

UCLA

UCLA Previously Published Works

Title

Fairness-Aware Kidney Exchange and Kidney Paired Donation

Permalink

<https://escholarship.org/uc/item/03g771qt>

Journal

Statistics in Biosciences

ISSN

1867-1764

Authors

Mingrui, Zhang

Xiaowu, Dai

Lexin, Li

Publication Date

2026

DOI

10.1007/s12561-026-09514-y

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at

<https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



Fairness-Aware Kidney Exchange and Kidney Paired Donation

Zhang Mingrui¹ · Dai Xiaowu² · Li Lexin¹ 

Received: 18 March 2025 / Revised: 24 December 2025 / Accepted: 17 January 2026
© The Author(s) under exclusive licence to International Chinese Statistical Association 2026

Abstract

The kidney paired donation (KPD) program provides an innovative solution to overcome incompatibility challenges in kidney transplants by matching incompatible donor-patient pairs and facilitating kidney exchanges. To address unequal access to transplant opportunities, there are two widely used fairness criteria: group fairness and individual fairness. However, these criteria do not consider protected patient features, which refer to characteristics legally or ethically recognized as needing protection from discrimination, such as race and gender. Motivated by the calibration principle in machine learning, we introduce a new fairness criterion: the matching outcome should be conditionally independent of the protected feature, given the sensitization level. We integrate this fairness criterion as a constraint within the KPD optimization framework and propose a computationally efficient solution using linearization strategies and column-generation methods. Theoretically, we analyze the associated price of fairness using random graph models. Empirically, we compare our fairness criterion with group fairness and individual fairness through both simulations and a real-data example.

Keywords Calibration · Column generation · Integer program · Kidney paired donation · Price of fairness · Random graph

✉ Li Lexin
lexinli@berkeley.edu

Zhang Mingrui
mingrui_zhang@berkeley.com

Dai Xiaowu
dai@stat.ucla.edu

¹ Division of Biostatistics, University of California, Berkeley, 2121 Berkeley Way, Berkeley, CA 94704, USA

² Department of Statistics and Data Science, University of California, Los Angeles, 520 Portola Plaza, Los Angeles, CA 90095, USA

1 Introduction

1.1 Kidney Paired Donation Programs

Kidney transplantation is the preferred treatment for end-stage renal disease (ESRD), offering significant improvements in both quality of life and survival compared to dialysis. However, a major obstacle is the incompatibility between donors and patients, often due to mismatches in blood type or human leukocyte antigens (HLA). According to the United Network for Organ Sharing (UNOS) and the Organ Procurement and Transplantation Network (OPTN), over 90,000 patients were on the kidney transplant waiting list at the end of 2023.

To overcome these incompatibility challenges, kidney paired donation (KPD) programs have been developed as an innovative solution. These programs match incompatible donor-patient pairs using a *virtual crossmatch*, a preliminary compatibility test, and facilitate kidney exchanges, allowing patients to receive kidneys from compatible donors through mutual exchange. These exchanges are called *exchange cycles* or simply *cycles*, with formal definitions provided in Sect. 2. The primary goal of KPD programs is to maximize the number of successful transplants or optimize generalized utilities based on predicted transplantation or survival outcomes.

Despite their promise, planned cycles may fail for various reasons, such as illness, pregnancy, or the death of a patient or donor, scheduling conflicts, or discrepancies between *virtual* and *laboratory crossmatch* results. These uncertainties necessitate the consideration of *recourse* strategies, which identify alternative transplant opportunities within the original cycle. Even when an entire cycle cannot proceed, smaller unaffected sub-cycles may still be viable. By accounting for these uncertainties, KPD programs can further maximize the expected number of successful transplants or optimize the expectation of some general utilities.

Following the framework of [1], KPD programs adopt three main recourse strategies. The first is *no recourse*, where a cycle either proceeds fully or fails entirely [e.g., 2]. The second, *internal recourse*, identifies the sub-cycle with the highest utility among those unaffected by the failure [e.g., 3]. The third, *subset recourse*, considers broader subsets that may include multiple cycles, enabling alternative arrangements when uncertainties arise [e.g., 4]. Each strategy computes expected utilities accordingly, with KPD programs aiming to maximize these expected utilities.

In addition to recourse strategies, we can integrate uncertainties into the optimization framework through other approaches, such as look-ahead strategy [5] and robust optimization [6].

1.2 Fairness Concerns in KPD Programs

Despite the success of KPD programs, important fairness concerns arise, particularly regarding unequal access to transplant opportunities. These disparities stem from two main factors. The first factor is differences in patients' HLA sensitization levels. A patient's HLA sensitization level is measured by their panel-reactive

antibody (PRA) score, which reflects the likelihood of HLA incompatibility with a random donor. Patients with high PRA scores, referred to as highly sensitized patients, are more difficult to match and, therefore, have fewer transplant opportunities. In contrast, patients with low PRA scores, known as lowly sensitized patients, are easier to match and benefit from more transplant opportunities. The second factor is asymmetric blood type compatibility. According to standard ABO blood type compatibility rules, patients with blood type O are harder to match because they can only receive kidneys from O-type donors. Conversely, patients with blood type AB are the easiest to match, as they can receive kidneys from donors of any blood type.

To address these disparities, KPD programs can incorporate fairness constraints to reduce unfairness caused by differences in HLA sensitization and blood type. Two widely used fairness criteria in KPD programs are group fairness and individual fairness. Group fairness focuses on ensuring that highly sensitized patients receive equitable consideration relative to lowly sensitized patients [6–8]. In contrast, individual fairness aims to provide balanced selection chances for each patient, ensuring that no one is unfairly disadvantaged [9]. While these two fairness criteria are central to addressing disparities in kidney exchange, other approaches have also been explored. For example, [10] incorporate the Nash standard of comparison (or proportional fairness) and Rawlsian justice principles. [11, 12], and [13] draw on game theory to address fairness within utility-maximization frameworks, considering the interests of stakeholders such as hospitals and regions.

