CS269I: Incentives in Computer Science
Lecture #3: Markets, Everywhere*

Tim Roughgarden†

October 1, 2018

1 Preamble

Last week we talked about several markets that you're familiar with: the market for rooms at Stanford, the market for admission to competitive colleges, and the market for medical residents. While some of our discussion had a computer science-y flavor—we did talk a fair amount about algorithms, like the Deferred Acceptance algorithm—the applications themselves were pretty old-school, predating computer science as an intellectual discipline. Today we'll discuss several markets that simply would not exist without the advances we've had in computer science over the past several decades. These are also the sorts of markets that some of you are likely to wind up working on, post-graduation.[1]

This lecture will be pretty wide-ranging, so here are a few themes you should be on the lookout for.

1. Markets are everywhere. After going through lots and lots of examples, you might see the world a little differently, and start recognizing markets where you never saw them before.

2. Many markets share some of the same features. Thus, when grappling with a market's issues in one domain, it's often useful to look to other domains for ideas and inspiration.[2]

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

[1]If you want to dig deeper into the topics of these first three lectures, the catchphrase you're looking for is "market design." There are entire courses devoted to market design in the economics department, the management science and engineering department, and the business school. Stanford is arguably the best place in the world to study or do research in the subject.

[2]For example, in Lectures #1 and #2, we starting looking at the market for doctors to organize our brainstorming for an alternative way of running college admissions in the U.S.

3. Each market has its own idiosyncratic quirks, and the details can have a big effect on the viability of different solutions.

4. Lots can go wrong in a market, but there are techniques for keeping them running relatively smoothly. The primary point of this and the next lecture is introduce you to some of these techniques.

2 Markets from (Almost) the 21st Century

Perhaps the most obvious example of a market that wouldn't exist without advances in computer science is Amazon.

Example #1: Amazon.

There would be no Amazon without the Internet. There would be no Amazon without the World Wide Web. Amazon would not have quite the same reach without modern wireless communication and smartphone technology.

What did Amazon bring to the table? After all, pre-Amazon, it's not like it was hard to buy a book—you'd just go to your local bookstore and buy it or ask them to order it. Similarly for clothes, for food, etc. Amazon's main contribution was to make an existing market better. For example:

- More convenience. Open 24 hours a day, no travel required, all genres of products in the same place.

- More efficiency. Ideally, because of its scale, Amazon should be able to supply a given set of products at a relatively low cost, hopefully translating to lower prices for consumers.

Switching from brick-and-mortar shops to Amazon wasn't all upside, however. On Amazon, you generally don't know much about the seller you're buying from, while you may have a relationship with the proprietors at your local shop. Shopping on Amazon was potentially less safe than traditional shopping.[3]

How about other examples of companies whose primary contribution was improving a market that already existed? Advertising companies like Google and Facebook are good examples.

Example #2: Google; Facebook.

Maybe you think of Google and Facebook as a search engine and a social network, respectively. But if you look at how they pay the bills, they're really massive advertising

[3]Reputation systems go some way toward mitigating this safety issue. More on reputation systems next lecture.

firms (with the vast majority of revenue coming from advertising).[4] Advertising has of course been a huge market forever—we've all seen *Mad Men*—but the rise of platforms like Google and Facebook allowed ads to be targeted at a far finer level of granularity than had ever been possible before.[5]

What about ride-sharing companies?

Example #3: Uber; Lyft.

To the extent that there was a ride-sharing market before Uber and Lyft, it was small, consisting of taxis plus various ad hoc arrangements.[6] Thus Uber and Lyft really created a new market, or at least grew an existing market by an order of magnitude or more. Uber started with limos, which made sense, as limo drivers could traditionally only be booked by appointment, and as a result often had long idle periods between rides. Many of these drivers were happy to earn some extra money through Uber during these idle times. Of course, it's now clear that the ride-sharing market is much bigger than was initially thought, with lots of amateur drivers happy to lend their time and car part-time to pick up some extra cash.[7]

There are many more examples of the same type.

