3

Market Design

Itai Ashlagi

Management Science and Engineering, Stanford University

3.1 Introduction

A marketplace enables economic transactions between the market participants according to a given set of rules. When a social planner has all the information, traditional operations and optimization can be applied directly to achieve optimal outcomes. But it is often the case that participants hold private information not directly accessible to the planner; engineering the rules (or the strategy sets) carefully is often crucial to allow the market to operate smoothly and reach desired outcomes. Another important piece of market design is paying attention to institutional details, which impose constraints on the planner and hence market outcomes, as in traditional operations.

Researchers have not only developed elegant theories in the area of market design, but also contributed to improving the operations of real marketplaces in collaboration with practitioners. This chapter provides a brief overview of two healthcare marketplaces, the National Residency Matching Program and Kidney Exchange, while emphasizing some of their design issues and ongoing challenges. These marketplaces might seem to be nonstandard due to the lack of monetary transfers or prices, but we shall see that they share some similar economic principles with classic markets.

3.2 Matching Doctors to Residency Programs

3.2.1 Early Days

Every year, thousands of doctors who graduate from medical school start a residency, specialty training, in the United States. In 2017 alone, more than 27,000 doctors began their first year of residency. Between the years 1900 and 1945,

Handbook of Healthcare Analytics: Theoretical Minimum for Conducting 21st Century Research on Healthcare Operations, First Edition. Edited by Tinglong Dai and Sridhar Tayur.

hospitals competed to hire doctors and the market in decentralized manner, and the market suffered from unraveling: offers to doctors were made earlier each year, until that doctors were often being hired more than a year before the completion of their training, which in turn led to inefficiency due to mismatching (Roth, 2008, 2002).[1]

In 1945, American medical schools agreed that all offers would be made on a specified day. Although this consensus solved the unraveling problem, it led to other frictions. Hospitals quickly noticed that when their first offer was rejected, other candidates had already accepted offers from other hospitals. This congestion further led hospitals to make exploding offers to which doctors had to reply immediately.[2] Finally, in the early '50s the American medical associations decided to use a centralized clearinghouse to match doctors to residency programs. The idea was that, within this new framework, after doctors interviewed with residency programs, both doctors and hospitals would submit ranking lists representing their match preferences, and an algorithm would determine matches based on these and on programs' capacities. The organization that runs the match is called the National Residency Matching Program (NRMP).

3.2.2 A Centralized Market and New Challenges

Before we describe the algorithm adopted by the NRMP, one natural question would be what objectives should such matching algorithm achieve? A common answer is to maximize some weighted function based on the rankings (an extreme version would be to maximize the number of doctors who are matched to their top choice). Such algorithms, however, may provide doctors misreporting their ranking of programs; that is a doctor who is interested in a highly demanded residency may be concerned that if she doesn't get that assignment, her second choice will be taken by another doctor who listed it as her first choice. Thus a doctor who does not get her first choice may well get a bad choice.

The NRMP adopted instead an algorithm that is similar to the *deferred acceptance* (DA) algorithm which outputs a *stable* matching. A matching of doctors to programs is stable if no doctor and hospital who are not matched to each other, prefer each other over their match (Roth, 1984; Gale and Shapley, 1962). The theory of stable matching was initiated by Gale and Shapely in their seminal paper "College admissions and the stability of marriage" where they also introduced the DA algorithm.

The DA algorithm takes the input rankings and simulates a natural process where agents on one side propose and agents on the other side evaluate offers.

1 There is much literature on unraveling in labor markets (see, e.g., Roth and Xing, (1994); Ünver (2001); Niederle and Roth, (2003)).
2 For congestion in labor markets, see, e.g., Avery et al., (2001); Roth and Xing (1997).

The *doctor-proposing DA algorithm* works as follows (the hospital-proposing version works similarly): at each stage of the algorithm, unmatched doctors propose to their favorite program, which they have not previously proposed to. Programs tentatively keep their favorite offers so far and reject all other offers. The algorithm stops if all doctors are matched, or, every unmatched doctor has applied to all the programs on her list. In addition to finding a stable matching, the doctor-proposing DA algorithm is strategyproof for doctors; in other words, it is always safe for doctors to report their true preferences. The original algorithm used by the NRMP was in fact similar to the hospital-proposing DA algorithm (Roth, 1984). It is interesting to note that the set of stable matchings has a lattice structure, and the DA algorithm generates the stable matching that is most (least) preferred by each agent on the proposing (courted) side.

But over the years, more women attended medical school, and the number of married doctors (couples) on the residency job market has grown. The NRMP began to face a real challenge because the DA algorithm allowed doctors to submit preferences only individually, and couples often had to find residencies outside the match (to avoid working in different cities).

Stability is arguably one of the keys to the success of the NRMP (Roth, 2002; Kagel and Roth, 2000). But in the case of couples who introduce complementary preferences, stability may not even exist (Klaus and Klijn, 2005; Roth, 1984). Despite this challenge, Roth and Peranson (1999) engineered a new algorithm, which has been in use by the NRMP since 1998. The key idea was to allow couples to apply together, by ranking pairs of hospitals. Roth and Peranson (1999) also switched from the hospital-proposing DA to a version based on the doctor-proposing DA.

The theory of stable matching hasn't provided guidelines for how to design the market with complementarities (a common issue in mechanism design), yet insights from the existing theory together with careful engineering have made this marketplace successful once again. Dozens of labor markets now use the stable matching mechanisms in different entry-level labor markets.[3]

3.2.3 Puzzles and Theory

Stability is arguably an important part of the success of the NRMP and other two-sided markets that adopt matching mechanisms (Roth, 2008). However, until recently various puzzles still remained.

