
# The Theory and Practice of Market Design

**Prize Lecture, December 8, 2012**

by Alvin E. Roth

Harvard University, Cambridge, MA, USA, Harvard Business School, Boston, MA, USA.

Preamble, to Lloyd Shapley:

Lloyd, when I began studying game theory, your work touched every part of it and shaped it and you were an inspiration—not just for me, but for the whole generation of game theorists who followed you. It is a great honor for me to share this prize with you.

I want to tell you about the theory and practice of market design, which is work that is still very much in progress. I should say personally that I am delighted to be recognized for work that we are still very much engaged with. Many of my colleagues are here in the audience and they are all waiting for me to get back to work.

The citation that the prize committee chose was "for the theory of stable allocations and the practice of market design." So the plan of my talk today is to tell you how stable allocations of the kind that Lloyd just spoke to us about connect to some of the most important markets that we are involved in—the matching markets that determine what schools we go to, and what jobs we get, and maybe who we are married to.

I also want to tell you about some additional theory that goes along with the theory of stable matchings, which helps us make game theory a practical tool for fixing markets when they are broken. And then I want to tell you about some of the applications. The ones I will talk about are job markets, school choice, and kidney transplantation of a certain sort.

Let us start by thinking about what markets and marketplaces do. How do they work, and how do they fail? And how can we fix them when they are broken?

When we think about markets, we often think about commodity markets. Commodity markets can be arm's length and anonymous. When you are buying 100 shares of stock on the New York Stock Exchange, you don't have to apply to buy, you don't have to convince the seller that you will take good care of those stocks. You don't have to worry about whether the seller took good took care of them when he had them. There is no courtship. The price does all the work. The New York Stock Exchange discovers a price at which supply equals demand.

But, in lots of markets, prices don't do all the work.

For example, I have been teaching at Harvard, and now I am teaching at Stanford. Those are both selective American universities. It is expensive to go to them, tuition costs a lot, but that is not what determines who gets to attend. Stanford doesn't raise its tuition until just enough students remain who want to attend. Even though tuition is expensive, it is low enough so that lots of people would like to be Stanford students. And then Stanford chooses the ones they would like.

So, universities don't rely on prices alone to determine who gets what. Labor markets also don't rely on prices alone, and labor markets and college admissions are more than a little like courtship and marriage. That is why the marriage metaphor is a good one because you can't just choose what you want. You also have to be chosen. You can't just tell Stanford that you're coming—they have to admit you.

You can't just tell Google that you are showing up for work. They have to hire you. And of course, it works that way on both sides. Yale can't just tell students to come to Yale. Yale has to compete for students with Harvard and Stanford. Google can't just hire who they want. Google has to compete for employees with Facebook.

So, matching markets are markets in which you can't just choose what you want (even if you can afford it), you also have to be chosen. College admissions and labor markets are two-sided matching markets, where both sides have preferences.

This is where we stand on the shoulders of Gale and Shapley's 1962 paper that Lloyd just told us about. My colleagues and I have also followed up on another paper, which I won't talk about since Lloyd didn't get to it—the Shapley and Scarf paper from 1974. Those two papers did something similar.

They defined a notion of stability. In the two-sided matching problem that Lloyd just told us about, a stable matching is one where there are not two people

who are not matched to each other but who would both prefer to be matched to each other than who they are matched to. That's a specialization of the more general idea of the *core* that Lloyd introduced in game theory. If an outcome is in the core, then there is no coalition of players that can go off on their own and do better than they could do in the marketplace.

That is going to be an essential idea for marketplace design. One of the tasks of a market designer is to create a marketplace that people want to come and transact in. They should not be able to do better by transacting outside of the marketplace. The core is a very important formulation of this idea. (Stability was explored explicitly in terms of the core in the Shapley and Scarf, 1974 paper.)

Let me remind you of the deferred acceptance algorithm, which Lloyd just talked to us about, because I want to emphasize that it gets an outcome in the core—as he told us—but I also want to talk about some of its other properties.

Here is a representation of that algorithm that is suitable for thinking of it as a centralized clearing house.

- Step 0: students and schools **privately** submit preferences to a clearinghouse
- Step 1: Each student "applies" to her first choice. Each school **tentatively** assigns its seats to its applicants one at a time in the order of the school's preferences/priorities over students. Any remaining applicants are rejected.
- . . .
- Step k: Each student who was rejected in the previous step applies to her next choice if one remains. Each school considers the students it has been holding together with its new applicants and **tentatively** assigns its seats to these students one at a time **in preference/priority order**[*]. Any remaining applicants are rejected.
- The algorithm terminates when no student application is rejected, and each student is (finally) given her current tentative assignment.

