

Incorporating Compatible Pairs in Kidney Exchange: A Dynamic Weighted Matching Model

ZHUOSHU LI*, Washington University in St. Louis

KELSEY LIEBERMAN^{†‡}, Truman State University

WILLIAM MACKE^{†‡}, University of Tulsa

SOFIA CARRILLO[†], Cornell University

CHIEN-JU HO, Washington University in St. Louis

JASON WELLEN, Washington University School of Medicine

SANMAY DAS[§], Washington University in St. Louis

Kidney exchange has been studied extensively from the perspective of market design, and a significant focus has been on better algorithms for finding chains and cycles to increase the number of possible matches. A more dramatic benefit could come from incorporating compatible pairs into the mechanism, but this possibility has been relatively understudied. In order to incentivize a compatible pair to participate in exchange, they must be offered a higher quality match for the recipient that can be performed without adding extra waiting time. In this paper, we make two main contributions to the study of incorporating compatible pairs in exchanges. First, we leverage the recently proposed Living Donor Kidney Profile Index (LKDPI) to measure match quality, and develop a novel simulator (based on data from a major transplant center) for the joint distribution of compatibility and quality across pairs. This simulator allows us to study the benefits of including compatible pairs under different models and assumptions. Second, we introduce a hybrid online/batch matching model with impatient (compatible) and patient (incompatible) pairs to capture the need for immediacy. We introduce new algorithms for matching in this model, including one based on online primal-dual techniques. Overall, our results indicate great potential in terms of both increased numbers of transplants of incompatible pairs (almost doubling the number transplanted) as well as improved match quality for recipients in compatible pairs (increasing expected graft survival by between 1 and 2 years). The results are also promising for hard-to-match subpopulations, including blood group O recipients.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design**; *Online algorithms*; • **Applied computing** → *Life and medical sciences*.

Additional Key Words and Phrases: Kidney exchange; matching; cardinal utility; online primal-dual methods

ACM Reference Format:

Zhuoshu Li, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das. 2019. Incorporating Compatible Pairs in Kidney Exchange: A Dynamic Weighted Matching Model. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19)*, June 24–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3328526.3329619>

*Now at Google.

[†]Lieberman, Macke, and Carrillo worked on this project as REU students at Washington University.

[‡]Equal contribution.

[§]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329619>

1 INTRODUCTION

In the end stages of renal failure, when a patient needs a transplant, one excellent option, when available, is to receive a kidney from a living donor who is willing to donate to them. Close to a third of kidney transplants annually in the US are living donor transplants.¹ As of today, most living donor kidney transplants match recipients directly with a donor willing to donate a kidney to that specific recipient. However, in some cases the donor and the recipient may not be medically compatible (due to ABO blood-type incompatibility or a positive crossmatch), in which case they can enter a *kidney exchange*, a type of barter market where incompatible donors donate to others with the understanding that their recipient will receive a medically compatible kidney from someone else [1, 41, 43, 45].

While paired kidney donation of this kind has had success in the United States, a raft of coordination problems and exchange fragmentation has prevented it from accounting for a truly significant fraction of transplants. One proposal to transplant more recipients from incompatible pairs has been to incorporate *compatible* pairs into exchanges [27]. While prior work has estimated the possible benefits of such donation in terms of the number of additional patients that could be transplanted from incompatible pairs, the recent development of new metrics for the quality of a living donor transplant [36] present an opportunity to reassess the possible benefits in the context of realistic models of compatible pair behavior, while also evaluating benefits in terms of both additional transplants made possible and improved outcomes from transplants. Further, it is reasonable to believe that compatible pairs may be more willing to enter exchanges if (1) their waiting times are kept low, and (2) they have a more precise idea of the potential benefit to doing so.

In this paper, we work towards a realistic model of kidney exchange with compatible pairs that incorporates these constraints. Our two main contributions are in explicitly incorporating match quality into the model, and in introducing a model of weighted hybrid online/batch matching with impatient and patient agents and analyzing algorithms for this model. We build and calibrate our models with data from a major transplant center, and by doing so we are able to better estimate the benefits of incorporating compatible pairs into kidney exchanges.

Modeling match quality. While kidney exchange has typically been modeled in a framework where the goal is to maximize the number of transplants, anecdotally it is clear that surgeons consider the quality of matching in their decision-making about which transplants to proceed with. Quantifying match quality in the context of kidney exchange could provide a sounder basis for decision-making, and, as argued above, provide a reason for compatible pairs to enter into exchanges. It would also contribute a real-world domain for research on matching with cardinal utilities [6].

Our approach here is to estimate expected survival of a graft from the recently proposed *Living Donor Kidney Profile Index (LKDPI)* [36], and use this as our measure of quality.² How beneficial could consideration of match quality be? The answer to this question depends largely on the heterogeneity of match qualities across pairs. At one extreme, LKDPIs across pairs could be completely independent of the original LKDPIs within the pairs. This would correspond to maximally heterogeneous match qualities and offer the highest possible benefits to recipients in compatible pairs of participating in the exchange. At the other extreme, LKDPIs could be completely determined by the characteristics

¹<https://www.kidney.org/news/newsroom/factsheets/Organ-Donation-and-Transplantation-Stats>

²LKDPI itself is a somewhat complex number to interpret. It is intended to be on the same scale as the KDPI for cadaveric kidneys, which is a percentile measure. Thus an LKDPI of 10 indicates that the kidney is comparable to the 10th percentile of cadaveric kidneys in terms of quality (with lower numbers being better). However, since some living donor kidneys can be better than *any* cadaveric kidney, LKDPI values can also be negative.

of the donor or the recipient in a pair, in which case there would be no social gains from trade [7]. We use detailed data on compatible pairs that have been transplanted at a major US transplant center over three years, and perform counterfactual simulation experiments confirming that the distribution of match quality (LKDPs) from “external” donors is far from independent of the match quality within a compatible pair, and this has significant implications on the possible gains from trade to the compatible pairs. We also provide evidence that the variability is largely driven by characteristics of the donor rather than the recipient, so there could be benefits from increasing the pool of possible donors, as has recently been suggested [26]. There are, however, some pairwise match characteristics that impact match quality (for example, HLA mismatches and body weight ratios), and lead to some potential “gains from trade.”

We use a novel dataset which contains data on 184 living donor kidney transplantations that took place at a major transplant center (henceforth “Center”) from 2014–2016 to construct a minimal simulator that produces realistic LKDPI / match quality values and closely matches the characteristics we observe in the data. Paired with the standard *Saidman simulator* [46] for generating recipient-donor pairs and compatibilities, this enables us to simulate realistic living donor kidney scenarios of any size with compatible and incompatible pairs.

A hybrid static-dynamic matching model. Compatible pairs may not be willing to wait even in order to find a potentially better match. In order to study a minimal departure from current practice, therefore, we introduce a hybrid matching model where the incompatible pool consists of *patient* agents who wait for a matching to be run on the pool (as is standard in the kidney exchange literature [2, 3, 11, 21–23, 48]), but compatible pairs are *impatient* and match greedily (with a pair from the incompatible pool if it improves the match for the compatible-pair recipient and directly from donor to recipient otherwise), as is assumed in truly dynamic markets like dating or rideshare [13, 29].³

In our model, there is a pool of agents that are waiting until a specified time T at which a batch matching or clearing (which can include cycles of length > 2) will be run, and a set of agents that enter dynamically between time 0 and time T . The pool of agents is embedded into a weighted compatibility graph, where edges exist between two agents if a match can be made between them, and weights w_{ij} reflect match quality. An agent entering dynamically is compatible with some subset of the pool; if it does not match through the pool, it receives a value w_i . While this directly models the kidney exchange setting with compatible pairs, the model itself is applicable to any matching scenario with patient and impatient agents.

