Scalable Robust Kidney Exchange

Duncan C McElfresh
Department of Mathematics
University of Maryland
College Park, MD 20742
dmcelfre@math.umd.edu

Hoda Bidkhori
Department of Industrial Engineering
University of Pittsburgh
Pittsburgh, PA 15260
bidkhori@pitt.edu

John P Dickerson
Department of Computer Science
University of Maryland
College Park, MD 20742
john@cs.umd.edu

Abstract

In barter exchanges, participants directly trade their endowed goods in a constrained economic setting without money. Transactions in barter exchanges are often facilitated via a central clearinghouse that must match participants even in the face of uncertainty—over participants, existence and quality of potential trades, and so on. Leveraging robust combinatorial optimization techniques, we address uncertainty in kidney exchange, a real-world barter market where patients swap (in)compatible paired donors. We provide two scalable robust methods to handle two distinct types of uncertainty in kidney exchange—over the *quality* and the *existence* of a potential match. The latter case directly addresses a weakness in all stochastic-optimization-based methods to the kidney exchange clearing problem, which all necessarily require explicit estimates of the probability of a transaction existing—a still-unsolved problem in this nascent market. We also propose a novel, scalable kidney exchange formulation that eliminates the need for an exponential-time constraint generation process in competing formulations, maintains provable optimality, and serves as a subsolver for our robust approach. For each type of uncertainty we demonstrate the benefits of robustness on real data from a large, fielded kidney exchange in the United States. We conclude by drawing parallels between robustness and notions of fairness in the kidney exchange setting.

1 Introduction

Real-world optimization problems face various types of uncertainty that impact both the quality and feasibility of candidate solutions. Uncertainty in combinatorial optimization is especially troublesome: if the *existence* of certain constraints or variables is uncertain, identifying a good—or even feasible—solution can be extremely difficult. Stochastic optimization approaches endeavor to maximize the *expected* objective value, under uncertainty. While sometimes successful, stochastic optimization relies heavily on a correct characterization of uncertainty; furthermore, stochastic approaches are often intractable—especially in combinatorial domains (Bertsimas, Brown, and Caramanis 2011). A complementary approach is *robust optimization*, which protects against worst-case outcomes. Robust approaches can be less sensitive to the exact characterization of uncertainty, and are often far more tractable than stochastic approaches (Ben-Tal, El Ghaoui, and Nemirovski 2009).

This paper addresses uncertainty in *kidney exchange*, a real-world barter market where patients with end-stage renal disease enter and trade their willing paired kidney donors (Rapaport 1986; Roth, Sönmez, and Ünver 2004). Kidney exchange is a relatively new paradigm for organ allocation, but already accounts for over 10% of living kidney donations in the United States, and is growing in popularity worldwide (Biró et al. 2017). Modern exchanges also include *non-directed donors* (NDDs), who enter the market without a paired patient and donate their kidney without receiving one in return. Computationally, kidney exchange is a packing problem: solutions (matchings) consist of cyclic organ swaps and NDD-initiated donation chains in a directed compatibility graph, representing all participants and potential transactions. Each potential transplant is given a numerical weight by policymakers; the objective is to select cycles and chains that maximize overall matching weight. In general, this problem is NP-hard (Abraham, Blum, and Sandholm 2007; Biró, Manlove, and Rizzi 2009); however, many efficient deterministic formulations exist that are fielded now and clear real exchanges (Abraham, Blum, and Sandholm 2007; Manlove and O'Malley 2015; Anderson et al. 2015; Dickerson et al. 2016; Dickerson, Procaccia, and Sandholm 2018).

Uncertainty in kidney exchange. Presently-fielded kidney exchange algorithms largely do not address uncertainty. Here, we consider two types of uncertainty in kidney exchange: over the *quality* of the transplant (weight uncertainty) and over the *existence* of potential transplants (existence uncertainty). Policymakers assign weights to potential transplants, which are (imperfect) estimates of transplant quality; weight uncertainty stems from both measurement uncertainty (e.g., medical compatibility and kidney quality) and uncertainty in the prioritization of some patients over others. Transplant existence is always uncertain: matched transplants "fail" before executing for a variety of reasons, severely impacting a planned kidney exchange. To address both cases, we propose *uncertainty sets* containing different realizations of the uncertain parameters. We then develop a scalable robust optimization approach, and demonstrate its success on data from a large fielded kidney exchange.

Robust optimization is a popular approach to optimization under uncertainty, with applications in reinforcement learning (Petrik and Subramanian 2014), regression (Xu, Caramanis, and Mannor 2009), classification (Chen et al. 2017), and network optimization (Mevissen, Ragnoli, and Yu 2013). Motivated by real-world constraints, we apply robust optimization to kidney exchange—a graph-based market clearing or resource allocation problem.

Our Contributions. To our knowledge, weight uncertainty has not been addressed in the kidney exchange literature. Our approach is similar to that of Bertsimas and Sim (2004) and Poss (2014), and uses some of their results. Several approaches have been proposed for existence uncertainty, primarily based on stochastic optimization (Dickerson et al. 2016; Anderson et al. 2015; Dickerson, Procaccia, and Sandholm 2018) or hierarchical optimization (Manlove and O'Malley 2015). The primary disadvantage of these approaches—in addition to tractability—is their reliance on, and sensitivity to, the explicit estimation of the probability of each particular potential transplant. This probability is extremely difficult to determine (Dickerson, Procaccia, and Sandholm 2018; Glorie 2012), and prevents the translation of those methods into practice. Our approach uses a simpler notion of edge existence uncertainty—an upper-bound on the number of non-existent edges—which is easier to interpret and estimate. Glorie (2014) proposed a related robust formulation that is exponentially larger than ours, and is intractable for realistically-sized exchanges.

In addition, we introduce a new scalable formulation for kidney exchange that combines concepts from two state-of-the-art formulations (Anderson et al. 2015; Dickerson et al. 2016), handles long or uncapped NDD-initiated chains without requiring expensive constraint generation, and ties into a developed literature on fairness in kidney exchange—thus addressing use cases that are becoming more common in fielded exchanges (Anderson et al. 2015).