Think of a centralized clearinghouse for matching students to schools, in which we ask students and schools to tell us their preferences. We ask the families and the students, "What is your first choice school, what is your second choice, what is your third choice? Write it down." And we say to the schools, "Rank order the students." Maybe the school principals have preferences as they do in New York, or maybe the schools have priorities for different students, established by the school district.

[*] Note that schools take no account of the step of the algorithm at which a student applied.

We take these preferences that have been submitted, and process them in a computer. We look at the students' preferences, and in the computer algorithm we have each student apply to her first choice. In this first step some schools get lots of applications, from the people who apply to those schools as their first choice.

But the schools don't immediately accept the people who applied. If they get more applications than they have spaces, the schools immediately reject those who they can't fit in, keeping—not yet rejecting—those that they prefer, i.e. not immediately rejecting those students who are highest on the school's preferences or priorities among those who have applied so far. That is, each school has its own preference, or priority list, and it keeps the highest priority students who have applied, and it rejects the rest.

But it doesn't accept those students yet, it just doesn't reject them, while it waits to see who else will apply. And every student who has been rejected goes ahead and applies to their next choice school. Every school looks at the new applications together with the old ones—the ones that have not been rejected yet. It orders them in terms of its preferences or priorities. It keeps the ones it likes best and rejects the rest without prejudice about when they applied.

So, you can be a highly preferred student, but you apply and get rejected by your first choice. Next you apply to some school as your second choice, and you can now bump someone who applied to it as their first choice and was not initially rejected—but now they will be rejected because there is not space for both of you, and you have higher priority.

That is how the algorithm works. At each step, any student who has been rejected applies to his next choice. Each school looks at the people who have applied so far, keeps the best ones who have applied so far, and rejects the rest. And the algorithm stops only when no student is rejected anymore, which must eventually happen because no student applies twice to any school. When the algorithm stops, the schools finally admit all the students whose applications they are holding. That is why Gale and Shapley ('62) called this a *deferred acceptance algorithm*, since acceptances are deferred until the end, when all applications have ceased.

What does it mean for the algorithm to result in a stable match? It means there isn't a student and a school, not matched to each other who would rather be matched to each other. Suppose I am a student, and I end up matched to my third choice school. The matching would be unstable if my second choice school would rather have me than someone who they have been matched with, because I would prefer my second choice to my third choice, which is what I have got. If they also preferred me, we would be a blocking pair. We would be

able to get together and produce a better match for ourselves, instead of accepting the match produced by the clearinghouse.

How do we know that doesn't happen? Well, the only time I get a chance to apply to my third choice school, according to this algorithm, is when my second choice school has already rejected me because they are full, and they have a full class of students they prefer to me.

So, if I am in my third choice school, it means that although I would prefer to be in my second choice, my second choice would not prefer to have me, and that is why the outcome is stable.

So the deferred acceptance algorithm finds a matching that is stable with respect to the preferences of the students and of the schools. How does the market find out those preferences? One of the questions you have to ask as a market designer is, "If you are going to set up a clearinghouse, and you are going to ask people to tell you their preferences, are they going to want to tell you their preferences, or are they going to want to tell you something else?"

That is, when you say to people, "Tell me your preferences," a natural reaction is, "What are you going to use that information for?" And if you are a parent dealing with a school system you might be reluctant to tell the school district your preferences if they are not going to make it safe for you to do so. When I started to look at two-sided markets in the 1980s, I showed that in fact it is impossible for a stable matching mechanism to always make it safe for everyone to reveal their true preferences.

But remember this is an algorithm with one side that proposes or applies, and one side that accepts and rejects. It turns out that it is completely safe for the side doing the proposing to reveal their true preferences—they can't come to harm by revealing their true preferences.

Let me give you an idea of how you could come to harm with a different algorithm.

When Atila Abdulkadiroğlu and Tayfun Sönmez first looked at the Boston school system, they found there an algorithm that tried to give as many people as possible their first choice. That sounds pretty good! That sounds like a sensible thing for a school system to try to do. But it turns out to make it unsafe for families to reveal their true preferences.

Think of an immediate acceptance algorithm instead of a deferred acceptance algorithm. The algorithm starts the same way—every student applies to their first choice school, and every school accepts as many of the kids who have applied as it can, rejecting the rest, using its own preference list to decide who to accept and who to reject. And then every student who was rejected applies to his or her second choice school. But, this is an immediate acceptance algorithm,

and the children who were not rejected on the first step were immediately accepted. So it could be that you have really high priority at your second choice school, but your second choice school is already full. It is filled with people who applied to it as their first choice. And there lies the problem.