The typical tradeoff in dynamic matching is between the known immediate value of a particular match and the unknown expected value of keeping agents around. In our case, that expected value now has two components, one based on the value of an agent to the set of dynamic arrivals, and one based on the value within the static set, which is known, but may itself change due to departures by other agents by the time the final static match takes place. We design a new algorithm for this type of hybrid online matching problem – ODASSE, for Online Dual Assignment Using Shadow Survival Estimates – motivated by the online dual assignment method [18, 32]. The algorithm is based on estimating values of agents in the static pool using both knowledge of the statistical properties of the population as well as information about the specific graph of static agents in any given problem.

Results and implications for incorporating compatible pairs in kidney exchange. The benefit from including directed donation (compatible) pairs in kidney exchanges can arise from two fronts: (1) an increase in the number of incompatible donors who find matches; (2) an increase in the quality

³Note that this differs from repeated batch matching [5], where agents arrive and a static matching takes place periodically.

of matches. We combine our compatibility/LKDPI simulator with the hybrid matching model in order to estimate the (minimal) potential benefits along both these fronts in a realistic manner.

We find that with compatible pairs joining the kidney exchange, the percentage of matched incompatible pairs increases substantially. Even in our pessimistic model where compatible pairs will not wait at all for a potential better match from the incompatible pool, more than 74% of incompatible pairs are matched, compared with 54% when the two- & three-cycle swaps only run within the incompatible pairs.

These results are similar to those of Gentry et al [27], who however focus only on compatible recipients gaining a donor age benefit. Our methods also allow us to estimate the benefits to recipients in compatible pairs, and we estimate an increase of 1.5 years in expected graft survival among recipients from compatible pairs.

We also look at the effects on the hard-to-match subpopulation of blood group O recipients. We estimate that the positive impacts on blood group O recipients are more substantial than in the general population (an increase from 32% to 62%).

By bringing quantitative estimates of these benefits into the light, we can inform policy debates. For example, how much expected benefit would be needed to convince compatible pairs to enter an exchange? The benefits could be even higher if compatible pairs were also patient – how long would they be willing to wait for a better match? These are all questions that can begin to be addressed from the foundation of the models and simulator we develop in this work.

2 KIDNEY EXCHANGE MODEL

In this section we describe the basics of kidney exchange. We largely follow Dickerson and Sandholm’s 2015 description. A kidney exchange can be represented as a directed compatibility graph $G = (V, E)$ [41–43]. Each vertex in the graph is a patient-donor pair in the pool. A directed edge e is constructed from vertex v_i to vertex v_j if the patient v_j is compatible with the donor kidney of v_i . Edges exist or do not exist due to medical characteristics (most importantly blood type, tissue antibodies and antigens) of the patient and the donor. There may also be other logistical constraints, but those are not relevant for our work here. In this pool, the donor of vertex v_i is willing to give her kidney if and only if the patient of v_i receives a kidney. A weight w_e can be assigned to an edge e . While this typically has been used in the literature to represent the priority of a transplantation (and therefore the utility to the *system* in some senses), we use it to represent the match quality when recipient v_j receives v_i ’s donor kidney (we discuss the function we use to determine match quality in the next section). In this graph, a sequence of transplants occurs when several vertices form a cycle c . A k -cycle refers to a cycle with exactly k pairs. In this paper, we only consider 2-cycles and 3-cycles, as is typical in fielded kidney exchange (incorporating cycles longer than 3 offers limited benefit given logistical constraints). Fielded exchanges also gain from chains, where an altruist donor without a paired patient enters the pool and start a directed path of transplants. We do not include chains in this work.

In this paper, a *matching* M is therefore a set of disjoint cycles in the compatibility graph G . The cycles must be disjoint because no donor can give more than one of her kidneys (some recent work explores multi-donor donation [25, 26] but we do not consider this here). Given a pre-defined utility function $u : \mathcal{M} \rightarrow \mathbb{R}$ and the set of all legal matchings \mathcal{M} , we are trying to find a matching which maximizes u , $M^* \in \arg \max_{M \in \mathcal{M}} u(M)$. Kidney exchanges typically find the maximum weighted cycle cover, formally, $u(M) = \sum_{c \in M} \sum_{e \in c} w_e$. In this paper, we consider two objectives, the number of matches (effectively $w_e = 1 \forall e$), and expected total graft survival (where w_e is defined as the expected graft survival for the recipient in edge e).

An integer programming (IP) solver is usually used to find the optimal solution [1, 9, 15, 21]. We use a formulation based on the position-indexed chain-edge formulation (PICEF) [21] method

to find the optimal solution when doing two- & three-cycle swap. Further details of our specific formulation are in Section 6.

3 MODELING MATCH QUALITY

Historically, much work on matching (and welfare economics broadly) has focused on ordinal preferences rather than cardinal utility [6]. This sidesteps the problem of having to make interpersonal comparisons of utility, and research has focused on outcomes in terms of objectives like stability and Pareto optimality [38]. However, with the increasingly important social roles played by matching mechanisms [19, 20, 39, 44], it is imperative to understand the outcomes of mechanisms in terms of overall social welfare (however this is defined for a given application) as well as distributional effects. There has been increasing interest in recent years in both the study of dynamic weight matching algorithms [8, 24] and consideration of specific models of utility [12, 28, 34, 37].

There is value in traditional parametric models that are used for utility, and these have been central to model development. Examples of such models include utilities that decay exponentially in waiting time [2, 4], and random utility models for specific match pairs [17]. However, a common criticism of such models is that it is unclear how general or valuable results are when the utility model itself is not grounded in reality. In our case, we are explicitly looking for a realistic model that can be used for decision-making. Further, in order to convince compatible pairs to enter kidney exchanges, we must be able to quantify the expected benefit to them in some meaningful manner, therefore, we need an individual model of match quality that can be reasoned about from the perspectives of agents in the market. One important consideration that we defer to future work is the waiting cost to agents in terms of cost and quality of life. For compatible pairs, this is a complex modeling problem from a practical standpoint, because the baseline waiting time is itself highly variable. The time from initial workup to transplantation for a compatible pair is at least several months long because of the barrage of necessary testing, and for part of this time the pair is not even sure that they will be judged compatible. Therefore, in this paper we focus on match quality, and subject our analyses to the pessimistic assumptions that compatible pairs must match immediately.⁴

Quantifying match quality. Transplant surgeons often have to make decisions on whether a proposed transplant is worthwhile to proceed with. The Kidney Donor Profile Index (KDPI) was developed as a means of assessing the quality of a cadaveric (deceased donor) kidney [40]. Recently, the Living Kidney Donor Profile Index (LKDPI) has been proposed as an analog for living donations [36]. LKDPI takes into account characteristics of both the donor and the recipient.