2 Preliminaries

Model for kidney exchange. A kidney exchange can be represented formally by a directed compatibility graph $G = (V, E)$. Here, vertices $v \in V$ represent participants in the exchange, and are partitioned as $V = P \cup N$ into P, the set of all patient-donor pairs, and N, the set of all NDDs (Roth, Sönmez, and Ünver 2004; Roth, Sönmez, and Ünver 2005; Abraham, Blum, and Sandholm 2007). Each potential transplant from a donor at vertex u to a patient at vertex v is represented by a directed edge $e = (u, v) \in E$, which has an associated weight $w_e \in \mathbf{w}$; weights are set by policymakers, and reflect both the medical utility of the transplant, as well as ethical considerations (e.g., prioritizing patients by waiting time, age, and so on). Cycles in G correspond to cyclic trades between multiple patient-donor pairs in P; chains, correspond to donations that begin with an NDD in N and continue through multiple patient-donor pairs in P. The kidney exchange *clearing problem* is to select a feasible set of transplants (edges in E) that maximize overall weight. Let $\mathcal{M}$ be the set of all feasible *matchings* (i.e., solutions) to a kidney exchange problem; the general formulation of this problem is $\max_{\mathbf{x} \in \mathcal{M}} \mathbf{x} \cdot \mathbf{w}$, where binary decision variables $\mathbf{x}$ represent edges, or cycles and chains. This problem is NP- and APX-hard (Abraham, Blum, and Sandholm 2007; Biró, Manlove, and Rizzi 2009).

Robust optimization. Robust optimization is a common approach to optimization under uncertainty, which is often more tractable and requires less accurate uncertainty information than other approaches (Bertsimas, Brown, and Caramanis 2011). This approach begins by defining an *uncertainty set* $\mathcal{U}$ for the uncertain optimization parameter; $\mathcal{U}$ contains different *realizations* of this parameter. Consider the example of edge weight uncertainty: we might design an edge weight uncertainty set $\mathcal{U}_w$ that contains the *realized* (i.e. "true") edge weights $\hat{\mathbf{w}}$ with high probability, $P(\hat{\mathbf{w}} \in \mathcal{U}_w) \geq 1 - \epsilon$, for $0 < \epsilon \ll 1$. The parameter ϵ is referred to as the *protection level*, and is often used to control the number of realizations in $\mathcal{U}$.

After designing $\mathcal{U}$, the robust approach finds the best solution, assuming the *worst-case* realization within $\mathcal{U}$. For kidney exchange (a maximization problem), this corresponds to a *minimization* over $\mathcal{U}$; for example, Problem (1) is the robust formulation with uncertain edge weights.

$$\max_{\mathbf{x} \in \mathcal{M}} \min_{\hat{\mathbf{w}} \in \mathcal{U}} \quad \mathbf{x} \cdot \hat{\mathbf{w}} \tag{1}$$

The robustness of this approach depends on the proportion of possible realizations contained in $\mathcal{U}$. If $\mathcal{U}$ contains all possible realizations, the approach may be too conservative; if $\mathcal{U}$ only contains one possible realization of $\hat{\mathbf{w}}$, the solution may be too myopic. The number of realizations in $\mathcal{U}$ is often controlled by a parameter: either an *uncertainty budget* Γ, or the protection level ϵ. Next we introduce the first type of uncertainty considered in this paper: edge weight uncertainty.

3 Optimization in the Presence of Edge Weight Uncertainty

Edge weights in kidney exchange represent the medical and social utility gained by a *single* kidney transplant. Weights are determined by policymakers, and are subject to several types of uncertainty.[1] Part of this uncertainty is due to insufficient knowledge of the future: a patient or donor's health may change, raising or lowering the "true" weight of their transplant edges. Another type of uncertainty stems from disagreement between policymakers regarding the social utility of a transplant. For example, some policymakers might prioritize young patients over older patients; other policymakers might prioritize the sickest patients above all healthier patients. Policymakers aggregate these value judgments to assign a single weight to each transplant edge, but this aggregation is a contentious and imperfect process (although recent work from the AI community has begun to address this using techniques from computational social choice and machine learning (Freedman et al. 2018; Noothigattu et al. 2018)). Still, there is no way to measure the "true" social utility of a transplant, and therefore this uncertainty is not easily measured.

Interval weight uncertainty. It is beyond the scope of this work to characterize these sources of uncertainty. We simply assume that the *nominal* edge weights $\mathbf{w}$, provided by

[1]The process used to set weights by the UNOS US-wide kidney exchange is published publicly (UNOS 2015).

policymakers, are an uncertain estimate of the *realized* edge weights $\hat{\mathbf{w}}$, i.e., the "true" value of each transplant. Next, we formalize edge weight uncertainty and our robust approach. This section focuses on edge weights, so we write our formulations with decision variables $x_e \in \mathbf{x}$ corresponding to individual edges.

We assume that realized edge weights $\hat{\mathbf{w}}$ are random variables with a partially known symmetric distribution, centered about the nominal weights $\mathbf{w}$. This assumption implies that $E[\hat{\mathbf{w}}] = \mathbf{w}$, thus a non-robust approach that maximizes $\mathbf{w}$ is equivalent to a stochastic optimization approach that maximizes *expected* edge weight. We refer to this edge uncertainty model as *interval uncertainty*.

Definition 1 (Interval Edge Weight Uncertainty)**.** Let $\hat{w}_e$ be the realized weight of edge e, with nominal weight w_e, and maximum discount $0 \leq d_e \leq w_e$. Let $\hat{w}_e \equiv w_e + d_e\alpha_e$, where α_e is the *fractional deviation* of edge e. Both α_e and $\hat{w}_e$ are continuous random variables, symmetrically distributed on $[-1, 1]$ and $[w_e - d_e, w_e + d_e]$ respectively.

Each discount factor d_e should reflect the level of uncertainty in w_e. If w_e is known exactly, then $d_e = 0$; if w_e is very uncertain, then we might set $d_e = w_e$, or higher.

