Personal Statement
Itai Ashlagi

My principal research is in market design, which studies how the rules and operations governing the marketplace affect economic outcomes. My research is both theoretical and applied and I use tools from operations, economics and computer science to study and (re)design marketplaces. One concrete example is kidney exchange, where my algorithms and recommendations have been adopted by numerous transplant centers and kidney exchange networks around the world.

Kidney exchange. Kidney exchange organizes swaps between patient-donor incompatible pairs and now accounts for about 12% of live donor transplants. The growth of kidney exchange has involved many operational and design challenges.

<u>How and when to match.</u> A kidney exchange platform must decide how and when to match pairs. While cyclic exchanges are limited to few pairs, chains, initiated by altruistic donors, can be arbitrarily long. We demonstrate the importance of chains in environments with many highly sensitized patients (Ashlagi et al., 2011b,a). We also developed a stylized dynamic stochastic matching model with a sparse environment and show that chains significantly reduce waiting times (Anderson et al., 2017) unless there are sufficiently many "easy-to-match" agents (Ashlagi et al., 2016).

I also study when to match, motivated by concerns that US platforms shifted towards frequent matching. We find that among batching policies, matching "greedily" is asymptotically optimal (Anderson et al., 2017), demonstrated also empirically (Ashlagi et al.). Increasing participation rate, however, has a significant impact on efficiency, unlike increasing matching frequency.

<u>The kidney exchange market.</u> After a decade since the launch of kidney exchange platforms and kidney exchange has become standard practice, we studied in Agarawal et al. (2017) the efficiency of the US market. We propose a simple price-theory model and identify two market failures that cause the loss of hundreds of transplants per year: hospitals do not internalize the benefits of their patients, and platforms do not provide hospitals adequate incentives. Indeed, more that half of the exchanges are facilitated within hospitals rather than through national platforms. Within-hospital exchanges often use easy-to-match pairs as predicted in Ashlagi and Roth (2014). We estimate the production function of a platform and find that individual hospitals operate below the efficient scale. We propose simple point-based mechanisms to encourage hospital participation based on their marginal contribution. While been controversial in the transplant community, recently such a system was implemented at the National Kidney Registry (NKR). We are now implementing such as system at the Alliance for Paired Donation (APD).

In Europe exchanges are typically arranged in national platforms, but recently I was invited to discuss local versus national exchanges at the Manchester International Living Donor meeting...

<u>Further impact.</u> As part of an interdisciplinary team associated with the APD, a multi-hospital program (with more than 60 hospitals), I was a Franz Edelman finalist. I developed matching algorithms (Anderson et al., 2015) and created software (pro-bono) with advanced histocompatibility capabilities that is now widely adopted. Just a couple of examples include Israel exchange program and Methodist Specialty and Transplant Hospital at San Antonio, which is the largest single hospital program in the US. I received the Terasaki Medical Innovation award from the NKR (largest US platform) for "significant impact in advancing paired exchange transplantation and saving the lives of those facing kidney failure".[1] This work demonstrates how market design, which draws on

[1] Also, the overview paper by Ashlagi and Roth (2012) was cited by the Nobel committee in the scientific background for the 2012 Economics Prize awarded to Alvin Roth.

both operations and game theory can lead to better outcomes.

Organ allocation We are now working on empirically evaluating allocation rules for cadaver waiting lists, in which patients may decline organs hoping to receive a better future offer (Agarwal et al., 2018). In the recent kidney allocation reform, we observed that when selecting between allocation rules, the acceptance-rejection simulation tool was invariant to the priority rule and hence ignored time preferences. We develop a methodology for estimating patients' preferences based on historical data and use it to evaluate the "price" of priorities on match quality and waste.

Two-sided matching markets. Stable matching theory has been instrumental in the study and design of two-sided markets. The literature traditionally studies whether good properties can be guaranteed for any preference realization. This line of research led to negative results and further provides almost no link between market characteristics and stable outcomes. By adopting a stochastic approach, I have addressed several fundamental issues in matching markets, establishing that desired properties hold with high probability.

In Ashlagi et al. (2017c), by considering random preferences, we find that matching markets generally have a small core, and realizations with multiple stable matchings are only a knife-edge case. This solves a longstanding puzzle explaining why empirically small cores are ubiquitous. Practically this means that in centralized marketplaces like the NRMP it is irrelevant what stable matching algorithm is implemented, and, it is also safe to recommend all agents to act sincerely. In this work we also find matching markets to be extremely competitive: agents on the short side match with very top choices and agents on the long side match almost with a random partner. [2]

Another issue is congestion, as forming matches often requires extensive and costly communication between participants. Segal (2007) finds that reaching a stable matching may require a high level of communication under arbitrary preferences. In Ashlagi et al. (2017b) we find that when the unobservable component of agents' preferences satisfies natural assumptions, one can recommend matches and use signals to reach a stable outcome with a low amount of communication.

Finally, stable matchings may not exist when agents have complementarities (as married doctors seek for pairs of jobs at the NRMP). In Ashlagi et al. (2014), we show constructively that a stable matching exists with high probability as long as the fraction of couples grows at a sub-linear rate with the size of the market.[3]

Service allocation: school choice and beyond. Many cities use school choice systems to assign students based on their preferences and priorities at different schools, often using a Gale-Shapley like algorithm. I have written on policy decisions that have a major impact on students' assignments.