One challenge is that when couples are part of the market, a stable matching may not exist. Roth and Peranson (1999) report that in every year they examined, there was a stable matching with respect to the reported preferences. Kojima et al. (2013) and Ashlagi et al. (2014) studied a large market model with

3 The deferred acceptance algorithm is also used in various cities around the world to assign a students to schools (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005a,b).

random preferences and showed that as long as the number of couples grows more slowly than the number of single doctors, a stable matching will be found with high probability.[4]

Another longstanding challenge has to do with the multiplicity of stable matchings in a two-sided matching market (without couples). Multiplicity of stable matching raises several issues. First, which stable matching should be implemented? The one that is best for the doctors, the one that is best for the hospitals, or something in between? Another related issue concerns incentives. Roth (1982) and Dubins and Freedman (1981) find that no stable matching mechanism is strategy-proof for all agents (so there are some instances in which an agent will benefit from misreporting her preferences). Demange et al. (1987) further find that opportunities to misreport preferences successfully will arise only if an agent has multiple stable partners.

Even though artificial markets with a large set of stable matchings can be easily constructed, empirical evidence from the NRMP suggests that this set is small (Roth and Peranson, 1999), with very few agents having more than one stable partner.[5]

When preferences are highly correlated, we would expect the core (the set of stable matchings) to be small (for example there is a unique stable matching if all doctors have the same preferences over programs). So the question is whether a large core can arise with uncorrelated preferences. Pittel (1992) and Knuth et al. (1990) study a two-sided matching market with n men and n women and find that almost every man has multiple stable partners as the market grows large, when preferences for each man and each woman over all the agents on the other side are drawn independently, at random from a uniform distribution.

Following simulations by (Roth and Peranson, 1999), Immorlica and Mahdian, (2005) find that if one side ranks (uniformly at random) only a constant fraction of agents on the other side, or alternatively when the ratio between supply and demand (competition) is large, then as the market grows large, almost all agents have a single stable partner.

More recently, Ashlagi et al. (2017) find that even with the slightest competition, the core is small: in a market with n men and $n + k$ women (for any $k \geq 1$) and preferences drawn independently and uniformly at random over all agents from the other side, almost all agents have a single stable partner as n grows large. This result provides a generic reason for why two-sided markets typically have a small core. Therefore, which side proposes in a the DA algorithm is not a real concern, and it is safe to recommend all participants to report their true preferences.

4 Recent studies established the existence of stability in more general large markets (Che et al., 2015; Azevedo and Hatfield, 2012).

5 A small core was found also in online dating markets with respect to estimated preferences (Hitsch et al., 2010).

3.3 Kidney Exchange

3.3.1 Background

Kidney transplantation is the preferred treatment for end-stage renal disease, increasing life expectancy by ten years on average (Wolfe et al., 1999). Transplantation also saves several hundred thousands of dollars over remaining on dialysis.[6] Currently, the vast majority of Medicare spending on kidney failure is directed to dialysis costs.

Unfortunately, there is a large shortage of organs, and there are currently more than 97,000 patients on the waiting list for cadaver kidneys in the US. Another source of transplants is live donation. In fact, transplantation from a live-donor kidney is also preferable to a cadaver kidney. However, not everyone who is healthy enough to donate a kidney can donate to her intended recipient because a successful transplant requires the donor and recipient to be blood-type and immunologically compatible. Incompatibility between a donor and her intended recipient creates the demand for *kidney exchange*: an incompatible patient-donor pair can donate a kidney to a compatible recipient and receive a kidney from a compatible donor.

Note that it is illegal to buy or sell organs for transplantation in almost all countries (see Roth (2007) and Leider and Roth (2010)).[7] Kidney exchange thus represents an attempt to organize a barter economy.

The first kidney exchange took place in Korea in 1991.[8] Until 2003, few exchanges have taken place in the US, but in 2016 the number of transplants from kidney exchanges had reached more than 640 (the number is higher because some transplants are recorded as anonymous donations rather than as kidney exchanges), and are more than 11% of all living donor transplants in the US.

Forms of exchanges: cycles and chains. One form in which exchanges are organized is a *cycle*, which involves only incompatible patient–donor pairs, with each patient receiving a kidney from a compatible donor of another pair. Another form is a chain, which is initiated by an altruistic donor (with no intended recipient) and donates to a pair, whose donor donates to the next pair, and so forth.[9] Exchanges are organized such that no patient–donor

6 In 2014, Medicare paid $87,638 per year per dialysis patient, but only $32,586 per year per transplant patient. Given a median waiting time of 3.5 years for a deceased donor kidney, the difference adds up to a cost savings of about $192,682 (United States Renal Data System, 2016).
7 In the US, the National Organ Transplant Act (NOTA 1984) makes it illegal to obtain organs for valuable consideration. For discussion in favor of compensation for donors see, e.g., Becker and Elias (2007).
8 See Rapaport (1986) who first raised the idea of kidney exchange and Ross et al. (1997) and Ross and Woodle (2000) for discussions regarding ethical considerations.
9 Chains typically end with a patient on the waiting list who has no intended donor.

pair donates a kidney prior to receiving a kidney. This means that cycles are conducted simultaneously, since the cost of a broken link would be high to a pair that first donated a kidney and later failed to receive one. Due to logistical barriers, cycles are typically short including or three pairs. But chains can be organized nonsimultaneously without breaking this requirement, and therefore can be longer (Saidman et al., 2006). The majority of kidney exchange transplants in the US are now conducted through chains.