If you are not careful which school you list as your first choice, and if you fail to be admitted to the school you say is your first choice, you may find that your second choice has already been filled with people who said it was their first choice. Therefore, it might not be safe for you to tell the school district what your true preferences are because if you don't get your first choice, there is a good chance you will not get your second choice. And then your third choice will be filled with people who listed it as their first or second choice, and now you might fall right through to the bottom.

So, it is unsafe to reveal your preferences in the immediate acceptance algorithm, because that algorithm makes it important to list as your first choice a school that you have a high chance of getting admitted to, and as your second choice a school that has a good chance of not having all its places filled immediately. The deferred acceptance algorithm avoids forcing families to bear these kinds of strategic uncertainties precisely by deferring the acceptances. What makes the immediate acceptance algorithm unsafe is that you can lose your place in a school if you don't list it first. But in the deferred acceptance algorithm, because the schools don't decide who to accept until they see everyone who is going to apply, if you don't get your first choice, you still have just as much chance of getting your second choice as if you had listed it as your first choice. And that is part of the proof that allows you to see why it becomes safe, a *dominant strategy*, for students to state their true preferences in the student-applying deferred acceptance algorithm.

So far we have been talking about theory. Let me tell you how this opens a window on how marketplaces work. I will have to tell you a little bit about how American doctors get their first jobs and what happened over the first fifty years of the 20th century.

Before 1900, doctors graduated from medical school and immediately began to practice medicine. Around 1900, the customary first job of doctors started to be a supervised clinical position in a hospital, called an internship or a residency. Medical graduates looked for their first job at around the time they graduated from their four years of medical school training.

These are very important jobs for doctors—a residency is a bit like graduate school—it shapes your future career. And they are very important jobs for hospitals, because the interns and residents are a big part of a teaching hospital's labor force. So, there started to be competition by hospitals to get good interns

and residents. One form this competition took was hospitals started to make offers a little earlier than their principal competitors. And pretty soon, instead of having lots of people trying to get jobs around June of the last year of medical school, medical students found that they were getting offers earlier in the year from hospitals, and the hospitals didn't give them a lot of time to decide. Instead of a thick market, with a lot of employers trying to hire at the same time, there started to be a lot of little thin markets in which a student would get an offer and, without knowing what other offers might come, have to decide yes or no, without being able to compare it to other offers.

These offers started getting earlier and earlier. By the 1930s, the standard time at which American doctors were getting their first jobs was around New Year of their last year of medical school, around the middle of the fourth year of medical school. This caused some distress to the hospitals because it meant hiring was taking place before information was available that would determine students' class standing, their grades in their fourth year of medical school, and things like that. But, the hospitals (and the students) could not stop themselves. Editorials in the medical journals from that time say, "Hiring is getting pretty early, let's not make it any earlier," and everyone agreed that was a good goal.

However by 1940, the standard time that American doctors were getting hired was two years before graduation. That's really inefficient because it turns out lots of information is missing. Not only can't hospitals tell who the good doctors are, the doctors can't tell which jobs they want.

After two years of medical school you might want to be a surgeon, because you got an A in anatomy. But you have not yet had any surgical experience. If the job market is going in the summer of your second year, you are a top student, you get a top surgery job—and only later in your third year, do you discover that you faint at the sight of blood, and that it will be just terrible to be a surgeon. So, there were big inefficiencies in going so early. The market was losing important information about match quality.

Eventually, the medical schools intervened. The medical schools are a third party: They are not the hospitals or the doctors. In 1945, they decided to help control the date at which offers were made, by not releasing any information about students before a certain date. No transcripts, no letters of reference. It had been risky to hire students two years before they graduated, but it was really risky to hire people without knowing any of their grades, or even if they were really medical students.

This succeeded in controlling the date at which offers were made, and then as it became clear that the date was under control, it was moved back from the second year, into the third year, and into the fourth year again.

But another problem developed, now that everyone was making offers at around the same time. There was congestion—there were lots of potential offers to be made, and not that much time to make them.

In 1945 the rules said that offers to students should remain open for ten days. What happened in 1945? There you are, a graduating medical student in 1945, and you get an offer from your third choice hospital. And your second choice hospital says to you, "You are on our waiting list, you are close to getting an offer. If someone turns us down, we will make you an offer. So, don't do anything in a hurry." And you are not in a hurry because you have ten days. So you wait, and, while everyone waits, the waiting lists don't move.

So on the tenth day, you have an offer from your third choice hospital and you are on the waiting list of your second choice. That was a formula for a chaotic tenth day. Some people accepted their third choice, and then later in the day, they accepted their second choice when they got that offer. And if they took a little while to tell their third choice that they were not coming, everybody on the third choice's waiting list had already accepted another position, and then that hospital was very unhappy.

So, to avoid the problems of the tenth day, in 1946 the rule was changed, so that it required that offers be open only for eight days.