Example #4: Airbnb. Airbnb opened up the market for spare rooms.[8]

Example #5: eBay. These days, you might regard eBay primarily as a smaller and less organized version of Amazon. But when eBay debuted back in 1995, it looked rather different, like the union of all the world's garage sales happening at the same time and place. Thus eBay first rose to prominence by opening up the market for second-hand bric-a-brac.

Example #6: StubHub. StubHub opened up the market for second-hand tickets to events like concerts and sporting contests. This was not a new market—there's always been ticket scalpers and ticket-resellers—but StubHub grew the market by making it more convenient and in many cases safer.

Example #7: Upwork, TaskRabbit, Postmates, Fiverr, etc. These companies brought parts of the labor market online (especially temp jobs and freelancing). Again, there's always been freelancing, but not at today's scale.[9]

---

[4] When I was a kid, I watched a lot of television. I thought that the point of a TV channel was to bring me cool programming. Only later on did I realize that the real point of a TV channel was to get my attention, so that it could be redirected to sponsored content.

[5] We'll say much more about how these systems actually work in a couple of weeks.

[6] Twenty-five or more years ago, at Tressider student union you could find a collection of index cards posted on a (physical) bulletin board, with offers and requests for rides.

[7] Like with Amazon, convenience was also a big factor in Uber's success. Most 21st-century success stories of new markets in the commercial sector both grew an existing market and also made it better in various ways.

[8] Craigslist served as an intermediate point in the evolution from traditional rental markets to Airbnb.

[9] A company like Monster.com is arguably more along the lines of Amazon, taking a well-established

Example #8: Tinder, etc. There's always been a market for dates, but it's never been as explicit as it is now.

3 Centralized vs. Decentralized Markets

3.1 Definitions and Examples

A market is *centralized* if the transactions that get executed are dictated by a third party. An extreme example would be many markets in the former Soviet Union. Two mechanisms from last week, the Draw and the Deferred Acceptance algorithm, correspond to centralized markets (all matches are centrally computed by the mechanism, and all participants can do is report their preferences).

Conversely, in a *decentralized* market, participants are free to transact directly. College admissions in the U.S. is decentralized, in that neither the actions of different applicants nor of different colleges are coordinated by any centralized mechanism.[10]

How about Examples #1–#8? Mostly, they are decentralized, but with one exception: Uber and Lyft. Both manage a two-sided market of drivers and riders, with every match centrally dictated by the platform. Drivers can to some extent pick which ride assignments to accept, but as a rider, you have to accept the driver assigned to you or face a cancellation penalty.[11,12]

3.2 Recommended Transactions

Markets are often somewhere in between centralized and decentralized. Even a seemingly decentralized market like Amazon influences the transactions you engage in via its (centralized) search functionality and user interface. More generally, an oft-used middle ground between the two extremes is for a third party (like the platform) to recommend several transactions, perhaps based on the data it has about a user and its possible transaction partners, while users remain free to pick any or none of them. (Uber/Lyft is the extreme case where only one transaction is recommended, and users cannot engage in a non-recommended transaction.)

---

market (the labor market for full-time employees) and making it more convenient.

[10] Warning: these days, when you use the word "decentralized" among computer scientists, they tend to assume that you're talking about blockchain. (Which we will talk about later, but not now.) We mean it in the older economic sense of the word, as the antithesis of a centralized economy. For example, one would call a capitalist economy like that of the U.S. a (primarily) decentralized market.

[11] Uber and Lyft do differ from a centralized mechanism like the Draw or the Deferred Acceptance algorithm in that they are asynchronous, which participants on each side of the market entering and exiting at whatever time they please.

[12] Why does Uber even bother to show you the rating of your driver? It's not like you're allowed to use it in choosing whom to transact with (like you can in Airbnb, Amazon, etc.). Presumably, Uber is collecting driver ratings anyways (with the important goal of identifying poor drivers), so why not pass them on to riders (and maybe make them feel a little better about the driver they're being forced to accept)?

3.3 Batched Transactions

A second way to interpolate between the extreme points of decentralization and centralization is to *batch* subsets of transactions for centralized processing. Let's look at two very different examples.

