CS269I: Incentives in Computer Science
Lecture #5: Market-Clearing Prices*

Tim Roughgarden†

October 8, 2018

1 Understanding Your Market

Markets come in different flavors. Here are three questions worth asking about a market, to clarify your thinking about it. While the questions are stated as binary choices, reality is often in between the two extremes.

Question #1: Does a participant care about the identity of her transaction partner?

When the answer is affirmative, the market is called a *matching market*. This was the case in all of our examples in the first two lectures (students care about the room they get, and colleges and students care who they are matched to).

What about in Amazon? Not really—you generally care only about what you get and the price you get it for, not the identity of the seller per se. Such markets are called *markets for goods*.[1] Most of our examples from last week fall into this category, including eBay and the NYSE. Uber and Lyft are also primarily markets for goods—you care about getting the ride and the price you pay for it, but not so much about which driver you are assigned.

Some markets resist easy classification. For example, in Airbnb, the identity of the seller and the good being purchased (a room rental) are usually tightly coupled. So Airbnb is a blend of sorts between a matching market and a market for goods. Some online labor markets (e.g., for hiring a freelancer) also have this character.

†Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

[1]Amazon's reputation system may make you pay a little attention to the seller's identity, but Amazon is still best regarded as a market for goods.

Question #2: Is money involved?

For this question, we see a big difference between the markets we discussed in Week 1 and those in Week 2. None of our Week 1 examples (room assignment, college admissions, resident matching) involved money, at least in the specific models that we studied. Zooming out, though, for the room assignment problem, all the participants are already paying Stanford tuition (or having it paid by someone else, or a scholarship, etc.). It's only after this "entry fee" that money does not play a role in the room assignment. Similarly, in college admissions, colleges cost money (and different amounts for different colleges), and admission offers can come with financial assistance. A more complete model of college admissions would also take these economic factors into account.

Meanwhile, almost all of our Week 2 examples—Amazon, eBay, Airbnb, Google/Facebook, etc.—involved money. The one big exception is dating platforms like Tinder, which are pure matching markets (without money).

Question #3: Are the goods in the market fungible, or idiosyncratic?

Here "fungible" means interchangeable, like shares of a stock or new copies of a book.[2] Lots of the sales on Amazon or eBay involve fungible goods (where all copies of a product sold across sellers are basically the same). By "idiosyncratic," we mean that every good in the market is different from one another, and accordingly buyers may value some more than others.

Sometimes goods are a blend of fungible and idiosyncratic. For example, consider the room assignment problem. All rooms of the same type (e.g., Toyon two-room doubles) are treated as interchangeable, and as a participant in the Draw you are not allowed to express any preference between them. Different room types are treated as idiosyncratic, and participants can express an arbitrary preference over them.

Today's lecture focuses on markets for goods, with prices, and it will segue naturally into the next three lectures (on auctions). The goods may or may not be fungible (we'll look at both cases). The two related questions that we're interested in are:

1. What prices do we expect to see?

2. What prices would we ideally like to see?

2 The Single Good Case

As a warm-up, let's think about the relatively easy case of a single fungible good, like shares of a stock or new copies of a book.[3] Since different copies of the goods are interchangeable, it makes sense to think about the case where all of them are priced the same.

[2] A related term is a "commodity," which typically refers to a raw material like wheat or oil.

[3] The discussion in this section might give you flashbacks to Econ 1, but we'll soon make things more general and interesting.