To vary the degree of uncertainty, we use an *uncertainty budget* Γ, which limits the total deviation from nominal edge weights. With our uncertainty model, it is natural to let Γ limit the total fractional deviation of each edge weight—i.e., sum of all α_e. This uncertainty set $\mathcal{U}^I_\Gamma$ is defined as:

$$\mathcal{U}^I_\Gamma = \left\{ \hat{\mathbf{w}} \mid \hat{w}_e = w_e + d_e\alpha_e, |\alpha_e| \leq 1, \sum_{e\in E} |\alpha_e| \leq \Gamma \right\}$$

For example if $\Gamma = 3$, there may be three edges with $|\alpha_e| = 1$, or one edge with $|\alpha_e| = 1$ and four edges with $|\alpha_e| = 1/2$, and so on.

Choosing an appropriate Γ is not straightforward. Matchings often use only a small fraction of the decision variables (e.g., transplant edges), and it is difficult to predict the size of the optimal matching. Intuitively, Γ should reflect the size of the final matching: for example if we assume that half of any matching's edges will be discounted, then we should set $\Gamma \simeq |\mathbf{x}|/2$. Generalizing this concept, we define a *variable*-budget uncertainty set $\mathcal{U}^I_\gamma$, with budget function $\gamma(|\mathbf{x}|)$.

$$\mathcal{U}^I_\gamma = \left\{ \hat{\mathbf{w}} \mid \hat{w}_e = w_e + d_e\alpha_e, |\alpha_e| \leq 1, \sum_{e\in E} |\alpha_e| \leq \gamma(|\mathbf{x}|) \right\}$$

Next, to define γ, we relate it to a much more intuitive parameter: the protection level ϵ.

3.1 Uncertainty Budget γ and Protection Level ϵ

The protection level ϵ mediates between a completely conservative approach, and the non-robust approach: as $\epsilon \to 0$ the approach becomes more conservative, and $\epsilon = 1$ corresponds to a non-robust approach. In this section we relate γ to ϵ, beginning with the following Theorem 1.

Theorem 1 (Adapted from Theorem 3 of (Bertsimas and Sim 2004))**.** *For a matching $\mathbf{x} \in \mathcal{M}$ with $|\mathbf{x}|$ edges, and uncertainty set $\mathcal{U}^I_\Gamma$, the probability that $\mathcal{U}^I_\Gamma$ contains the realized edge weights for $\mathbf{x}$ is bounded below by*

$$P(\hat{\mathbf{w}} \in \mathcal{U}^I_\Gamma) \geq 1 - B(|\mathbf{x}|, \Gamma),$$

with

$$B(|\mathbf{x}|, \Gamma) = \frac{1}{2^{|\mathbf{x}|}} \left((1-\mu) \binom{|\mathbf{x}|}{\lfloor \eta \rfloor} + \sum_{l=\lfloor \eta \rfloor + 1}^{|\mathbf{x}|} \binom{|\mathbf{x}|}{l} \right),$$

with $\eta = (\Gamma + |\mathbf{x}|)/2$ and $\mu = \eta - \lfloor \eta \rfloor$.

That is, for some ϵ, if Γ is chosen such that $\epsilon = B(|\mathbf{x}|, \Gamma)$, then the inequality $P(\hat{\mathbf{w}} \in \mathcal{U}^I_\Gamma) \geq 1 - \epsilon$ holds by Theorem 1. Next, we use this result to define a variable uncertainty budget function γ, using the intuitive definition introduced by Poss (2014): for matching $\mathbf{x} \in \mathcal{M}$ and protection level ϵ, we find the minimum Γ such that $B(|\mathbf{x}|, \Gamma) \leq \epsilon$. If this is not possible (i.e., the matching is too small, or ϵ is too small), then $\gamma = |\mathbf{x}|$. This budget function is defined as:

$$\beta(|\mathbf{x}|) = \begin{cases} |\mathbf{x}| & \text{if } \min_{\Gamma>0} \{\Gamma \mid B(|\mathbf{x}|, \Gamma) \leq \epsilon\} \text{ is infeasible,} \\ \min_{\Gamma>0} \{\Gamma \mid B(|\mathbf{x}|, \Gamma) \leq \epsilon\} & \text{otherwise.} \end{cases}$$

It may not be clear how to solve the edge weight robust problem with this variable uncertainty budget. We use the approach of Poss (2014), which solves the variable-budget robust problem by solving several instances of the *constant*-budget robust problem; details of this approach can be found in Appendix A.4[2]. Thus, to solve the variable-budget robust problem we first solve the constant-budget robust problem.

3.2 Constant-Budget Edge Weight Robust Approach

We now describe our approach to the constant-budget edge weight robust problem; a full discussion and derivation can be found in Appendix A. We need to solve Problem (1) with edge weight uncertainty set $\mathcal{U}^I_\Gamma$. This requires a minimization of the objective, over $\hat{\mathbf{w}} \in \mathcal{U}^I_\Gamma$, followed by a maximization over matchings in $\mathcal{M}$.

First we *directly minimize* the objective of Problem (1) over $\mathcal{U}^I_\Gamma$. That is, for any matching $\mathbf{x} \in \mathcal{M}$, we find the minimum objective value for any realized edge weights in $\mathcal{U}^I_\Gamma$, denoted by $Z(\mathbf{x})$:

$$Z(\mathbf{x}) = \min_{\hat{\mathbf{w}} \in \mathcal{U}^I_\Gamma} \mathbf{x} \cdot \hat{\mathbf{w}} \tag{2}$$

Thus, solving the robust problem corresponds to maximizing $Z(\mathbf{x})$ over all feasible matchings. Our approach to doing so is as follows. First, we linearize $Z(\mathbf{x})$ using several new variables and constraints; we then add these to an existing kidney exchange formulation (Dickerson et al. 2016). The complete linear formulations of $Z(\mathbf{x})$ and Problem (1) are given in Appendix A.2, but are omitted here for space. Our robust formulation is scalable—it has a polynomial count of variables and constraints, regardless of finite chain cap; on realistic exchanges it takes only a few seconds to solve. We demonstrate our method's impact on match composition in Section 5, and show how it effectively controls for the impact of robustness using protection level ϵ.