Hinge. Hinge is a new-ish dating app. Users can browse profiles, initiate a conversation, and upvote different parts of others' profiles. In addition to this, once per day, Hinge does the following centralized computation using everybody currently active in the system:[13]

1. For each user, infer a preference ranking over all other users from the user's past actions (i.e., which parts of which profiles they looked at and upvoted).
2. Run the—wait for it—Deferred Acceptance algorithm, using the inferred preference lists! A user's match in this algorithm is displayed in their feed as a recommended "match of the day."[14]

Thus batching enabled Hinge to shoehorn a seemingly centralized idea (stable matching) into a decentralized environment.

NYSE. For a very different example, consider the New York Stock Exchange (NYSE). This market is about as decentralized as it gets, and anyone can place a buy or sell order (at a given price) on any stock at any time (during opening hours). Trade occurs whenever there is a buy order that quotes a price at least that of some sell order, and matching transactions are executed immediately, on a first-come first-serve basis.

An issue with the current implementation of the NYSE is *high-frequency trading*. To give one example of how this comes up, imagine that you submit an order to buy one share of Facebook stock at a price of $200. Suppose the order remains outstanding, say because all standing sell orders are at prices $210 or higher. Now imagine that the Cambridge Analytica scandal hits, and you suspect that Facebook's stock price is about to drop dramatically. The race is on: you want to cancel your standing buy order, while opportunistic sellers want to beat you to it and clear your order (at the now-lucrative price of $200). Whether or not the transaction gets executed depends on who gets to the market first. (Remember that transactions are executed serially, in first-come first-serve fashion.)

This is just one of many scenarios where speed translates directly into dollars on the NYSE. In fact, in the U.S. alone, a conservative estimate is that the high-frequency trading industry earns several billion dollars a year. From this, we can conclude that there are

[13] See `https://techcrunch.com/2018/07/11/hinge-employs-new-algorithm-to-find-your-most-compatible-match-for-you/`.

[14] Actually, this isn't quite right. What are the two sides of the market? For the purely heterosexual subset of the market, they would be the men and the women, and the Deferred Acceptance algorithm makes perfect sense. For the rest of the market, Hinge uses an algorithm for the *stable roommates* problem, where anyone can be paired with anyone else. Unlike in stable matching, there may not be any stable solution to a stable roommates problem. Still, there are Deferred Acceptance-style proposal algorithms that compute useful solutions, including a stable solution whenever one exists [4].

significant arbitrage opportunities at the millisecond time scale (e.g., between essentially identical index funds on the Chicago and New York exchanges).

With this background, maybe it's not surprising that, in 2010, a company called Spread Networks invested 300 million dollars into covertly building a high-speed fiber-optic cable connecting New York and Chicago (the latter also hosts a major stock exchange). Of course, such cables already existed. Their innovation? Rather then circuitously following railroad tracks or dodging mountains, Spread Networks's cable was on a straighter line than any of the existing ones. The benefit? A Chicago-New York round-trip time of 13.3 milliseconds, shaving a whole *3 milliseconds* off of the previous record. (The blink of an eye takes at least 100 milliseconds.) This improvement was enough to make a lot of money.[15] The joke at the time was that the next innovation would be a cable that tunneled directly through the Earth, thereby avoiding the Earth's pesky curvature. In reality, by 2011 competitors were already switching from fiber-optic cable to microwaves to get even better round-trip times (the refractive index of air is a bit less than that of glass).

Intuitively, there should be better uses of resources than building infrastructure to participate in the high-frequency trading rat race.[16] Could we redesign the NYSE to mitigate the incentive to engage in high-frequency trading?

Budish, Cramton, and Shim [2] propose an elegant solution: to collect transactions into *batches* and compute all the trades in a batch at the same time. For example, imagine that we break time into one-second intervals. At the end of each second, the exchange considers all the standing buy and sell orders (possibly submitted during the preceding second, or possibly submitted earlier and not yet executed or canceled), and removes any orders that were canceled in the preceding second. With the remaining buy and sell orders, the exchange executes as many trades as possible, all at a common price, subject to the buyers paying no more than their bid and sellers receiving at least their ask.[17] With one batch per second, an extra three milliseconds of speed is no longer very valuable.[18]

Budish et al.'s suggestion has not yet been adopted. As you can imagine, there's a lot of inertia to make any changes to the NYSE, and also certain parties are lobbying against the change (like those that financially benefit from high-frequency trading). However, because the change is conceptually simple and the authors are putting in the work to drum up support for the idea, it's conceivable that some version of it will be implemented in the next 5–10 years.[19]

[15] What matters is being a tiny bit faster than your competitors—as long as you get to the market first, you're good.

