attain at the current prices, which is simply his current price, p_j.

Proof that the Algorithm Stops

To prove that the algorithm stops, we define a non-negative quantity that starts at a finite level and decreases by at least 1 with each iteration of the algorithm.

- Since it can only decrease finitely many times, the algorithm stops after finitely many iterations.

Given the current price vector p, define the **aspiration** of buyer i as the best payoff she can attain at the current prices:

$$\max_j v_{ij} - p_j.$$

Similarly, define the **aspiration** of seller j as the best payoff he can attain at the current prices, which is simply his current price, p_j.

We show that the sum of all agents' aspirations (both buyers' and sellers') starts at a finite level and decreases by at least 1 with each iteration.

- The first part is obvious: the sum of aspirations starts at $\sum_{i \in I} \max_j v_{ij}$, which is finite.

Proof that Aspirations Decrease with Each Iteration

Aspirations change only at two steps of the algorithm: Step 5 (where prices of all goods in $N(S)$ increase by 1, for some constricted set S), and Step 6 (where all prices are reduced by the same non-negative constant).

The change in Step 6 reduces the aspiration of each seller by the constant, and increases the aspiration of each buyer by the same constant.

- This has no effect on the sum of everyone's aspirations.

The change in Step 5 increases the aspiration of each seller in the set $N(S)$ by 1 and decreases the aspiration of each buyer in the set S by 1 (as the prices of **all** of their most-preferred goods increase by 1), and has no effect on anyone else's aspiration.

- The effect on the sum of everyone's aspirations equals $|N(S)| - |S|$. Since S is a constricted set, this is at most -1.

Thus, the sum of everyone's aspirations decreases by at least 1 with each iteration of the algorithm. This completes the proof.

Intuition for the Last Step

Intuitively, by gradually increasing the price of over-demanded goods, the algorithm gradually reduces the "aspirational" total value available to the agents, until it reaches the maximum actually-attainable value V^*, at which point the market clears.

In economics, such a process of gradually increasing the prices of over-demanded goods is called **tâtonnement**. It captures how "market forces" can adjust prices to equate supply and demand.

Summary

- **Matching markets** or **assignment games** capture situations where each seller sells a single (heterogeneous) good and each buyer wants at most one good.
- A **competitive equilibrium** is an assignment of goods to buyers together with prices for all goods that lead each buyer to select her assigned good.
- A competitive equilibrium always exists, and all competitive equilibria generate the maximum possible total value.
- A competitive equilibrium can be viewed a perfect matching in the **preferred-seller graph**. Our proof of its existence takes this perspective and applies **Hall's theorem** on perfect matchings in bipartite graphs.