NEW CHALLENGES FOR MARKET DESIGN‡

New Challenges in Multihospital Kidney Exchange†

By Itai Ashlagi and Alvin E. Roth*

A kidney transplant is the treatment of choice for end-stage renal disease, but over 90,000 patients are waiting for a cadaver kidney in the United States, and fewer than 11,000 such transplants are performed annually. Live donation is also possible, and there are now more live than deceased donors each year in the US, although they still account for fewer than 6,500 transplants a year (since living donors only donate one kidney). And having a healthy, willing live donor is not enough: sometimes a donor's kidney is incompatible with the intended recipient, either because of blood type or immunological incompatibilities. Incompatibility between donor and intended recipient creates the demand for *kidney exchange* (a.k.a. *kidney paired donation*): an incompatible patient-donor pair can donate a kidney to a compatible recipient and receive a kidney from a compatible donor.

The first kidney exchange was in Korea (Kwak et al. 1999), where the high frequency of blood types A and B make exchanges due to blood-type incompatibility readily available (an A-B pair exchanging with a B-A pair, where X-Y denotes a patient of blood type X and donor of blood type Y), more than in the US, where blood type B is relatively rare. The first kidney exchange in the US, in New England in 2000, also involved two blood type–incompatible pairs

(see Wallis et al. 2011 for history and references). And for most patients on the waiting list for cadaver kidneys, blood type determines compatibility with a given donor. Only 10 percent of those 90,000 patients are "highly sensitized," meaning they are *immunologically* incompatible with more than 80 percent of donors with compatible blood type.

But the patients enrolled in the most active kidney exchange networks are much more highly sensitized: in the 4 exchange networks with which we have worked, the percentage of highly sensitized patients is from 50–80 percent of those enrolled (Ashlagi, Gamarnik, and Roth 2011). The present paper considers why this is the case, and its consequences.

While the first proposal for organizing kidney exchange on a large scale involved exchanges organized as both cycles and chains, logistical constraints required that the initial exchanges conducted by the New England Program for Kidney Exchange, the Alliance for Paired Donation (APD), and other networks were between just two patient-donor pairs (Roth, Sönmez, and Ünver 2004, 2005a,b).[1] Subsequent work suggested that as patient pools grew larger, expanding the infrastructure to allow only slightly larger, three- and four-way exchanges, would be efficient (Roth, Sönmez, and Ünver 2007a). But the prevalence of highly sensitized patients among those enrolled in kidney exchange has brought long chains back into the picture in an important way, after the

[1] In addition to those two large kidney exchange clearinghouses, kidney exchange today is practiced by a growing number of hospitals and consortia. Computer scientists have become involved, and an algorithm of Abraham, Blum, and Sandholm (2007) designed to handle large populations was briefly used in the APD and has been used in the United Network for Organ Sharing (UNOS) pilot program for a national exchange.

introduction of nonsimultaneous chains initiated by a nondirected donor (Roth et al. 2006; Rees et al. 2009). Chains now contribute many of the kidney exchanges performed by all of the largest multihospital networks except the UNOS pilot program. The usefulness of chains turns out to be closely related to the highly sensitized patient population. And one of several causes of the high percentage of highly sensitized patients is that many large transplant centers are withholding their easy-to-match patient-donor pairs, and only enrolling their hard-to-match pairs. This reduces the total number of transplants that can be achieved, particularly for the most highly sensitized patients.[2] (Another reason the patient pool becomes highly sensitized is that sensitized patients remain unmatched longer, and build up in the patient pool, but the data show even the initially enrolled patients are highly sensitized.)

There are several reasons why hospitals wish to conduct their own exchanges for easy-to-match pairs, including *compatible pairs* in which the donor can donate to the intended patient.[3] These include the difficult logistics of coordinating with other hospitals, paying for the tests that establish donor-recipient compatibility before knowing who the recipient will be (Rees et al. forthcoming), and, finally, the fact that presently used matching algorithms do not make it individually rational (IR) for hospitals to enroll all their pairs, since they do not guarantee each hospital that enrolling all its patients will allow it to perform as many transplants as it could get by enrolling only some patients and doing some exchanges internally among its own patients.