[16] This intuition can be turned into a forceful economic argument, see [2].

[17] For example, imagine there are five sell orders (one share each) for prices 1, 3, 5, 7, and 9, and five buy orders for prices 2, 4, 6, 8, and 10. Three trades will be executed, involving the three buyers willing to pay at least 6 and the three sellers willing to accept at most 5. The common selling price determined by the exchange will be somewhere between 5 and 6—an example of a *market-clearing* price.

[18] Speed still helps, since there's a chance that you'll sneak in at the end of one time interval while your competitor's order arrives at the beginning of the next one. But the chance of this goes to zero quickly as the time interval grows large.

[19] Perhaps the computerized version of the NYSE should have been implemented this way in the first place. In effect, the strong incentives for high-frequency trading are caused by a bug in the original implementation,

4 Types of Market Failure

Lots of things can go wrong in a market. In this section we'll catalog four of them, and there will be more in future lectures. The purpose of this discussion is to give you tools and language to diagnose poorly functioning markets that you come across, and alert you to some pitfalls to avoid if you're creating your own market.

4.1 Timing Issues

We've already seen some examples of how timing issues can roil a market. In Lecture #1 we talked about how markets can unravel (with offers being made earlier and earlier), and about the chaos caused by exploding offers in the pre-1952 market for medical doctors (and in the market for law graduates to this day). High-frequency trading in the NYSE is another good example, a byproduct of the current implementation's over-emphasis on the exact order of operations.

Imagine if Stanford and other top schools switched to admitting students on a first-come first-serve basis, say with applications accepted starting at the stroke of midnight on March 15th. What would happen? After one year? After five years? How would the schools' makeup and reputation change? (Seriously, think about it.)

4.2 Safety Issues

A market has a safety issue when participants do not feel safe engaging in transactions in the market. A black market would be a canonical example. Online platforms that eventually lead to face-to-face meetings (like dating platforms) obviously have to pay attention to safety issues. Another interesting example is Craigslist—over the years, the fraction of fraudulent listings seems to have grown, and as a consequence the site increasingly suffers from safety issues.

4.3 Not Thick Enough

A market is *thick* if, from the perspective of a participant, it's easy to find a desirable trading partner. Otherwise, the market is *thin*. Returning to our running examples, Amazon and eBay would certainly qualify as thick. For Uber/Lyft, Airbnb, or Tinder, the corresponding market may be thick or thin, depending on the location (e.g., compare San Francisco vs. Death Valley). Thin advertising markets can also show up in Google and Facebook, because the targeting of users can be so fine-grained.

The most obvious reason for a thin market is that not enough people show up (see: countless failed startups). Craigslist is again a good example. Back in the late 90s and early 00s, it was *the* place to find housing; now, it seems to get thinner by the day (presumably

and are an artifact of the specific mechanism used. We should be reminded here of a lesson learned in Lecture #1: small changes to a mechanism can make a big difference in what users do and how well the system functions!

because of the rise of Airbnb and the aforementioned safety issues). Fixing this problem is extremely important, but it has as much to do with marketing strategies as technology, and we won't say much more about it.

Even if there are lots of participants, there is a thinness problem if a participant can't easily *find* someone to transact with. Here, technology and your design and business decisions can help thicken a market. For example:

1. Good search functionality is obviously important—a user has to be able to quickly and easily find promising trading partners.

2. More generally, keep transaction costs (in terms of time and money) low.

3. Make it easy for a user to offer a transaction to multiple parties at once. For example, in U.S. college admissions, the introduction of the Common Application made it much easier for students to apply to lots of colleges, and it had the effect of thickening the college admissions market. Or, in keyword-based online advertising, most systems allow an advertiser to enable "broad match," where by default an advertiser's bid for clicks on one keyword is automatically extended to apply to closely related keywords and phrases, as well.[20]

4. When Apple first announced its mobile operating system, it restricted access to its API for apps to approved app developers (i.e., no third-party developers allowed). Android, meanwhile, made its API open and welcomed all comers. Apple suffered some serious blowback and quickly changed its mind. This decision made the market for iPhone apps thicker, with more developers and consequently more apps available to users.