[2]All appendices are included in the full version of this paper, available on arXiv: https://arxiv.org/abs/1811.03532

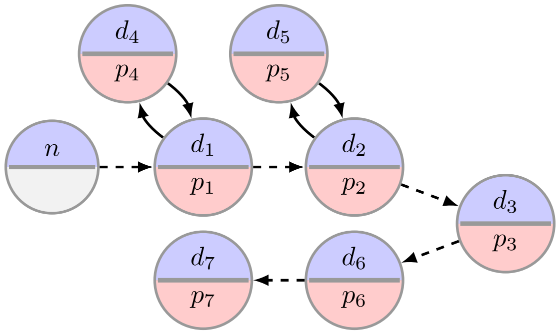

Figure 1: Sample exchange graph with a 5-chain and two 2-cycles. The NDD is denoted by n, and each patient (and her associated donor) is denoted by p_i (d_i). A maximum-cardinality matching algorithm would select the 5-chain, denoted with dashed edges; however, the smaller matching consisting of two disjoint 2-cycles, shown with solid edges, may be more robust to edge failure.

4 Optimization in the Presence of Edge Existence Uncertainty

In this section we consider *edge existence uncertainty*, where an algorithmic match must be chosen before the full realization of edges is revealed. Algorithmically-matched transplants in a kidney exchange can fail before transplantation for a variety of reasons: a patient may become too ill to undergo transplantation, or pre-transplantation testing may reveal that a patient is incompatible with her planned donor kidney. Furthermore, some edges are more likely to fail than others (e.g., edges into particularly sick patients). Edge failure significantly impacts fielded exchanges–with failure rates often above 50% (Dickerson, Procaccia, and Sandholm 2018; Anderson et al. 2015; Ashlagi, Jaillet, and Manshadi 2013).

For illustration, consider the simple exchange in Figure 1 with two potential matchings: single 5-chain initiated by the NDD, or two 2-cycles (with pairs $\{1, 4\}$ and $\{2, 5\}$). The 5-chain matches the most patient, but is less robust to edge failures. Consider the *worst-case* outcome for each matching, when 1 edge is *guaranteed* to fail: with the 5-chain, in the worst-case the *first* edge fails, causing the entire chain to fail; with the 2-cycles, a single edge failure only causes a *single* cycle to fail, leaving the other cycle complete. With this notion of edge existence uncertainty (which we define later), the 2-cycles are more robust than the 5-chain.

Managing edge failure in kidney exchange has been addressed in the AI and optimization literature in application-specific (Manlove and O'Malley 2015; Chen et al. 2012) or stochastic-optimization-based (Dickerson, Procaccia, and Sandholm 2018; Dickerson et al. 2016; Anderson et al. 2015; Klimentova, Pedroso, and Viana 2016) ways. These *failure-aware* approaches associate with each edge a pre-determined failure probability p_e; these probabilities are used to then maximize *expected* matching score, possibly subject to some recourse actions. This stochastic approach is tractable when p_e is identical for each edge. Our work addresses two major drawbacks of the failure-aware approach. First, when each edge has a unique p_e, those models require enumerating every cycle and chain, which is intractable for large graphs or long chains. Second, the failure-aware approach is very sensitive to p_e (as discussed in, e.g., §4.4 of Dickerson, Procaccia, and Sandholm (2018)). In practice, precise values of p_e are not known, thus the failure-aware approach can easily produce unreliable results. We use a simpler notion of edge existence uncertainty, which assumes that in any matching, the number of edges is *bounded* by a constant (Γ). This parameter is intuitive and simple to estimate from past exchanges.

To our knowledge, ours is the first *scalable* robust optimization approach to edge existence uncertainty in kidney exchange. Glorie (2014) develops several elegant robust methods for edge existence uncertainty, but requires that all cycles and chains are found during pre-processing and stored in memory. The number of chains grows exponentially in both the number of edges and the maximum chain length; thus, these approaches are intractable for exchanges involving more than a few dozen patient-donor pairs and NDDs.

Edge existence uncertainty. Here we briefly describe our robust approach to edge existence uncertainty; a full discussion and derivation can be found in Appendix B. For ease of exposition, in this section, decision variables $x_c \in \mathbf{x}$ correspond to cycles and chains rather than edges. We use the following model of edge existence uncertainty.

Definition 2 (Γ-Failures Edge Existence Uncertainty)**.** Up to Γ edges may fail in any matching. After failures occur, the realized exchange graph is $\hat{G} = (V, \hat{E})$, such that edges $\hat{E} \subseteq E$ succeed and remain in existence, while all other edges fail and do not exist.

With this notion of uncertainty, without regard to computational or memory constraints, a stochastic-optimization-based approach could identify the best matching over all possible realizations $\hat{G}$ (Anderson et al. 2015). This is clearly intractable, as the number of realized graphs is exponential in $|E|$. Instead, we take a robust optimization approach by maximizing the worst-case (minimum) matching score over a set of realizable graphs $\hat{G}$ in an uncertainty set $\mathcal{U}$. Like the stochastic approach, the robust approach considers a huge number of realizations $\hat{G}$; however the robust approach is far more tractable, as it need only find the worst-case realization and need not represent all realizable graphs explicitly.