In the context of fair machine learning, a *protected feature* (or sensitive attribute) refers to a characteristic legally or ethically recognized as needing protection from discrimination. Our key question is how to establish a fairness criterion that ensures equal access to transplant opportunities across patient groups defined by protected characteristics—and how to achieve this in practice. This specific focus has not yet been explored in the KPD literature, as existing fairness criteria do not account for protected patient features.

Some protected features, such as race and gender, are associated with differences in sensitization levels and blood types. For instance, studies have shown that parous women are more likely to develop high sensitization to HLA antigens [14], making them less compatible with most donors in a KPD program and harder to match. This leads to unequal access to transplant opportunities between females and males. A simplistic approach might aim to balance overall selection rates between genders, but this could inadvertently disadvantage highly sensitized male patients, as a higher number of highly sensitized female patients would need to be matched to achieve gender balance. A more equitable approach would balance selection rates within subgroups, such as highly sensitized females versus males and lowly sensitized females versus males.

Motivated by this example, we propose a new fairness criterion: the matching outcome should be conditionally independent of the protected feature, given the sensitization level. The randomness associated with this fairness criterion is determined by a randomization policy, as proposed for individual fairness in kidney exchange [9, 10] and general matching problems [15, 16]. This approach provides guarantees for average selection rates within protected groups across each sensitization level.

1.3 Fairness in General Decision-Making Problems

Our fairness criterion in KPD programs is defined based on the conditional outcome given protected features, drawing on similar concepts from the literature on general decision-making problems.

Demographic parity ensures fairness by requiring that the rate of positive decisions is consistent across groups defined by protected features, promoting equality in outcomes regardless of group membership. *Equalized odds*, introduced by [17], aligns predictive performance such that the false positive and false negative rates are similar across groups, leading to a fair distribution of errors. *Predictive parity*, discussed by [18], ensures fairness by equalizing the positive predictive value (PPV) across groups, thereby making positive predictions equally reliable and trustworthy for all groups. *Calibration within groups*, explored by [19], requires that individuals with the same predicted probability have consistent actual outcome rates across groups, ensuring well-calibrated predictions.

These fairness concepts emphasize different priorities: overall outcome equality, error distribution, prediction reliability, or probability calibration. They are frequently incorporated as fairness constraints in statistical optimization problems [e.g., 20].

Our fairness criterion in KPD programs is closely aligned with calibration within groups; see a more detailed discussion in Sect. 3.1.

1.4 Our Contributions

This paper makes several contributions to the field of kidney exchange and fairness in allocation. First, we propose a new fairness criterion based on a protected feature, which has not been explored in the kidney exchange literature. We integrate this fairness criterion as a constraint within the optimization framework commonly used in kidney paired donation (KPD) programs. This flexible structure can accommodate other fairness criteria and potential recourse strategies. Furthermore, we propose a computationally efficient solution to the resulting optimization problem, achieving scalability by avoiding explicit enumeration of all exchange plans and instead relying on an iterative cutting-plane procedure.

Second, we investigate the *price of fairness* associated with our proposed criterion, defined as the relative loss in system efficiency when a fair allocation is prioritized over an optimal (unconstrained) allocation [21]. Theoretically, we derive an upper bound on the asymptotic price of fairness using random graph models that incorporate ABO blood type distributions. Empirically, through simulation studies, we show that the efficiency loss from implementing our fairness criterion is relatively low.

Our findings align with prior studies on the tradeoff between efficiency and fairness in resource allocation. For example, [7] examine the efficiency loss associated

with group fairness in kidney exchange, while [11] analyze the price of ensuring individual rationality in multi-hospital kidney exchanges, both employing random graph models with ABO blood types. [10] utilize the Nash Social Welfare Program to address the tradeoff between fairness and efficiency. Similarly, [22] propose a framework for fair policy targeting that balances fairness and efficiency using Pareto optimal treatment allocation rules, offering theoretical guarantees and practical solutions applicable to social welfare contexts.

2 Review of KPD Program in an Optimization Framework

In this section, we review the notations, terminologies, and optimization framework used in KPD programs. Section 2.1 focuses on the classical optimization problem without incorporating fairness constraints, while Sect. 2.2 reviews group fairness and individual fairness, introducing an additional constraint to integrate fairness into the framework.

2.1 KPD Program Without Fairness

2.1.1 Classical Formulation

We represent a KPD program as a directed graph $G = (V, E)$, where $V = \{v_1, \dots, v_N\}$ denotes the vertex set and E denotes the edge set. The vertex set V is the set of N incompatible donor-patient pairs. The edge set E consists of all ordered pairs (v_i, v_j) that the donor in vertex $v_i \in V$ is compatible with the patient in vertex $v_j \in V$. An *exchange cycle*, or simply *cycle*, is defined as a sequence of distinct incompatible donor-patient pairs. We denote a cycle c as an ordered sequence of vertices $[c_1, \dots, c_{|c|}]$ in V , where $|c|$ is the cycle length of c , satisfying that $(c_i, c_{i+1}) \in E$ for $1 \leq i \leq |c| - 1$, and $(c_{|c|}, c_1) \in E$. To execute the cycle, the patient in c_{i+1} will receive the donor kidney of c_i for $1 \leq i \leq |c| - 1$, and the patient c_1 will receive the donor kidney of $c_{|c|}$. Figure 1 illustrates how a two-way cycle and a three-way cycle work in KPD programs. An *exchange plan* is a collection of vertex-disjoint cycles in the graph.




Fig. 1 Illustration of exchange cycles shown with solid arrows. Transplantations along the dashed arrows cannot proceed due to incompatibility

In fielded KPD programs, we usually restrict cycles to be no longer than 3. We use $\mathcal{C}$ to denote the set of all cycles under the restrictions on the length of the cycles. We assign a utility u_{v_i, v_j} to each edge $(v_i, v_j) \in E$. Based on the edge utility, we define the cycle utility $u(c) = \sum_{i=1}^{|c|-1} u_{c_i, c_{i+1}} + u_{c_{|c|}, c_1}$. In a classical KPD program, we aim to select a disjoint packing of cycles in $\mathcal{C}$ with the maximum sum of utilities, which can be formulated as the following integer program:

$$\max_{x_c \in \{0,1\}} \sum_{c \in \mathcal{C}} x_c u(c) \quad \text{subject to} \quad \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \leq 1, \quad \forall v \in V. \quad (1)$$

The binary variable x_c indicates whether c is selected or not. The objective function in (1) is the sum of utilities of the selected cycles. The constraint in (1) requires that the selected cycles are vertex-disjoint.