KDPI is a percentile score. For example, a score of 4 implies that the kidney is in the “top 4%” of cadaveric kidneys. LKDPI is intentionally designed to be on the same scale (as mentioned in the Introduction, since living donor kidneys can

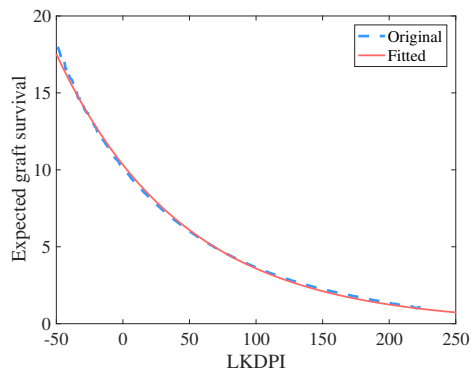


Fig. 1. An exponential curve fits the graft half-life as a function of LKDPI.

⁴The question of how to analyze waiting cost from the perspective of the matching market is also complex; however, one reasonable way to think about it is as costs to the healthcare *system*. For example, dialysis costs \$70,000-\$100,000 per year [31, 33, 35], and this is a cost that must be borne by some agent (individuals, private insurance, hospitals, or the government). Incorporating this can be useful when the modeling task is to assess matching policies and how they change costs over the entire system, rather than from the perspective of individual agents, hospitals, and so on.

be better than any cadaveric kidney, the LKDPI often takes negative values as well). Therefore, optimizing for LKDPI, while a useful proxy, is semantically ill-founded. However, since LKDPI is computed based on a survival model (Cox regression [16]), one can translate the model to a model of expected graft survival (or graft half-life), the survival time of the transplanted organ in the donor [30].⁵ We have found that an exponential curve fits the graft half-life as a function of LKDPI almost perfectly (see Figure 1), and can thus estimate expected graft survival as $14.78e^{-0.01239x}$ where x is the LKDPI. We can use this measure in place of LKDPI where it is more appropriate.⁶ Thus the edge weight w_e in each cycle is defined as the estimated expected graft survival of the recipient.

4 DRIVERS OF MATCH QUALITY: ANALYZING TRANSPLANT CENTER DATA

In this section we analyze actual data from transplants in terms of match quality, to gain an initial empirical understanding of the potential for “gains from trade” because of heterogeneity in pairwise match quality. In order to do so, we perform a counterfactual analysis using detailed data from only compatible pairs.

4.1 Data description

Massie *et al* 2016 come up with the LKDPI measure based on several important characteristics for determining graft survival. We gathered de-identified data on all donor and recipient characteristics that are used in computing LKDPI from all directed living-donor transplants performed at the center in a three year period (2014-2016). There were 166 such transplants with complete characteristics for calculating LKDPI and graft survival; 121 of them also include complete HLA antibody and antigen information.⁷ The distribution of each characteristic is shown in Table 1. We also analyze the correlation of every pair of characteristics, shown in Figure 2, which serves as a fundamental building block for designing the simulator in Section 5.

4.2 Counterfactual analysis within the center

Typically, if a donor and recipient are deemed medically compatible, a directed transplant is performed, with the donor’s kidney going to the recipient. However, there may be cases where the match quality is low even if they are compatible, and perhaps the recipient could receive a better kidney through an exchange; for example, they may be able to receive a kidney from a younger donor, or avoid an immunologically risky donor/recipient combination, like child to mother or husband to wife [27]. Such scenarios are hypothetical, and may seem unlikely at first glance. To validate our conjecture, for these donor-recipient pairs, we computed the expected graft survival (EGS) of each pair, and then performed counterfactual simulations to assess the potential to improve outcomes. The two counterfactual simulations share a Pareto improvement restriction—no recipient may receive a kidney with a worse quality (shorter EGS) for them than that of the kidney from their original paired donor.

Optimal. In the first simulation, we find the best matching, allowing arbitrary length cycles; this can also be treated as a bipartite matching problem (with the restriction that the matching must be perfect) between donors on one side and recipients on the other.

⁵After graft failure, the donor typically needs another transplant.

⁶One could also use expected graft survival as input to an expected “Quality Adjusted Life Year” (QALY) [14, 47, 49] computation over the lifetime of the recipient.

⁷We do not have data for the remaining 45 pairs simply because of a technical change in the software system; there is no selection bias.

	Mean	s.d.
Donor Age	48.22	12.68
Donor eGFR	98.11	15.08
Donor Systolic BP	124.14	13.11
Donor BMI	27.78	4.46
Recipient Weight (Female)	180.7	42.26
Recipient Weight (Male)	190.34	39.9
Donor Weight (Female)	160.75	30.06
Donor Weight (Male)	200.8	32.8
Donor Sex	F: 0.7	M: 0.3
Rec Sex	F: 0.35	M: 0.65
Donor African-American	Y: 0.05	N: 0.95
Donor Cigarette Use	Y: 0.32	N: 0.68
Donor/Rec Related	Y: 0.50	N: 0.50
Donor Blood Type	O: 0.6, A: 0.3, B: 0.07, AB: 0.03	
Rec Blood Type	O: 0.46 A: 0.39 B: 0.12 AB: 0.03	
Donor/Rec ABO compatible	Y: 0.88	N: 0.12
	Donor/Rec related	Donor/Rec unrelated
Donor/Rec HLA-B Mismatches	0: 0.18, 1: 0.32, 2: 0.5	0: 0.01, 1: 0.1, 2: 0.89
Donor/Rec HLA-DR Mismatches	0: 0.13, 1: 0.06, 2: 0.81	0: 0.01, 1: 0.06, 2: 0.93
	Counterfactual Matrix: all unrelated	
Donor/Rec HLA-B Mismatches	0: 0.009, 1: 0.091, 2: 0.9	
Donor/Rec HLA-DR Mismatches	0: 0.02, 1: 0.04, 2: 0.94	

Table 1. Distribution of each characteristic of the center’s data. F/M means Female/Male, Y/N represents Yes/No, and Rec is a shortening of Recipient.

Two and Three-cycle swap. In the second simulation, we only allow either a direct donation from the donor to the recipient or through a two- and three-cycle kidney exchange, to more closely approximate realistic logistical constraints.

We consider two subsets of the data. The “complete” 166-pair subset, assuming no HLA incompatibilities, and the “restricted” subset of 121 pairs for which we have complete antibody/antigen information and can determine all incompatibilities and rule out such transplants. The distribution of EGS and corresponding LKDPI among the real pairs and in the results of our counterfactual simulations are shown in Figure 3. The mean and median EGS and LKDPIs are given in Tables 2 and 3.

We can see there is a median improvement of 1.93 years of expected graft survival for the Optimal and 1.38 years for the two- & three-cycle swap (over a median half-life of 10.84 years). We also see that including compatibility constraints itself does not have a huge effect on the results (some of the improvement in the larger set is simply due to having a thicker market). Beyond the specific results, it is surprising to see the high number of transplants that were performed with LKDPIs above 50, since these indicate that the average *cadaveric* kidney would have been better for the recipient, in contrast to the conventional wisdom that living donor kidneys are always better. The optimized matches from the counterfactual “exchange” are much better, with many fewer “bad” matches and many more with LKDPI of 20 or lower, predictive of excellent outcomes.




Fig. 2. Correlation matrix of each pair of characteristics.

	Optimal	Two&Three -cycle Swap	Original		Optimal	Two&Three -cycle Swap	Original
Mean	11.44	10.88	9.80	Mean	11.58	11.14	9.67
Median	11.39	10.84	9.46	Median	11.67	11.18	9.34

(a) Restricted subset (121 pairs) with HLA anti-
gen/antibody compatibility constraints

(b) Complete subset (166 pairs) without HLA anti-
gen/antibody compatibility constraints

Table 2. Mean and median EGS for two counterfactual simulations, compared to reality over the last three years at the center.