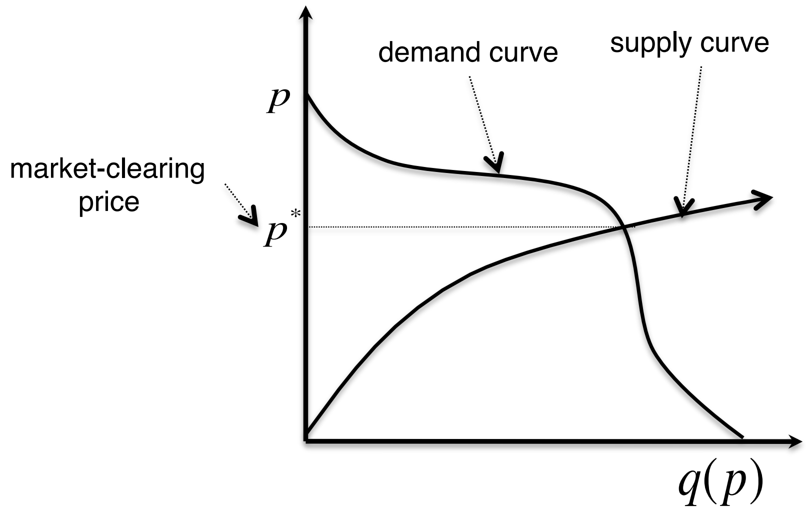

Figure 1: A market-clearing price for interchangeable goods.

Next is one of the most basic primitives in economics.

Vocabulary Lesson
demand curve (n.): a function that specifies the quantity $q(p)$ of a good that buyers are willing to purchase at a per-unit price of p.

Typically, a demand curve is a decreasing function—as the good becomes more expensive, fewer people are willing to buy it. See also Figure 1.[4] A simple example is a linear demand curve, say $q(p) = \{0, 100 - 5p\}$. In this case, when the good is free the demand is 100, and the demand drops to 0 once the price hits 20.

For example, the outstanding buy orders for a stock in the NYSE can be thought of as a demand curve. If there are outstanding buy orders at prices 2, 4, 6, 8, and 10, then

$$q(p) = \begin{cases} 5 & \text{if } p < 2 \\ 4 & \text{if } 2 \leq p < 4 \\ 3 & \text{if } 4 \leq p < 6 \\ 2 & \text{if } 6 \leq p < 8 \\ 1 & \text{if } 8 \leq p < 10 \\ 0 & \text{if } p \geq 10. \end{cases}$$

A demand curve tells you the quantity purchased if you happen to know the going price. But where does this price come from?

[4]By convention, economists plot price on the y-axis and quantity on the x-axis. Thus the figure actually plots the *inverse* demand curve (specifying price as a function of quantity), which contains the same information.

Vocabulary Lesson

supply curve (n.): a function that specifies the quantity $s(p)$ that suppliers are willing to produce if sold at a per-unit price of p.

Typically, $s(p)$ is increasing with p—the more lucrative it is to sell the good, the more suppliers are willing to produce.

For example, the outstanding sell orders for a stock on the NYSE provide a supply curve. If there are outstanding orders with offer prices 1, 3, 5, 7, and 9, then

$$s(p) = \begin{cases} 0 & \text{if } p < 1 \\ 1 & \text{if } 1 \leq p < 3 \\ 2 & \text{if } 3 \leq p < 5 \\ 3 & \text{if } 5 \leq p < 7 \\ 4 & \text{if } 7 \leq p < 0 \\ 5 & \text{if } p \geq 9. \end{cases}$$

A *market-clearing price* is then a price p^* that equalizes supply and demand: $q(p^*) = s(p^*)$. For example, in Figure 1, the market-clearing price is the y-value at which the two curves cross. In our NYSE example, any price between 5 and 6 is market-clearing (with supply and demand equal to 3).

A market-clearing price is more or less what we expect to see in the case of a single fungible good with lots of buyers and sellers. (Is it what we want to see? Hold that thought.) In some markets, these are the only concepts you need to have a first-order understanding of what's going on. It's exactly what's happening in the NYSE, simultaneously in parallel across thousands of stocks. Large swaths of Amazon act in a similar way (in parallel across many products).

What about a ride-sharing platform like Uber and Lyft? One difference to note is that the price in Uber/Lyft is dictated centrally by the platform, while in a decentralized market like the NYSE it emerges organically from the prices quoted by buyers and sellers. So the question is: what kinds of prices do we expect Uber/Lyft to choose? These companies are famously secretive about their pricing strategies and algorithms, but presumably it's not a million miles away from trying to balance supply and demand.[5] For example, when does surge pricing get triggered? When the supply (of drivers) falls too far below the demand (from riders). What does surge pricing accomplish? A higher price results in both increased supply and decreased demand, bringing supply and demand closer together.[6]

[5] There are of course additional complications, for example caused by the spatial and temporal dimensions of the market.