2.1.2 Extension: Incorporating Recourse Strategies

To account for uncertainty, we consider the expected utilities in (1), depending on the chosen recourse strategy. Different strategies lead to different calculations of expected utilities.

No-recourse strategy

Under the no-recourse strategy, the expected utility is determined without adapting to failures. Given the failure probabilities of each vertex in V and each edge in E , we can compute the probability of each cycle in $\mathcal{C}$ being executable, where all vertices and all edges in the cycle do not fail. An explicit formula is available to compute the expected utility for this strategy [1, 2].

Internal-recourse strategy

The internal-recourse strategy considers adapting to failures within a given cycle. For a cycle c , let $\mathcal{M}(c)$ represent the set of all exchange plans involving vertices in c . Without loss of generality, assume the elements of $\mathcal{M}(c)$ are ordered by descending utility. Specifically, the utility of $M_i(c)$, the i -th element in $\mathcal{M}(c)$, is no less than that of $M_j(c)$, the j -th element, for $i < j$. If all vertices and edges in $M_1(c)$ do not fail, $M_1(c)$ will be executed. If any vertices or edges in $M_k(c)$ fail but those in $M_{k+1}(c)$ do not fail, $M_{k+1}(c)$ will be executed ($k \geq 1$). While there is no explicit formula for this strategy, efficient algorithms are available for computing the expected utility [1–3].

Subset-recourse strategy

The subset-recourse strategy expands the search space to disjoint *relevant subsets* rather than disjoint cycles, providing more flexibility in adapting to uncertainty. Formally, a relevant subset of size (k, q) is the set of at most $(k + q)$ vertices in graph G inducing a strongly connected subgraph, such that any edge of the paths that provide the strong connectivity belongs to some cycle of size at most k , whose vertices are in the subset. We follow the same approach as the internal-recourse strategy to compute the expected utility of each relevant subset. To solve the optimization problem in (1), we must enumerate the new set $\mathcal{C}$, with algorithms given in [1] and [23].

2.2 KPD Program with Fairness Constraint

2.2.1 Group Fairness

The group fairness aims to ensure fairness towards highly sensitized patients. We partition the vertex set V into $V_h \cup V_e$, where V_h denotes the set of incompatible donor-patient pairs with highly sensitized patients, and V_e denotes the set of incompatible donor-patient pairs with lowly sensitized patients. Following [7], we formulate the group fairness as a constraint in the optimization problem below

$$\max_{x_c \in \{0,1\}} \sum_{c \in \mathcal{C}} x_c u(c) \quad \text{subject to} \quad \sum_{v \in V_h} \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \geq \alpha, \quad \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \leq 1, \quad \forall v \in V. \quad (2)$$

Specifically, we consider a fairness constraint that the number of highly sensitized patients involved in the matching is no less than some threshold α .

2.2.2 Individual Fairness

The individual fairness utilizes a randomization policy to ensure that every patient has a similar chance of being matched. Let $\mathcal{F} \subseteq 2^{\mathcal{C}}$ be the set of possible exchange plans that consist of disjoint cycles in $\mathcal{C}$. We assign a probability distribution δ on $\mathcal{F}$, which indicates the probability of selecting each exchange plan in $\mathcal{F}$. Based on δ , let x_c be the overall probability that the cycle $c \in \mathcal{C}$ is selected as one of the disjoint subsets. Further, we can compute $\pi_v = \sum_{c \in \mathcal{C}} x_c 1_{v \in c}$, which is the probability of each vertex v to be selected. Let $\pi = (\pi_v)_{v \in V} \in \mathbb{R}^N$ be the vector of selection probability of all patients and let $\bar{\pi} = 1^T \pi / N$ be the average selection probability of all patients. The quantity π_v is central in defining individual fairness and our new fairness criterion in Sect. 3.

Recall that x_c in (1) and (2) is a binary variable indicating whether c is selected or not, but here x_c is a continuous variable bounded between 0 and 1. Following [9], we formulate the individual fairness as a constraint in the optimization problem below

$$\max_{\delta} \sum_{c \in \mathcal{C}} x_c u(c) \quad \text{subject to} \quad \sum_{v \in V} |\pi_v - \bar{\pi}|^p \geq \beta^p. \quad (3)$$

Specifically, the individual fairness promotes a similar chance of being selected for each patient. We choose the L_p norm of the vector $(\pi - \bar{\pi})$ to measure the variation of the selection probability among all the patients. We consider a fairness constraint that the L_p norm of the vector $(\pi_v - \bar{\pi})$ is no less than some threshold β .

3 A new Fairness Criterion Based on a Protected Feature

3.1 Fairness Formulation

In this section, we introduce a new fairness criterion based on a protected feature A of patients in V , motivated by the concept of *calibration within group* in

the context of machine learning. In essence, our new fairness criterion requires that the probability of a patient being selected for a transplant (the matching outcome) should not depend on the protected feature, such as race or gender, once we account for the patient's sensitization level. This means that, within each sensitization category, patients with different protected attributes should have similar chances of being matched.

In binary classification, a score function $R(X)$ satisfies calibration within groups if $\Pr(Y = 1 \mid R(X) = r, A = a) = r$ for all score values r and group level a , where Y is the outcome variable and X is the feature variables. A slightly weaker definition [24] only requires that Y is conditionally independent of A given the score $R(X)$. Our fairness definition is based on this weaker definition. In KPD programs, we view Y as the selection indicator and R as the sensitization level of a patient. Therefore, we define our new fairness criterion as satisfying that the selection indicator is conditionally independent of the protected feature given the sensitization level. In other words, at each sensitization level, the matching outcome is independent of the protected feature. We can view this fairness criterion in KPD programs as a reverse problem of that in machine learning. In machine learning, the randomness is due to the underlying population model, and our goal is to construct a score $R(X)$ satisfying the fairness condition. In KPD programs, we need to determine the randomness such that the observed sensitization level satisfies the fairness condition.