	Optimal	Two&Three -cycle Swap	Original		Optimal	Two&Three -cycle Swap	Original
Mean	23.99	27.69	36.10	Mean	23.46	25.50	37.15
Median	21	25	36	Median	19	22.5	37

(a) Restricted subset (121 pairs) with HLA anti-
gen/antibody compatibility constraints

(b) Complete subset (166 pairs) without HLA anti-
gen/antibody compatibility constraints

Table 3. Mean and median LKDPI for two counterfactual simulations, compared to the reality over the last three years at the center.

4.3 Discussion

This is a proof-of-concept for the potential of improving quality of matching. One immediate question arises from the fact that we are using three years worth of data on recipients and donors in a static setting; this is obviously unrealistic. However, the main point is to estimate realistic distributions from data; we can use projections to then analyze differently-sized markets (from

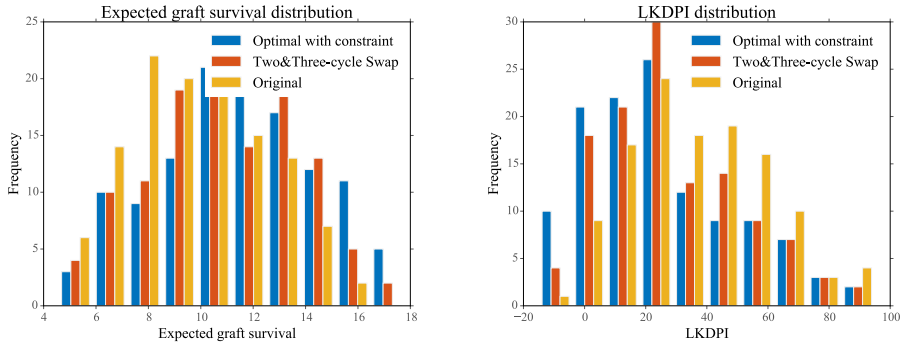


Fig. 3. Distribution of the expected graft survival (left) and LKDPI (right) of the original matched pairs and matched pairs in the two counterfactual simulations, using the restricted subset of real data with HLA antigen/antibody compatibility constraints from the center over the last three years.

smaller ones to larger ones that could be realized through regional pooling or already-functioning national exchanges). We turn to these questions and beyond in the next section.

5 DESIGNING A COMPATIBILITY+LKDPI SIMULATOR

A central goal of incorporating compatible pairs in exchanges would be to improve the *number* of matches, in addition to improving match quality. This could also lower costs for transplant centers by allowing for more internal matches where the transplant center does not need to go to a regional or national exchange to find a match for an incompatible pair.⁸ In order to estimate the possible benefits more systematically over different possible population sizes, we need to efficiently and correctly simulate LKDPIs over donor and recipient populations.

This would be simple if LKDPIs were distributed in a manner that was easy to correctly estimate, for example, independently, or independently conditional on the LKDPI of the original compatible pair. Unfortunately, this turns out not to be the case. To get a simple benchmark of how much this may affect the results, we can simulate different distributions based on data from the center.

We first build a counterfactual matrix of estimated graft survival based on the complete (166-pair) data by calculating LKDPI values for each of these 166 pairs. We then investigate the expected graft survival of donor-patient pairs under the Optimal and two- & three-cycle swapmatching algorithms when resampling the matrix in different ways. To simulate independent LKDPIs, we resample individual LKDPIs from the whole matrix. To simulate donor-dependent LKDPIs, we shuffle all donors for a given recipient, and to simulate recipient-dependent LKDPIs, we shuffle all recipients for a given donor. The results are shown in Table 4. The first row shows statistics from the original compatible matching. As we see, most of the methods for generating LKDPIs vastly overestimate the possible gains, and the evidence is consistent with the observation that the determination of LKDPI/expected graft survival is largely based on the donor’s characteristics [36]. These results demonstrate the need for a good simulator.

The central empirical facts that allow us to construct an efficient simulator are analyses of the joint distributions of variables involved in determining compatibility (PRA and ABO compatibility, based on the simulator of Saidman et al. (2006)) and computing LKDPI (See Table 1), and analysis of the possible underlying mechanisms of dependence. In particular, compatibility is solely a function

⁸This could have positive and negative effects overall, by perhaps increasing fragmentation, but lowering costs. However, many centers choose not to participate in broader exchanges much of the time in practice, for a variety of reasons.

of blood type and antibodies, while LKDPI considers many other factors, most of which have limited relationship to those (Figure 2). Since the state of practice in kidney transplantation has been to always assume that any living donor is excellent (a practice called into question by our results above), it is unlikely that there is any selection bias in the characteristics we sample for typical compatible pair arrivals. We first generate a donor-recipient pair, with all LKDPI-related characteristics generated sequentially in a manner that respects the data distributions in Table 1 and the correlation structure shown in Figure 2. We then generate the PRA (percentage reactive antibodies) and compatibility based on the Saidman model. Details of our simulator are available in an online appendix. The last line of Table 4 shows that the simulator produces results very close to the real data.

	EGS original	EGS 2&3 swap	EGS Optimal	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	9.67	11.14	11.58	37.15	25.50	22.46
Sample from the whole matrix	9.23	14.40	15.30	40.51	2.67	-2.5
Shuffle all donors per recipient	9.19	14.16	14.94	40.92	4.11	-0.47
Shuffle all recipients per donor	9.21	11.74	12.50	40.70	20.6	15.49
Sample from the simulator	9.38	11.40	11.80	39.21	24.50	20.09

Table 4. EGS and LKDPI comparison for different sampling methods and different market clearing algorithms.

6 INCLUDING COMPATIBLE PAIRS IN KIDNEY EXCHANGES

We now turn to exploring the benefits of including compatible pairs in kidney exchange. In particular, we introduce a hybrid online/batch matching model with patient agents and impatient agents. Patient agents represent incompatible pairs. They are waiting in the market and can be matched at any time. Impatient agents represent compatible pairs who arrive sequentially. Since compatible pairs can always opt for internal transplant, we consider the setting in which they are not willing to wait in the market and need to be matched immediately upon arrival. Moreover, a compatible pair is only willing to participate in the match if the recipient of the pair receives a kidney of higher quality.⁹

6.1 Problem formulation

There is a pool of *patient agents* (modeling incompatible pairs) who can be matched at any time, and a set of *impatient agents* (modeling compatible pairs) arrive sequentially and need to be matched

⁹In practice, incompatible pairs also may arrive and depart dynamically; however, since they cannot demand immediacy and the incompatible market clears periodically, this model is a reasonable approximation.

immediately upon arrival. We denote a matching problem as a k -matching problem if up to k agents can be involved in a match (which corresponds to a k -cycle in kidney exchanges). Moreover, impatient agents can perform “self-matches” and obtain utility without matching with other agents (this could of course be set to 0 for modeling other applications where patient agents must match immediately with some other agent). For ease of exposition, in this section, we focus on 2-matching problems (e.g., only directed (compatible) transplant and 2-cycles are allowed in kidney exchanges). We discuss the generalization to k -cycles in Section 6.3.

For any positive integer K , let $[K]$ denote the integer set $\{1, \dots, K\}$ and $[K]^*$ denote the set $\{0, \dots, K\}$. There are I patient agents, indexed by $i \in [I]$, and T impatient agents, indexed by $t \in [T]$. Let $N = T + I$; we use $n \in [N]$ to index all agents, with the first T agents being the impatient agents and the rest being the patient agents.

