Empirical Price of Fairness in Failure-Aware Kidney Exchange

John P. Dickerson Ariel D. Procaccia Tuomas Sandholm

Computer Science Department
Carnegie Mellon University
{dickerson,arielpro,sandholm}@cs.cmu.edu

ABSTRACT

Fielded kidney exchanges typically use utilitarian or near-utilitarian matching rules, potentially at great cost to certain classes of hard-to-match patients. Dickerson, Procaccia, and Sandholm [6] recently adapted the *price of fairness*, a measure of the tradeoff between fairness and efficiency, to kidney exchange; they showed that the price of fairness is small in theory but often non-negligible in practice. We extend their work by directly comparing fairness in two models of kidney exchange—where edges either can or cannot fail after an algorithmic match but before transplantation. On real and simulated data, even matching under strict fairness constraints in the model with edge failure results in significantly more expected transplants than matching *efficiently* in the deterministic model.

1. INTRODUCTION

The preferred treatment for kidney failure is transplantation. Successful transplantation of a kidney relies on tissue-type compatibility between the donor organ and patient, among other factors. Compatibility is determined through a tissue-type *crossmatch* between a potential donor and patient's blood. If the two types differ substantially, the patient's body will reject the donor organ.

Some patients are *highly-sensitized*, with a very low probability of passing a crossmatch test with a random organ. For these patients, finding a kidney is difficult [10]. Roughly 17% of the adult patients on the waiting list for *deceased* donor kidneys are highly-sensitized [7]. Recently, an allocation policy was designed for deceased donors that balances fairness and efficiency while working within the currently fielded priority-based framework [4].

Complementing deceased donation is kidney exchange, which allows patients with willing but medically incompatible *living* donors to swap donors with other patients. The percentage of highly-sensitized patients in fielded kidney exchanges is quite high; over 60% of the United Network for Organ Sharing (UNOS) nationwide kidney exchange is highly-sensitized, as shown in Figure 1.

Currently fielded kidney exchanges tend to match using utilitarian or near-utilitarian rules. Intuitively, maximizing social welfare may come at the cost of marginalizing highly-sensitized patients. Bertsimas, Farias, and Trichakis [3] introduced a general measure of this cost as the *price of fairness*—the relative loss in total welfare from using a "fair" objective (in our case, a "fair" matching rule), instead of an overall utility-maximizing one. A recent paper by Dickerson, Proccacia, and Sandholm [6] adapts this concept to kidney exchange, finding that the theoretical price of fairness (with respect to highly-sensitized patients) in kidney exchange is low. They then formalize two natural "fair" utility functions and show how to optimize either of these functions in two models of kidney exchange. They find that, empirically, the price of fairness is frequently high on real data.

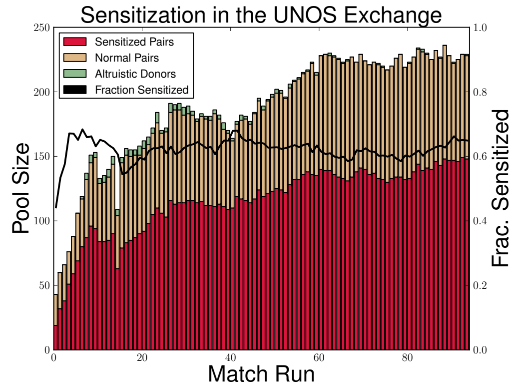

Figure 1: Composition of the UNOS national kidney exchange from inception. For each of 94 match runs (x-axis), the raw number of highly-sensitized patients, non-highly-sensitized patients, and altruists are plotted (left y-axis), as well as the percentage of patients who are highly-sensitized as a percentage of the pool size (right y-axis).

In this paper, we extend their work by explicitly comparing the impact of fairness constraints on match efficiency in either of two models of kidney exchange. One is deterministic, where a matched edge certainly results in a transplant; the other is failure-aware, where a matched edge may fail before transplantation. Under utilitarian matching rules, the failure-aware model will result in more transplants in expectation in both theory and practice [5]. However, we show that failure-aware matching results in more expected transplants than maximum cardinality matching, *even when strict fairness constraints are enforced*. We conclude that there is a huge "price of using the wrong model" that is potentially more harmful to *all* patients than ignorance of fairness constraints.