Compatibility between a donor and recipient. For a transplant to take place, a patient must be both blood-type (ABO) and tissue-type compatible with a donor. Thus, an O donor is valuable because O is ABO compatible with any other patient. Tissue-type compatibility means that the patient has no antibodies to the donor's antigens. The common measure for patient sensitivity is the Panel Reactive Antibody (PRA), which captures the likelihood that, based on her antibodies, the patient is tissue-type incompatible with a random donor in the population.

We survey some important steps in the progress of kidney exchange in the US, while emphasizing the economic and operational perspectives. Due to the shortage of space, we elaborate on only about handful of (more recent) issues.

3.3.2 Creating a Thick Marketplace for Kidney Exchange

The first proposal for organizing kidney exchange on a large scale involved integrating cycles of patient–donor pairs while considering patients' preferences (Roth et al., 2004). However, in the early days, only pairwise exchanges were conducted.[10] Subsequent work suggested that allowing only slightly larger, three- and four-way exchanges, would increase efficiency (Saidman et al., 2006; Roth et al., 2007). Common to these studies is taking a centralized approach (clearinghouse) to kidney exchange and in particular using *optimization*.[11]

Efficiency in Large Markets

Roth et al. (2007) characterize efficient allocations using cycles in a large market and find no need for cycles longer than four. Their large market assumption assumes that compatibility between patient and donors depends only on blood types.[12]

The characterization follows from the blood-type structure. To get some intuition consider, for example, the set of A-O and O-A patient–donor pairs. A

10 Roth et al. (2005) proposed a mechanism for conducting pairwise exchanges.
11 See also Segev et al. (2005).
12 This prediction is true even when patients' PRAs are included, by building on the Erdős-Renyi random graph theory (Ashlagi and Roth, 2014) because in a sufficiently large market, the PRA will not be a barrier, and the blood-type structure will determined the efficient allocation (In fact, there is no need for cycles longer than 3.)

kidney exchange pool is likely to have fewer A-O pairs than O-A pairs because many A-O pairs are compatible and select to go through a direct live-donor transplant.[13] If there were only A-O pairs and O-A pairs, in a sufficiently large market, it would be efficient (and possible) to match every A-O pair with a different O-A pair in a two-way cycle.[14] In particular, some pairs are *overdemanded* and some are *underdemanded*; the former will all match whereas some fraction of the latter will remain unmatched in an efficient outcome.

It is not hard to extend the characterization to show that under large market assumptions also chains need not be long.[15] However, the experience of kidney exchange platforms suggested that chains in fact play a crucial role.

Optimization

Kidney exchange platforms often use optimization to find matches through cycles and chains. The optimization problem of maximizing the number of matches is NP-complete (Abraham et al., 2007) and various algorithms have been developed to help programs with this task, following Roth et al. (2007) and Abraham et al. (2007).[16]

The Need for Chains

After the first nonsimultaneous chain (Rees et al., 2009), which involved more than ten transplants, chains became common. Today the average chain length is between four and five. The two longest chains so far contained 30 and 35 transplants in 2012 and 2014, respectively. The pairs in a chain, especially longer ones, are not all identified at once. The last donor in a chain segment either ends the chain by donating to a patient on the waiting list or becomes a *bridge donor* and initiates a chain segment in a future period.

One reason that chains have become ubiquitous is the large fraction of (very) highly sensitized patients in kidney exchange networks (Ashlagi et al., 2012). Another reason, which we elaborate in our discussion on incentives (Section 3.3.4), is that hospitals often match internally easy-to-mach pairs and enroll the pairs they cannot match. As a result, the compatibility graphs induced from the patient–donor pairs in real kidney exchange pools are sparse.[17] Data-driven simulations by Ashlagi et al. (2011a,b) and Dickerson

13 An A-O pair is incompatible if the A patient is tissue-type incompatible with the O donor.
14 A-O pairs can potentially match with each other, but this is a waste, as each such pair could potentially help a different O patient.
15 Altruistic donors can initiate chains that can be at most of length three, including at most two underdemanded pairs and one patient on the cadaver waiting list.
16 Abraham et al. (2007) develop an algorithm that can identify an optimal solution in a large pool using cycles up to length three. Researchers have expanded this line of work to include chains as well as various objectives (e.g., Biro et al. (2009), Glorie et al. (2012), Anderson et al. (2015), Constantino et al. (2013), Klimentova et al. (2014), Dickerson et al. (2012c)).
17 See Ashlagi et al. (2012, 2013) for a theoretical models based on Erdős Renyi graphs (in static and dynamic settings), in which the the key ingredient is the sensitivity of the patient (and blood

et al. (2012a) reveal that nonsimultaneous chains increase efficiency in a dynamic environment.[18]

3.3.3 Dynamic Matching

Early theoretical papers took a static approach and focused on the importance of the *matching technology*. But kidney exchange pools are dynamic, with pairs arriving and being matched over time; thus, it is natural to ask how the platform should match in this dynamic environment.

The *policy* employed by the clearinghouse, which determines which exchanges to implement and when, also affects the efficiency of the marketplace. One natural policy is a *greedy* policy, where the clearinghouse forms exchanges as soon as an opportunity arises. Another possibility is to adopt a *batching* policy, which identifies a (weighted) optimal allocation within the pool every number of periods. More sophisticated policies may take into account both the compatibility graph and the future.