There are three types of matches: matches between impatient agents and patient agents, matches between patient agents and patient agents, and self-matches of impatient agents. We use $x_{n,i} \in \{0, 1\}$ as the variable to indicate whether a match is made between agent n and agent i (we use $i = 0$ to denote self-match), where $n \in [N]$ indexes the set of impatient agents and patient agents, and $i \in [I]^*$ indexes self-matches (i.e., $i = 0$) and the set of patient agents (e.g., $i \in [I]$).

- When $i=0$, $x_{n,0}$ represents a self-match of agent n .
- When $i > 0$ and $n \leq T$, $x_{n,i}$ represents a match between impatient agent n and patient agent i .
- When $i > 0$ and $n > T$, $x_{n,i}$ represents a match between patient agent $j = n - T$ and patient agent i .

We let $w_{n,i}$ denote the value (or weight) for the match between agent n and agent i . Our framework can deal with any model of cardinal utilities, but for our purposes in this paper we describe how to set the values of $w_{n,i}$ for the kidney exchange application. When $i \neq 0$, we set $w_{n,i}$ to be the sum of the expected graft survivals of the recipients for pair n and pair i if the two pairs exchange their kidneys. We set $w_{n,i} = 0$ if there is no acceptable kidney exchange between two pairs (including the case where impatient agent n is medically compatible with patient agent i , but the internal transplant is higher quality for the recipient in n). When $i = 0$, we set $w_{n,0}$ to be the utility generated by internal transplant. Naturally $w_{n,0} = 0$ if $n > T$ since $n > T$ indexes incompatible pairs who cannot perform internal transplant.

Integer program formulation. We are now ready to formulate the 2-matching problem (up to two agents are involved in a match). We first consider an offline formulation, in which we assume we know which impatient agent is arriving at each time t . Our problem can be formulated as the following integer program (IP):

$$\begin{aligned}
& \max \sum_{n=1}^N \sum_{i=0}^I w_{n,i} x_{n,i} \\
& \text{s.t.} \sum_{i=0}^I x_{n,i} \leq 1, \forall n \in [T] \\
& \sum_{n=1}^N x_{n,i} + \sum_{j=1}^I x_{T+i,j} \leq 1, \forall i \in [I] \\
& x_{n,i} \in \{0, 1\}, \forall n \in [N], \forall i \in [I]^*
\end{aligned}$$

The objective is to maximize the total utility, which can be written as $\sum_n \sum_i w_{n,i} x_{n,i}$. The first constraint ensures each impatient agent can either self-match (e.g, $x_{n,0} = 1$) or be matched to at

most one patient agent. The second constraint ensures each patient agent can only be matched to at most one agent, either an impatient agent or a patient agent.

6.2 Matching Algorithms

In practice, the integer program above cannot be directly solved since we do not have knowledge of future impatient agent arrivals. Hence we must derive an algorithm that can make matching decisions online using only knowledge of the static pool, a single arriving impatient agent, and overall distributional properties. One natural, albeit computationally expensive, choice is to set up and solve the full IP after each arrival of an impatient agent, but only perform a match the impatient agent is involved in at that time. We term this algorithm *Online Assignment via Exhaustive Search* or OAES. Below we propose another algorithm, ODASSE (Online Dual Assignment Using Shadow Survival Estimates), based on online primal-dual techniques.

6.2.1 The ODASSE algorithm. In ODASSE, instead of working with the integer program, we will work with the relaxed linear program that allows the $x_{n,i}$ variables to take on fractional values in $[0, 1]$. This relaxation enables us to derive the dual and utilize primal-dual techniques to obtain an online assignment algorithm (note that our proposed algorithms only make integer assignments).

The dual of the above linear program can be written as

$$\begin{aligned} \min \quad & \sum_{t=1}^T \alpha_t + \sum_{i=0}^I \beta_i \\ \text{s.t.} \quad & w_{t,i} - \alpha_t - \beta_i \leq 0, \forall t \in [T], i \in [I]^* \\ & w_{t+j,i} - \beta_j - \beta_i \leq 0, \forall i \in [I], j \in [I] \\ & \alpha_t, \beta_i \geq 0, \forall t \in [T], i \in [I] \\ & \beta_0 = 0 \end{aligned}$$

where α_t, β_i are the dual variables that can be interpreted as the estimated values of compatible pairs and incompatible pairs respectively. β_0 comes from self-match and is always set to 0. Analogous to *shadow prices* in constrained optimization, we call β_i the *shadow survival* of patient agent i .

Online Assignment. Given the primal-dual formulation, we now turn to formulating the online assignment rule. We first state an additional assumption that is easy to satisfy in practice and is useful in our discussion. We assume that for each $w_{n,j} \neq 0$, we can add a small perturbation to the weight value, which can help remove potential tie-breaking issues. Since this perturbation can be arbitrarily close to 0, the influence on the optimal solution can also be arbitrarily small.

Online dual assignment using optimal $\{\beta_i\}$. First we assume an idealized scenario in which there exists an oracle which provides us with optimal values of β_i^* in the dual problem. From complementary slackness, we know that, for $1 \leq t \leq T$, the optimal assignment $x_{t,i}^*$ should satisfy $x_{t,i}^*(w_{t,i} - \alpha_t^* - \beta_i^*) = 0$.

From the constraint, we have $\alpha_t^* \geq w_{t,i} - \beta_i^*$. Since the objective is attempting to minimize $\sum_t \alpha_t$, therefore $\alpha_t^* = \max_i \{w_{t,i} - \beta_i^*\}$. Since we have the small perturbation assumption, there is a unique i such that $w_{t,i} - \alpha_t^* - \beta_i^* = 0$. Denote it as i^* . We get $\alpha_t^* = \{w_{t,i^*} - \beta_{i^*}^*\}$. We can conclude $i^* = \arg \max_i \{w_{t,i} - \beta_i^*\}$.

Therefore, in the optimal primal, we should set $x_{t,i^*}^* = 1$ and set $x_{t,i}^* = 0$ for all $i \neq i^*$ to satisfy complementary slackness. We can apply the same arguments for the constraints $w_{t+j,i} - \beta_j - \beta_i \leq 0$ and get the assignment decision for patient agent $n = T+1, \dots, T+I$, as summarized in Algorithm 1. Note that it is easy to check that this assignment algorithm satisfies both primal and dual feasibility. Furthermore, all of the assignments $x_{n,i}^*$ are either set to 1 or 0. Therefore, ODASSE generates the

optimal assignment $x_{n,i}^*$ when given optimal β_i^* for the 2-matching problem as specified in the integer program.

Estimating $\{\beta_i\}$. In the above algorithm, assume optimal β_i^* is given, we can make optimal online assignments without knowing which impatient agent will arrive in the future (i.e., we do not need information about α_t). In practice, we do not have access to the true β_i . We propose two approaches in estimating β_i values:

- Learning from past populations. By solving the dual on simulated populations, we are able to gather training data that can be used to approximate β_i without knowing the sequence of arriving compatibles, using a machine learning approach. We train off two major information sources: 1) the demographic and medical information of an incompatible pair (all the individual features needed to compute LKDPI) and 2) the initial graph state of the incompatible pairs. The latter of these 2 takes the form of the β_i value when solving the dual on just the incompatible pool. We find that a linear model is a good predictor of β_i , performing as well as more complex approaches like random forests. It is also worth noting that the initial graph state of the incompatible pairs is an important feature for prediction.
- Simulating future arrivals. Since we have access to historical populations, we are able to simulate the arrivals of impatient agents and solve the optimal β_i values for the simulated population. We can estimate the β_i values using the average of many simulations. We note here that we use a method that re-simulates the future at every impatient agent arrival, since the current state of the graph provides meaningful information. This is again computationally expensive.