You are starting to be market designers already, and you can see that that is not going to solve the problem. There is a problem, but that won't be the solution to it.

By 1950, the hospitals rejected twelve hours as too long to leave offers open, so there was an explicit exploding offer environment in which you would get an offer on the phone and have to say yes or no immediately, because the hospital would say to you, "we have to move down our waiting list, we know that it moves fast."

And then the participants in this market did a remarkable thing. They got together and developed, partly by trial and error, a centralized clearinghouse, with the outward form I have told you about. They asked people for their preferences.

Hospitals submit preferences over the students they have interviewed. Students submit their preferences over the hospitals, more specifically over the residency programs at which they interviewed. And. . .it worked. It was a voluntary program, but pretty soon, everyone was getting matched through the clearinghouse.

When I started to study this in a 1984 paper, I found that this 1950s algorithm was different from but equivalent to Gale and Shapley's deferred acceptance algorithm with hospitals proposing.

So, this market went from being a really chaotic market for the first fifty years of the twentieth century, to being a very orderly market for the next twenty-five years at least. What did the trick was the adoption of a centralized clearinghouse that used something equivalent to the deferred acceptance algorithm.

This was a first among many observations that started to make me think about how marketplaces work, and what they are supposed to do, and how they can fail. This market had been failing. And the way it had been failing from 1900 to 1945, was it had failed to provide thickness. It had failed to bring everyone to the marketplace at the same time so that there would be lots of opportunities available.

In 1945, it proved possible to assemble a thick market—everyone was available at the same time—but there was congestion. It took time to process offers, and people were waiting. Therefore if you were a hospital, you could not make as many offers as you might like. You wanted to make offers and get them accepted or rejected so you could move down your waiting list, but people were holding your offers. You could not move down your waiting list, so you could not make as many offers as you wanted to and as you needed to get to a good match. So, to work well, a marketplace not only has to be thick, it has to be able to deal with congestion. And finally, a marketplace has to be safe and simple to participate in.

In a clearinghouse organized by the deferred acceptance algorithm, if you can get people to wait for and join the central clearinghouse, then it solves the thickness problem without running into congestion, because you can make all those waiting lists move really fast. In the 1950s they ran the medical resident clearinghouse on card-sorting machines; today it is run on computers.

It is safe and simple in the sense that it is safe to submit your true preferences if you are on the proposing side of this market. (It turns out, because the market is a large one, it is quite safe to submit your true preferences whichever side of the market you are on.)

One way you can tell that the stability of the final outcome is important is by looking at clearinghouses around the world. This is not a very common form of market organization, but it is not non-existent, so we can learn from the experience of a number of labor market clearinghouses.

So I started studying different clearinghouses, and whether or not they produced stable matchings in terms of the stated preferences. One of the things you see on the list here is that often, when they produce stable matchings, they also succeed: they are still in use and they stopped unraveling. The first line is the National Resident Matching Program, the American market I just told you about.

| **Market** | **Stable** | **Still in use (stopped unraveling)** |
|---|---|---|
| NRMP | yes | yes (new design in '98) |
| *Edinburgh ('69)* | *yes* | *yes* |
| *Cardiff* | *yes* | *yes* |
| *Birmingham* | *no* | *no* |
| *Edinburgh ('67)* | *no* | *no* |
| *Newcastle* | *no* | *no* |
| Sheffield | no | no |
| Cambridge | no | yes |
| London Hospital | no | yes |
| Medical Specialties | yes | yes (~30 markets, **1 failure**) |
| Canadian Lawyers | yes | yes (Alberta, no BC, Ontario) |
| Dental Residencies | yes | yes (5 ) (no 2) |
| Osteopaths (< '94) | no | no |
| Osteopaths (≥ '94) | yes | yes |
| Pharmacists | yes | yes |
| Reform rabbis | yes (first used in '97–98) | yes |
| Clinical psych | yes (first used in '99) | yes |
| Lab experiments | yes | yes |
| (Kagel&Roth *QJE*, 2000) | no | no |

It turns out that in various regions of the British National Health Service the market for new medical graduates unraveled in the 1960s much as the American market for new doctors had in the 1940s. A Royal Commission was formed and said, "Do what the Americans do." But the only thing that was described in the American medical literature was that there was a clearinghouse. So every region of the British National Health Service adopted a different kind of clearinghouse, and some were stable, and some were not. They provided something like a small natural experiment.

Those markets are part of a small data set; it is certainly not as comprehensive as we might like. But by and large, when we see clearinghouses that produce stable matchings, they succeed. And when we see unstable matchings, they mostly fail. There are some exceptions, in particular the two markets at

Cambridge and London hospitals, which each matched students from a single medical school to a single hospital and succeeded despite being unstable.