Uncertainty set. Let $F \subseteq E$ be the subset of failed edges for a realized graph $\hat{G}$; thus, $\hat{E} = E \setminus F$ is the set of realized edges. Equation (3) defines uncertainty set $\mathcal{U}^{ex}_{\Gamma}$ in this way: up to Γ edges may fail (i.e., $|F| \leq \Gamma$).

$$\mathcal{U}^{ex}_{\Gamma} = \left\{ \hat{G} = (V, \hat{E}) \mid \hat{E} = E \setminus F, |F| \leq \Gamma \right\} \qquad (3)$$

In kidney exchange, one edge failure can cause other edge failures: if one cycle edge fails, all edges in the cycle also fail; edge failure in a chain causes all *subsequent* chain edges to also fail. This leads to a notion of weight uncertainty for cycles and chains, where the realized weight of a cycle or chain $\hat{w}_c$ may be smaller than nominal weight w_c. Let α_c be a discount parameter for cycle or chain c, such that $\hat{w}_c = w_c(1 - \alpha_c)$. For example, if any edge fails in cycle

c, then the entire cycle fails and $\alpha_c = 0$. We define the cycle/chain weight uncertainty set $\mathcal{U}^w_\Gamma$ in this way:

$$\mathcal{U}^w_\Gamma = \left\{ \hat{\mathbf{w}}_\mathbf{c} \mid \hat{w}_c = w_c(1-\alpha_c), \alpha_c \in [0,1], \sum_{c \in X} \alpha_i \leq \Gamma \right\}$$

This uncertainty set is less intuitive than $\mathcal{U}^{ex}_\Gamma$, but more suited to the robust approach. In Appendix B we show that $\mathcal{U}^w_\Gamma$ and $\mathcal{U}^{ex}_\Gamma$ are equivalent for integer Γ, and thus can be used for our robust approach.

4.1 Robust Optimization Approach

In this section we briefly describe our robust approach; for a full discussion and derivation, please see Appendix B. Our robust formulation for uncertainty set $\mathcal{U}^w_\Gamma$ follows a similar approach to Section 3. First, we directly minimize the kidney exchange objective over $\mathcal{U}^w_\Gamma$, for some feasible solution $\mathbf{x} \in \mathcal{M}$. We express this minimization as a function $Z(\mathbf{x})$: in effect, $Z(\mathbf{x})$ discounts the Γ largest-weight cycles and chains. We then linearize $Z(\mathbf{x})$ using several variables and constraints—this requires a formulation with variables tracking individual total chain weights—which is not possible in any existing compact kidney exchange formulations. For this purpose, we introduce a new kidney exchange formulation.

The PI-TSP formulation. We propose the position-indexed TSP formulation (PI-TSP); for details, please see Appendix B. Our formulation combines innovations from the two leading kidney exchange clearing approaches: PICEF (Dickerson et al. 2016) and PC-TSP (Anderson et al. 2015). PICEF introduced an indexing schema that enables a more compact formulation in the context of long chains; our formulation builds on this to allow tracking of individual chain weights, a necessity that PICEF could not do. PC-TSP builds on techniques from the prize-collecting travelling salesperson problem (Balas 1989) to provide a tight linear programming relaxation; in general, the PC-TSP formulation has exponentially many constraints and thus requires constraint generation to solve. Our formulation uses an efficient version of position indexing that also requires only $O(|E|)+O(|V|\cdot|N|)$ constraints. Unlike PICEF, our formulation does not grow with the chain cap L: PICEF uses $O(|V|^3)$ variables (when $L \to |V|$); for large graphs, the PICEF model becomes too large to fit into memory (Dickerson et al. 2016). Our formulation uses a fixed number of variables—$O(|V|^2)$—for any chain cap, alleviating this memory problem. This is especially relevant to existing exchanges, as long chains can significantly increase efficiency in kidney exchange (Ashlagi et al. 2012). PI-TSP uses the following parameters:

- a kidney exchange graph G with cycles C
- L: chain cap (maximum number of edges used in a chain),
- w_e: edge weights for each edge $e \in E$,
- w^C_c: cycle weights for each cycle $c \in C$,

and the following decision variables:

- $p_{\{e,v\}} \geq 1$: the position of edge e or vertex v in any chain,
- $\hat{p}_e \geq 0$: equal to p_e if e is used in a chain, and 0 otherwise,
- $z_c \in \{0,1\}$: 1 if cycle c is used in the matching,
- $y_e \in \{0,1\}$: 1 if edge e is used in a chain, and 0 otherwise,
- $y^n_e \in \{0,1\}$: 1 if edge e is used in a chain starting with NDD n, and 0 otherwise,
- w^N_n: total weight of the chain starting with NDD n,
- f^i_v and f^o_v: chain flow into and out of vertex $v \in P$,
- $f^{i,n}_v$ and $f^{i,n}_v$: chain flow into and out of vertex $v \in P$, from the chain starting with NDD $n \in N$.

The PI-TSP formulation with chain cap L is given in Problem 4. We use the notation $\delta^-(v)$ for the set of edges into vertex v and $\delta^+(v)$ for the set of edges out of v.

$$\max \quad \sum_{n \in N} w^N_n + \sum_{c \in C} w^C_c z_c \tag{4a}$$

$$\text{s.t.} \quad \sum_{e \in E} w_e y^n_e = w^N_n \qquad n \in N \tag{4b}$$

$$\sum_{n \in N} y^n_e = y_e \qquad e \in E \tag{4c}$$

$$\sum_{e \in \delta^-(v)} y_e = f^i_v \qquad v \in V \tag{4d}$$

$$\sum_{e \in \delta^+(v)} y_e = f^o_v \qquad v \in V \tag{4e}$$

$$\sum_{e \in \delta^-(v)} y^n_e = f^{i,n}_v \qquad v \in V, n \in N \tag{4f}$$

$$\sum_{e \in \delta^+(v)} y^n_e = f^{o,n}_v \qquad v \in V, n \in N \tag{4g}$$

$$f^o_v + \sum_{c \in C : v \in c} z_c \leq \quad f^i_v + \sum_{c \in C : v \in c} z_c \leq 1 \qquad v \in P \tag{4h}$$