Algorithm 1 Online Dual Assignment ($\{\beta_i\}$)

```

1: Input:  $\{\beta_i\}$ 
2: Initialize  $\mathbb{I} = [I]^*$ 
3: for  $t = 1, \dots, T$  do
4:   Set  $i^* \leftarrow \arg \max_{i \in \mathbb{I}} \{w_{t,i} - \beta_i\}$ 
5:   if  $i^*$  equals to 0 then
6:     Self-match impatient agent  $t$ 
7:   else
8:     Match patient agent  $i^*$  with impatient agent  $t$ 
9:     Set  $\mathbb{I} \leftarrow \mathbb{I} \setminus \{i^*\}$ 
10:  end if
11: end for
12: for  $j = 1, \dots, I$  do
13:  if  $j$  is in  $\mathbb{I}$  then
14:    Set  $i^* \leftarrow \arg \max_{i \in \mathbb{I} \setminus \{j\}} \{w_{T+j,i} - \beta_i\}$ 
15:    if  $w_{T+j,i^*}$  equals to 0 then
16:      No match is found
17:    else
18:      Match patient agent  $j$  with  $i^*$ 
19:      Set  $\mathbb{I} \leftarrow \mathbb{I} \setminus \{i^*\}$ 
20:    end if
21:  end if
22: end for

```

6.3 Extensions

So far, we have focused the discussion on 2-matching problems. However, our framework can extend to other settings, including k -matching problems. Below we discuss a few extensions.

6.3.1 k -matching problems. Our framework can easily handle k -matching problems for any k (i.e., including up to k cycles in kidney exchanges). Below we first describe how to model 3-matching problems and provide discussion on general k -matching problems.

In 3-matching problems (in kidney exchanges, this corresponds to allowing internal-transplant, 2-cycles, and 3-cycles in a match), we can add an additional dimension $j \in [I]^*$ to our variables to expand the formulation to perform matches between 2 or 3 agents. Our variable indicates the match involving different agents:

- $x_{n,0,0}$: self match of impatient agent n
- $x_{n,i,0}$: impatient agent n and patient agent i
- $x_{n,i,j}$: impatient agent n and patient agents i and j
- $x_{T+k,i,j}$: patient agents i, j , and k

We also need to calculate all weights of $w_{n,i,j}$ with any given utility models. For example, in kidney exchange, if $i, j \neq 0$, $w_{n,i,j}$ is the sum of the utility for the 3 kidneys exchanged in the 3-cycle. With this new notations, we can again write down the primal and dual formulations. In particular, the dual can be written as:

$$\begin{aligned} \min \quad & \sum_{t=1}^T \alpha_t + \sum_{i=1}^I \beta_i \\ \text{s.t.} \quad & w_{n,i,j} - \alpha_t - \beta_i - \beta_j \leq 0, \forall(n, i, j) \\ & w_{T+k,i,j} - \beta_i - \beta_j - \beta_k \leq 0, \forall(i, j, k) \\ & \alpha_t, \beta_i \geq 0, \forall(t, i) \\ & \beta_0 = 0 \end{aligned}$$

Note that the constraints are very similar to 2-matching problems, except we need to subtract an additional β value on the left-hand side of the first two constraints. Looking at the first constraint, using complementary slackness, when impatient agent t arrives, we choose $(i^*, j^*) = \arg \max_{i,j} \{w_{t,i,j} - \beta_i - \beta_j\}$ to match with t . We can write down similar assignment rules for the second constraint. The same derivations can easily extend to k -matching problems. There will be $k-1$ and k β variables to be subtracted on the left-hand side of the first two constraints of the dual. And we can derive the optimal assignment rules using complementary slackness.

6.3.2 Including chains. We note that fielded kidney exchanges often gain from chains, where an altruist donor without a paired patient enters the pool and starts a directed path of transplants [10]. Technically, it is relatively straightforward to include chains in our framework. The formulation is the same as the matching problem, the main change needed is to calculate weights based on the edges on a chain (instead of on a cycle), and to estimate a value for the “leftover” kidney at the end of a chain. This estimation will interact with the online formulation, and thus needs careful modeling.

6.3.3 Dynamic pool of patient agents. In practice, pools of incompatible pairs are not static, but are dynamically changing with pairs arriving and departing at each time interval. Our framework can be expanded to account for this dynamic environment by setting $w_{n,i,j} = 0$ unless there exists a time when all pairs are present. Again, careful modeling of arrival and departure will be critical to ensuring that the results are meaningful in this case.

6.4 Experimental Results

We now investigate the impacts of including compatible pairs in exchanges. The Oracle model (or ODASSE with perfect information about β values) is an upper bound on what can be achieved under the constraint that compatible pairs must be matched immediately, while OAES and the variants of ODASSE with estimation of β values are specific algorithms that can make online decisions without knowledge of the future. We compare our online matching algorithms with the baseline, in which compatible pairs are always matched with their original donors, and incompatible pairs are matched among themselves using two- & three-cycle swap.

Kidney matching environment. We use the simulator as described in Section 5. The simulator generates donor-recipient pairs based on known blood type, tissue type and demographic feature distributions in the population. Based on these, it can determine compatibility using the approach of Saidman 2006, as well as the LKDPI score [36], between pairs. The LKDPI score can be mapped directly to expected graft survival time, which provides a meaningful measure of match quality. Each experiment we report consists of 500 simulations of 50 arriving compatible pairs with a pool size of 100 incompatible pairs.¹⁰

Evaluation metrics. We evaluate performance on two main dimensions: (1) Overall match quality, the expected number of years of total graft survival of all transplants; (2) The number of transplants performed.

	Baseline	OAES	ODASSE	Oracle
Matched proportion of incompatible pairs	54.4%	74.6%	70.6%	76.0%
Expected graft survival of compatible pairs	9.6	11.1	11.2	11.4
Expected graft survival of incompatible pairs	10.4	9.8	9.6	10.0

Table 5. The matched proportion of incompatible pairs (Row 1) and the expected graft survival of compatible (Row 2) and incompatible (Row 3) pairs, where Row 3 is conditional on being matched. The baseline in Rows 1 and 3 represents running two- & three-cycle swap with only incompatible pairs and in Row 2 represents the case where recipients from compatible pairs are always matched with their original donors. The ODASSE algorithm reported here uses machine learning estimates for β values.

Results. Our headline results are in Table 5, which shows that including compatible pairs in exchange, even with the immediacy constraint, can improve outcomes dramatically, matching 37% more incompatible pairs, a 20 percentage point improvement in the match rate. For compatible pairs, there is an increase of 1.5 years in the expected survival time of the transplanted kidney. There is a slight decrease in expected graft survival for incompatible pairs, but this is conditional on being transplanted. This is because of a “rival” effect. In order to incentivize compatible pairs to participate, “better” quality kidneys have to be offered to those pairs. This is presumably more than balanced by the substantially increased probability of receiving a transplant, given that you are incompatible.