Kidney exchange platforms in the US have gradually moved to small batches and thus frequently identify exchanges. The Alliance for Paired Donation (APD) and the National Kidney Registry (NKR) identify exchanges on a daily basis, and the UNOS program identifies matches twice a week. These are national programs, in which multiple hospitals participate in. A major concern is that this behavior is driven by platform competition. However, Methodist at San Antonio (MSA), which is a single center program and faces no competition, also matches on a daily basis. In contrast, national exchange programs in several countries such as Canada, United Kingdom, Netherlands, and Australia, search for exchanges every three or four months (Ferrari et al., 2014; Malik and Cole, 2014; Johnson et al., 2008). We briefly discuss recent research on this front.

In a simulation study, Ashlagi et al. (2018) looked at the effect of batching policies on efficiency (measured by the fraction of matched pairs and waiting times), using empirical data from the APD and MSA programs, which have different pool compositions, partially because participating hospitals in national pools often match easy-to-match pairs internally. Pairs in the simulation arrive according to a Poisson process and depart according to an exponential random variable, unless they are matched earlier. (Ashlagi et al. also model various frictions such as delays due to blood shipping and match cancellations.) They find that there is essentially almost no harm in matching frequently. Figure 3.1 plots the fraction of matched pairs and the average waiting time under different

types are ignored); the models explain the relationships between the fraction of highly sensitized patients and the need for long chains.

18 Ashlagi et al. (2011a,b) have been also part of an ongoing debate regarding the importance of chains (see also Gentry and Segev (2011) and Gentry et al. (2009)).

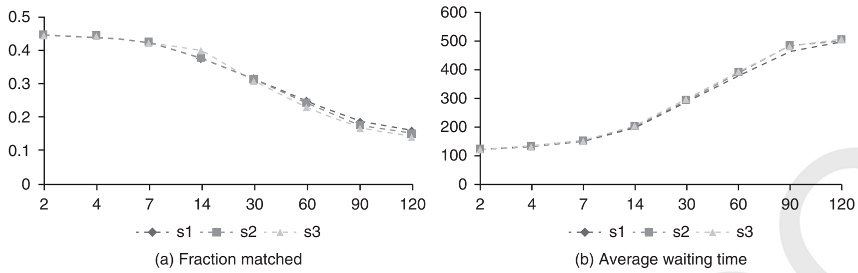

Figure 3.1 Statistics using APD data (taken from Ashlagi et al.). The x-axis represents the time interval between two match-runs. Each plot, S1-S3, stands for different prioritizations/weights assigned to pairs based on the patients' PRA.

matching frequencies and different sets of weights. Assigning high priority to highly sensitized patients increases their match rate but at the expense of other pairs.

Intuitively, matching frequently does not harm the fraction of transplanted patients because both underdemanded pairs and highly-sensitized patients accumulate in the pool. For instance, when an A-O patient-donor arrives, if the A patient is not too sensitized, there is likely an immediate match with an O-A pair. But if the A-O cannot match with an O-A pair, a match with any pair arriving in the near future is also unlikely; thus delaying other pairs from matching is also unlikely to help this pair. In other words, when the departure rate is low, many hard-to-match pairs accumulate in the pool, and waiting with a newly arriving, easy-to-match pair is unnecessary, because it is likely to match a hard-to-match pair. When the departure rate is high, matching infrequently will result in many departures of easy-to-match pairs.

Artificially thickening the market does not increase the fraction of matched pairs, but Ashlagi et al. (2018) also find that increasing the arrival rate increases the fraction of matched pairs up to a certain point (which is the fraction matched in a large market). Figure 3.2 plots the fraction of matched pairs under different arrival rates. Note that there is a diminishing return to scale. Note also, however, that the waiting time will keep decreasing as the arrival rate increases, even when the fraction of matched pairs does not increase anymore.

Consequently, increasing participation rate is much more important than artificially thickening the market. Rough intuition is that, at a low arrival rate, some O donors may match A patients, whereas at a large arrival rate, such A patients can be matched by other A donors.

Theoretical frontiers

Ünver (2010) first studied the problem of dynamic matching under large market assumptions. He found that a greedy algorithm that uses two-way and

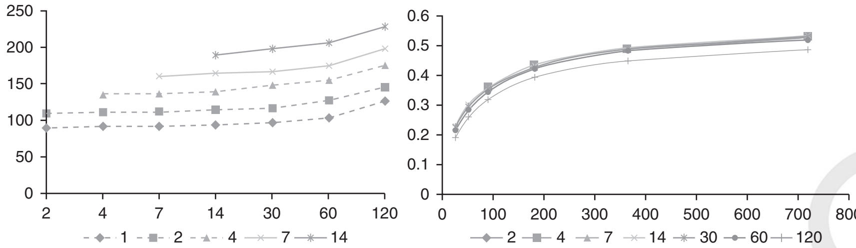

Figure 3.2 Varying arrival rates using APD data (taken from Ashlagi et al.). (Left) The x-axis represents the time interval between two match-runs, and each plots represents a different arrival rate. (Right) The x-axis represents the number of arrivals per year and y-axis the fraction of matched pairs.

three-way cycles is almost optimal and minor gains would be achieved by holding up some pairs.

Motivated by the sparsity of kidney exchange pools, various authors have studied dynamic matching in random graph-based models while abstracting away from blood types (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2014).

