Organ Exchanges: A Success Story of AI in Healthcare

John P. Dickerson Tuomas Sandholm

January 11, 2016

1 Description of proposed tutorial

Goal of the tutorial. Kidney exchanges represent a truly fielded example of technology at the intersection of artificial intelligence and economics—an example that is both recent and an active research area. Kidney exchanges are saving lives on an ongoing basis and growing annually, both in the United States and worldwide. They provide a wealth of ongoing opportunities for theoretical and experimental research in both fields. Prof. Sandholm's lab has been actively involved in kidney exchange research and fielding of the techniques since 2005. For example, the algorithms and software from this lab run the United Network for Organ Sharing (UNOS) kidney exchange, which now includes 143 transplant centers—over 60% of the centers in the US. That practical experience fielding a large-scale healthcare application of AI will also be reflected in the tutorial.

Our tutorial targets a general AI audience and will be of interest to both theoreticians and practitioners—and those in between. Kidney exchange research appears in conferences and journals that are of interest to AAAI attendees; an incomplete list is given below.

Computer Science. AAAI [5,17,21,22,28,31], AAMAS [18,20,29], EC [1, 2, 8, 12, 13, 19, 23], IJCAI [11, 32], SODA [3], Theoretical Computer Science [14]

Operations Research. Management Science [48], MSOM [27], EJOR [15]

Economics. GEB [6, 45], AER [9, 39, 41], QJE [38], JET [40, 43, 47], RES [46], Theoretical Economics [10]

The audience of this tutorial will walk away with knowledge of the following.

- Tree search and discrete optimization techniques used to clear large kidney exchanges in practice (particularly branch-and-price-based approaches to solving very large (mixed) integer programs);

- Dynamic optimization techniques used to learn how to balance competing dimensions in healthcare, like equity versus efficiency, while maintaining high-level qualitative control by humans;
- Theoretical random graph models used to inform policy decisions in markets like kidney exchange;
- Knowledge of the interplay between economics, computer science, operations research, and healthcare research, and how the intersection of those four areas informed the creation and informs the evolution of a large-scale healthcare application;
- Current challenges in kidney exchange and, more general, challenging aspects of fielding AI technology in healthcare; and
- An end-to-end perspective of the fielding and evolution of kidney exchange programs from a practitioner who was, and two practitioners who are, at the forefront of that nascent field.

Content. Please see Section 2 for a detailed outline of the proposed tutorial.

Duration. 3.5 hours

Tutorial description. This tutorial covers past and current research in kidney exchange, a method by which patients in need of a kidney can swap willing but incompatible donors. Throughout, it also gives a higher-level overview of the steps taken to translate a purely academic idea into a large fielded healthcare system. The tutorial focuses on the computational aspects of kidney exchange, starting by introducing past research that has now been implemented in real-world exchange (or deemed impractical and not yet fielded) and then by covering the current research problems available at the intersection of AI, optimization, and economics. It especially dives into the computational methods developed and used to solve extremely large discrete optimization problems that reflect kidney exchange, along with the interplay between modeling decisions, computational tractability, exchange efficiency, equity, dynamism in matching, and a variety of other real-world constraints and considerations. Research toward exchanges for other organs such as livers and lungs, as well as cross-organ exchanges, will also be covered.

Prerequisite knowledge. This tutorial will be accessible to the general attendee of AAAI and should not require specialized knowledge in a specific subfield of AI. One of our main goals is to showcase a recent fielded AI "success story," thus motivating practitioners in the *general* AAAI community to take their ideas outside of academia and into practice.

2 Outline of proposed tutorial

The following is an outline for a 3.5 hour tutorial on kidney exchange. It assumes an intermission at the 1.75 hour mark. When appropriate, we include a reference or two to a sample paper related to a specific point in the outline.

2.1 Introduction & preliminaries

- Kidney disease and the state of organ transplantation in the US and abroad; medical and legal constraints in organ transplantation
- Genesis of the kidney exchange idea from a medical point of view [36]
- Economics & matching research on large-scale kidney exchange [38–42]
- Modern generalizations such as altruist-initiated chains—as opposed to traditional kidney exchange where only cycles were used. We will cover both closed [34] and open [37] chains. Nowadays most kidney exchange transplants are conducted using open chains (i.e., never-ending altruist-initiated chains), which Prof. Sandholm helped invent [37].