¹⁰While the arrival rates of compatible and incompatible pairs in the system are similar, because they match less often, there are more incompatible pairs waiting at any given time interval than the number of compatible pairs that arrive.

6.4.1 *Fairness.* An important question is how the introduction of a new mechanism could impact different subpopulations. While aggregate effects may appear strongly positive, it would be problematic if they came at the expense of specific subgroups. In kidney exchange, the focus is often on two hard-to-match subpopulations: (1) Blood type O recipients, who can only receive kidneys from Type O donors, and are thus matched at lower rates than those with other blood types; (2) Highly sensitized patients, who have antigens likely to make them reject many donor kidneys ($PRA \geq 90\%$). We focus on the former here, because of changing practice (surgeons now believe the appropriate hard to match PRA threshold should be 98%) and the complex interaction between HLA, blood group, and PRA.

Since compatible pairs are always being matched, we only consider incompatible pairs here. Figure 4 shows our main results. Using the OAES algorithm increases the proportion of blood Type O recipients (in incompatible pairs) matched from 32% to 62%, a proportionally much higher increase than for the general population of incompatible recipients. The increased proportion of matches demonstrates that the gains in social welfare do not come at the cost of fairness for these harder-to-match patients.

Another question that could arise about fairness relates to the result in Table 5 showing that the expected graft survival (EGS) of incompatible pairs decreases in both algorithms compared to the baseline. This is a complex issue; Table 5 shows EGS conditional on transplantation. More incompatible pairs are getting transplants, however, so unconditional EGS increases. The decline is for recipients with higher quality donors (because of the “rival” effect introduced by the need to incentivize compatible pairs to participate with better donor kidneys).

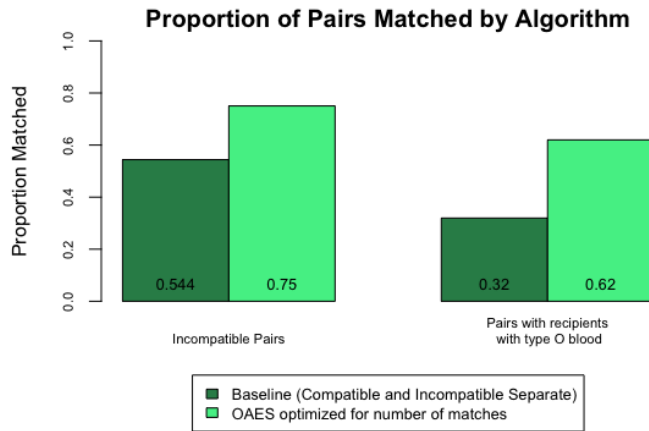


Fig. 4. Change in the proportion of the entire population and the subpopulation of pairs with blood type O recipients in the incompatible pool that are matched when compatible pairs are included, using the OAES algorithm to benchmark.

6.4.2 *Comparison of online matching algorithms.* Figure 5 compares the different algorithms for matching. It is interesting to note that ODASSE with learned estimates of β achieves performance better than the version with expensive repeated simulation, and that the performance of ODASSE and OAES is so close to the oracle. This is also promising in the sense that better estimation of β could perhaps get even closer to optimal performance.

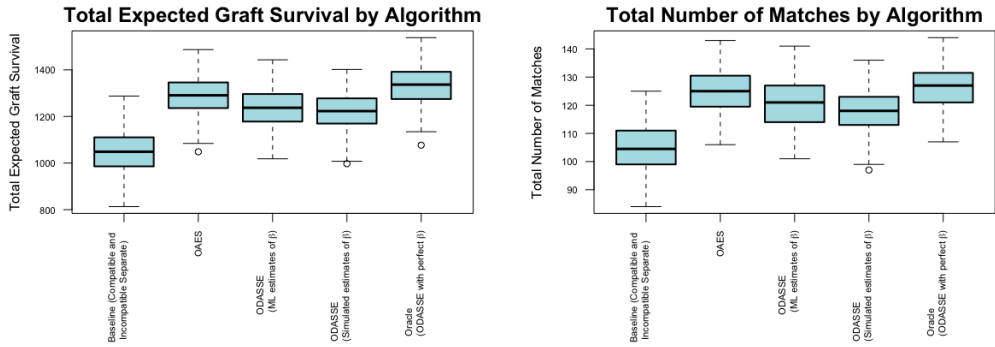


Fig. 5. Box plots for the expected total years of graft survival and number of matches for the different matching algorithms. The oracle algorithms solve the integer program with complete knowledge of future arrivals of compatible pairs – one version (count) optimizes for the number of transplants, and the other (quality) for the expected graft survival.

It is also worth commenting on run-time. ODASSE requires many initial runs to generate training data and estimate β values, but is then very fast in the online phase. OAES is expensive: one run with 50 compatible arrivals and an incompatible pool of size 100 takes 2-3 minutes on commodity hardware, and does not scale well to larger instances. Nevertheless, OAES is likely sufficient for the scale of problems dealt with in kidney exchange, since transplants cannot actually be conducted immediately. ODASSE is likely a better choice in settings requiring close-to-real-time immediacy.

7 CONCLUSION AND FUTURE RESEARCH

Living donor kidney transplantation has proven to be an important domain for the development of matching theory and algorithms. It is becoming increasingly important to study cardinal utilities in kidney exchange, and we believe this could open up more fertile avenues for research. Our main goal in this paper is to develop a robust framework for analyzing match quality in models of kidney exchange. Our framework for modeling match quality is based on real donor and recipient data from a major transplant center, and we introduce a hybrid online-batch matching model to capture compatible pairs' expected desire for immediacy. We use the model to estimate the benefits of including compatible pairs in exchanges, in terms of both quantity and quality of transplants. We find that if we were able to induce compatible pairs to join kidney exchanges, the percentage of matched incompatible pairs would increase dramatically, and there would also be a substantial increase in expected graft survival for recipients in compatible pairs. Quantifying the potential quantitative benefits of participating through LKDPI may also make compatible pairs more likely to join.

The development of our simulator for match quality should allow for future work on weighted matching to use realistic weight distributions from a real application. In addition, our new model of matching with patient and impatient agents also opens up interesting algorithmic questions. Finally, there are many interesting modeling and policy questions for kidney exchange specifically that can be pursued – particularly interesting ones include consideration of the tradeoffs between waiting costs and improved match quality, and possible systemic effects of including compatible pairs in exchanges, for example, changes in incentives for centers.

ACKNOWLEDGMENTS

We are grateful for support from NSF awards 1527037 and 1560191.

