

Market Design: Understanding markets well enough to fix them when they're broken

"Grand Challenge" White Paper for Future Research in the Social, Behavioral & Economic Sciences

By Alvin E. Roth, Harvard University

Abstract

In the past fifteen years, the emerging field of [Market Design](#) has solved important practical problems, and clarified both what we know and what we don't yet know about how markets work. The challenge is to understand complex markets well enough to fix them when they're broken, and implement new markets and market-like mechanisms when needed.

Among markets that economists have helped design are multi-unit auctions for complementary goods such as spectrum licenses; computerized clearinghouses such as the National Resident Matching Program, through which most American doctors get their first jobs; decentralized labor markets such as those for more advanced medical positions and for academic positions; school choice systems; and kidney exchange, which allows patients with incompatible living donors to exchange donor kidneys with other incompatible patient-donor pairs.

These markets differ from markets for simple commodities, in which, once prices have been established, everyone can choose whatever they can afford. Most of these markets are *matching* markets, in which you can't just choose what you want, you also have to be *chosen*. One of the scientific challenges is to learn more about the workings of complex matching markets, such as labor markets for professionals, college admissions, and marriage.

1. Fundamental questions

“Market design” is the term used to refer to a growing body of work that might also be called microeconomic engineering, and to the theoretical and empirical research in economics, computer science and other disciplines that supports this effort and is motivated by it.

At its heart is the fundamental question: How do markets work? For competitive commodity markets, economists have a good grasp of some of the basic elements. When price discovery and adjustment operate smoothly, agents choose what they want at the prices they see. But many markets are more complicated than that; you can’t simply choose what you want, even if you can afford it; you also have to be chosen.

Examples of such “matching markets” abound: colleges don’t select their entering classes by raising tuition until just enough students remain interested, rather they set tuition so that lots of students would like to attend, and then they admit some fraction of those who apply. (And colleges also can’t just choose their students, they have to woo them, since many students are admitted to multiple colleges.) Neither do employers of professionals reduce wages until just enough applicants remain to fill the positions; there’s courtship on both sides (e.g. many new economics Ph.D.s would like to work for Stanford at the wages they offer, but Stanford receives many applications and makes only a few offers, and then has to compete with other top universities to actually hire those they make offers to).

Particularly for entry-level professionals, wages are often rather impersonal (e.g. many new assistant professors of economics, or new associates at large law firms, earn around the same wage, just as many students are offered the same tuition packages). Prices seem to play a different role in clearing matching markets than in markets for commodities. Labor markets and college admissions are thus more than a little like the marriage market; each is a two sided matching market that involves searching and courting on both sides.

Among markets that economists have helped design are multi-unit auctions for complementary goods such as spectrum licenses; computerized clearinghouses such as the [National Resident Matching Program](#), through which most American doctors get their first jobs; decentralized labor markets such as those for [more advanced medical positions](#), and for [academic positions](#); the [school choice systems](#) used to assign children to big city schools; and [kidney exchange](#), which allows patients with incompatible living donors to exchange donor kidneys with other incompatible patient-donor pairs. For surveys, see Milgrom (2004) and Roth (2002, 2008).

Auction design is perhaps the part of market design most closely connected to the traditional function of commodity markets: price discovery and efficient allocation. However, when multiple, heterogeneous goods are offered, and buyers may want to consume packages of complementary goods, matching bidders to packages becomes necessary, and recent research motivated by FCC auctions of radio spectrum has drawn close parallels between auctions and matching markets. Many open questions

remain, on the interface of economics and computer science, about the design and conduct of auctions that will make it safe and simple for bidders to bid on packages of goods.

Matching markets sometimes suffer persistent market failures, and this has opened a door through which economists have engaged in market design. For example, a number of labor markets have lost thickness due to the [unraveling of transaction dates](#): e.g. presently big law firms often hire new associates while they still have a year remaining of law school, and appellate judges hire [law clerks](#) via exploding offers that don't allow them to compare offers. Failures associated with phenomena like these caused the market for new doctors to explore various forms of centralized clearinghouse. In 1995 I was asked to direct the [redesign](#) of the clearinghouse for new doctors (the National Resident Matching Program), to address a number of issues, including the fact that there are a growing number of [married couples](#) in that labor market who seek two positions in the same vicinity. Each of these issues raises questions whose further answers will be important for understanding and designing complex markets:

- How does the timing of transactions influence market clearing? In particular, what is needed to create a marketplace in which sufficiently many transactions are available at the same time to achieve the benefits of a thick market? (Economists have devoted great effort to understanding the price of transactions, but much less is known about other features of transactions.) The timing of transactions concerns not just when they are made, but also their duration, as in e.g. the case of [“exploding” offers](#).
- How does the growing number of two-career households influence the labor market? How does it influence the marriage market? How are these related (e.g. in migration to cities, and in spousal hiring policies of firms such as universities located outside of cities, and labor law involving what kinds of questions applicants can be asked about their marital status)? Some of these are questions that will involve collaboration among economists, demographers, and sociologists.

A marketplace that successfully becomes thick by attracting many participants may face a problems of *congestion* resulting from all the transactions that can potentially be considered, since in many markets such consideration take time (e.g. interviews in labor markets, time between offers and acceptances, etc.) Congestion was the problem that led to the redesign of the high school assignment process in New York City, and it has led to the redesign of a number of other markets, such as the market for [clinical psychologists](#). Many open questions remain about the management of congestion.