Anderson et al. (2017) looked at homogenous agents and considered three settings distinguished by the types of feasible exchanges: two-way cycles, two- and three-way cycles, and a single chain. They found that a greedy algorithm is approximately optimal. Furthermore, allowing three-way cycles results in significantly lower waiting times than only two-way cycles, and a single unrestricted chain reduces the waiting times even more. This result sheds light on the importance of chains in sparse pools. Akbarpour et al. (2014) looked at a similar model with departures and found that if departure times are known to the planner, matching patients just before they depart reduces the loss rate significantly in comparison to a greedy algorithm, but without this information, greedy matching is approximately optimal.

Ashlagi et al. (2013) allow for two types of agents, hard- and easy-to-match, and ask how the market thickness, as determined by arrival rates, affects efficiency under a class of myopic policies. They found a tight connection between the market thickness and the desired matching technology. When easy-to-match agents arrive more frequently to the market, two-way cycles are approximately optimal; otherwise, using chains is important.

More sophisticated policies, which consider the future, have been developed in the computer science literature. Dickerson et al. (2012b) learn the potential of different nodes in the compatibility graph to determine whether to match them or not (see also Dickerson and Sandholm (2014)).

In practice, myopic algorithms (which optimize periodically) are ubiquitous, and there is no evidence that such sophisticated algorithms lead to significant benefits. Moreover, when the market is sufficiently thick, simple matching rules are essentially optimal (Ünver, 2010). Otherwise, as figure 3.2 suggests, increasing arrival rates is arguably more important than designing the perfect algorithm.

3.3.4 The Marketplace for Kidney Exchange in the United States

In the last decade several national platforms emerged in the US, which seek to organize kidney exchange at a large scale. Participation, however, is not mandatory, and hospitals may decide to engage only partially in these platforms by enrolling some of their pairs and internally matching others. It is therefore natural to ask: How common is this behavior? Is efficiency harmed, and if so by how much? How can the platform align incentives to increase efficiency? Can full efficiency even be achieved?

The current national platforms vary according to size, operations, and algorithms. The major national platforms, the Alliance for Paired Donation (APD), United Network for Organ Sharing kidney exchange program (UNOS), and the National Kidney Registry (NKR) all involve the participation of many hospitals. Other platforms involve fewer hospitals and even single hospitals, such as Methodist at San Antonio (MSA). Typically, large scale programs use optimization algorithms to identify exchanges and may vary in how they prioritize patients or how frequently they search for exchanges.

Why Should Hospitals' Incentives Matter?

By and large, patients rely on hospitals (surgeons), and hospitals benefit from conducting more transplants (Sönmez and Ünver, 2013; Ashlagi and Roth, 2014). Let us see how hospitals may benefit from participating only partially in kidney exchange platforms. Consider a hospital A with two pairs, $a1$ and $a2$, which it can match internally through a two-way cycle. Suppose A enrolls both pairs to the platform, and the platform can either match $a1$ with $a2$ or match $a1$ with some other pair $b1$ that does not belong to A. If $b1$ has the highest priority, than $a1$ will remain unmatched. Thus hospital A is better off matching internally $a1$ and $a2$. Suppose there is another pair $b2$ that can only match with $a2$. If A would match $a1$ and $a2$ internally, only two transplants will happen instead of four.

Roth and Sonmez and Unver find that there is there is no efficient strategyproof mechanism ((Sönmez and Ünver, 2013; Roth, 2008)). Ashlagi and Roth (2014) studied extensively the free riding problem and found that under large market assumptions the cost from requiring allocations to be individually rational is low (see also Toulis and Parkes (2015)). They further suggest to adopt a "frequent flyer program" that will encourage hospitals to enroll their easy-to-match pairs. Therefore, optimizing without considering hospitals' incentives may potentially result in a large loss due to hospitals' behavior. But how big of an issue is this in practice?

Market Failure in Kidney Exchange

In a recent paper, Agarawal et al. (2017) accumulated data from various sources to study the kidney exchange marketplace in the US, quantified the efficiency loss, and offer solutions for how to fix it. They proposed a stylized model based on producer theory that predicts the sources of inefficiency and pursues these objectives empirically.

The first dataset they used is from the NKR, the largest platform in the US and includes all submissions (medical data of all incompatible pairs and altruistic donors) and transplants conducted between the years 2012–2014. Another dataset, obtained from UNOS, is the list of all live-donor transplants ever conducted in the US, and whether they are due to an exchange or not.

Agarawal et al. (2017) find that the US kidney exchange market is *highly fragmented*, with more than 60 percent of transplants conducted through internal matches (that are not facilitated by the NKR). Furthermore, larger hospitals are more likely to participate at the NKR, though hospitals vary in their level of participation. Importantly, they find that the NKR is *selected to have harder-to-match pairs*: the smaller the fraction of pairs a hospital enrolls, the more sensitized the patients within these enrolled pairs. They further provide smoking-gun evidence for large efficiencies: **while** *2.5 percent of the transplanted O kidneys go to non-O recipients at the NKR, more than 11 percent of such transplants occur within internal matches.*[19] Recall that in a large market it is efficient to transplant O kidneys in O recipients.

Agarawal et al. (2017) propose to model the kidney exchange as a platform that receives submissions and produces transplants. They take a steady-state approach (where submissions and transplants are per time period, say per year), but do not adhere to the previous large market assumptions. The platform is associated with a *production function* f, which receives vectors of submissions and generate transplants. The platform rewards hospitals with transplants, either immediately or in the future. Agarawal et al. (2017) considered the problem of maximizing welfare subject to two constraints. First, hospitals submit pairs optimally to maximize their own utility, and second, the platform is constrained from promising more transplants than it can generate.