For simplicity of presentation, we assume the protected feature A is binary with two levels $\{0, 1\}$ and the sensitization R has M levels $\{r_1, r_2, \dots, r_M\}$. In practice, it is common that $M = 2$, where we partition all patients into highly sensitized and lowly sensitized patients, or $M = 3$, where we partition all patients into highly sensitized, moderately sensitized and lowly sensitized patients. The variables A and R partition all patients into $2M$ subgroups, denoted as $\{V_{ij}\}_{0 \leq i \leq 1, 1 \leq j \leq M}$, where $V_{ij} = \{v \in V : A(v) = i, R(v) = r_j\}$. As in Sect. 2.2.2, we assign a probability distribution δ on the set of exchange plans $\mathcal{F}$, which indicates the probability of selecting each exchange plan in $\mathcal{F}$. Moreover, π_v is the probability of each vertex v to be selected in the end, and thus $|V_{ij}|^{-1} \sum_{v \in V_{ij}} \pi_v$ is average selection rate in V_{ij} . Given the sensitization level $R = r_j$, our fairness criterion restricts the average selection rates in V_{0j} and V_{1j} to be close. It is natural to impose constraints that the absolute difference in these average selection rates is bound by a constant l_j . Therefore, we formulate the problem under our fairness criterion as

$$\max_{\delta} \sum_{c \in \mathcal{C}} x_c u(c) \quad \text{subject to} \quad \left| |V_{0j}|^{-1} \sum_{v \in V_{0j}} \pi_v - |V_{1j}|^{-1} \sum_{v \in V_{1j}} \pi_v \right| \leq l_j, \quad \forall 1 \leq j \leq M. \quad (4)$$

For whichever choice of parameters $l_1, \dots, l_M \geq 0$, there always exists a solution to (4). The strength of the fairness constraint depends on the parameters $l_1, \dots, l_M$. In practice, we suggest two candidate values for l_j 's: $l_j = 0.5 / \min\{|V_{1j}|, |V_{2j}|\}$ and $l_j = 0.5 / \max\{|V_{1j}|, |V_{2j}|\}$. We can interpret these two candidate values as the desired precision based on the larger or smaller subgroup. The former represents a weaker fairness constraint, and the latter represents a stronger one.

3.2 Algorithm for Solving (4)

This subsection presents a linear optimization procedure for solving (4) without enumerating all exchange plans. The approach follows a master-subproblem structure similar to column generation. The master problem optimizes over the marginal probabilities of selecting each cycle and each vertex, while an auxiliary subproblem verifies whether the proposed solution is implementable and, if necessary, adds new linear constraints that refine the master problem.

Specifically, we express (4) as an optimization problem in terms of the cycle-selection variables x_c , rather than the probability distribution δ over exchange plans. Recall that the selection probability of each vertex v is

$$\pi_v = \sum_{c \in \mathcal{C}} 1_{v \in c} x_c.$$

Both the objective and the fairness constraints are linear in x_c . A direct constraint of variables x_c is $0 \leq x_c \leq 1$ for all $c \in \mathcal{C}$ and $0 \leq \pi_v \leq 1$ for all $v \in V$. However, not every such vector $x = (x_c)_{c \in \mathcal{C}}$ satisfying the above constraint can be realized by a valid probability distribution over disjoint exchange plans.

To enforce implementability, we introduce a scalar variable $\eta \in \mathbb{R}$ that represents the best achievable expected utility at the current selection profile and iteratively add linear inequalities, called *cuts*, to refine the feasible region. Let $\mathcal{B}^{\text{opt}}$ and $\mathcal{B}^{\text{feas}}$ denote the sets of indices for optimality and feasibility cuts generated so far; both sets start empty and expand monotonically. The restricted master problem is formulated as

$$\begin{aligned} & \max_{x, \pi, \eta} \quad \eta \\ & \text{subject to} \quad \pi_v = \sum_{c \in \mathcal{C}} 1_{v \in c} x_c \leq 1, \quad \forall v \in V, \\ & \quad 0 \leq x_c \leq 1, \quad \forall c \in \mathcal{C}, \\ & \quad \left| |V_{0j}|^{-1} \sum_{v \in V_{0j}} \pi_v - |V_{1j}|^{-1} \sum_{v \in V_{1j}} \pi_v \right| \leq l_j, \quad j = 1, \dots, M, \\ & \quad \pi^\top \theta^{(\ell)} + \beta^{(\ell)} \geq \eta, \quad \forall \ell \in \mathcal{B}^{\text{opt}}, \\ & \quad \pi^\top \hat{\theta}^{(k)} + \hat{\beta}^{(k)} \leq 0, \quad \forall k \in \mathcal{B}^{\text{feas}}, \end{aligned} \quad (\text{MP})$$

where $(\theta^{(\ell)}, \beta^{(\ell)}) \in \mathbb{R}^N \times \mathbb{R}$ and $(\hat{\theta}^{(k)}, \hat{\beta}^{(k)}) \in \mathbb{R}^N \times \mathbb{R}$ are coefficients provided by the auxiliary subproblem.

Given a candidate π from (MP), the auxiliary subproblem checks whether π can be implemented by a valid distribution over exchange plans and determines the best achievable expected utility under that profile. For each plan $F_t \in \mathcal{F}$, let $u(F_t)$ denote its total utility and let $\delta(F_t)$ denote the probability that plan F_t is selected. The auxiliary configuration problem is

$$\begin{aligned}
 & \max_{\delta} \quad \sum_{F_t \in \mathcal{F}} u(F_t) \delta(F_t) \\
 & \text{subject to} \quad \sum_{F_t \in \mathcal{F}} 1_{v \in F_t} \delta(F_t) = \pi_v, \quad \forall v \in V, \\
 & \quad \sum_{F_t \in \mathcal{F}} \delta(F_t) = 1, \quad \delta(F_t) \geq 0, \quad \forall F_t \in \mathcal{F},
 \end{aligned} \tag{SP}$$

with the corresponding dual formulation

$$\begin{aligned}
 & \min_{\theta \in \mathbb{R}^N, \beta \in \mathbb{R}} \quad \pi^\top \theta + \beta \\
 & \text{subject to} \quad \beta + \sum_{v \in V} \theta_v 1_{v \in F_t} \geq u(F_t), \quad \forall F_t \in \mathcal{F}.
 \end{aligned} \tag{SP-D}$$

In practice, we solve the dual problem (SP-D) rather than the primal (SP), since the primal has one variable for each exchange plan and cannot be handled directly. The dual problem is of manageable size and can be solved efficiently using column generation. The feasibility of (SP) is inferred from the behavior of (SP-D): if the dual is bounded, then (SP) is feasible, and if the dual is unbounded, then (SP) is infeasible.