2. PRELIMINARIES

In this section, we briefly overview the standard graph-based model of kidney exchange and discuss two optimization models—deterministic and failure-aware—adapted in Dickerson, Procaccia, and Sandholm [6] to include fairness constraints.

2.1 Kidney Exchange Model

The standard model for kidney exchange encodes an n-patient kidney exchange as a directed *compatibility graph* $G = (V, E)$ by constructing one vertex for each patient-donor pair. An edge e from v_i to v_j is added if the patient in v_j is compatible with the donor kidney of v_i. A donor is willing to give her kidney if and only if the patient in her vertex v_i receives a kidney. The weight w_e of an edge e represents the utility to v_j of obtaining v_i's donor kidney.

A cycle c in the graph G represents a possible kidney swap, where each vertex in c obtains the kidney of the previous vertex. In fielded kidney exchange, cycles of length at most only some small constant L are allowed. In most fielded kidney exchanges, including the UNOS kidney exchange, $L = 3$ (i.e., only 2- and 3-cycles are allowed). Fielded kidney exchanges also gain great utility through the use of *chains*. Chains start with an altruistic donor donating his kidney to a patient, whose paired donor donates her kidney to another patient, and so on.

A *matching* M is a collection of vertex-disjoint cycles and chains in the graph G. Given the set of all legal matchings $\mathcal{M}$, the *clearing problem* in kidney exchange is to find a matching M^* that maximizes some utility function $u : \mathcal{M} \to \mathbb{R}$. Formally:

$$M^* = \operatorname*{argmax}_{M \in \mathcal{M}} u(M)$$

In fielded kidney exchanges, one typically finds the maximum weighted cycle cover (i.e., $u(M) = \sum_{c \in M} \sum_{e \in c} w_e$). This *utilitarian* objective can favor certain classes of patient-donor pairs while marginalizing others.

2.2 Weighted Fairness in Kidney Exchange

One simple method to emphasize a certain class of patient-donor pairs—for us, those in the set of highly-sensitized vertices V_H—is to increase the weight of edges with a sink in V_H. This definition generalizes the policy UNOS currently applies to highly-sensitized patients in their fielded kidney exchange.

Building on the standard (deterministic or probabilistic) kidney exchange integer programming formulation, we rewrite the objective as follows:

$$\max \quad \textstyle\sum_c v_\Delta(c) x_c$$

Here, $v_\Delta(c)$ is the value of a cycle or chain c (either in the deterministic or probabilistic model) such that the weight of each edge $e \in c$ is adjusted by some re-weighting function $\Delta : E \to \mathbb{R}$.

A simple example re-weighting function is multiplicative:

$$\Delta^\beta(e) = \begin{cases} (1+\beta)w_e & \text{if } e \text{ ends in } V_H \\ w_e & \text{otherwise} \end{cases}$$

Intuitively, for some $\beta > 0$, this function scales the weight of edges ending in highly-sensitized vertices by $(1+\beta)$. For example, if $\beta = 0.5$, then the optimization algorithm will value edges that result in a highly-sensitized patient receiving a transplant at 50% above their initial weight (possibly scaled by other factors like failure probability and chain position, as in the probabilistic model).

For any $M \in \mathcal{M}$, let M' be the matching such that every cycle $c \in M$ has augmented weight $v_\Delta(c)$. Then define the *weighted* fairness rule u_Δ in terms of the utilitarian rule u applied to the augmented matching M', such that $u_\Delta(M) = u(M')$.

3. EXPERIMENTAL RESULTS

In this section, we compare failure-aware matching against deterministic matching. Our metric is the expected number of transplants. We draw from real match runs and simulated UNOS data [9] (see [5] for a derivation of failure rates).[1] We also draw from aggregate failure rates published by the Alliance for Paired Donation (APD) [1]. Our experiments unilaterally support the hypothesis that failure-aware matching, even with strict fairness constraints implemented, results in more actual transplants than even fully utilitarian deterministic matching.

[1]Please get in touch for data. All code is available at `https://github.com/JohnDickerson/KidneyExchange`.

3.1 Constant Failure Probability

We begin by assuming every edge fails with the same constant probability. For example, a failure probability of 0.2 in the figures below corresponds to a compatibility graph in which every edge in an algorithmic match will fail 20% of the time and succeed 80% of the time. This assumption, while not likely to hold in practice, is easily parameterized and allows us to explore the differences in models as matchings become less reliable.