$$f^o_v \leq 1 \qquad v \in N \tag{4i}$$

$$p_e = 1 \qquad e \in \delta^+(N) \tag{4j}$$

$$\hat{p}_e = p_e y_e \qquad e \in E \tag{4k}$$

$$p_v = \sum_{e \in \delta^-(v)} \hat{p}_e \qquad v \in P \tag{4l}$$

$$p_e = p_v + 1 \qquad v \in P, e \in \delta^+(v) \tag{4m}$$

$$\sum_{e \in E} y^n_e \leq L \qquad n \in N \tag{4n}$$

$$f^{o,n}_v \leq f^{i,v} \leq 1 \qquad v \in V, n \in N \tag{4o}$$

$$y_e \in \{0,1\} \qquad e \in E \tag{4p}$$

$$z_c \in \{0,1\} \qquad c \in C \tag{4q}$$

$$y^n_e \in \{0,1\} \qquad e \in E, n \in N \tag{4r}$$

The ability to express individual chain weights as decision variables has applications beyond robustness. For particularly valuable NDDs (such as those with so-called "universal donor" blood-type O), exchanges may enforce a *minimum* chain length or chain weight, to ensure that these rare NDDs are not "used up" on short chains; such a policy was formerly used by the United Network for Organ Sharing (Dickerson, Procaccia, and Sandholm 2012), using a much less scalable form of optimization—that also does not consider uncertainty—than our approach (Abraham, Blum, and Sandholm 2007). Such a policy can be implemented efficiently with PI-TSP, inefficiently with PC-TSP, and not with PICEF, where decision variables do not indicate from which NDD a chain originated. In Appendix B we show–using real kidney exchange data–that PI-TSP can enforce a minimum chain length, and that this restriction has *almost no* impact on overall matching score.

5 Experimental Results

In this section, we compare each robust formulation against the leading non-robust formulation, PICEF (Dickerson et al. 2016), with varying levels of uncertainty. These experiments use real exchange graphs collected from the United Network for Organ Sharing (UNOS)—a large US-wide kidney exchange with over 160 participating transplant centers—between 2010 and 2016, as well simulated exchanges generated from known patient statistics using the standard method (Dickerson, Procaccia, and Sandholm 2018).[3]

For each exchange, we calculate the optimal non-robust matching M_{OPT} (with total score $|M_{\text{OPT}}|$), and the robust matching M_{R}, for varying uncertainty budgets. We then draw many *realizations* of the exchange graph, based on the uncertainty model, and calculate the realized scores of the robust matching $|M_{\text{R}}|$ and non-robust matching $|M_{\text{NR}}|$. We are primarily interested in the fractional difference from $|M_{\text{OPT}}|$, calculated as $\Delta OPT\left(M_{\{\text{R,NR}\}}\right) = \left(|M_{\text{OPT}}| - |M_{\{\text{R,NR}\}}|\right)/|M_{\text{OPT}}|$.

We calculate $\Delta OPT\,(M_{\text{R}})$ and $\Delta OPT\,(M_{\text{NR}})$ for $N = 400$ realizations, and compare the robust and non-robust approaches. In rare cases the optimal matching is empty (i.e., there is no solution, or the uncertainty budget exceeds the matching size), we exclude these exchanges from the results.

Edge Weight Uncertainty We begin by exploring the impact on match utility of robust approaches to managing edge weight uncertainty. Here, each edge is randomly labeled as *probabilistic* (P) or *deterministic* (D). P edges receive weight 0 or 1 with probability 0.5, while D edges always receive weight 0.5; thus, expected edge weight is 0.5. The non-robust approach maximizes *expected* edge weight, making this a kind of stochastic approach. The robust approach considers the discount value (0 or 0.5) of each edge, and avoids edges with a positive discount value. We vary the level of uncertainty with the fraction of P edges (α). Realizations are drawn by assigning the P edges to have weight either 0 or 1.

We compute M_{R} for protection levels $\epsilon \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 0.5\}$, and then calculate both $\Delta OPT\,(M_{\text{R}})$ and $\Delta OPT\,(M_{\text{NR}})$. Figure 2 shows ΔOPT on realistic 64-vertex simulated graphs (left) and larger (typically 150–300-vertex) real UNOS graphs (right); these figures show results for each protection level ϵ and for various α. Note that M_{NR} does not depend on ϵ, and thus the non-robust results are shown as (constant) dashed lines.

The robust approach guarantees a better worst-case (minimum) ΔOPT, but results in a lower median ΔOPT. The protection level ϵ controls the robustness of our approach; smaller ϵ protects against more uncertain outcomes, but at greater cost to nominal behavior. As $\epsilon \to 1$, the robust approach protects against fewer bad outcomes, and approaches the behavior of non-robust.

Edge Existence Uncertainty We now address edge existence uncertainty, and compare the robust and non-robust approaches with Γ edge failures, for $\Gamma \in \{1, 2, 3, 4, 5\}$. Each Γ corresponds to a different notion of uncertainty, such that exactly Γ edges fail.[4] For each Γ, we calculate M_{R}, and draw $N = 400$ realizations by failing Γ edges in the matching.

We calculate ΔOPT for each realization, and compare these results for the robust and non-robust matchings. With edge existence uncertainty, the worst-case outcome is almost always an empty matching ($\Delta OPT\,() = -1$). Thus, rather than compare the worst-case ΔOPT, we compare the *distribution* of ΔOPT for each approach: we treat ΔOPT as a random variable, and use three simple statistical tests to demonstrate that—as expected—the robust approach produces more conservative and predictable results.

First, we use the Wilcoxon signed-rank test to determine that the robust and non-robust approaches produce a different distribution of ΔOPT. For each Γ, this test produces p-values well below 10^{-3}, indicating that the distributions of ΔOPT *are different* for the robust and non-robust approach. Second, for all exchanges and all Γ, the mean ΔOPT is typically 1% *higher*, and the standard deviation 1–2% *lower* for the robust approach. That is, the robust approach more consistently produces higher-weight solutions.

Third, we visualize the difference between these distributions using their histograms. Figure 3 shows the bin-wise difference between the histograms of ΔOPT (robust minus non-robust), with mean ΔOPT for non-robust shown as a dotted line. In these plots, the height of the bars indicate the change in probability density due to robustness. On all plots, the bars are *negative* for high and low values of ΔOPT, meaning that the robust matching is *less likely* to have an abnormally high or low ΔOPT. The bars are *positive* when ΔOPT is near its mean non-robust value—meaning that the robust matching is *more likely* to have a ΔOPT near the mean non-robust value. This is exactly the desired behavior: robustness produces more predictable and less varied results. In this application robustness exceeds expectations: the robust approach achieves a lower variance, *and* slightly improves nominal performance.