We would like better evidence than this. We use a lot of tools in market design, because to understand existing markets and marketplaces, design new ones, and actually get things implemented, you have to attack the problem from all directions. One of the tools we use is laboratory experimentation.

So, on this list there is a market with twenty-five thousand positions a year (the NRMP); and other markets with a couple of hundred positions, such as the regional markets in Britain. And at the bottom of the list are some experimental markets with a dozen positions, in the laboratory. In the lab we can look at a stable algorithm and an unstable algorithm and see that the difference really affects the success of the clearinghouse.

So, experiments fit nicely on this list; they are one of the tools of market design. They would not carry the day alone: we would not convince medical administrators to implement a stable algorithm just because we found it worked well in the lab. But experiments amplify and help us understand what we are seeing in the field data, and they also help us communicate it.

The way I got involved in redesigning the medical match was I got a telephone call in 1995. Elliott Peranson, who is here today, was involved—he is a matchmaker, he was involved in helping administer the match and I think he put the doctors on to me. They were having some problems with new features in the medical environment—not just married couples, but let me tell you about married couples.

In the 1950s there were almost no women medical students in the United States. Today fifty percent of American medical students are women. In the 1970s, when the figure was around 10%, the market started to run into some trouble, with married couples not always taking the assignments they got in the match. The reason was that the clearinghouse was treating married couples as if they were not married.

That was not working well for married couples because they would get positions that were far apart, or even when they got positions that were close together they might not be getting assignments they liked, since they were not asked what *pairs* of positions they wanted. This means the algorithm being used, which found stable matchings when applicants were all single, was not finding stable matchings any more.

If my wife and I get two jobs in Boston, but one is a good job, and one is a bad job, we are not going to be happy. (As those of you who are married know, the iron law of marriage is that you cannot be happier than your spouse.) We would be happier in New York, with two pretty good jobs. So, until the clearinghouse

started asking married couples for their preferences over pairs of jobs, it didn't have a chance of getting stable matchings, in terms of the preferences of married couples. The clearinghouse eventually did allow couples to express their preferences over pairs of jobs, but there remained other problems.

Elliott Peranson and I redesigned the algorithm so that it dealt more gracefully with married couples, and with other kinds of situations in which individuals might need two positions, and other features of the medical market that made some choices complements rather than substitutes. Because Elliott is such a matchmaking entrepreneur, there are dozens of healthcare markets that today use our algorithm in labor clearinghouses. And Muriel Niederle and I have worked on clearinghouses to help fix problems that arise in markets that come later in the career of doctors. She is here today too.

Another recent application is that school districts are starting to use clearinghouses that employ deferred acceptance, and this is close to what Lloyd and David thought of in their 1962 article. Just as Elliott is the "Johnny Appleseed" of clearinghouses in healthcare, Neil Dorosin, who is also here, is the "Johnny Appleseed" of school choice algorithms that are stable. Neil was the director of high school operations in New York City when we first met him, when we were involved in helping to redesign the way New York City high school students are assigned to schools. That was with Atila Abdulkadiroğlu and Parag Pathak, who are also here.

New York City had a semi-centralized system that handled admissions and waiting lists through the mail. For example they would send to a high school student a letter that said, "Congratulations, you have been admitted to two high schools, tell us which one you want." And only when they got that letter back could they look at the vacancies that had been created and then admit some more people.

That was a congested process. There are almost 90,000 students a year who enter New York City high schools, and just before school began about 30,000 of them had not been admitted anywhere and had to be assigned administratively to schools they had not expressed a preference for. We helped New York develop a stable clearinghouse and almost immediately that 30,000 number went down to 3,000.

Different school districts have had different problems. Boston schools used an immediate acceptance algorithm of the kind I told you about that made it unsafe for parents to reveal their true preferences to the school district. And so we helped Boston Public Schools implement a deferred acceptance algorithm as well, that made it safe for families to list their true preferences. That is a technology that has now started to spread to other school districts. Atila and Parag

and Neil and I and lots of education professionals in each of these cities have started to make that a technology that is accessible to more and more families in the United States.

Let me move to a very different kind of market and tell you about kidney exchange. It is one of the exciting markets that we deal with. You each have two kidneys. And if you are as healthy as you look, you could remain healthy with just one. But kidney disease is a deadly disease, and transplantation is the treatment of choice. So you might know someone who has kidney disease, and because you're healthy, you could donate a kidney to someone you love.

But sometimes you are healthy enough to donate a kidney, but you can't donate it to the person you love, because of some kind of incompatibility. The figure shows a simple blood type incompatibility.

Donor 1 would love to give her kidney to Recipient 1, who needs one, but Donor 1 has blood type A and Recipient 1 has blood type B, and that is not going to work. There is another patient-donor pair who have the same problem in reverse.