If (SP) is feasible, strong duality ensures that the optimal values of (SP) and (SP-D) coincide. The resulting optimal dual pair (θ, β) defines a valid optimality cut $\pi^\top \theta + \beta \geq \eta$, which is added to (MP). If (SP) is infeasible, then, by Farkas' lemma, there exists a certificate $(\hat{\theta}, \hat{\beta})$ such that $\hat{\beta} + \sum_{v \in V} \hat{\theta}_v 1_{v \in F_t} \leq 0$ for all $F_t \in \mathcal{F}$ and $\pi^\top \hat{\theta} + \hat{\beta} > 0$. The separating inequality $\pi^\top \hat{\theta} + \hat{\beta} \leq 0$ is then added as a feasibility cut.

The dual problem (SP-D) contains one constraint for each feasible plan F_t , which makes it impossible to construct explicitly. Column generation avoids this by starting with a small subset of constraints and iteratively adding only those that are necessary for optimality. At each iteration, the algorithm identifies the most violated dual constraint under the current dual solution (θ, β) .

Formally, given (θ, β) , we search for a plan F_t whose dual constraint

$$\beta + \sum_{v \in V} \theta_v 1_{v \in F_t} \geq u(F_t)$$

is most violated. This search is equivalent to solving a maximum-weight packing problem over vertex-disjoint cycles:

$$\max_{y_c \in \{0,1\}} \sum_{c \in \mathcal{C}} \left(u(c) - \sum_{v \in c} \theta_v \right) y_c - \beta \quad \text{subject to} \quad \sum_{c \in \mathcal{C}} y_c 1_{v \in c} \leq 1, \quad \forall v \in V.$$

Each binary variable y_c indicates whether cycle c is selected, and the constraint ensures that selected cycles are vertex-disjoint. If the optimal value of this problem is positive, the corresponding plan violates a dual constraint and is added to a restricted version of (SP-D). If the optimal value is nonpositive, all constraints in the full dual problem are satisfied. Because (θ, β) is already optimal for the restricted dual problem, and adding satisfied constraints cannot further decrease its objective value, (θ, β) must also be optimal for the full dual (SP-D).

The algorithm alternates between solving the master problem (MP) and invoking the auxiliary subproblem. Each iteration produces a candidate (x, π, η) that satisfies the fairness constraints in (4). The auxiliary subproblem then determines whether this profile π is implementable. If it is, the resulting dual solution provides an optimality cut that tightens the objective bound η in (MP); otherwise, a feasibility cut removes the unattainable profile π from the feasible region.

The process continues until no violated constraint can be found by the auxiliary subproblem. In practice, this occurs when the optimal value of the separation problem in the column-generation step is nonpositive, indicating that all dual constraints are satisfied. At this point, the restricted master problem is equivalent to the full master problem, and no further cuts or columns can improve the solution.

From a theoretical standpoint, each iteration introduces a valid linear inequality that progressively refines the feasible region, and only finitely many distinct cuts can be generated since the set of feasible plans $\mathcal{F}$ is finite. Hence, the procedure terminates after a finite number of iterations. Upon termination, the current π is guaranteed to be implementable by a valid distribution over exchange plans, the fairness constraints are satisfied, and the corresponding objective value η equals the optimum of (4). The final primal solution of (SP) explicitly provides this distribution, yielding an optimal randomized policy that achieves both fairness and efficiency.

3.3 Prediction of Individual Selection Probability

The selection probability π_v is a central quantity in defining both individual fairness and our new fairness criterion. Since a KPD program is not static but dynamic, a natural statistical question to ask is how to predict individual selection probabilities before more incompatible donor-patient pairs enter the pool for the next round of exchange allocation. We provide a solution based on sample splitting. The following discussion can accommodate any fairness criteria within our general framework.

Our method utilizes historical data of incompatible donor-patient pairs independent of the current and future pairs. Suppose the historical pool consists of incompatible donor-patient pairs denoted as $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$, and the current pool consists of pairs denoted as $\{v_1, \dots, v_{N_1}\}$. Assume that the exchange allocation occurs when the size of the incompatible pairs pool reaches N , where $N_1 < N < N_0 + N_1$. The historical pool $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$ aggregates incompatible donor-patient pairs collected across potentially multiple past KPD rounds rather than a single pool. In each iteration b , we sample a subset of these historical vertices to simulate a feasible past pool of size $N - N_1$. The edge set $E^{(b)}$ is then determined over the combined vertex set $V^{(b)} = \{v_1, \dots, v_{N_1}\} \cup \tilde{V}^{(b)}$, treating unobserved edges between vertices that never coexisted in historical rounds as absent. To construct $E^{(b)}$, we recompute donor-patient compatibilities among all vertices in $V^{(b)}$ using the same ABO and HLA matching criteria as in the real KPD graph. Each edge (v_i, v_j) is included in $E^{(b)}$ if donor v_i is compatible with patient v_j . The prediction procedure can be described in Algorithm 1 below.