6 Robustness as Fairness

Balancing efficiency and fairness is a classic economic problem; recently, a body of literature covering fairness in kidney exchange has developed in the AI/Economics (Dickerson, Procaccia, and Sandholm 2014; McElfresh and Dickerson 2018; Ashlagi, Jaillet, and Manshadi 2013; Ding et al. 2018) and medical ethics (Gentry, Segev, and Montgomery 2005) communities; Appendix C presents a more thorough discussion. Though seemingly unrelated, fairness and robustness share a key characteristic: the balance between two competing properties. Fairness rules in kidney exchange often mediate between a fair and efficient outcome, using a parameter to set the balance. Similarly, robustness mediates between a good nominal outcome with the worst-case outcome, using an uncertainty budget or protection level to set that balance.

[3]All experiments were implemented in Python and used Gurobi (Gurobi Optimization, Inc. 2018), a state-of-the-art industrial combinatorial optimization toolkit, as a sub-solver. Our code is available on GitHub: https://github.com/duncanmcelfresh/RobustKidneyExchange.

[4]This is slightly more conservative than the notion of uncertainty introduced previously; in Section 4, *up to* Γ edges may fail, while in the experiments *exactly* Γ edges fail.

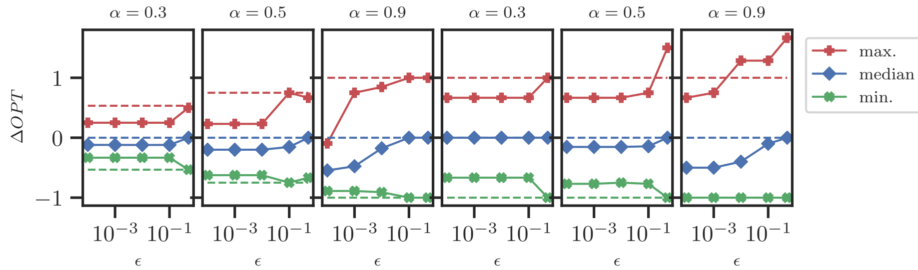

Figure 2: ΔOPT for non-robust (dashed lines) and edge weight robust (solid lines) matchings, for 64-vertex simulated exchange graphs (3 left plots) and real UNOS exchanges (3 right plots).

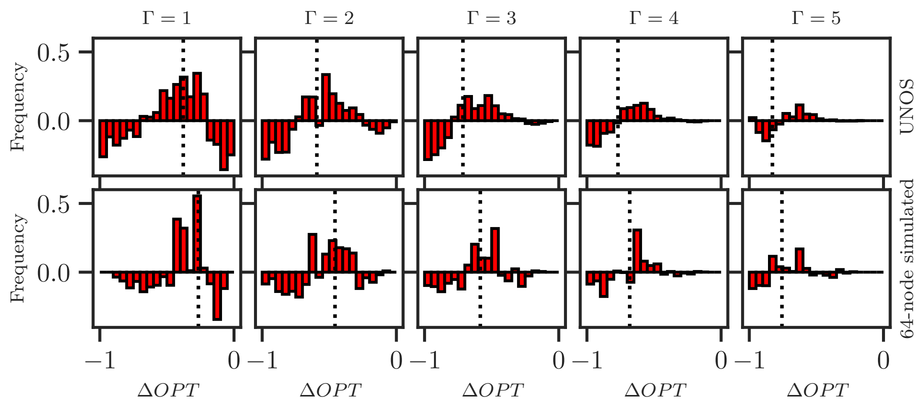

Figure 3: Difference between the robust and non-robust histograms of ΔOPT (robust minus non-robust) for real UNOS (top) and simulated exchanges (bottom), for various Γ. Dotted line: mean ΔOPT for non-robust.

In kidney exchange, fairness often refers to the prioritization of pediatric or disadvantaged (*highly-sensitized*) patients, who are unlikely to find a compatible donor. In the weighted fairness approach, edges that represent transplants to prioritized patients receive additional edge weight, making them more likely to be matched by standard algorithms; versions of this prioritization scheme are used by most exchanges, including UNOS. To generalize weighted fairness, let each edge have a *priority weight* $\hat{w}_e \in [0, \infty)$, equal to the nominal weight w_e multiplied by a factor $(1 + \alpha_e)$, with $\alpha_e \in [-1, \infty)$. For example, we might set $\alpha_e > 0$ for all edges *into* prioritized patients; this will help prioritized patients, but will likely lower overall efficiency (a tradeoff often described as the *price of fairness* (Caragiannis et al. 2009; Bertsimas, Farias, and Trichakis 2011; Dickerson, Procaccia, and Sandholm 2014; McElfresh and Dickerson 2018)).

To balance fairness with efficiency, policymakers limit the degree of prioritization. Let $\mathcal{P}_\Gamma$ be a *budgeted prioritization set*, which bounds the sum of absolute differences between each w_e and $\hat{w}_e$; this prioritization set is given as:

$$\mathcal{P}_\Gamma = \left\{ \hat{\mathbf{w}} \mid \hat{w}_e = w_e(1 + \alpha_e), \alpha_e \geq -1, \sum_{e \in E} \alpha_e w_e \leq \Gamma \right\}$$

As with edge weight uncertainty, the budget Γ balances between fairness and efficiency. If Γ is large, the algorithm might sacrifice matching size in order to match prioritized patients—but the maximum amount of efficiency sacrificed will be predictable, given Γ, which is attractive to policymakers. In Appendix C we further develop this concept, propose fairness rules that use $\mathcal{P}_\Gamma$, and present some theoretical results regarding the balance between fairness and efficiency.

7 Conclusions & Future Research

In this paper, we presented the first *scalable* robust formulations of kidney exchange. Our methods address both uncertainty over the *quality* and the *existence* of a potential transplant. On real and simulated data from a large, fielded kidney exchange, we showed that our methods (i) clear the market within seconds and (ii) result in more predictable and better quality matchings than the status quo.