[6] For all its bad publicity, surge pricing seems natural and even inevitable to an economist. (What's in a name? A lot. "Happy hours" at bars serve essentially the same function as surge pricing (if not as dynamically), but they never get bad publicity.) Of course, the pure economist is baffled to ever wait for a table at a restaurant—if they're so busy, why don't they raise their prices? As always, reality is much more complex than our toy models (but there is no shame in this!).

3 Idiosyncratic Goods

Our model so far is clearly inadequate for analyzing, say, Airbnb. The goods for sale in Airbnb (spare rooms) all differ from one another, and buyers will care about which one they get (and also the price they get it at). What would be an analog of a market-clearing price? What prices might we expect to arise in a decentralized market, and what prices would make sense to impose in a centralized market? Presumably, there are multiple prices involved, with more popular goods going for higher prices.

3.1 The Model

Here's the model we'll use for Airbnb and other similar types of markets, where each buyer is buying at most one of the goods for sale.[7]

- There are n potential buyers (e.g., potential renters).
- There m different sellers (e.g., hosts), each with one (idiosyncratic) good (e.g., a spare room).
- For each buyer i and good j, let v_{ij} denote the maximum price that i would be willing to pay for j; this is often called the buyer's *value* or *valuation* for good j. Because the goods are idiosyncratic, a buyer can have a different valuation for each of them (as would be the case in Airbnb).[8]
- For each good j, r_j is the minimum price that its seller is willing to accept for it; this can be thought of as the seller's *reserve price* (e.g., the expected cost to the seller of renting their spare room for a night).[9]

Next we make two simplifying assumptions.

1. To save on notation, we'll assume that $r_j = 0$ for all goods j for the rest of the lecture. Everything we'll say extends to the case of general r_j's (see Homework #3).
2. We'll assume that each buyer has an "outside option" of not buying any of the goods. Each buyer has value 0 for their outside option. (Again, the rest of the lecture would extend to outside options with arbitrary values.)

[7]Another example would be an online labor market, where the buyers are firms, the goods are (heterogeneous) workers, and prices correspond to wages.

[8]The very definition of a valuation refers to a price. This is why we never talked about them in our applications without prices, like the room assignment and college admission problems. Without valuations, all we had to work with was relative preferences. With prices and valuations, we also know *how much* a participant favors one outcome over another (as measured in dollars).

[9]By allowing only a single reserve price for each seller (independent of the buyer), the model assumes that sellers treat buyers as interchangeable and only care about the price they get. To first order, this appears to be the case in a system like Airbnb.

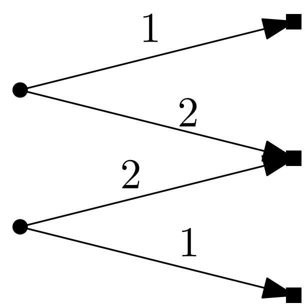

Figure 2: Visualizing a market with idiosyncratic goods as a bipartite graph. Edges are annotated with valuations. Zero-value edges and outside options are not shown.

The model can be usefully visualized as a bipartite graph (Figure 2), with buyers on one side, goods on the other side, and each edge (between a buyer and a good) annotated with the buyer's valuation for the good. In Figure 2, zero-value edges and the outside options are omitted from the picture.

3.2 Competitive Equilibrium

An *outcome* in this model has two ingredients, specifying who gets what, and who pays what. Formally, the first ingredient is a *matching* M, where $M(i)$ denotes the good assigned to buyer i (or to i's outside option, as appropriate). In M, every good is assigned to at most one buyer, and every buyer is assigned either a good or her outside option. The second ingredient is a price vector $\mathbf{p}$, where $p_j \geq 0$ specifies the (nonnegative) selling price of the jth good. (Outside options always have price 0.)

What kind of outcomes do we want or expect?