Algorithm 1 Prediction procedure with selection probabilities

-
- 1: **Input:** Historical pool of vertex set $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$, current pool of vertex set $\{v_1, \dots, v_{N_1}\}$, and number of repetitions B .
 - 2: **Output:** Prediction of π_v for each $v \in \{v_1, \dots, v_{N_1}\}$.
 - 3: **for** $b = 1$ to B **do**
 - 4: Sample $\{\tilde{v}_1^{(b)}, \dots, \tilde{v}_{N-N_1}^{(b)}\}$ from $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$ without replacement.
 - 5: Determine the edge set $E^{(b)}$ in $V^{(b)} = \{v_1, \dots, v_{N_1}, \tilde{v}_1^{(b)}, \dots, \tilde{v}_{N-N_1}^{(b)}\}$.
 - 6: Based on the graph $(V^{(b)}, E^{(b)})$, solve (4) to obtain the selection probability of $\pi_v^{(b)}$ for each $v \in \{v_1, \dots, v_{N_1}\}$.
 - 7: **end for**
 - 8: **Return:** mean and quantiles of $\{\pi_v^{(1)}, \dots, \pi_v^{(B)}\}$ as the mean prediction and interval prediction of π_v , for each $v \in \{v_1, \dots, v_{N_1}\}$.
-

In practice, the computational cost can be high if we directly enumerate all cycles or relevant subsets within the simulated pool $\{v_1, \dots, v_{N_1}, \tilde{v}_1^{(b)}, \dots, \tilde{v}_{N-N_1}^{(b)}\}$ for each of the B replications (e.g., $B = 1000$). When the number of protected feature groups M is small, it is more efficient to first enumerate the feasible cycles or subsets once using the entire combined vertex set $\{v_1, \dots, v_{N_1}, \tilde{v}_1, \dots, \tilde{v}_{N_0}\}$, and then, for each replication b , retain only those cycles whose vertices appear in the sampled subset $\{v_1, \dots, v_{N_1}, \tilde{v}_1^{(b)}, \dots, \tilde{v}_{N-N_1}^{(b)}\}$. This strategy avoids repeated enumeration while preserving correctness. When N_0 is extremely large, however, full enumeration over all historical vertices may still be infeasible. In such cases, the historical pool $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$ can be partitioned into several disjoint subsets of manageable size, and the enumeration of feasible cycles or relevant subsets can be performed separately within each subset. The results are then combined across subsets as needed.

4 Price of our new Fairness Criterion Under Random Graph Models

The *price of our new fairness criterion* quantifies the efficiency loss incurred when enforcing fairness constraints compared to the unconstrained optimal allocation. Formally, it is defined as the relative reduction in the total utility (e.g., expected number of transplants or any general utility) achieved under our fairness-constrained optimization versus the utility achieved by the unconstrained solution. This concept follows the standard “price of fairness” framework in the literature [e.g., 21], here adapted specifically to our new fairness criterion. In this section, we establish theoretical guarantees for the price of our fairness criterion using a random graph model that incorporates ABO blood types. We describe the model assumptions below.

Recall that a donor and a patient are compatible if they match in both blood type and HLA. We assume blood type compatibility follows standard medical guidelines: AB patients can receive kidneys from donors of any blood type, A and B patients can receive from donors of their own type or type O, while O patients can only receive from type O donors. We assume HLA compatibility

follows binomial distributions. Specifically, we randomly assign each patient a PRA score, representing the probability of being HLA incompatible with any donor. We assume the PRA scores can take discrete values $\{r_1, \dots, r_M\}$, which also determine sensitization levels $\{r_1, \dots, r_M\}$; and we assume HLA compatibility between different donor-patient pairs is independent. The vertex set V is formed by independently drawing donor-patient pairs from an underlying population, keeping only incompatible pairs until a total of N incompatible pairs is reached. Let $\mu_{b_1, b_2, r, a}$ denote the probability of sampling an incompatible donor-patient pair with donor blood type b_1 , patient blood type b_2 , patient sensitization level r , and patient protected group level a . Here, the parameters $\{\mu_{b_1, b_2, r, a}\}$ are fixed but unknown. The edge set E is determined by the compatibility between donor-patient pairs in V . Regarding the utility assignment, we assume $u(c)$ only depends on the induced subgraph of vertices in c , independent of the protected feature of vertices in c . We also assume $u(c)/|c|$ is bounded by a constant.

In our random graph model, the PRA levels $\{r_1, \dots, r_M\}$ are treated as fixed parameters representing different sensitization categories rather than random draws from a continuous distribution. Specifically, each vertex v_i (representing a donor-patient pair) is assigned a PRA level r_i from this finite set according to empirical proportions observed in real KPD data [e.g., 25]. Hence, the PRA levels are not sampled uniformly over $[0, 1]$ but follow a categorical distribution reflecting realistic sensitization patterns.

In our theoretical analysis, we adopt a simplified random graph model where each edge is independently present with a constant probability $p \in (0, 1)$. This assumption is primarily technical and facilitates the use of classical random graph results [e.g., 26] to establish asymptotic matching properties. In real kidney exchange programs, however, the probability that an edge exists between two vertices depends on medical compatibility, particularly blood type and PRA sensitization level, so that the effective edge probabilities are heterogeneous across vertex pairs. Nevertheless, the assumption of a constant p captures the key structural property required in our proofs: the existence of a sufficiently dense graph where each vertex has an expected number of compatible partners growing linearly in N . Under any model where the minimal edge probability remains bounded away from zero across vertex types (as is empirically the case in our simulations), the asymptotic results and efficiency bounds derived here continue to hold qualitatively. Therefore, while the assumption of a constant p simplifies the theoretical exposition, it does not materially affect the validity of the conclusions regarding the asymptotic price of fairness.

Similar to group fairness that prioritizes highly sensitized patients, our new fairness criterion needs to prioritize some subgroups of patients based on the protected feature, and further balance their average selection probabilities. In Sect. 4.1, we derive a general result quantifying the efficiency loss due to subgroup prioritization in KPD programs. Specifically, we focus on prioritizing patients with either $A = 1$ or $A = 0$, given their blood type and sensitization level. In Sect. 4.2, we apply the result to establish theoretical guarantees for the price of fairness.