2.2 Optimization models for the basic kidney exchange problem

- Two basic formulations of the kidney exchange clearing problem as an integer programming search problem [1, 41]
  - Tightness of different formulations
  - Complexity [1, 33] and runtime tradeoffs in more expressive models (e.g., traditional matching versus more complex matching with cycles and chains, different utility functions)
- A custom branch-and-price-based integer programming solver built to clear large kidney exchange pools (of the type that existed in the nascency of kidney exchange) [1]
- Two recent integer program-based formulations built to handle long chains in kidney exchange [4, 27], an adaptation to the original branch-and-price-based solver [35], and runtime comparisons between the three approaches

2.3 State of, and challenges in, practical implementation

At this point in the tutorial, the audience will know about the high-level kidney exchange problem, its initial formulation as a theoretically and empirically hard optimization problem, and will have in-depth knowledge of the original techniques and AI/OR methods used to overcome that computational hardness. The presenters have worked extensively with very large kidney exchanges to implement these clearing algorithms in practice. We will cover some of the early challenges to translating AI-style research into practice, such as the following.

- Convincing practitioners that "hardcore" optimization is necessary, when compared to more understandable (to a layperson) heuristic methods
- Feedback loop between healthcare professionals and CS/Econ researchers
- Legality and ethics, especially with respect to moving beyond simple kidney swaps to more elaborate constructs like non-directed donor-initiated kidney chains
- Other challenges

We will discuss how these challenges were overcome, with learnings that will hopefully help AI researchers get AI systems fielded in the large in other applications as well.

— Intermission —

Following the intermission, we will deep dive into different dimensions of fielded kidney exchange that are most pertinent to today's and the future's fielded exchanges, currently unsolved, and would benefit immensely from AI research.

2.4 Post-match failure

In all fielded kidney exchanges, many algorithmic matches fail to move to transplant [7,19]. This is not explicitly taken into account in any fielded optimization algorithms, but is an active area of research.

- Reasons for post-match failure in kidney exchange, and learning failure rates from data [26]
- Approaches for incorporating failure rates into the optimization problem [4,19,30,33]
- Tractability versus expressiveness in stochastic optimization formulations

2.5 Fairness

Prioritization of otherwise-marginalized candidates is a contentious topic and ongoing discussion in healthcare and, specifically, kidney exchange. We will cover the following.

- How should we define fairness in general [41]
- Two different ways to define fairness in an optimization model, and theoretical bounds on the "price of fairness" relative to an efficient matching in these models [20]
- Open questions in balancing equity and efficiency in healthcare applications

2.6 Dynamic exchange

Healthcare applications are often dynamic, with patients and other agents arriving and departing over time, and the interactions between the various agents changing over time as well. In kidney exchange, patients and donors arrive over time, and the exchange can choose when, how, and if to match these actors. We will cover general dynamic optimization techniques developed and then applied to kidney exchange [11,17], along with theoretical approaches to dynamic kidney exchange and matching [2,3,46]. This is an open problem; we will also cover promising research directions in dynamic exchange, with an emphasis on AI approaches to the problem.

2.7 Game-theoretic analysis

Here, we view the kidney exchange problem as a multi-agent system with transplant centers as agents, and their set of patient-donor pairs (and possibly altrustic donors) as a private type. The issue is that the transplant centers can match easy-to-match pairs locally and only report hard pairs to the exchange because transplant centers gain financially from conducting surgeries locally and local transplants are logistically easier. Such hiding of the pairs from the exchange compromises the global efficiency of the system. Prof. Sandholm in collaboration with UNOS has empirically shown that the problem is rampant: today transplant centers not only hide their easy-to-match pairs but all their internally matchable pairs [44]. We will discuss mechanism design approaches to this problem, such as the following.

- Mechanism design impossibility results in a dense model [10]
- Better mechanisms for a reduced model of kidney exchange [6,45]
- Ongoing ideas and techniques to circumvent the impossibility results [28]

2.8 Research toward exchanges of other organs and cross-organ exchanges

We will briefly overview nascent theoretical and practical pushes to field new types of organ exchanges—liver exchange [21,25], lung exchange [24,32], multi-organ exchange [21]—with a focus on the modeling differences between those and kidney exchange.