The initial papers on kidney exchange focused on incentives for patients and their surgeons, but the current problems facing kidney exchange arise from the fact that hospitals have become the main players, and have different strategy sets than individuals, since directors of transplant centers deal with multiple patient-donor pairs. Section I considers how the current algorithms

fail to make it safe for hospitals to enroll all their pairs, and how this could be fixed.

Section II considers why long chains play such an important role. Exchange pools are modeled as compatibility graphs whose vertices are incompatible pairs with directed links indicating compatibility between donors and patients. Previous studies focused on the relatively dense compatibility graphs that would arise if blood-type incompatibilities were dominant, rather than the sparse graphs corresponding to many highly sensitized patients.

I. Individual Rationality and Incentives for Hospitals to Participate Fully

Most kidney exchange clearinghouses try to maximize the (weighted) number of transplants without attention to whether some incompatible pairs can be matched internally by the hospitals that entered them into the database. Thus, it may not be individually rational (IR) for a hospital to contribute those pairs it can match internally (see, e.g., Roth 2008).[4] For example, consider a hospital A with two pairs, a_1 and a_2, that it can match internally. Suppose it enters those two pairs in a centralized exchange. It may be that the weighted number of transplants is maximized by including a_1 in an exchange but not a_2, in which case only one of hospital A's patients will be transplanted, when it could have performed two transplants on its own.

An allocation (set of disjoint exchanges) is IR if no hospital can match more pairs internally than the number of its pairs matched in the allocation. In Ashlagi and Roth (2011), we show that efficiency and individual rationality cannot always be satisfied simultaneously.[5] We also show constructively, however, that this is not an issue in large exchange pools; under minor assumptions, as the number of hospitals grows (the compatibility graph grows), with probability tending to one there exists an ϵ-efficient allocation that is (i) IR and (ii) doesn't use exchanges of size more than three.[6] Our model builds on Erdos-Renyi probabilistic graphs, which allows

---

[2] For example, the UNOS kidney exchange program was begun in 2010, and through the end of 2011 it had accomplished only 17 transplants, from a small pool of enrolled patient-donor pairs, despite having many affiliated hospitals.

[3] Failure to enroll compatible pairs in kidney exchange not only deprives those pairs of the possibility of a better matched kidney, it also exacerbates the shortage of blood type O donors who are blood type compatible with all patients.

[4] Some weighted matching algorithms put some weight on internal exchanges, but this does not solve the problem.

[5] Roth, Sönmez, and Ünver (2007b) first showed that no strategyproof mechanism is efficient.

[6] Toulis and Parkes (2010) show a similar result assuming exchanges are limited to size 2.

TABLE 1

| Hospitals | Pairs | Max-match | | | IR-match | |
		H's gain	Truthful reporting No. matches	Withholding No. matches	H's gain	Truthful reporting No. matches
10	100	1.055	52 (12.9)	43 (8)	0.994	51.6
12	120	1.066	65 (17.2)	54 (11)	0.995	64 (16.6)
15	150	1.061	84 (23.8)	71 (16)	0.988	83.5 (23)
20	200	1.048	117 (36)	101 (26)	0.986	116.6 (34.9)

Note: Withholding internal matches versus reporting truthfully in Max match and IR match.

us to show that in the limit as the graph grows large and dense, an efficient allocation needs exchanges of size no more than three.[7]

Limit theorems, however, do not address patient populations of clinically relevant size, and so we report simulations that show the limit results are achieved in populations of the size we presently see. These simulations suggest that considerable gains could be achieved by adopting an IR mechanism.

Simulations.—For each iteration we generate compatibility graphs as follows. According to blood type and sensitivity (Panel Reactive Antibody (PRA)) distributions consistent with the UNOS population (see, e.g., Ashlagi and Roth 2011), a patient and one to three related donors are drawn uniformly. We then test tissue type compatibility between the patient and her donors using the patient's PRA (which is a probability of incompatibility). A patient and one of her related donors join the pool (as an incompatible pair) if the patient is incompatible with all her related donors. In addition, a nondirected donor is generated randomly. To complete the graph, tissue type tests are conducted between donors and patients of different pairs. Finally, we associate each incompatible pair and each nondirected donor to a random hospital.