This is what opens up the possibility of exchange. The B blood type patient could get the B blood type kidney from the donor in Pair 2. And at the same time, the donor in Pair 1 can give an A kidney to Recipient 2. So, that is an exchange – that is a kind of thing economists are good at, thinking about exchange, and how it can be organized.

The question was how to develop a marketplace to allow transplantation to occur when otherwise it would not have. This is work done with a number of people, with my colleagues Tayfun Sönmez, and Utku Ünver, and Itai Ashlagi

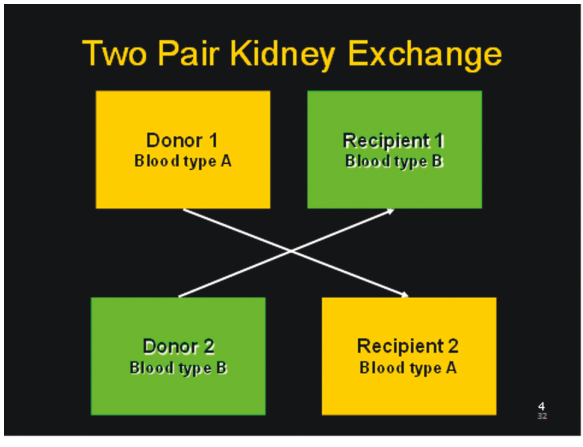


**FIGURE 1.**

who is here, and with many medical professionals, including Mike Rees—a surgeon whom I will tell you more about and who is here today, too, and Frank Delmonico, who we helped to start the New England Program for Kidney Exchange.

There is congestion in kidney exchange. Here is a picture taken during a pair of surgeries—one nephrectomy—taking a kidney out, and one transplant—putting a kidney in—in Cincinnati, Ohio.

I am in the yellow gown, keeping my hands out of the way, so no one hands me anything. In the bucket is a kidney. Behind me is another operating room you can't see, just steps away, where the nephrectomy took place, moments before. A surgeon and kidney exchange pioneer named Steve Woodle is in the striped cap, and at the same time these surgeries are going on, this donor's patient, and this recipient's donor are at Mike's hospital in Toledo, Ohio undergoing the same set of operations.

When I say at the same time, I mean literally at the same time. The surgical teams got the patients anesthetized, they did the initial incisions, and then they got on their cell phones and they said, "We're ready to go here. Are you ready?"

**FIGURE 2.**

The reason they do the surgeries simultaneously is that you cannot contract on a kidney in the United States. It is illegal to give valuable consideration for an organ for transplantation. That mostly means you can't buy and sell organs, but it also means it is hard to enforce any kind of agreement such as 'today we give you a kidney and tomorrow you give us a kidney.'

So, these operations go on at the same time. That causes congestion because to do this simple two-way exchange you need four operating rooms and four surgical teams, because you have to do two nephrectomies and two transplants at the same time. So, it is already logistically hard to do even the simplest kind of two-pair exchange.

We could do more transplants if we could do larger exchanges. A three-way exchange skirts the "double coincidence" that Jevons told us makes barter exchanges so difficult. That is, sometimes patients can be included in a three-way exchange even if there were no-two way exchanges in which they could take part. To do a three-way exchange simultaneously requires six operating rooms and six surgical teams.

When we initially started, we could only muster four operating rooms and surgical teams at a time. Nowadays, six is not so hard, but you can see that if we were trying to get an even longer exchange, it might be really hard to muster the resources.

But kidney exchange can be done not just in shorter and longer cycles. It can also be done in chains that begin with non-directed donors. Deceased donors are non-directed, they don't have a particular patient in mind, and there is an increasing number of living non-directed donors who want to donate a kidney and don't have a particular patient in mind. One question is, "How many transplants can a non-directed donor facilitate?" When our surgical colleagues were doing all the surgeries simultaneously, we would only have six people in such a chain—three donors and three recipients. But Mike Rees took the initiative in conducting the first non-simultaneous chain, that was reported in an unusual article in the *New England Journal of Medicine* that was a collaboration between surgeons and market designers. Mike led the way in what has become standard practice in kidney exchange. The idea is that a chain initiated by a non-directed donor can be arranged so that each patient-donor pair gets a kidney before they give one. So the cost of a broken link isn't nearly as great as it would be if some pair underwent a nephrectomy and then failed to receive a kidney themselves. The cost of that kind of breach would be not only that they had undergone a surgery that they didn't need and that didn't help them, but that they would no longer have a kidney to exchange in the future, and to avoid that is why we still do cyclic exchanges simultaneously.

Of course in a non-simultaneous chain you have to rely on people who say that they are going to give a kidney next Tuesday to do so. You cannot compel them. But it turns out that this has worked pretty well in practice.