2.9 Combining human value judgments and optimization to learn to match

We present a holistic method to combine all of the dimensions mentioned above, along with high-level human-provided value judgments, into a unified framework for learning to match in a general dynamic model [22]. The framework takes as input a high-level objective (e.g., “maximize graft survival of transplants over time") decided on by experts, then automatically (i) learns based on data how to make this objective concrete and (ii) learns the "means" to accomplish this goal—a task, in our experience, that humans handle poorly. It uses data from all live kidney transplants in the US since 1987 to learn the quality of each possible match; it then learns the *potentials* of elements of the current input graph offline (e.g., potentials of pairs based on features such as donor and patient blood types), translates these to weights, and performs a computationally feasible batch matching that incorporates dynamic, failure-aware considerations through the weights.

This method, shown in Figure 1, provides sensitivity analysis for the UNOS exchange. We will discuss challenges in implementation and the disconnect between general optimization methods and a fielded application, and how the given framework attempts to unify the two different viewpoints.

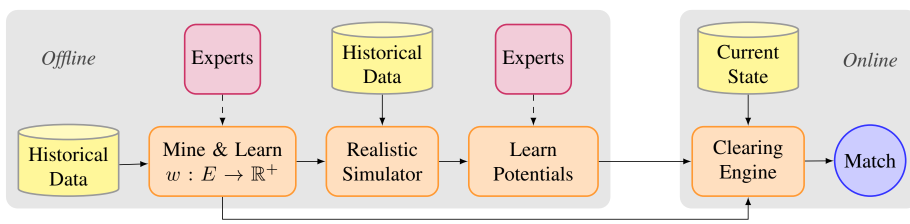

Figure 1: The FUTUREMATCH framework.

2.10 Conclusion & recap of research directions

We will provide insight into where we think fielded kidney exchanges are—and where they should be—going, and touch on the different directions in this area where AI researchers could make an impact.

- Failure-aware, fairness-aware, dynamic kidney exchange models, and the optimization techniques to solve real-world instances of the kidney exchange problem, as discussed above
- Mechanism design aspects of kidney exchange and their effect on fielded exchange performance
- Aspects of the kidney exchange model that have *not been touched* by the optimization literature—mainly, money and the difficulties of international exchange
- Difficulties in fielding an AI application in healthcare, and ways this could be streamlined in the future

3 About the presenters

John P. Dickerson

Mailing address. 9219 Gates-Hillman Center
Carnegie Mellon University
Pittsburgh, PA 15213

Phone. +1 240 715 2514

Email. dickerson@cs.cmu.edu

John is a final-year Ph.D. candidate in the Computer Science Department at Carnegie Mellon University. His Ph.D. thesis research focuses on kidney exchange. He has published extensively on kidney exchange in top-tier AI conferences [12, 16–20, 22]. His work has also set policy at the United Network for Organ Sharing (UNOS) nationwide kidney exchange, an exchange with which he works closely; it now includes 143 transplant centers in the United States. With Tuomas Sandholm, he created FutureMatch, a general framework for learning to match subject to human value judgments. FutureMatch won a 2014 HPCWire Supercomputing Award and now provides sensitivity analysis for matching policies at UNOS. He is the winner of a 2012–2015 NDSEG Fellowship and a 2015–2017 Facebook Fellowship, and is a 2015–2016 Siebel Scholar.

John has also lectured extensively on kidney exchange. For a "tutorial-style" presentation, please see a 90-minute talk on kidney exchange created and given by John in a Ph.D.-level course at CMU: http://www.cs.cmu.edu/~arielpro/15896s15/docs/896s15-15.pdf.

Tuomas Sandholm

Mailing address. 9205 Gates-Hillman Center
Carnegie Mellon University
Pittsburgh, PA 15213