One insight from the model is that to maximize welfare, optimal rewards should equal marginal products minus some adjustment term (the adjustment terms is zero when the platform operates at a constant returns to scale regime). In particular, hospitals should be rewarded based on their marginal contribution to the platform. This provides a simple explanation why current algorithms, which essentially attempt to maximize the number of transplants don't provide hospitals with the right incentives.

19 Large gaps remain even when restricting the non-O recipient to be low sensitized.

The finding by Agarawal et al. (2017) suggests that platform incentives can be solved by simply using point systems. Hospitals will maintain a point balance of transplants, and the platform will break ties in favor of hospitals that have a larger amount. For platforms that operate at constant returns to scale regime, the exact rewards can be computed by estimating the derivatives of the production function. The idea is that hospitals will earn points based on their marginal contribution to the platform, whereas the platform will favor favor hospitals with higher balances during the matching process (for example through tie-breaking). For implementation details and challenges, see Agarawal et al. (2017).

Agarawal et al. (2017) further quantify the inefficiency in the data. Estimating the production function reveals that the NKR is operating at the constant returns to scale regime, but many hospitals that match internally operate at an inefficient scale. They find that misaligned incentives account for around 200–400 transplants per year.

Finally, Agarawal et al. (2017) find that platform incentives do not account for all inefficiency. The remaining inefficiency is due to agency problems. This prediction, also generated by their model, is also supported by significant efforts in recent years to organize financial agreements between insurance companies and hospitals (Rees et al., 2012; Held et al., 2016; Irwin et al., 2012; DHHS, 2016). One challenge was to engage private insurance companies in a standard acquisition charge to reimburse for expenses prior to transplants.

As Agarawal et al. (2017) point out, kidney exchange, which is a seemingly unusual market, faces classic market failures, which can be addressed using market and nonmarket tools. It will be interesting to see how this market evolves in the next few years, now that some of these agency problems have been alleviated and platforms such the NKR have adopted point systems.

3.3.5 Final Comments on Kidney Exchange

Kidney exchange is now responsible for more than 11 percent of live-donor transplants in the United States. Despite this success, many challenges still remain in order to allow platforms to operate more smoothly, as well as to grow this marketplace. Although optimization plays a role in kidney exchange, increasing participation is a first-order consideration in order to increase the number of of transplants. In the US creating one national pool is more challenging than in other countries, arguably due to institutional structure. However, small countries that seek to work together in order to increase matching opportunities may face similar concerns. For a thorough survey of kidney exchange practices in European countries, see Biro et al. (2017).

There are ongoing efforts towards new innovations, such as the Global Kidney Exchange that aims to overcome medical and financial incompatibilities by matching pairs in developed countries that lack transplantation facilities with pairs in developed countries (Rees et al., 2016). But similar to other innovations in the field of transplantation, ethical considerations are an integral part of the process toward implementation. Indeed, Global Kidney Exchange sparked a loud debate in the transplantation community (Delmonico and Ascher, 2017; Roth et al., 2017; Rees et al., 2017). It remains to be seen how this potential new market will develop.

Marketplaces such as platforms for kidney exchanges are usually part of a larger market. Little is still known about the design of marketplaces that face outside competition.

References

Abdulkadiroğlu, A., Pathak, P. and Roth, A. (2005a). The New York City High School Match. *American Economic Review*, 95(2), pp. 364-367.

Abdulkadiroğlu, A. and Sönmez, T. (2003). School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3), pp. 729-747.

Abdulkadiroğlu, A., Pathak, P., Roth, A. and Sönmez T. (2005b). The Boston Public School Match. *American Economic Review*, pp. 368–371.

Abraham, D.J., A. Blum, and T. Sandholm. (2007). Clearing Algorithms for Barter Exchange Markets: Enabling Nationwide Kidney Exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, pp. 295–304, 2007.

Agarawal, N., I. Ashlagi, E. Azavedo, C. Featherstone, and O. Karaduman. (2017). "Market Failure in Kidney Exchange." Working paper.

Akbarpour, M., S. Li, and S. Oveis Gharan. (2014). "Dynamic matching market design." Working paper.

Anderson, R., I. Ashlagi, Y. Kanoria, and D. Gamarnik. (2017). Efficient dynamic barter exchange. *Operations Research* 56(6):1446–1459.

Anderson. R., I. Ashlagi, D. Gamarnik, and A.E. Roth. (2015). Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences*, 112(3):663–668.

Ashlagi, I., and A.E. Roth. (2014). Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863.

Ashlagi, I., D.S. Gilchrist, A.E. Roth, and M.A. Rees. (2011a). Nonsimultaneous Chains and Dominos in Kidney Paired Donation—Revisited. *American Journal of Transplantation*, 11:984–994.

Ashlagi, I., D. S. Gilchrist, A. E. Roth, and M. A. Rees. NEAD Chains in Transplantation. *American Journal of Transplantation*, 11:2780–2781, 2011b.

Ashlagi, I., Adam Bingaman, Maximilien Burq, Vahideh Manshadi, David Gamarnik, Cathi Mur- phey, A.E. Roth, Marc L. Melcher, and Michael A. Rees. The effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. (2018). *American Journal of Transplantation.* 18(5): 1177–1186.

Ashlagi, I., David Gamarnik, Michael A Rees, and Alvin E Roth. The need for (long) chains in kidney exchange. Technical report, National Bureau of Economic Research, 2012.