So the chain reported in that paper had twenty people in it—ten nephrectomies and ten transplants at the time it was reported. That chain continued subsequently. The reason there can be a lot of people in it and not just a few is that the congestion has been avoided because these chains can now be done non-simultaneously. The marketplace now allows more exchanges, more transplants to be done because you don't have to assemble all of the operating rooms at the same time.

There have been larger chains, such as one with sixty people in it, thirty nephrectomies and thirty transplants. The same transplants could not be done via a set of small exchanges. The reason has to do with the fact that compatibility graphs that record which patients can take which kidneys have become sparse, as we have a growing proportion of patients who are "highly sensitized," which is to say that they have developed antibodies to many human proteins, and so are immunologically incompatible with most kidneys. Sparse graphs contain very few cycles, but many long chains. And that is why these long chains we are seeing are proving so useful. Itai Ashlagi, who is here today, has been leading the study of those long chains and why they occur and how they can best be managed.

So, why do we have to do kidney exchange, instead of just buying kidneys from willing sellers? My colleague Gary Becker at the University of Chicago has argued that there is no shortage of kidneys, since everyone has two, so there is a surplus. The problem he says, and he is far from alone in this position, is that we are just not mobilizing them properly: the prices, set by law at zero, are too low.

But it turns out a lot of people think it is a terrible idea to buy and sell kidneys, the kind of bad idea that only bad people have. It is against the law here in Sweden. It is against the law in the United States. It is officially illegal in just about every country in the world except for the Islamic Republic of Iran to buy and sell kidneys. There is a legal market in Iran, and there are black markets in many places, and "gray markets" in some where the laws are less than clear.

Which things can and should be bought and sold is an interesting subject that I have written a bit about. I call a transaction "repugnant" if some people would like to do it, and other people don't think they should. So kidney sales are a repugnant transaction. We could talk about repugnant transactions for a long time. But I won't talk about them here, except to say it is an important subject. It turns out there are lots of things we are reluctant to buy and sell. We are reluctant to use money as the tool to decide who gets what. As social scientists,

The NEW ENGLAND JOURNAL of MEDICINE

VOL. 360 NO. 11

ESTABLISHED IN 1812 MARCH 12, 2009 NEJM.ORG

# A Nonsimultaneous, Extended, Altruistic-Donor Chain

Michael A. Rees, M.D., Ph.D., Jonathan E. Kopke, B.S., Ronald P. Pelletier, M.D., Dorry L. Segev, M.D., Matthew E. Rutter, M.D., Alfredo J. Fabrega, M.D., Jeffrey Rogers, M.D., Oleh G. Pankewycz, M.D., Janet Hiller, M.S.N., Alvin E. Roth, Ph.D., Tuomas Sandholm, Ph.D., M. Utku Ünver, Ph.D., and Robert A. Montgomery, M.D., D.Phil.

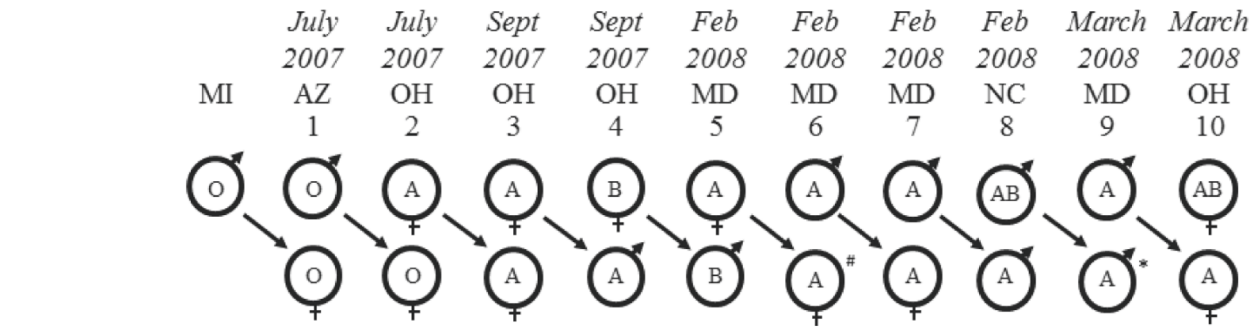


* This recipient required desensitization to Blood Group (AHG Titer of 1/8).
# This recipient required desensitization to HLA DSA by T and B cell flow cytometry.

FIGURE 3.

I think we need to understand that better. There are things about people's perceptions of markets that economists have not yet spent enough time trying to understand.