Phone. +1 412 268 8216

Email. sandholm@cs.cmu.edu

Tuomas Sandholm is Professor at Carnegie Mellon University in the Computer Science Department, with affiliate professor appointments in the Machine Learning Department, Ph.D. Program in Algorithms, Combinatorics, and Optimization (ACO), and CMU/UPitt Joint Ph.D. Program in Computational Biology. He is the Founder and Director of the Electronic Marketplaces Laboratory. He has published over 450 papers. He has 25 years of experience building optimization-powered electronic marketplaces, and has fielded several of his systems. In parallel with his academic career, he was Founder, Chairman, and CTO/Chief Scientist of CombineNet, Inc. from 1997 until its acquisition in 2010. During this period the company commercialized over 800 of the world's largest-scale generalized combinatorial multi-attribute auctions, with over \$60 billion in total spend and over \$6 billion in generated savings. Dr. Sandholm's algorithms also run the UNOS kidney exchange, which includes 60% of the transplant centers in the US. He is Founder and CEO of Optimized Markets, Inc., a startup that is bringing a new paradigm to advertising campaign sales and scheduling - in TV, radio, Internet display, streaming, mobile, game, and cross-media advertising. He also served as the redesign consultant of Baidus sponsored search auctions and display advertising markets; within two years Baidus market cap increased 5x to \$50 billion due to better monetization. He has served as consultant, advisor, or board member for Yahoo!, Google, Chicago Board Options Exchange, swap.com, Granata Decision Systems, and others. He has developed the leading algorithms for several general classes of game. They won the 2014 world championships in computer Heads-Up No-Limit Texas Hold'em. He holds a Ph.D. and M.S. in computer science and a Dipl. Eng. (M.S. with B.S. included) with distinction in Industrial Engineering and Management Science. Among his many honors are the NSF Career Award, inaugural ACM Autonomous Agents Research Award, Sloan Fellowship, Carnegie Science Center Award for Excellence, Edelman Laureateship, and Computers and Thought Award. He is Fellow of the ACM, AAAI, and INFORMS.

Tuomas published the first computer science paper on kidney exchange [1], and is widely seen as the leading expert on the computational aspects of kidney exchange. He has widely lectured on his research in fielding kidney exchanges in academia and industry.

Tuomas has extensive teaching experience, and typically teaches graduate AI and undergraduate AI at CMU, as well as Foundations of Electronic Marketplaces, a course he designed. He has also given 35 tutorials on market clearing and expressive commerce at academic conferences and in industry. These include tutorials at AAAI-{98,99,00,02,04,05}, AAMAS/AGENTS-{97,98,99,00,01,02,04}, EC-{00,04,05}, and IJCAI-{97,99,01,03}.

References

[1] David Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 295–304, 2007.

[2] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Dynamic matching market design. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, page 355, 2014.

[3] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. A dynamic model of barter exchange. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.

[4] Ross Anderson, Itai Ashlagi, David Gamarnik, and Alvin E Roth. Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences*, 112(3):663–668, 2015.

[5] Elliot Anshelevich, Meenal Chhabra, Sanmay Das, and Matthew Gerrior. On the social welfare of mechanisms for repeated batch matching. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 60–66, 2013.

[6] Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.

[7] Itai Ashlagi, Duncan S. Gilchrist, Alvin E. Roth, and Michael Rees. Nonsimultaneous chains and dominos in kidney-paired donation—revisited. *American Journal of Transplantation*, 11(5):984–994, 2011.

[8] Itai Ashlagi, Patrick Jaillet, and Vahideh H. Manshadi. Kidney exchange in dynamic sparse heterogenous pools. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 25–26, 2013.

[9] Itai Ashlagi and Alvin E. Roth. New challenges in multihospital kidney exchange. *American Economic Review*, 102(3):354–59, 2012.

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

[11] Pranjal Awasthi and Tuomas Sandholm. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 405–411, 2009.

[12] Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2015.

[13] Avrim Blum, Anupam Gupta, Ariel D. Procaccia, and Ankit Sharma. Harnessing the power of two crossmatches. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 123–140, 2013.

[14] Ioannis Caragiannis, Aris Filos-Ratsikas, and Ariel D. Procaccia. An improved 2-agent kidney exchange mechanism. *Theoretical Computer Science*, 589:53–60, 2015.

[15] Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research*, 231(1):57–68, 2013.

[16] Sanmay Das, John P. Dickerson, Zhuoshu Li, and Tuomas Sandholm. Competing dynamic matching markets. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, 2015.

[17] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1340–1346, 2012.

[18] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 711–718, 2012.

[19] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Failure-aware kidney exchange. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 323–340, 2013.

[20] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Price of fairness in kidney exchange. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1013–1020, 2014.

[21] John P. Dickerson and Tuomas Sandholm. Multi-organ exchange: The whole is greater than the sum of its parts. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1412–1418, 2014.

[22] John P. Dickerson and Tuomas Sandholm. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[23] Yichuan Ding, Dongdong Ge, Simai He, and Christopher Ryan. A non-asymptotic approach to analyzing kidney exchange graphs. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2015.

[24] Haluk Ergin, Tayfun Sönmez, and M Utku Ünver. Lung exchange, 2014. Working paper.