Ashlagi, I., P. Jaillet, and V.H. Manshadi. (2013). Kidney exchange in dynamic sparse heterogenous pools. In *In ACM Conference on Electronic Commerce*. Citeseer, 2013.

Ashlagi, I., M. Braverman, and A. Hassidim. (2014). Stability in large matching markets with complementarities. *Operations Research*, 62(4):713–732.

Ashlagi, I., Y. Kanoria, and J.D. Leshno. (2017). Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125(1):69–98.

Avery, C., C. Jolls, R.A. Posner, and A.E. Roth. (2001). The market for federal judicial law clerks. *The University of Chicago Law Review*, 68(3):793–902.

Azevedo, E.M. and J.W. Hatfield. (2012). Complementarity and multidimensional heterogeneity in matching markets. *Unpublished mimeo*.

Becker, G.S. and J.J. Elias. (2007). Introducing incentives in the market for live and cadaveric organ donations. *The Journal of Economic Perspectives*, 21(3):3–24.

Biro, P., L. Burnapp, H. Bernadette, A. Hemke, R. Johnson, J. van de Klundert, and D. Manlove. (2017). First Handbook of the COST Action CA15210: European Network for Collaboration on Kidney Exchange Programmes (ENCKEP).

Biro, P., D.F. Manlove, and R. Rizzi. (2009). Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Mathematics, Algorithms and Applications*, 1(04):499–517.

Che, Y-K, J. Kim, and F. Kojima. (2015). Stable matching in large economies. *Department of Economics, Columbia University*. Working paper.

Constantino, M., Xenia Klimentova, A.V., and A. Rais. (2013). New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research*, 231(1):57–68.

Delmonico, F.L. and N.L. Ascher. (2017). Opposition to irresponsible global kidney exchange. *American Journal of Transplantation*, 17(10):2745–2746.

Demange, G., D. Gale, and M. Sotomayor. (1987). A further note on the stable matching problem. *Discrete Applied Mathematics*, 16(3):217–222.

DHHS. Medicare Provider Reimbursement Manual, Part 2. (2016). www.cms.gov/Regulations-and-Guidance/Guidance/Transmittals/Downloads/R7PR241.pdf.

Dickerson, J.P., A.D. Procaccia, and T. Sandholm. (2012a). Optimizing Kidney Exchange with Transplant Chains: Theory and Reality. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*.

Dickerson, J.P., A.D. Procaccia, and T. Sandholm. (2012b). Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI Workshop: Modern Artificial Intelligence for Health Analytics*.

Dickerson, J.P., A.D. Procaccia, and T. Sandholm. (2012c). Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of the 11th*

International Conference on Autonomous Agents and Multiagent Systems—Volume 2, pp. 711–718. International Foundation for Autonomous Agents and Multiagent Systems.

Dickerson, J.P. and T. Sandholm. (2014). Balancing efficiency and fairness in dynamic kidney exchange. In *AAAI Workshop: Modern Artificial Intelligence for Health Analytics*.

Dubins, L.E. and D.A. Freedman. (1981). Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494.

Ferrari, P., Willem Weimar, Rachel J Johnson, Wai H Lim, and Kathryn J Tinckam. (2014). Kidney paired donation: principles, protocols and programs. *Nephrology Dialysis Transplantation*, 30 (8):1276–1285.

Gale, D. and L.L. Shapley. (1962). College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 69:9–15.

Gentry, S.E., R. A. Montgomery, B. J. Swihart, and D. L. Segev. (2009). The Roles of Dominos and Nonsimultaneous Chains in Kidney Paired Donation. *American Journal of Transplantation*, 9: 1330–1336.

Gentry, S.E. and D.L. Segev. (2011). The honeymoon phase and studies of nonsimultaneous chains in kidney paired donation. *American Journal of Transplantation*, 11:2778–2779.

Glorie, K., A. Wagelmans, and J. van de Klundert. (2012). Iterative branch-and-price for hierarchical multi-criteria kidney exchange. Technical report.

Held, P.J., F McCormick, A Ojo, and John P Roberts. (2016). A cost-benefit analysis of government compensation of kidney donors. *American Journal of Transplantation*, 16(3):877–885.

Hitsch, G.J., A. Hortaçsu, and D. Ariely. (2010). Matching and sorting in online dating. *American Economic Review*, 100(1):130–163.

Immorlica, N. and M. Mahdian. (2005). Marriage, honesty, and stability. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 53–62. Society for Industrial and Applied Mathematics.

Irwin, F.D., AF Bonagura, SW Crawford, and M Foote. (2012). Kidney paired donation: a payer perspective. *American Journal of Transplantation*, 12(6):1388–1391.

Johnson, R.J., J.E. Allen, S.V. Fuggle, J.A. Bradley, C. Rudge et al. (2008). Early experience of paired living kidney donation in the United Kingdom. *Transplantation*, 86(12): 1672–1677.

Kagel, J.H. and A.E. Roth. (2000). The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment. *The Quarterly Journal of Economics*, 115(1): 201–235.

Klaus, B. and F. Klijn. (2005). Stable matchings and preferences of couples. *Journal of Economic Theory*, 121(1):75–106.

Klimentova, X., F. Alvelos, and A. Viana. (2014). A new branch-and-price approach for the kidney exchange problem. In *International Conference on Computational Science and Its Applications*, pp. 237–252. Springer.

Knuth, D.E., R. Motwani, and B. Pittel. (1990). Stable husbands. *Random Structures & Algorithms*, 1(1):1–14.