Is thickness always a good thing? For a platform whose raison d'être is to enable transactions, the answer is mostly yes. (Modulo congestion issues, discussed in the next section.) What about for participants? For example, suppose you're trying to sell a concert ticket on StubHub. Are you happy if the market grows thicker? It depends: you're very happy for the *buyer* side to get thicker, since it increases your chance of a sale; but you're unhappy for the seller side to get thicker, since your chance of losing a sale to a competitor increases.

4.4 Congestion

Congestion occurs in markets in which there's too much of a good thing, with participants overwhelmed by the number of options.

Let's start with a centralized clearinghouse, like the Draw or the Deferred Acceptance algorithm. Is congestion a big deal? Does anything get harder as the market grows bigger and thicker? Formulating and communicating preferences definitely gets harder for the participants—the more options you have to rank, the more onerous the task. When the market becomes sufficiently large, participants' rankings need to be communicated implicitly

[20]For example, a bid placed for the keyword "algorithms" might be automatically extended to "algorithms book," "algorithms textbook," and possibly even "computer programming."

or perhaps just inferred by the system (as in Hinge, described in Section 3). But once these rankings are in hand, scale poses no problem (unless its really massive)—both the Draw and stable matching are efficient algorithms that scale easily to millions of users.

In a decentralized market with congestion, users can't possibly contemplate all of the available options. Sometimes, the solution is as simple as incorporating good search functionality or a good recommendation system. Other times, the problem is trickier. In dating platforms, for example, the most popular profiles receive far too many messages to reply to, and it's not always easy (for the user or an algorithm) to figure out which ones to focus on.

A natural idea is to push back against thickness and low transaction costs, so that an offer to transact becomes meaningful again—hunting for a sweet spot between the thinness and congestion problems. Hinge's "match of the day" suggestion is one approach. But is it possible for the users to directly differentiate themselves, with no platform intervention needed? Section 6 discusses one possible solution: costly signals.

5 Digression: Network Effects

5.1 Network Effects

Before discussing solutions to the congestion problem, we digress to explain why the thinness and congestion problems are fundamental and ubiquitous. The reason is "network effects."

> **Vocabulary Lesson**
>
> *network effect* (n.): the effect that one user of a good or service has on the value of that good or service to its other users.

Unless otherwise noted, network efforts are generally assumed to be a good thing, with the value of the good or service increasing with the number of users.[21]

For example, should we be surprised that there is a very small number of big social networks? No: most people want to use the social network that everyone else is using, resulting in one or a few networks having almost all the users.

Similarly, should we be surprised that so few operating systems are dominant? No, for the same reasons. As a user, you generally want to use the OS with the deepest pool of developers; as a developer, you want to write apps or software for an OS with lots of users.

There are strong network effects in many of the types of markets we've been discussing, and these inevitable lead to one or a small number of markets that suffer from congestion, and a large number of markets that suffer from thinness. Note that, given the choice between these two challenges, congestion is the problem that you want!

[21]A related idea is "Metcalfe's law," which points out that, given that the number of trading partners for a fixed participant grows linearly with the number of other users, the *total* number of possible pairings grows super-linearly (in fact, quadratically).

5.2 Digression within a Digression: Antitrust Regulation

You sometimes hear Silicon Valley leaders griping about the prospect of government regulation. But why would anyone want to regulate the big tech companies, anyway? Networks effects are often part of the concern. The worry is that, when there are strong network effects, a platform with many users can exploit its users' high switching costs (in terms of the lost network effects) by failing to innovate, raising prices, or engaging in anticompetitive behavior. Such worries lead some to call for increased regulation—at the extreme, for breaking up a company like Google into multiple smaller independent companies (as was done with AT&T/Bell in the 1980s, and Standard Oil seven decades earlier).