[25] Haluk Ergin, Tayfun Sönmez, and M Utku Ünver. Liver exchange, 2015. Working paper.

[26] Kristiaan M. Glorie. Estimating the probability of positive crossmatch after negative virtual crossmatch. Technical report, Erasmus School of Economics, 2012.

[27] Kristiaan M. Glorie, J. Joris van de Klundert, and Albert P. M. Wagelmans. Kidney exchange with long chains: An efficient pricing algorithm for clearing barter exchanges with branch-and-price. *Manufacturing & Service Operations Management (MSOM)*, 16(4):498–512, 2014.

[28] Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[29] Jian Li, Yicheng Liu, Lingxiao Huang, and Pingzhong Tang. Egalitarian pairwise kidney exchange: Fast algorithms via linear programming and parametric flow. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 445–452, 2014.

[30] Yijiang Li, Jack Kalbfleisch, Peter Xuekun Song, Yan Zhou, Alan Leichtman, and Michael Rees. Optimization and simulation of an evolving kidney paired donation (KPD) program. Department of biostatistics working paper 90, University of Michigan, May 2011.

[31] Yicheng Liu, Pingzhong Tang, and Wenyi Fang. Internally stable matchings and exchanges. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1433–1439, 2014.

[32] Suiqian Luo and Pingzhong Tang. Mechanism design and implementation for lung exchange. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

[33] David Manlove and Gregg O'Malley. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM Journal of Experimental Algorithmics*, 19(1), 2014.

[34] Robert Montgomery, Sommer Gentry, William H Marks, Daniel S Warren, Janet Hiller, Julie Houp, Andrea A Zachary, J Keith Melancon, Warren R Maley, Hamid Rabb, Christopher Simpkins, and Dorry L Segev. Domino paired kidney donation: a strategy to make best use of live non-directed donation. *The Lancet*, 368(9533):419–421, 2006.

[35] Benjamin Plaut, John P. Dickerson, and Tuomas Sandholm. A fast and optimal clearing engine for barter exchanges with capped chains, 2015. Working paper.

[36] F. T. Rapaport. The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proceedings*, 18:5–9, 1986.

[37] Michael Rees, Jonathan Kopke, Ronald Pelletier, Dorry Segev, Matthew Rutter, Alfredo Fabrega, Jeffrey Rogers, Oleh Pankewycz, Janet Hiller, Alvin Roth, Tuomas Sandholm, Utku Ünver, and Robert Montgomery. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine*, 360(11):1096–1101, 2009.

[38] Alvin Roth, Tayfun Sönmez, and Utku Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.

[39] Alvin Roth, Tayfun Sönmez, and Utku Ünver. A kidney exchange clearinghouse in New England. *American Economic Review*, 95(2):376–380, 2005.

[40] Alvin Roth, Tayfun Sönmez, and Utku Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151–188, 2005.

[41] Alvin Roth, Tayfun Sönmez, and Utku Ünver. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review*, 97:828–851, 2007.

[42] Alvin Roth, Tayfun Sönmez, Utku Ünver, Frank Delmonico, and Susan L. Saidman. Utilizing list exchange and nondirected donation through 'chain' paired kidney donations. *American Journal of Transplantation*, 6:2694–2705, 2006.

[43] Tayfun Sönmez and Utku Ünver. Altruistically unbalanced kidney exchange. *Journal of Economic Theory*, 152(1):105–129, 2014.

[44] Darren Stewart, Ruthanne Leishman, Elizabeth Sleeman, Catherine Monstello, Guy Lunsford, Jude Maghirang, Tuomas Sandholm, Sommer Gentry, Richard Formica, John Friedewald, and Kenneth Andreoni. Tuning the OPTN's KPD optimization algorithm to incentivize centers to enter their easy-to-match pairs. In *American Transplant Congress (ATC)*, 2013. Talk abstract.

[45] Panos Toulis and David C. Parkes. Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior*, 91(0):360–382, 2015.

[46] Utku Ünver. Dynamic kidney exchange. *Review of Economic Studies*, 77(1):372–414, 2010.

[47] Özgür Yılmaz. Kidney exchange: An egalitarian mechanism. *Journal of Economic Theory*, 146(2):592–618, 2011.

[48] S. A. Zenios. Optimal control of a paired-kidney exchange program. *Management Science*, 48(3):328–342, 2002.