REFERENCES

- [1] David Abraham, Avrim Blum, and Tuomas Sandholm. 2007. Clearing Algorithms for Barter Exchange Markets: Enabling Nationwide Kidney Exchanges. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 295–304.
- [2] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. 2017. Thickness and information in dynamic matching markets. Working paper. Initial version appeared at ACM Conference on Economics and Computation (EC-14).
- [3] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. 2015. A dynamic model of barter exchange. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1925–1933.
- [4] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. 2017. Efficient dynamic barter exchange. *Operations Research* 65, 6 (2017), 1446–1459.
- [5] Elliot Anshelevich, Meenal Chhabra, Sanmay Das, and Matthew Gerrior. 2013. On the Social Welfare of Mechanisms for Repeated Batch Matching. In *AAAI Conference on Artificial Intelligence (AAAI)*. 60–66.
- [6] Elliot Anshelevich and Sanmay Das. 2010. Matching, cardinal utility, and social welfare. *ACM SIGecom Exchanges* 9, 1 (2010), 4.1–4.7.
- [7] Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. 2013. Anarchy, Stability, and Utopia: Creating Better Matchings. *Autonomous Agents and Multi-Agent Systems* 26, 1 (2013), 120–140.
- [8] Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. 2018. Maximum weight online matching with deadlines. *arXiv preprint arXiv:1808.03526* (2018).
- [9] Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. 2015. Mix and Match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91 (2015), 284–296.
- [10] Itai Ashlagi, David Gamarnik, Michael Rees, and Alvin E. Roth. 2017. The Need for (Long) Chains in Kidney Exchange. Working paper. Initial version appeared at ACM Conference on Electronic Commerce (EC-12).
- [11] Itai Ashlagi, Patrick Jaillet, and Vahideh H. Manshadi. 2013. Kidney Exchange in Dynamic Sparse Heterogenous Pools. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. ACM, New York, NY, USA, 25–26. <https://doi.org/10.1145/2482540.2482565>
- [12] Itai Ashlagi and Alvin E Roth. 2014. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics* 9, 3 (2014), 817–863.
- [13] Joseph H Bragdon, Ali Hortaçsu, and Dan Ariely. 2010. Matching and sorting in online dating. 100, 1 (2010), 130–163.
- [14] R Scott Braithwaite, David O Meltzer, Joseph T King Jr, Douglas Leslie, and Mark S Roberts. 2008. What does the value of modern medicine say about the \$50,000 per quality-adjusted life-year decision rule? *Medical Care* 46, 4 (2008), 349–356.
- [15] Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. 2013. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research* 231 (2013), 57–68.
- [16] David Roxbee Cox and David Oakes. 1984. *Analysis of survival data*. Vol. 21. CRC Press.
- [17] Sanmay Das and Emir Kamenica. 2005. Two-Sided Bandits and the Dating Market. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*. Edinburgh, UK, 947–952.
- [18] N. Devanur and T. Hayes. 2009. The AdWords problem: Online keyword matching with budgeted bidders under random permutations. In *Proceedings of ACM Conference on Electronic Commerce (EC)*.
- [19] Nicholas Diakopoulos. 2015. Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism* 3, 3 (2015), 398–415.
- [20] Nicholas Diakopoulos. 2016. Accountability in algorithmic decision making. *Commun. ACM* 59, 2 (2016), 56–62.
- [21] John P. Dickerson, David Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. 2016. Position-Indexed Formulations for Kidney Exchange. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [22] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. 2012. Dynamic Matching via Weighted Myopia with Application to Kidney Exchange. In *AAAI Conference on Artificial Intelligence (AAAI)*. 1340–1346.
- [23] John P. Dickerson and Tuomas Sandholm. 2015. FutureMatch: Combining Human Value Judgments and Machine Learning to Match in Dynamic Environments. In *AAAI Conference on Artificial Intelligence (AAAI)*. 622–628.
- [24] Yuval Emek, Shay Kutten, and Roger Wattenhofer. 2016. Online Matching: Haste Makes Waste!. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*. 333–344.
- [25] Haluk Ergin, Tayfun Sönmez, and M. Utku Ünver. 2017. Multi-Donor Organ Exchange. Working paper.
- [26] Gabriele Farina, John P. Dickerson, and Tuomas Sandholm. 2017. Operation Frames and Clubs in Kidney Exchange. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. AAAI Press, 199–205.
- [27] SE Gentry, DL Segev, Mary Simmerling, and RA Montgomery. 2007. Expanding kidney paired donation through participation by compatible pairs. *American Journal of Transplantation* 7, 10 (2007), 2361–2370.
- [28] Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. 2015. Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*. 921–928.
- [29] Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle. 2017. Labor Market Equilibration: Evidence from Uber. Working paper Revision requested from American Economic Review.

- [30] Sundaram Hariharan, Christopher P Johnson, Barbara A Bresnahan, Sarah E Taranto, Matthew J McIntosh, and Donald Stablein. 2000. Improved graft survival after renal transplantation in the United States, 1988 to 1996. *New England Journal of Medicine* 342, 9 (2000), 605–612.
- [31] Philip J Held, 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 (2016), 877–885.
- [32] Chien-Ju Ho and Jennifer Wortman Vaughan. 2012. Online Task Assignment in Crowdsourcing Markets. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI'12)*. AAAI Press, 45–51.
- [33] Kimberly D Krawiec and Michael A Rees. 2014. Reverse transplant tourism. *Law & Contemporary Problems* (2014), 145–173.
- [34] Jian Li, Yicheng Liu, Lingxiao Huang, and Pingzhong Tang. 2014. Egalitarian Pairwise Kidney Exchange: Fast Algorithms via linear Programming and Parametric Flow. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS '14)*. 445–452.
- [35] Thaminda Liyanage, Toshiharu Ninomiya, Vivekanand Jha, Bruce Neal, Halle Marie Patrice, Ikechi Okpechi, Ming-hui Zhao, Jicheng Lv, Amit X Garg, John Knight, et al. 2015. Worldwide access to treatment for end-stage kidney disease: A systematic review. *The Lancet* 385, 9981 (2015), 1975–1982.
- [36] Allan B Massie, Joseph Leanza, LM Fahmy, EKH Chow, Niraj M Desai, X Luo, EA King, MG Bowring, and DL Segev. 2016. A risk index for living donor kidney transplantation. *American Journal of Transplantation* 16, 7 (2016), 2077–2084.
- [37] Nicholas Mattei, Abdallah Saffidine, and Toby Walsh. 2017. Mechanisms for Online Organ Matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. 345–351.
- [38] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press Cambridge.
- [39] Cathy O'Neil. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.
- [40] Panduranga S Rao, Douglas E Schaubel, Mary K Guidinger, Kenneth A Andreoni, Robert A Wolfe, Robert M Merion, Friedrich K Port, and Randall S Sung. 2009. A comprehensive risk quantification score for deceased donor kidneys: The kidney donor risk index. *Transplantation* 88, 2 (2009), 231–236.
- [41] Alvin Roth, Tayfun Sönmez, and Utku Ünver. 2004. Kidney exchange. *Quarterly Journal of Economics* 119, 2 (2004), 457–488.
- [42] Alvin Roth, Tayfun Sönmez, and Utku Ünver. 2005. A kidney exchange clearinghouse in New England. *American Economic Review* 95, 2 (2005), 376–380.
- [43] Alvin Roth, Tayfun Sönmez, and Utku Ünver. 2005. Pairwise Kidney Exchange. *Journal of Economic Theory* 125, 2 (2005), 151–188.
- [44] Alvin E Roth. 2015. *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.
- [45] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. 2007. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *The American Economic Review* (2007), 828–851.
- [46] Susan L. Saidman, Alvin Roth, Tayfun Sönmez, Utku Ünver, and Frank Delmonico. 2006. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation* 81, 5 (2006), 773–782.
- [47] George W Torrance and David Feeny. 1989. Utilities and quality-adjusted life years. *International Journal of Technology Assessment in Healthcare* 5, 4 (1989), 559–575.
- [48] Utku Ünver. 2010. Dynamic kidney exchange. *Review of Economic Studies* 77, 1 (2010), 372–414.
- [49] Yolanda Bravo Vergel and Mark Sculpher. 2008. Quality-adjusted life years. *Practical Neurology* 8, 3 (2008), 175–182.