What are some next steps for market design? Well, it was easy for us to study markets that used centralized clearinghouses, and in many cases, those are the markets we have helped build, centralized clearinghouses for kidney exchange, for school choice, for labor market clearinghouses. But most markets are organized in a less centralized way than via a clearinghouse, and so part of market design is thinking about rules of engagement. What should the rules be in various kinds of decentralized marketplaces? How should offers be made? How long should they be open? What happens if you accept an offer and later you change your mind? How do you get information about preferences? When we talked about clearinghouses, I said to you that it is possible to organize a clearinghouse so that it is safe for people to reveal their preferences. But in most markets no one will ever ask you your preferences.

You are not American doctors, so you are probably not going to go through a clearinghouse when you look for a job; you will go through application and selection procedures. No one will ever really ask you, "Are we your first choice?" And if they did, you might say, "Of course you are," and say that to everyone who asks you, but that would not be critical in deciding who gets what.

One of the markets that is just getting underway right now in December is the market for new Ph.D. economists in North America. One bit of market design that my colleagues and I did in that market is institute a signaling mechanism—a mechanism by which applicants can send a limited number of signals of interest to employers, to help break through the congestion that arises because it is so cheap to apply nowadays. Most employers looking for certain kinds of economists get applications from all those economists who are on the market.

So, for each job, employers get many, many applications, and the question is how to sort among them, which ones should you pay attention to? The problem is similar to those that you might see on dating websites, which some of my students and colleagues have studied as well.

The problem on a dating website is that women with attractive pictures get many, many emails from men, saying, "How about me?" If they get too many to look at, then they just ignore most of them. And if the men are finding that most of their emails are ignored, then they start sending more emails, reflecting more superficial investigations and less information about match quality.

Muriel Niederle and Soohyung Lee and some colleagues intervened in a dating website on which people were sending lots of emails: they gave

people just a few of what they called 'virtual roses'. You could send many emails, but you couldn't send as many roses as you wanted, because you had a strictly limited supply, so if someone got an email that had a rose attached, they knew it was special, and worth paying attention to. Interventions like that helped convey information about matches, just as we found in the economics job market.

Similarly, in a lot of decentralized markets we have looked at, we find that helping convey information, and helping preserve options, for example by allowing people who accept early offers to change their mind, helps the market work better.

A market designer has to think a lot about what constitutes a free market. We just went through an election in the United States and our politicians like to talk about free markets or regulated markets as if these are entirely different things.

I think a useful way to think about free markets is that a free market is a market with rules and institutions that let it operate freely. When we talk about a wheel that can rotate freely, we don't mean a wheel that is unconnected to anything else. We mean a wheel that has an axle, and well-oiled bearings. I think that is a good metaphor for a free market. A free market needs institutions that let it work well. Markets like the New York Stock Exchange have lots of rules—they're not laissez-faire markets, they have rules about when they open in the morning, and when they close in the afternoon, and these help keep the market thick. And they have rules that have the force of law, like rules against insider trading that you can go to jail for violating, and those rules are intended to make it safe for ordinary investors to participate in the market.

So, what makes the New York Stock Exchange work well is a good set of rules. And that is how we should be thinking about free markets—markets with a set of rules that allow them to work freely.

Market design is an ancient human activity, but only recently have contemporary economists focused attention on it and started to take an active role. As economists have started to work as engineers, we find ourselves using a collection of tools to investigate particular markets and marketplace designs.

Game theory is the central tool, and it is a combination of strategic and coalitional models that used to be called "non-cooperative" and "cooperative" game theory, but I think that distinction is not useful anymore. We don't use coalitional and strategic models to study different kinds of games, we use them to ask different questions about the same game, like, "Is the outcome stable?" and "Is it safe to reveal your preferences?"

We need to make lots of careful observations of actual markets and how they work and how they fail. If you are a game theorist, *rules are data*. When we look at markets, we want to know what are their rules and how are those rules changing over time, and what kinds of behavior are people trying to make rules against?

And often, when we are trying to design a market, there's a deadline by which time the design has to be completed and implemented, for example by the time students need to go to school next year. So sometimes we are working beyond our deeply reliable scientific knowledge, and then we use tools such as computation to help us look at the data and try to understand what is going to happen if we implement a particular market design. As I already indicated, controlled experiments are useful in this effort too—especially when you are looking at a new design for which there is not as yet any field evidence about how it would work. Controlled experiments are also useful when you are looking at field data and trying to understand what they mean.

Let me close by saying that market design is a team sport. And it is a team sport in which it is hard to tell who are theorists or practitioners because it blurs those lines.


**FIGURE 4.**

Here are some of my colleagues, with whom I have worked on designing, getting adopted and implementing various marketplaces. Some of them are economists and some of them are not. There are kidney surgeons there and educators, and a matchmaker, and some of them are here to join in this wonderful party with us this week.

Thank you.

Portrait photo of Alvin E. Roth by photographer Ulla Montan.