Some markets fail to reach efficient outcomes because it isn't safe for market participants to reveal the necessary private information. This was what led to the [redesign](#) of the school choice system for Boston: the [old Boston algorithm](#) made it risky for families to reveal what schools they wished their children to attend, since a family that failed to get the choice it listed first would likely drop far down in the rankings. The new assignment mechanism makes it safe—a “dominant strategy”—for families to state their true preferences. However in many cases it can be shown to be impossible to make safe participation a dominant strategy, and so many questions remain about how to make participation safe. Recent [results](#) in economics and computer science suggest that some of these problems may become more tractable as markets grow large.

Developing kidney exchange in the United States involved many people working to overcome each of the problems mentioned above. First, a thick marketplace had to be made possible by establishing databases of incompatible patient donor pairs. Then, congestion had to be overcome, in the form of the number of operating rooms and surgical teams who could be assembled simultaneously to carry out larger exchanges. (The recent development of [non-simultaneous chains](#) has helped.) Presently, kidney exchange programs are grappling with the problem of how to make it safe for transplant centers to participate fully, by revealing all of their incompatible patient-donor pairs to the exchange.

At each step of the process, there has been collaboration between economists and doctors, and lately also computer scientists (about which more in a moment). A characteristic of market design is that it is going to require a great deal of collaboration among all sorts of people to design appropriate markets, and get them adopted and implemented.

In addition, kidney exchange and the shortage of transplantable organs also make clear that not every kind of market transaction is welcomed: some kinds of market transactions are viewed as [repugnant](#). In particular, it is against the law in the U.S. and in most developed nations to buy or sell organs for transplantation. More broadly, market solutions are not welcomed for a variety of transactions. Understanding the sociology and psychology of repugnant transactions and markets is a big task, which is likely to illuminate many aspects of markets and market design. At present, this is work that brings together economists, psychologists, sociologists, legal scholars, and philosophers.

Many of the market designs referred to above involve computer-assisted markets. Computers can assist markets in a number of ways, some of them more profound than others. Markets can be run on computers, so that transactions are recorded and processed in an orderly way. Markets can be accessed over the internet, so that many more people can participate than could at a marketplace in the physical world. Markets can use computers as trusted intermediaries, to accomplish something more or more cheaply than could be done without computers (for example the computer can hold a reserve price without revealing it unnecessarily, or job applicants can send a certified number of signals (as in the signaling mechanism now used in the [market for new economists](#)). Finally, computers can add computational intelligence to the market; instead of just reporting bids and asks, market outcomes can be determined by (possibly computationally intensive) algorithms that process market information in ways that couldn't be done, or done quickly, without computers. In this latter category, note that finding optimal kidney exchanges of constrained size is an NP hard problem solved by integer programming, while many labor market clearinghouses take as input rank order lists and use a deferred acceptance algorithm to find a stable matching.

Note that market design is not just about computerized or even centralized marketplaces, but also about the rules, procedures and customs of decentralized markets, what might be called their market culture. For example, in helping repair an unraveled market for [gastroenterologists](#), an essential feature was changing the rules about whether applicants could change their minds about offers received before a specified time.

To summarize, the last fifteen years have increased our understanding of how markets fail and how they can sometimes be fixed. The theory and practice of market design are deeply intertwined, and each particular market design brings economists into close contact with experts in the particular domain, and in other academic disciplines. For economics as a discipline, market design provides a fresh source of theoretical problems and empirical data, about the most fundamental questions of economics, concerning how markets work, and how they can be fixed when they fail.

2. What are the implications for advancing the domain? For building capacity? And for providing infrastructure?

As we understand more about markets (and perhaps about repugnant transactions) we'll know more about where and in what ways better markets can improve welfare, and perhaps also more about where we might pause to look for alternatives before instituting simple or unregulated or monetary markets or relaxing the restrictions against them.

As market design grows, it will become more like an engineering discipline, demanding both design knowledge and knowledge of particular domains of application. Right now, design papers mostly are judged by the journals as theory papers. But frontier design papers might not necessarily have the same focus on theory, or empirical work, that standard papers do, they might derive their value from how those things are combined in novel ways on some new domain of application. So, as market design develops, we'll have to nurture a market design literature that judges and recognizes frontier work in appropriate ways.

3. Who is doing provocative research?

Market designers are starting to be too numerous for a short list (here's a [link to a very partial list](#) of mostly economists and computer scientists), but there are big active groups in the Boston area and at Stanford. The Stanford group includes Milgrom, Bulow, Levin; Niederle, Ostrovsky, Hatfield, and Kojima, and the Boston group includes Roth, Athey, Parkes, Edelman, and Coles at Harvard; Pathak and Ashlagi at MIT; Sonmez and Unver at Boston College. Other centers include Maryland: Ausubel and Cramton; Michigan: Chen, Resnick, and Leider; Chicago: Budish; CMU: Sandholm...

References

Milgrom, Paul Putting Auction Theory to Work, Cambridge U. Press, 2004

Roth, Alvin E. "[The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics](#)," *Econometrica*, 70, 4, July 2002, 1341-1378.

Roth, Alvin E. "[What have we learned from market design?](#)" Hahn Lecture, *Economic Journal*, 118 (March), 2008, 285–310.

Creative Commons License: <http://creativecommons.org/licenses/by-nc-sa/3.0/>  <http://i.creativecommons.org/l/by-nc-sa/3.0/88x31.png>
 xmlns:dc="http://purl.org/dc/elements/1.1/" href="http://purl.org/dc/dcmitype/Text" property="dc:title" rel="dc:type">Market Design by [property="cc:attributionName" rel="cc:attributionURL">Alvin E. Roth is licensed under a \[Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.\]\(http://creativecommons.org/licenses/by-nc-sa/3.0/\)](http://kuznets.fas.harvard.edu/~aroth/papers/NSF%20Grand%20Challenge.SBE%202020.%20Market%20Design.pdf)