We compare an algorithm like those currently in use, *max-match*, to an IR mechanism, *IR-match*. Two behaviors by hospitals are considered: (i) truth-telling—the hospital reports all its pairs; and (ii) *withhold internal matches*—the hospital withholds a maximum set of pairs it can match internally.

Each row of Table 1 corresponds to a scenario. The first scenario includes 10 hospitals, and 100 pairs. The ratio between the number of matches a given hospital, *H*, obtains from withholding to the number of matches *H* obtains from reporting truthfully, given that all other hospitals report truthfully, is given in the third column (under the max-match algorithm) and the sixth column (under the IR-match algorithm). Thus, it is beneficial for *H* to withhold its internal matches under max-match, but not under IR-match.[8] The maximum number of matches that can possibly be obtained is given in the fourth column (under max-match when all hospitals report truthfully). The number of matches obtained when all hospitals withhold their internal matches is given in the fifth column. Finally, the number of matches obtained when all hospitals report truthfully under IR-match is given in the last column.

To summarize, (i) a hospital profits from withholding its internal matches under max-match but not under IR-match; (ii) more than ten percent more matches will be achieved under truthful reporting under IR-match than under max-match assuming hospitals withhold internal matches; and (iii) the cost of using

---

[7] This refines a result by Roth, Sönmez, and Ünver (2007b) that in sufficiently large graphs exchanges of size more than 4 are not needed for efficiency.

[8] When all other hospitals withhold their pairs, the gain under max-match for hospital *H* increases to more than 10 percent.
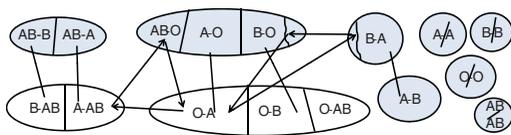
FIGURE 1. THE STRUCTURE OF AN EFFICIENT ALLOCATION
WITHOUT ALTRUISTIC DONORS

Notes: All self-demanded pairs are matched to each other. All B-A pairs are matched to A-B (assuming more B-As than A-Bs); the remainder of the A-B pairs are matched in three-way exchanges using O-As and B-Os. AB-O are matched in three-ways, each using two overdemanded pairs, and every other overdemanded pair is matched to an under-demanded pair.

IR-match is very small even assuming that all hospitals report truthfully under max-match, and the gain from using IR-match is substantial if hospitals withhold under max-match (assuming they report truthfully under IR-match).

Individual rationality alone is insufficient to guarantee full participation: hospitals may have incentives to withhold *overdemanded* pairs other than internal matches (for example, an A-O pair is overdemanded since there are fewer such pairs than O-A pairs due to blood type compatibility). Intuitively, if hospital *A* has an overdemanded pair a_1 that it can internally match to pair a_2 or a_3 that are each underdemanded, individual rationality constraints alone may still leave hospital *A* better off withholding a_1 and waiting to see if one of its underdemanded pairs, a_2 or a_3, is unmatched by the mechanism and then internally matching that pair to a_1. In Ashlagi and Roth (2011) we introduce an "almost" efficient mechanism that makes truthful reporting an ϵ-Bayes-Nash equilibrium as the number of participating hospitals grows large.

II. The Need for Long Chains

For large, dense compatibility graphs, there is almost always an efficient allocation as in Figure 1 (see Roth, Sönmez, and Ünver 2007a and Ashlagi and Roth 2011). Exchanges of size greater than three are not needed. Furthermore, Roth, Sönmez, and Ünver (2007a) showed via simulations that even in small exchange pools, exchanges of size greater than three do not add many transplants. These simulations show (as does the related theory) that each nondirected

donor can increase the number of transplants by at most three. But current exchange programs have many highly sensitized patients, and thus have sparser graphs than graphs generated from the statistics of the general patient population.