Kojima, F., P.A. Pathak, and A.E. Roth. (2013). Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics*, 128(4):1585–1632.

Leider, S. and A.E. Roth. Kidneys for sale: Who disapproves, and why? (2010). *American Journal of transplantation*, 10(5):1221–1227.

Malik, S. and Edward Cole. (2014). Foundations and principles of the Canadian living donor paired exchange program. *Canadian Journal of Kidney Health and Disease*, 1(1):6, 2014.

Niederle, M. and A.E. Roth. (2003). Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match. *Journal of Political Economy*, 111(6):1342–1352.

Pittel, B. (1992). On likely solutions of a stable matching problem. In *Proceedings of the third annual ACM-SIAM symposium on Discrete algorithms*, pp. 10–15. Society for Industrial and Applied Mathematics.

Rapaport, F.T. (1986). The case for a living emotionally related international kidney donor exchange registry. *In Transplantation proceedings, volume* 18(3), pp. 5–9.

Rees, M., T. Dunn, S. Rees, J. Rogers, L. Reece, A. Roth, C. Kuhr, O. Ekwenna, D. Fumo, K. Krawiec, et al. (2016). Global kidney exchange. In *American Journal of Transplantation*, volume 16, pp. 352–352. Wiley Online Library.

Rees, M.A., J.E. Kopke, R.P. Pelletier, D.L. Segev, M.E. Rutter, A.J. Fabrega, J. Rogers, O.G. Pankewycz, J. Hiller, A.E. Roth, T. Sandholm, M.U. Ünver, and R.A. Montgomery. (2009). A non-simultaneous extended altruistic donor chain. *New England Journal of Medicine*, 360: 1096–1101.

Rees M.A., S. Paloyo, A.E. Roth, K.D. Krawiec, O. Ekwenna, C.L. Marsh, A.J. Wenig, and T.B. Dunn. (2017). Global kidney exchange: Financially incompatible pairs are not transplantable compatible pairs. *American Journal of Transplantation*.

Rees, M.A., M.A. Schnitzler, E.Y. Zavala, J.A. Cutler, A.E. Roth, F.D. Irwin, S.W. Crawford, and A.B. Leichtman. (2012). Call to develop a standard acquisition charge model for kidney paired donation. *American Journal of transplantation*, 12(6):1392–1397.

Ross, L.F. and E.S. Woodle. (2000). Ethical issues in increasing living kidney donations by expanding kidney paired exchange programs. *Transplantation*, 69(8):1539–1543.

Ross, L.F., D.T. Rubin, M. Siegler, M.A. Josephson, J.R. Thistlethwaite Jr., and E.S. Woodle. (1997). Ethics of a paired-kidney-exchange program.

New England Journal of Medicine 336:1752–1755. DOI: 10.1056/NEJM199706123362412.

Roth, A.E. (2008). What have we learned from market design? *Economic Journal*, 118:285–310.

Roth, A.E. and E. Peranson. (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review*, 89:748–780.

Roth, A.E., T. Sönmez, and M. U. Ünver. (2004) Kidney exchange. *Quarterly Journal of Economics*, 119: 457–488.

Roth, A.E., T. Sönmez, and M. U. Ünver. (2005). Pairwise kidney exchange. *Journal of Economic Theory*, 125:151–188.

Roth, A.E, T. Sönmez, and M. U. Ünver. (2007). Efficient kidney exchange: coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97:828–851.

Roth, A.E. (2002). The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica*, 70:1341–1378.

Roth, A.E. (1982). The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628.

Roth, A.E. (1984) The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy*, 92(6):991–1016.

Roth, A.E. (2007). Repugnance as a constraint on markets. *The Journal of Economic Perspectives*, 21 (3):37–58.

Roth, A.E. and Xiaolin Xing. (1994). Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review*, pp. 992–1044.

Roth, A.E. and X. Xing. (1997). Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of Political Economy*, 105(2):284–329.

Roth, A.E., K.D. Krawiec, S. Paloyo, O. Ekwenna, C.L. Marsh, A.J. Wenig, T.B. Dunn, and M.A. Rees. (2017). People should not be banned from transplantation only because of their country of origin. *American Journal of Transplantation*, 17(10):2747–2748.

Saidman, S.L., A.E. Roth, T. Sönmez, M.U. Ünver, and F.L. Delmonico (2006). Increasing the Opportunity of Live Kidney Donation by Matching for Two and Three Way Exchanges. *Transplantation*, 81:773–782.

Segev, D.L., S.E. Gentry, D.S. Warren, B. Reeb, and R.A. Montgomery (2005). Kidney paired donation and optimizing the use of live donor organs. *JAMA*, 293(15):1883–1890.

Sönmez, T. and M.U. Ünver. (2013). Market design for kidney exchange. *The Handbook of Market Design*, pp. 93–137.

Toulis, P. and D.C. Parkes. (2015). Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior*, 91:360–382.

United States Renal Data System. *2016 USRDS annual data report: Epidemiology of kidney disease in the United States*, volume 3. (2016). Bethesda, MD: National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases.

Ünever, M.U. (2010). Dynamic Kidney Exchange. *Review of Economic Studies*, 77(1):372–414.

Utku, M.U. (2001). Ünver. Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets. *Journal of Economic Dynamics and Control*, 25(6):1039–1080.

Wolfe, R.A., V.B. Ashby, E.L. Milford, A.O. Ojo, R.E. Ettenger, L.Y.C. Agodoa, P.J. Held, and F.K. Port. (1999). Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England Journal of Medicine*, 341(23):1725–1730.