Figure 2 compares the weighted fairness rule u_Δ applied to the failure-aware model—which takes edge failure into account *during the optimization process*—against the utilitarian rule u applied to the deterministic model, which computes a maximum cardinality disjoint cycle cover without regard for edge failure. The efficient failure-aware matching always results in at least as many (typically more) expected transplants as the efficient deterministic matching; however, of interest, even matchings under the fair rule u_Δ in the failure-aware model often result in significant overall gains when compared to the utilitarian deterministic matching.

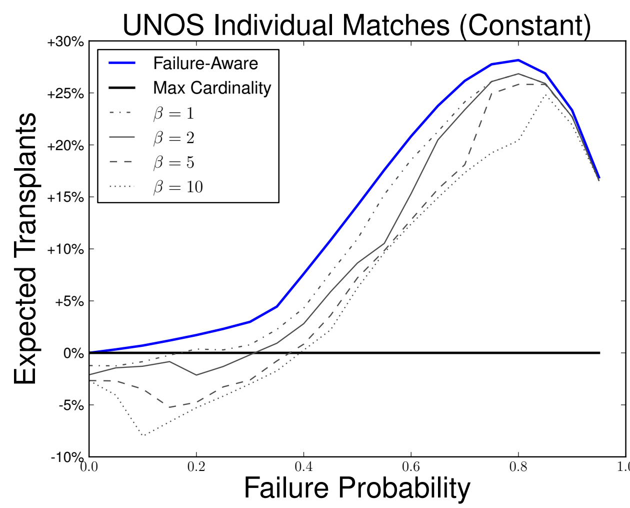

Figure 2: Percentage change in the expected number of transplants for *actual* UNOS match runs when using failure-aware matching—possibly with fairness constraints—instead of maximum cardinality matching. The x-axis varies the constant edge failure probability from 0 to 1.

For example, for $\beta = 1.0$—that is, when highly-sensitized patients are valued at twice that of lowly-sensitized patients—we see a drop of only a couple of percentage points of expected transplants when there is *no* probability of edge failure. When the probability of edge failure is at least 20%, this fairer failure-aware matching beats the efficient deterministic matching. In fact, when the probability of edge failure is at least 40%, valuing highly-sensitized transplants at 11x ($\beta = 10.0$) that of a lowly-sensitized patient results in more expected transplants overall than the "unfair" deterministic matching! This general behavior is supported in Figure 3, which performs the same experiments on generated data that mimics the UNOS distribution, for pools of size 50 and 250.

3.2 Bimodal Failure Probability

In reality, not all potential transplants are created equal. There is evidence that matches to highly-sensitized patients are more likely to fail than matches to lowly-sensitized patients [1, 5]. In this section, we perform similar experiments to Section 3.1, only now drawing edge failure probabilities from such a *bimodal* distribution.

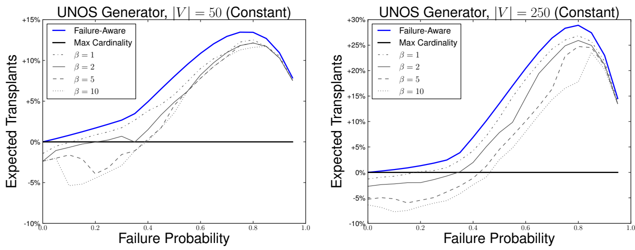

Figure 3: Percentage change in the expected number of transplants for *generated* UNOS match runs when using failure-aware matching—possibly with fairness constraints—instead of maximum cardinality matching. The x-axis varies the constant edge failure probability from 0 to 1.

We draw from failure rates estimated from the UNOS exchange [9] and from from aggregate failure rates published by the APD [1].