Definition 3.1 (Competitive Equilibrium) An outcome is a *competitive equilibrium (CE)*[10] if:

(a) whenever a buyer i is assigned a good or outside option j,

$$\underbrace{v_{ij} - p_j}_{\text{utility of } j\text{, given } \mathbf{p}} \quad \geq \quad \underbrace{v_{ij'} - p_{j'}}_{\text{utility of } j'\text{, given } \mathbf{p}} \tag{1}$$

for all possibilities j' (a good or i's outside option).[11]

(b) If the jth good is unassigned, then $p_j = 0$.

This is our definition of "market-clearing prices" when there are heterogeneous goods (and every buyer wants only one good). Part (a) says that, given the prices, all buyers are as

[10] Also known as a *Walrasian equilibrium*.

[11] When j or j' is i's outside option, we interpret both the value and price to be zero.

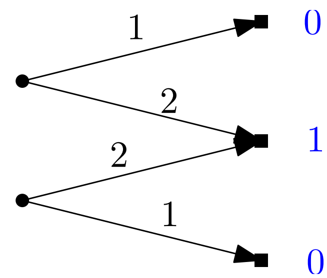

Figure 3: A price vector that participates in two competitive equilibria. Edges are labeled with valuations, goods with prices.

happy as they could be. This property by itself can be satisfied in a trivial and undesirable way, by setting $p_j = +\infty$ for every good j and assigning all buyers to their outside options. Part (b) of the definition requires that the outcome is "market clearing," in the sense that every good which is desired is sold.

To start understanding this definition, first note that part (a) implies that every buyer has nonnegative utility (since when j' is i's outside option, the right-hand side of (1) is 0). Conversely, if the ith buyer actually wants some good j' (i.e., $v_{ij'} > p_{j'}$), then i is not assigned to her outside option (otherwise the left-hand side of (1) would be 0 while the right-hand side would be strictly positive).

3.3 Discussion

The requirements of a competitive equilibrium are strong.[12] In effect, we are putting a price tag p_j on each good, and letting each buyer i independently pick whichever good she wants (or her outside option). Magically, there are no conflicts and every buyer simultaneously gets what she wants.[13,14]

Returning to our example (Figure 3), suppose we label the two buyers A and B and the three goods 1, 2, and 3 (in both cases, from top to bottom). Set $p_1 = 0$, $p_2 = 1$, and $p_3 = 0$. Is it possible to define a matching M so that the pair $(M, \mathbf{p})$ constitutes a competitive equilibrium? We can't assign A to 3 or B to 1 in a CE, since this would give the buyer 0 utility despite the presence of a good (e.g., good 2) for which the buyer has positive utility (violating property (a) of CE). Similarly, we can't assign A or B to her outside option. So we need to assign A to either 1 or 2 and B to either 2 or 3. There are three ways to do this (as B can only be assigned to one of them). If we assign A to 1 and B to 3, then we violate requirement (b) of Definition 3.1 and do not get a CE. The other two matchings (A to 1 and B to 2, or A to 2 and B to 3) are CE.

[12] If you're worried about whether they actually exist, see Section 3.5.

[13] There might be a tie between goods for being a buyer i's favorite, in which case we allow these ties to be broken in a coordinated way.

[14] This property may remind you of the student-optimality property of the Deferred Acceptance algorithm (Lecture #2).

How do we feel about the definition of a competitive equilibrium? There's a plausible narrative about why we might expect prices in a market like Airbnb to resemble a competitive equilibrium. In an outcome $(M, \mathbf{p})$ where property (b) is violated, we might expect the seller of good j to decrease her price. If property (a) is violated, then there's a good j which is the favorite of two or more different buyers, and we might expect the seller of good j to increase her price. In other words, an outcome that is not a competitive equilibrium is unlikely to persist for long. Meanwhile, if an outcome *is* a competitive equilibrium, then all buyers are happy (given the prices) and there is no upward or downward price pressure on any good, so we might expect the outcome to persist.

3.4 First Welfare Theorem

Should we be happy with the prices at a competitive equilibrium? In at least one sense, yes. In our example in Figure 3, we observed that one matching ($A \to 1$ and $B \to 3$) is not a CE, while the two others ($A \to 1$ and $B \to 2$, and $A \to 2$ and $B \to 3$) are CE. Notice any other differences between the first matching and the latter two? One difference is that the total value of the first matching is $1 + 1 = 2$, while that of the second and third are $1 + 2 = 3$. More generally, it's easy to see that the matchings in our two CEs are exactly the matchings with maximum-possible total value. This is just a simple example, however; could it be true in general?