4.1 Efficiency Loss of Subgroup Prioritization

Recall the definition of V_{ij} in Sect. 3.1. We further write V_{ij} as the union

$$V_{ij} = \cup_{b_1, b_2 \in \{O, A, B, AB\}} \{V_{b_1, b_2, ij}\},$$

where $V_{b_1, b_2, ij}$ is the subset of vertices in V_{ij} with donor blood type b_1 and patient blood type b_2 . For random graph G , consider again the optimization problem in 1:

$$\max_{x_c \in \{0,1\}} \sum_{c \in \mathcal{C}} x_c u(c) \quad \text{subject to} \quad \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \leq 1, \quad \forall v \in V. \quad (5)$$

But here, we allow for a general utility u and a set of cycles or relevant subsets $\mathcal{C}$ with some length limits. The term “general utility” refers to a flexible specification of the objective beyond the traditional “number of transplants” criterion. Specifically, $u(c)$ can represent any valid utility function that incorporates additional clinical or policy considerations, such as predicted graft survival, donor-recipient age matching, or sensitization adjustments, while still assigning a scalar utility value to each feasible exchange cycle c . Regardless of how $u(c)$ is defined, all utilities enter the optimization problem through binary decision variables, so the resulting formulation remains an integer programming problem.

Let $\mathcal{P}$ denote the set of indices (b_1, b_2, i, j) , where the subgroup $V_{b_1, b_2, ij}$ should be prioritized over $V_{b_1, b_2, 1-ij}$. For fixed $\epsilon \geq 0$, we consider the optimization problem under the constraint of prioritizing these subgroups in $\mathcal{P}$:

$$\max_{x_c \in \{0,1\}} \sum_{c \in \mathcal{C}} x_c u(c) \quad (6)$$

$$\begin{aligned} \text{subject to} \quad & \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \leq 1, \quad \forall v \in V, \\ & \sum_{v \in V_{b_1, b_2, ij}} \sum_{c \in \mathcal{C}} x_c 1_{v \in c} - |V_{b_1, b_2, ij}| \leq \epsilon N \quad \text{or} \\ & \sum_{v \in V_{b_1, b_2, 1-ij}} \sum_{c \in \mathcal{C}} x_c 1_{v \in c} \leq \epsilon N, \quad \forall (b_1, b_2, i, j) \in \mathcal{P}. \end{aligned} \quad (7)$$

The problems (5) and (6)–(7) share the same objective function. The constraint (7) means that at most ϵN in $V_{b_1, b_2, 1-ij}$ can be matched or at most ϵN patients in $V_{b_1, b_2, ij}$ can be unmatched, if the subgroup $V_{b_1, b_2, ij}$ is prioritized over $V_{b_1, b_2, 1-ij}$. We refer to (5) as the unconstrained problem, aside from the vertex-disjointness condition, and to (6)–(7) as the constrained problem.

The following Proposition 1 is the main result of this subsection.

Proposition 1 *Consider the random graph model in Sect. 4, and fix any sequence $\epsilon_N \downarrow 0$. For each N , let $\text{Opt}(G_N)$ be the optimum of the unconstrained problem (5), and let $\text{Opt}_{\epsilon_N}^{\text{prio}}(G_N)$ be the optimum of the constrained problem (6)–(7) with $\epsilon = \epsilon_N$. Then, with probability one as $N \rightarrow \infty$,*

$$\text{Opt}(G_N) - \text{Opt}_{\varepsilon_N}^{\text{prio}}(G_N) = o(N).$$

Proposition 1 shows that we can prioritize patients with either $A = 1$ or $A = 0$, given their blood type and sensitization level, with ignorable relative efficiency loss when the random graph is large. The result is useful to derive upper bounds for the price of our new fairness criterion, defined as the relative overall utility loss due to the fairness constraint in (4). We present these results in the next subsection.

4.2 Upper Bounds for Price of our new Fairness Criterion

4.2.1 Optimizing some General Utilities

First, we apply Proposition 1 to the scenario with general utilities, allowing for potential recourse strategies. We can obtain a crude upper bound for the price of fairness in Proposition 2 below.

Proposition 2 *The price of fairness due to the fairness constraint in (4) is no greater than*

$$\max_{b_1, b_2, r} \max \left\{ \frac{\mu_{b_1, b_2, r, 1} \bar{\mu}_{r, 0} - \mu_{b_1, b_2, r, 0} \bar{\mu}_{r, 1}}{(\mu_{b_1, b_2, r, 1} + \mu_{b_1, b_2, r, 0}) \bar{\mu}_{r, 0}}, \frac{\mu_{b_1, b_2, r, 0} \bar{\mu}_{r, 1} - \mu_{b_1, b_2, r, 1} \bar{\mu}_{r, 0}}{(\mu_{b_1, b_2, r, 1} + \mu_{b_1, b_2, r, 0}) \bar{\mu}_{r, 1}} \right\}$$

almost surely as $N \rightarrow \infty$, where $\bar{\mu}_{r, 1} = \sum_{b_1, b_2} \mu_{b_1, b_2, r, 1}$ and $\bar{\mu}_{r, 0} = \sum_{b_1, b_2} \mu_{b_1, b_2, r, 0}$.

If the blood type distributions are balanced across all subgroups defined by different levels of A and R , i.e., $\mu_{b_1, b_2, r, 0} / \bar{\mu}_{r, 0} = \mu_{b_1, b_2, r, 1} / \bar{\mu}_{r, 1}$ for all b_1, b_2, r , then the upper bound in Proposition 2 is 0. If the blood type distributions are not balanced within a specific subgroup level $A = a$ and $R = r$, and the optimal solution only matches patients in this subgroup, then the upper bound in Proposition 2 is attainable.

4.2.2 Maximizing the Number of Transplants Without Recourse Strategies

Then, we apply Proposition 1 to the scenario that maximizes the number of transplants without any recourse strategies. In this scenario, an explicit optimal allocation is explicitly available in [27].

We introduce the following assumptions to simplify the decomposition of the four-way probability $\mu_{b_1, b_2, r, a}$. First, we assume that the patient and donor in each incompatible pair share the same protected feature level. Second, we assume that the distributions of sensitization levels are consistent across protected feature levels. While these assumptions are not strictly necessary, they facilitate the decomposition of $\mu_{b_1, b_2, r, a}$ into more manageable terms.