Optimal assignment. A major concern in Boston was unsustainable busing costs, representing 10% of the total school board budget. Prior to 2014, the city was divided into three "zones" and families could rank schools within their zone. Together with my student Peng Shi we developed theory for allocating *optimally* heterogeneous services to agents under incomplete information (Ashlagi and Shi, 2015). [4] We solve a large market approximation and translate the solution into a finite market mechanism. We find that incentive compatible and Pareto efficient mechanisms are simple: planners can optimize only over priorities and *menus* of services, which agents can rank.[5] Applying

[2]See a blog post in Combinatorics and More (by Gil Kalai).

[3]Elliott Peranson, who developed the matching algorithm for the NRMP jointly with Alvin Roth, writes me: "based on your work we are modifying the procedures we use to attempt to find a stable match.."

[4]Little is known about optimal allocation without money when agents have multi-dimensional preferences.

[5]This characterization is applied in Ashlagi and Shi (2014) to study when it is possible to improve community cohesion without sacrificing choice; the key tool is to use correlated lotteries that maintain incentive constraints.

to Boston data, the solution outperforms the existing design with respect to efficiency, equity, and predictability. This paper bridges assortment planning with markets without money and was awarded the MSOM best publication in service management.

In 2014, Boston reformed its school choice, based on a simplification of this work, to assign families from each geocode a personalized menu of schools to rank from. This impacts around 4000 families annually.

Designing tie-breaking rules in school choice. Often lotteries are used to resolve ties between students in schools with excess demand. Motivated by debates in NYC, Chile and Amsterdam, my student Afshin Nikzad and I compare two common tie-breaking rules (Ashlagi and Nikzad, 2016): either all schools use a common lottery, or each school uses a separate lottery. Separate lotteries seem fairer but may create artificial inefficiencies. We find that in *popular schools* a single lottery stochastically dominates multiple lotteries. This removes the confusion in early discussions and suggests that popular schools should use a common lottery.[6]

Dynamic matching. Motivated by pooling problems in ride-sharing, I am interested in dynamic matching when matches yield heterogenous values. Matching greedily is sub-optimal, and we develop algorithms that perform well under worst-case assumptions (Ashlagi et al., 2017a, 2018).

Mechanism design. I have written theoretical papers on mechanism design and highlight two of them. Towards designing an optimal auction, Myerson characterizes incentive compatibility via monotonicity. This does not extend to multi-dimensional settings. In Ashlagi et al. (2010) we characterize the domains of multi-dimensional preferences, in which monotonicity is necessary and sufficient for implementing an allocation rule.

In Ashlagi et al. (2012), I address an open problem posed in the seminal paper by Nisan and Ronen (1999), to assign tasks to workers who privately know how long it takes them to complete each task. I find that anonymous incentive compatible mechanisms for minimizing the makespanmust result in far-from-optimal outcome. (This won the Outstanding Paper award in EC.)

References

N. Agarawal, I. Ashlagi, E. Azavedo, C. Featherstone, and O. Karaduman. Market failure in kidney exchange. Working paper, 2017.

N. Agarwal, I. Ashlagi, P. Somaini, and D. Waldinger. Dynamic incentives in waitlist mechanisms. *American Economic Review, Papers and Proceedings*, 108, 2018. forthcoming.

R. Anderson, I. Ashlagi, D. Gamarnik, and A.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.

R. Anderson, I. Ashlagi, D. Gamarnik, and Y. Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.

I. Ashlagi and A. Nikzad. What matters in tie-breaking rules? how competition guides design. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 767–768. ACM, 2016.

[6] I was invited to talk about this with policymakers prior to the recent school choice reform in Chile.

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

I. Ashlagi and A.E. Roth. New Challenges in Multi-hospital Kidney Exchange. *American Economic Review: Papers and Proceedings*, 102(3):354–359, 2012.

I. Ashlagi and P. Shi. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research*, 62(6):1247–1264, 2014.

I. Ashlagi and P. Shi. Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097, 2015.

I. Ashlagi, A. Bingaman, M. Burq, V. Manshadi, D. Gamarnik, C. Murphey, A.E. Roth, M.L. Melcher, and M.A. Rees. The effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. *American Journal of Transplantation*. fothcoming.

I. Ashlagi, M. Braverman, A. Hassidim, and D. Monderer. Monotonicity and implementability. *Econometrica*, 78(5):1749–1772, 2010.

I. Ashlagi, D. S. Gilchrist, A. E. Roth, and M. A. Rees. Nonsimultaneous chains and dominos in kidney paired donation - revisited. *American Journal of Transplantation*, 11:984–994, 2011a.

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

I. Ashlagi, S. Dobzinski, and R. Lavi. Optimal lower bounds for anonymous scheduling mechanisms. *Mathematics of Operations Research*, 37(2):244–258, 2012.

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

I. Ashlagi, M. Burq, P. Jaillet, and V. Manshadi. On matching and thickness in heterogeneous dynamic markets. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 765–767. ACM, 2016.

I. Ashlagi, Y. Azar, M. Charikar, A. Chiplunkar, O. Geri, H. Kaplan, R. Makhijani, Y. Wang, and R. Wattenhofer. Min-cost bipartite perfect matching with delays. 81, 2017a.

I. Ashlagi, M. Braverman, Y. Kanoria, and P. Shi. Communication requirements and informative signaling in matching markets. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 263–263. ACM, 2017b.

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

I. Ashlagi, M. Burq, P. Jaillet, and A. Saberi. Maximizing efficiency in dynamic matching markets. working paper, 2018.