However, having a big market share is not in and of itself a crime, and nobody wants to simply penalize success on principle. To justify government action, it's necessary that a company also engages in anticompetitive behavior. This was the government's charge in *U.S. vs. Microsoft, 2001*; they claimed that Microsoft was exploiting its dominance in the operating systems market to also dominate the Internet browser market (by bundling Internet Explorer with Microsoft's operating system), without necessarily offering a superior product. You may know that Google is currently embroiled in a similar situation, and stands accused of exploiting its dominance in search to conquer other sectors. The European Union fined Google about 5 billion USD for anticompetitive behavior, while the U.S. has not taken any action (yet).[22]

6 Signaling

We now return to the problem of congestion in a market, where participants are overwhelmed by the number of options. One approach to exporting more information and guidance to participants is to allow a limited amount of *signaling*.

There are different kinds of signals. For example, you can signal your qualifications in a college application (grades, test scores, letters, etc.), a job application, a dating profile, and so on. (More on this next lecture, when we talk about mitigating adverse selection.)

Signals can also be used to indicate exceptional *interest* in a transaction. For example, in college admissions, applying early (binding or non-binding) acts as such a signal. Crucially, you cannot apply early everywhere—you usually have to choose at most one college to apply to early, and as a result the signal is meaningful. (And early applicants are admitted at a higher rate than regular applicants, at least at the top schools [1].)

What about the market for professors? Departments usually get hundreds of applications per open position, so there's no way to interview everybody. At Stanford, the strategy is pretty straightforward: invite the 12 (say) strongest-looking applicants for interviews, and make offers to some subset of them.

Why wouldn't this strategy be optimal for a somewhat lower-ranked department? The

[22]When determining whether behavior is anticompetitive, law in the U.S. tends to focus on whether consumer prices were driven upward (and of course, Google is more-or-less free). The EU uses a broader definition of anticompetitive behavior, and on these grounds decided that Google's behavior qualified.

worry is that all 12 of the strong applicants would wind up accepting more prestigious offers, leaving the department with no hires. Thus it makes sense to interview at least some applicants who are a little less strong than the strongest ones. But which ones?

One negative consequence of this uncertainty is that there are always some applicants who slip through the cracks, with no offers from any departments that they interviewed at, and with no interviews from lower-ranked schools who incorrectly guessed that the applicant was out of their league.

The market for economics professors took a cue from the college admissions system, and it is now possible for each applicant to send a special signal (through a third party) to at most two different departments. The intended use is to signal your interest in a department that might otherwise think that you're not interested in them. (The signals are generally not sent to the very top departments, who generally focus entirely on qualifications.) It's not obvious how to do a rigorous empirical study of the effects of this signaling scheme, but the anecdotal evidence seems largely positive [3].

Could something similar work in the dating market, where owners of popular profiles receive far too many messages to reply to? Lee and Niederle [5] ran a field study in collaboration with a Korean dating site. In the experiment, users were given two "roses" per week. A user could send a large number of messages, but could attach a rose to at most two of them.

The dating company gave the researchers all of the relevant data about users, including an internally computed "desirability score" for each. This provided a measure of how effective the rose signals were, by comparing the probability of a message reply with and without a rose, controlled for the desirability score of both parties. The roses were especially effective when sent to someone with a similar or lower desirability score (who might otherwise think the sender was out of their league). Overall, roses increased the probability of message acceptance (from 15% to 18%, averaged over all cases), and also increased the overall number of messages. Perhaps someday we'll see a similar use of limited signals in production on a major dating platform.

References

[1] C. Avery and J. D. Levin. Early admissions at selective colleges. *American Economic Review*, 100(5):2125–2156, 2010.

[2] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621, 2015.

[3] P. Coles, J. Cawley, P. B. Levine, M. Niederle, A. E. Roth, and J. J. Siegfried. The job market for new economists: A market design perspective. *Journal of Economic Perspectives*, 24(4):187–206, 2010.

[4] R. W. Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6(4):577–595, 1985.

[5] S. Lee and M. Niederle. Propose with a rose? Signaling in internet dating markets. *Experimental Economics*, 18:731–755, 2015.