Theorem 3.2 (First Welfare Theorem) *If $(M, \mathbf{p})$ is a competitive equilibrium, then M is a matching with maximum total value. That is,*

$$\sum_{i=1}^{n} v_{iM(i)} \geq \sum_{i=1}^{n} v_{iM'(i)}$$

for every matching M'.

Proof: Let's follow our nose. Consider some matching M^* with the maximum-possible total value. All we have going for us is our assumption that $(M, \mathbf{p})$ is a CE. The first condition of Definition 3.1 feels like the stronger one, but how can we use it? In particular, how should we choose j' on the right-hand side of (1)? The only thing we know other than M is M^*, so one natural option is to take $j' = M^*(i)$. That is, we'll use the fact that each bidder i prefers the good she received in M (at prices $\mathbf{p}$) to the good she received in M^* (again at prices $\mathbf{p}$) to obtain:

$$v_{iM(i)} - p_{M(i)} \geq v_{iM^*(i)} - p_{M^*(i)}.$$

(As usual, if $M(i)$ or $M^*(i)$ is i's outside option, then the valuation and price are both interpreted as zero.) Denote the sum $\sum_{j=1}^{m} p_j$ of all prices by P. Summing the inequality above over all bidders i gives us

$$\underbrace{\sum_{i=1}^{n} v_{iM(i)}}_{\text{total value of } M} - \underbrace{\sum_{i=1}^{n} p_{M(i)}}_{= P \text{ by CE property (b)}} \geq \underbrace{\sum_{i=1}^{n} v_{iM^*(i)}}_{\text{total value of } M^*} - \underbrace{\sum_{i=1}^{n} p_{M^*(i)}}_{\leq P},$$

where we are using the fact that $\sum_{i=1}^{n} p_{M(i)}$ sums over all of the items that have a non-zero price (and hence the sum equals P), and the fact that M^* can only assign each good to one buyer (and thus $\sum_{i=1}^{n} p_{M^*(i)}$ is at most P). Rearranging terms shows that the total value of M is at least that of M^*, and hence M is also an optimal matching. ■

Several comments.

1. It's kind of amazing that a competitive equilibrium automatically solves a non-trivial computational problem, namely computing a maximum-weight matching in a bipartite graph. (The problem is polynomial-time solvable, but the algorithms for it are too advanced for CS161 and are taught only in CS261.)

2. If we're happy with the objective of maximizing the total value of the matching, then every competitive equilibrium gives us the best-possible matching we could have hoped for. In this sense, we should be happy if a market reaches a competitive equilibrium.

3. Let's compare the First Welfare Theorem to the properties we proved about stable matchings (Lecture #2). In stable matching, there was no money, so it didn't make sense to define valuations for anybody, and so there was no way to talk about the "total value" of a matching. In other words, we didn't even have the vocabulary to state a result like Theorem 3.2 back in Week #1, and had to settle for proving that every stable matching is Pareto optimal. Here, because we have a quantitative measure of buyers' preferences (via their valuations), we can also aggregate across them (by summing). This implies, in particular, that every competitive equilibrium is Pareto optimal (see Homework #3).[15]

4. When you hear an economist say that "markets are efficient," they usually mean a statement along the lines of Theorem 3.2 (which shows up in different forms in different models).

3.5 Existence and Computation

Given the strength of the CE requirements and the First Welfare Theorem, you'd be right to wonder whether they are guaranteed to exist.[16,17] In the current model, with at most one good assigned to each buyer, they do.

[15]Remember that Pareto optimality is a necessary but not sufficient condition for an outcome to be "good," and in particular that there can be lots of inequity in a Pareto optimal outcome.

[16]Remember, any random person can write down a definition, you can't stop them. When confronted with a new definition, always ask: (i) are there any interesting examples that satisfy the definition? and (ii) are there any interesting consequences of satisfying the definition? (Here, the answers are yes and yes: a CE exists in every market of the type we're considering in this lecture, and the First Welfare Theorem is an interesting property of CE.)