Ashlagi, Gamarnik, and Roth (2011) provide empirical evidence that longer exchanges and long chains increase efficiency, and provide a theoretical framework based on sparse Erdos-Renyi graphs. Intuitively, if a patient *p* is highly sensitized and can receive a kidney from very few donors, the chance that *p* will be part of a short exchange is small (and including a small exchange, even if one exists, might be inefficient). Previous papers explicitly or implicitly ignored tissue type compatibility, and Ashlagi, Gamarnik, and Roth (2011) suggest that this type of compatibility cannot be neglected. They show analytically that for sparse graphs, longer exchanges and longer chains increase the number of transplants linearly as the graph grows. They use simulations to show that their results give a good approximation in small graphs as well.[9]

Below, we report simulations of the effectiveness of long chains. Graphs are generated as in the previous simulations, except we generate more highly sensitized patients by assuming each patient has 3–7 potential related donors; each patient is tested for compatibility with each related donor (this generates a pool with approximately 60 percent high-PRA patients). The advantage of longer exchanges and chains is clear from Table 2 (see also Dickerson, Procaccia, and Sandholm 2011 who also conduct simulations to confirm the utility of long chains).

Each row in the table corresponds to a different scenario. In the first 3 scenarios there are 100 pairs (with an average of 62.5 high-PRA patients) and the number of nondirected donors is either 2 or 6. The third through seventh columns each describe the number of matches obtained (and in parentheses the number of highly sensitized matches) under different matching algorithms— (k, l) means we search for the maximum number of transplants allowing exchanges (cycles) up to size *k* and chains up to size *l*. Thus, with 100

---

[9] Ashlagi et al. (2011a, b) use simulations to show that long chains increase efficiency in a dynamic setting. Dickerson, Procaccia, and Sandholm (2012) support these findings with extensive simulations.

Table 2

Size	NDDS	(2, 3)	(3, 3)	(3, 4)	(3, ∞)	(∞, ∞)
100 (62.5)	1	27.72 (11.78)	42.58 (22.19)	43.18 (22.52)	46.44 (25.66)	56.78 (35.62)
	2	29.44 (12.71)	44.48 (23.48)	45.64 (24.22)	54.25 (32.17)	58.24 (36.75)
	6	39.08 (18.58)	52.52 (28.82)	55.21 (31.32)	63.1 (38.89)	63.52 (39.98)
150 (94.11)	1	47.49 (21.32)	76.11 (42.86)	76.78 (43.24)	76.72 (43.46)	98.08 (65.07)
	2	49.88 (22.14)	79.26 (44.66)	80.62 (45.76)	90.16 (55.13)	101.34 (67.8)
	6	60.71 (29.02)	87.47 (50.19)	91.04 (53.62)	104.53 (66.81)	105.46 (69.57)

Note: Average number of matches and high-PRA matches (in parentheses) with different number of nondirected donors (NDDS) and different-sized pools.

pairs and 2 nondirected donors, an average of 44.48 matches were found when searching for matches using cycles of length at most 3 and chains of length 3.

III. Conclusion

Kidney exchange programs are maturing, yet progress is still slow. This note describes some major issues facing kidney exchange today and suggests solutions that may significantly increase the number of transplants. First, transplant centers withhold their easy-to-match patient-donor pairs, and we suggest modifying the commonly used matching mechanisms to make it IR for hospitals to participate with all their patients. This applies to compatible as well as incompatible pairs. Blood types cause some pairs to be overdemanded and others to be underdemanded. Roth, Sönmez, and Ünver (2005a) showed that a significant increase in the number of kidney exchanges could be achieved by allowing compatible pairs to participate (see also Sönmez and Ünver 2011, who analyze graphs with compatible pairs).

Second, due to the highly sensitized pools and their sizes, allowing long chains will significantly increase the number of transplants, especially for highly sensitized patients. While many exchange networks have begun to do this, the national pilot program has yet to adopt this approach.