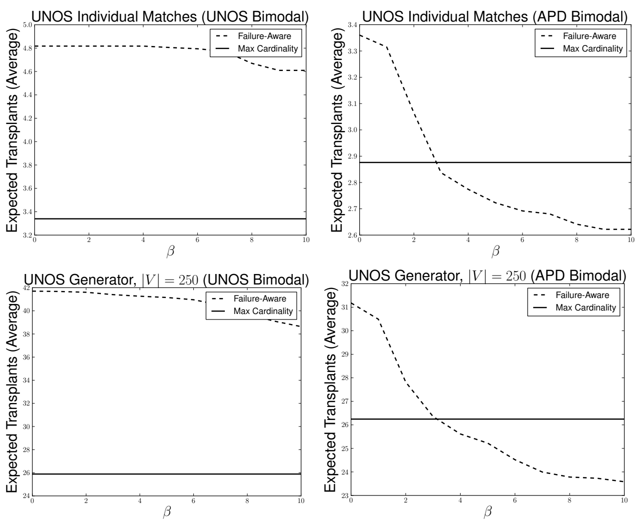

Figure 4: Change in the expected number of transplants on average for *actual* (top row) and *generated* (bottom row) UNOS match runs when using failure-aware matching instead of maximum cardinality matching, assuming bimodal edge failure rates derived from UNOS (left column) and APD (right column). The x-axis varies the β fairness factor applied to the failure-aware matching algorithm.

Figure 4 compares the expected number of transplants resulting from a match under $u_\triangle$ for increasing values of $\beta \in \{0, 1, \ldots, 10\}$ against a baseline number of transplants from the efficient deterministic matching under u (shown as a thick horizontal line on the plots). Again, we see that even applying fairly strict fairness bumps to hard-to-match patients still results in significant overall gains in expected transplants under the failure-aware model when compared to a utilitarian maximum cardinality matching.

4. DISCUSSION

Balancing efficiency and equity in general allocation of resources in healthcare, not just kidney exchange, is a contentious issue. Formally, full resolution of the issue has proven itself to be difficult (if not impossible) across the board; however, recent progress has been made in areas like deceased donor kidney and liver allocation. We believe similar progress can and should be made in kidney exchange, and that this work can be formalized mathematically in adapted models of kidney exchange.

While the price of fairness in kidney exchange *within* a model can be quite high for strict definitions of "fairness", we showed that there is often an even *higher* price paid for optimizing in the wrong model, regardless of fairness. In this paper, we considered the loss in expected transplants from using an efficient deterministic model against both efficient and various versions of "fair" failure-aware matching. However, future work should consider other aspects of real kidney exchange that are not presently modeled in fielded exchanges, like dynamic matching and game-theoretic considerations from the viewpoint of participating hospitals and other legal entities. Some research has been done in this area already, but has primarily focused on utilitarian and not fair matching rules.

Moving forward, the kidney exchange community would benefit immensely from *combined* approaches to handling not just dynamic matching, match failures, and fairness in the optimization problem, but also game-theoretic and legal considerations in the design of the matching mechanism itself. Our research group plans to draw on work from the operations research and economics literature to move in this direction: for example, Hooker and Williams present a general methodology for balancing a particular form of fairness (that we feel would not be the criterion of choice in kidney exchange) and efficiency [8]; Bertsimas, Farias, and Trichakis formalized a proposal for balancing fairness and efficiency in *deceased* kidney allocation [4]; and Ashlagi, Jaillet, and Manshadi theoretically address dynamic exchange in a reduced model [2]. A general parameterized model of kidney exchange will increase the efficacy of fielded exchange and aid in the widespread adoption of new exchanges in differing legal and political environments.

5. REFERENCES

[1] I. Ashlagi, D. S. Gilchrist, A. E. Roth, and M. Rees. Nonsimultaneous chains and dominos in kidney-paired donation—revisited. *American Journal of Transplantation*, 11(5):984–994, 2011.

[2] I. Ashlagi, P. Jaillet, and V. H. Manshadi. Kidney exchange in dynamic sparse heterogenous pools. In *ACM Conference on Electronic Commerce (EC)*, pages 25–26, 2013.

[3] D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.

[4] D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

[5] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Failure-aware kidney exchange. In *ACM Conference on Electronic Commerce (EC)*, pages 323–340, 2013.

[6] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of fairness in kidney exchange. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2014.

[7] HHS/HRSA/HSB/DOT. OPTN/SRTR annual data report, 2011.

[8] J. N. Hooker and H. P. Williams. Combining equity and utilitarianism in a mathematical programming model. *Management Science*, 58(9):1682–1693, 2012.

[9] Kidney Paired Donation Work Group. OPTN KPD pilot program cumulative match report (CMR) for KPD match runs: Oct 27, 2010 – Apr 15, 2013, 2013.

[10] United Network for Organ Sharing (UNOS). http://www.unos.org/.