[17]For guaranteed existence, it's important that at most one good can be assigned to each buyer (as is typical in Airbnb, say). Next lecture we'll talk about spectrum auctions, where a buyer might be assigned multiple items. In this more general setting, competitive equilibria are not guaranteed to exist.

Theorem 3.3 (Existence of Competitive Equilibria) *In every market of the above type, there is at least one competitive equilibrium.*

The proof is constructive, and uses a variant of the Deferred Acceptance algorithm from Lecture #2. That might sound nonsensical, since competitive equilibria involve prices and the Deferred Acceptance algorithm used no prices. So we need to extend the algorithm so that it includes prices. Buyers will propose to sellers, with a price attached to the proposal. (Assume that there is a finite number of possible prices, like all multiples of \$10 between 0 and \$1000.) Think of a buyer i as having a ranked list of all the possible (good, price) pairs (j, p), listed in decreasing order of $v_{ij} - p$. Think of a seller as having a ranked list over the same set, listed in decreasing order of price. (In both cases, break ties arbitrarily.) This induces an instance of stable matching of the type studied in Lecture #2, and it typechecks to run the Deferred Acceptance algorithm on it. The argument that the original Deferred Acceptance algorithm terminates with a stable matching translates to the present algorithm and shows that it terminates with a competitive equilibrium (up to a small discretization error).

In more detail, consider the following variant of the Deferred Acceptance algorithm:

Deferred Acceptance Algorithm (with Prices)

while there is at least one unassigned buyer **do**
  every unassigned buyer i "proposes" a price p to a seller j, where j and p are chosen to maximize $v_{ij} - p$ over all pairs (j, p) that have not already been rejected[18]
  every seller retains the highest-price offer ever received, and rejects all other proposals
all unrejected proposals are made final

Because every buyer is always tentatively assigned to at most one seller and every seller to at most one buyer, the algorithm terminates with a matching M. It also terminates with a price vector $\mathbf{p}$, where p_j is the final (and hence highest) offer than seller j accepted. (If nobody ever proposed to a seller j, then we define p_j as 0.)

A seller only rejects a proposal in favor of one with a higher price, and buyers never withdraw proposals. Thus once a seller is matched, it is matched forevermore, and at ever-higher prices. Hence, the only way a good can go unassigned is if its seller was never proposed to, in which case the final price of the good is 0. We conclude that the output $(M, \mathbf{p})$ of this algorithm satisfies property (b) of a competitive equilibrium.

To argue property (a), fix a buyer i and a good j'. If $j' = M(i)$ then equality holds in (1) and we're done, so assume that $M(i)$ and j' are different. Why wasn't i assigned to j' by our extended Deferred Acceptance algorithm? One possibility is that i never proposed to j', not even at a price of 0. Since buyer i worked her way down her (good,price) options in

[18]If nothing else, the buyer can "propose" to her outside option at a price of 0.

decreasing order of $v_{ij} - p$, and since her final proposal was to the good $M(i)$ with a price of $p_{M(i)}$ (rather than to j' with a price of 0), we have

$$v_{iM(i)} - p_{M(i)} \geq v_{ij'} - 0 \geq v_{ij'} - p_{j'},$$

as required.

The other possibility is that all of i's proposals to j' were rejected in favor of better (i.e., higher-priced) offers. Suppose i's final proposal to j' was at the price p, and so i never made an offer to j' at a price of $p + \epsilon$. (We're assuming that all prices are restricted to be multiples of ϵ.) Since i instead made an offer to $M(i)$ at a price of $p_{M(i)}$, it must have been that

$$v_{iM(i)} - p_{M(i)} \geq v_{ij'} - (p + \epsilon).$$

Since the price of good j' only increases over the course of the algorithm, this inequality also holds at termination:

$$v_{iM(i)} - p_{M(i)} \geq v_{ij'} - p_{j'} - \epsilon.$$

This verifies property (a) of a competitive equilibrium, up to the discretization error ϵ.

References