As kidney exchange has grown, the set of players has changed. To foster further growth,

the design of kidney exchange clearinghouses must respond.

REFERENCES

Abraham, David J., Avrim Blum, and Tuomas Sandholm. 2007. "Clearing Algorithms for Barter Exchange Markets: Enabling Nationwide Kidney Exchanges." In *Proceedings of the 8th ACM Conference on Electronic Commerce*, 295–304. San Diego, CA: Association for Computing Machinery.

Ashlagi, Itai, and Alvin E. Roth. 2011. "Individual Rationality and Participation in Large Scale, Multi-Hospital Kidney Exchange." National Bureau of Economic Research Working Paper 16720.

Ashlagi, Itai, David Gamarnik, and Alvin E. Roth. 2011. "The Need for (Long) NEAD Chains." Unpublished.

Ashlagi, Itai, D. S. Gilchrist, Alvin E. Roth, and Michael A. Rees. 2011a. "NEAD Chains in Transplantation." *American Journal of Transplantation* 11: 2780–81.

Ashlagi, Itai, D. S. Gilchrist, Alvin E. Roth, and Michael A. Rees. 2011b. "Nonsimultaneous Chains and Dominos in Kidney Paired Donation—Revisited." *American Journal of Transplantation* 11 (5): 984–94.

Dickerson, John P., Ariel D. Procaccia, and Tuomas Sandholm. 2011. "Optimizing Kidney Exchange with Transplant Chains: Theory and Reality." In *Proceedings of the Eleventh*

International Conference of Autonomous Agents and Multiagent Systems (AAMAS). Valencia, Spain. June 4–8.

Kwak, J. Y., O. J. Kwon, K. S. Lee, C. M. Kang, H. Y. Park, and J. H. Kim. 1999. "Exchange-Donor Program in Renal Transplantation." *Transplantation Proceedings* 31 (1–2): 344–45.

Rees, Michael A. et al. 2009. "A Nonsimultaneous, Extended, Altruistic-Donor Chain." *New England Journal of Medicine* 360: 1096–1101.

Rees, Michael A., M. A. Schnitzler, J. A. Cutler, Alvin E. Roth, F. D. Irwin, S. W. Crawford, and Alan B. Leichtman. Forthcoming. "Call to Develop a Standard Acquisition Charge Model for Kidney Paired Donation." *American Journal of Transplantation*.

Roth, Alvin E. 2008. "What Have We Learned from Market Design?" *Economic Journal* 118 (527): 285–310.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2004. "Kidney Exchange." *Quarterly Journal of Economics* 119 (2): 457–88.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2005a. "A Kidney Exchange Clearinghouse in New England." *American Economic Review* 95 (2): 376–80.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2005b. "Pairwise Kidney Exchange." *Journal of Economic Theory* 125 (2): 151–88.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2007a. "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences." *American Economic Review* 97 (3): 828–51.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2007b. "Notes on Forming Large Markets from Small Ones: Participation Incentives in Multi-Center Kidney Exchange." Unpublished.

Roth, Alvin E., Tayfun Sönmez, M. Utku Ünver, Francis L. Delmonico, and Susan L. Saidman. 2006. "Utilizing List Exchange and Nondirected Donation through Chain Kidney Paired Donations." *American Journal of Transplantation* 6 (11): 2694–705.

Sönmez, Tayfun, and M. Utku Ünver. 2011. "Altruistic Kidney Exchange." Unpublished.

Toulis, Panagiotis, and David C. Parkes. 2010. "A Random Graph Model of Kidney Exchanges: Optimality and Incentives." Unpublished.

Wallis, C. Bradley, Kannan P. Samy, Alvin E. Roth, and Michael A. Rees. 2011. "Kidney Paired Donation." *Nephrology Dialysis Transplantation* 26 (7): 2091–99.

This article has been cited by:

1. Judd B. Kessler,, Alvin E. Roth. 2012. Organ Allocation Policy and the Decision to Donate. *American Economic Review* **102**:5, 2018-2047. [Abstract] [View PDF article] [PDF with links]