Adapting automated ethical decision-making frameworks that aggregate noisy human value judgments (Noothigattu et al. 2018; Freedman et al. 2018; Bonnefon, Shariff, and Rahwan 2016) into our robust formulation is a natural way to handle uncertainty in the weights determined by a committee of stakeholders. Approaching *dynamic* kidney exchange, where participants arrive and depart over time, via robust reinforcement learning methods would be fruitful (Lim, Xu, and Mannor 2013; Xu and Mannor 2010).

References

Abraham, D.; Blum, A.; and Sandholm, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 295–304.

Anderson, R.; Ashlagi, I.; Gamarnik, D.; and Roth, A. E. 2015. Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences* 112(3):663–668.

Ashlagi, I.; Gamarnik, D.; Rees, M.; and Roth, A. E. 2012. The need for (long) chains in kidney exchange. NBER Working Paper No. 18202.

Ashlagi, I.; Jaillet, P.; and Manshadi, V. H. 2013. Kidney exchange in dynamic sparse heterogenous pools. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 25–26.

Balas, E. 1989. The prize collecting traveling salesman problem. *Networks* 19(6):621–636.

Ben-Tal, A.; El Ghaoui, L.; and Nemirovski, A. 2009. *Robust Optimization*. Princeton University Press.

Bertsimas, D., and Sim, M. 2004. The price of robustness. *Operations Research* 52(1):35–53.

Bertsimas, D.; Brown, D. B.; and Caramanis, C. 2011. Theory and applications of robust optimization. *SIAM Review* 53(3):464–501.

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

Biró, P.; Burnapp, L.; Haase, B.; Hemke, A.; Johnson, R.; van de Klundert, J.; and Manlove, D. 2017. Kidney exchange practices in Europe. First Handbook of the COST Action CA15210: European Network for Collaboration on Kidney Exchange Programmes.

Biró, P.; Manlove, D. F.; and Rizzi, R. 2009. Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Mathematics, Algorithms and Applications* 1(04):499–517.

Bonnefon, J.-F.; Shariff, A.; and Rahwan, I. 2016. The social dilemma of autonomous vehicles. *Science* 352(6293):1573–1576.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. The efficiency of fair division. International Workshop on Internet and Network Economics (WINE).

Chen, Y.; Li, Y.; Kalbfleisch, J. D.; Zhou, Y.; Leichtman, A.; and Song, P. X.-K. 2012. Graph-based optimization algorithm and software on kidney exchanges. *IEEE Transactions on Biomedical Engineering* 59:1985–1991.

Chen, R. S.; Lucier, B.; Singer, Y.; and Syrgkanis, V. 2017. Robust optimization for non-convex objectives. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 4708–4717.

Dickerson, J. P.; Manlove, D.; Plaut, B.; Sandholm, T.; and Trimble, J. 2016. Position-indexed formulations for kidney exchange. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 711–718.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2014. Price of fairness in kidney exchange. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1013–1020.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2018. Failure-aware kidney exchange. *Management Science*. To appear; earlier version appeared at EC-13.

Ding, Y.; Ge, D.; He, S.; and Ryan, C. T. 2018. A non-asymptotic approach to analyzing kidney exchange graphs. *Operations Research*. To appear; earlier version appeared at EC-15.

Freedman, R.; Schaich Borg, J.; Sinnott-Armstrong, W.; Dickerson, J.; and Conitzer, V. 2018. Adapting a kidney exchange algorithm to align with human values. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Gentry, S.; Segev, D.; and Montgomery, R. A. 2005. A comparison of populations served by kidney paired donation and list paired donation. *American Journal of Transplantation* 5(8):1914–1921.

Glorie, K. M. 2012. Estimating the probability of positive crossmatch after negative virtual crossmatch. Technical report, Erasmus School of Economics.

Glorie, K. 2014. *Clearing barter exchange markets: Kidney exchange and beyond*. PhD dissertation, Erasmus University Rotterdam.

Gurobi Optimization, Inc. 2018. Gurobi reference manual.

Klimentova, X.; Pedroso, J. P.; and Viana, A. 2016. Maximising expectation of the number of transplants in kidney exchange programmes. *Computers & Operations Research* 73:1–11.

Lim, S. H.; Xu, H.; and Mannor, S. 2013. Reinforcement learning in robust markov decision processes. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.

Manlove, D., and O'Malley, G. 2015. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM Journal of Experimental Algorithmics* 19(1).

McElfresh, D. C., and Dickerson, J. P. 2018. Balancing lexicographic fairness and a utilitarian objective with application to kidney exchange. *AAAI Conference on Artificial Intelligence (AAAI)*.

Mevissen, M.; Ragnoli, E.; and Yu, J. Y. 2013. Data-driven distributionally robust polynomial optimization. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 37–45.

Noothigattu, R.; Gaikwad, S. N. S.; Awad, E.; D'Souza, S.; Rahwan, I.; Ravikumar, P.; and Procaccia, A. D. 2018. A voting-based system for ethical decision making. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Petrik, M., and Subramanian, D. 2014. RAAM: The benefits of robustness in approximating aggregated MDPs in reinforcement learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 1979–1987.

Poss, M. 2014. Robust combinatorial optimization with variable cost uncertainty. *European Journal of Operational Research* 237(3):836–845.

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

Roth, A.; Sönmez, T.; and Ünver, U. 2004. Kidney exchange. *Quarterly Journal of Economics* 119(2):457–488.

Roth, A.; Sönmez, T.; and Ünver, U. 2005. A kidney exchange clearinghouse in New England. *American Economic Review* 95(2):376–380.

UNOS. 2015. Revising kidney paired donation pilot program priority points. OPTN/UNOS Public Comment Proposal.

Xu, H., and Mannor, S. 2010. Distributionally robust Markov Decision Processes. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*.

Xu, H.; Caramanis, C.; and Mannor, S. 2009. Robust regression and LASSO. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 1801–1808.