Specifically, let μ_a represent the frequency probability of the protected feature level $A = a$. Define $\mu_{O|a}$, $\mu_{A|a}$, $\mu_{B|a}$, and $\mu_{AB|a}$ as the frequency probabilities of blood types O, A, B, and AB, respectively, within the protected feature level $A = a$. Under

these assumptions, there exists a constant c such that $\mu_{b_1, b_2, r, a} = cr\mu_a\mu_{b_1|a}\mu_{b_2|a}$ for all $b_1, b_2 \in \{O, A, B, AB\}$, $r \in \{r_1, \dots, r_M\}$, and $a \in \{0, 1\}$.

Moreover, let $\bar{\mu}_O, \bar{\mu}_A, \bar{\mu}_B, \bar{\mu}_{AB}$ denote the frequency probability of blood types O, A, B, AB, respectively, among the whole population. Let $\bar{\gamma}$ be the average PRA score among the whole population. We can obtain a more precise upper bound for the price of fairness in Proposition 3 below.

Proposition 3 Assume $1.5\bar{\mu}_A > \bar{\mu}_O > \bar{\mu}_A > \bar{\mu}_B > \bar{\mu}_{AB}$ and $\bar{\gamma} < 0.4$. Let $\phi_{b_1, b_2} = \sum_{k=0}^1 \mu_k \bar{\gamma} \mu_{b_1|k} \mu_{b_2|k}$, for $b_1, b_2 \in \{O, A, B, AB\}$. Let

$$\begin{aligned} T_a(r) &= \mu_a r (\mu_{O|a} + \mu_{AB|a} - \mu_{O|a} \mu_{AB|a} + \mu_{A|a}^2 + \mu_{B|a}^2) + 2\mu_a \mu_{A|a} \mu_{B|a}, \\ S_a(r) &= T_a(r) + \mu_a \{ \mu_{O|a} (1 - \mu_{O|a}) + \mu_{A|a} \mu_{AB|a} + \mu_{B|a} \mu_{AB|a} \}, \\ Q(r) &= T_1(r) + T_0(r) + \phi_{B,AB} + \phi_{O,AB} + \phi_{A,AB} + \phi_{A,O} + \phi_{AB,O} + \phi_{O,B}, \end{aligned}$$

for $a = 0, 1$ and $r \in \{r_1, \dots, r_M\}$, and

$$\begin{aligned} R_a &= \min \{ \phi_{B,AB}, 2\mu_a \mu_{B|a} \mu_{AB|a} - \phi_{B,AB} \} \\ &\quad + \min \{ \phi_{O,AB} + \phi_{A,AB}, 2\mu_a \mu_{A|a} \mu_{AB|a} - \phi_{O,AB} - \phi_{A,AB} \} \\ &\quad + \min \{ \phi_{A,O} + \phi_{AB,O}, 2\mu_a \mu_{O|a} \mu_{A|a} - \phi_{A,O} - \phi_{AB,O} \} \\ &\quad + \min \{ \phi_{O,B}, 2\mu_a \mu_{O|a} \mu_{B|a} - \phi_{O,B} \}, \end{aligned}$$

for $a = 0, 1$. Then, the price of fairness due to the fairness constraint in (4) is no greater than

$$\max_{r \in \{r_1, \dots, r_M\}} \max \left\{ \frac{S_1 T_0 - S_0 T_1 - S_0 R_1}{S_1 Q}, \frac{S_0 T_1 - S_1 T_0 - S_1 R_0}{S_0 Q}, 0 \right\}$$

almost surely as $N \rightarrow \infty$.

Intuitively, the quantities in Proposition 3 summarize how much matching potential each protected group has, and how much of that potential can be reallocated when we impose our fairness constraint. For a fixed protected group level $A = a$ and sensitization level $R = r$, $T_a(r)$ is the expected fraction of incompatible pairs in group a (at level r) that are fully matched in the optimal unconstrained allocation. It aggregates all blood-type combinations that are essentially “always used” in the optimal construction (e.g., X–X pairs, A–B/B–A pairs, and the core part of AB–O, A–O, B–O exchanges) and therefore contribute deterministically to the match count. $S_a(r)$ is the expected total fraction of incompatible pairs in group a (at level r) that belong to blood-type classes that can potentially participate in exchanges under the Ashlagi–Roth optimal structure. It includes both the fully matched categories counted in $T_a(r)$ and the partially matched categories (such as B–AB, A–AB, O–A, O–B). Thus, the ratio $T_a(r)/S_a(r)$ represents the limiting selection probability for group a at sensitization level r in the unconstrained optimum.

At the population level, $Q(r)$ is the expected total fraction of incompatible pairs at sensitization level r that are matched in the unconstrained optimal allocation,

aggregated over both protected groups and all relevant blood-type combinations. It is the denominator in the expression for the asymptotic price of fairness, i.e., the benchmark efficiency against which we compare the fair allocation. Finally, R_a captures the *additional matching mass that can be shifted* toward group $A = a$ within the partially matched blood-type classes while preserving feasibility. Each $\min\{\cdot, \cdot\}$ term in the definition of R_a corresponds to a specific partially matched blood-type class (B–AB, A–AB, O–A/O–AB, O–B), and ensures we do not reassign more matches to group a than either (i) the total number of matches available in that class, or (ii) the supply of compatible pairs in group a . In this sense, R_a measures the “room to maneuver” for reallocating matches across protected groups without changing the blood-type structure of the exchange.

In combination, $T_a(r)$ and $S_a(r)$ describe the baseline matching rates for each group and sensitization level, $Q(r)$ encodes the total unconstrained efficiency, and R_a quantifies how much we can adjust the allocation in favor of each group when enforcing our fairness constraint. The bound in Proposition 3 then compares the best fair allocation achievable under these structural limits with the unconstrained optimum, yielding an explicit asymptotic upper bound on the price of fairness.

The condition $1.5\bar{\mu}_A > \bar{\mu}_O > \bar{\mu}_A > \bar{\mu}_B > \bar{\mu}_{AB}$ imposes a mild and empirically realistic constraint on the population-level blood type distribution, while the condition $\bar{\gamma} < 0.4$ implies that most patients are not highly sensitized. Both assumptions are standard in the kidney exchange literature and appear in Ashlagi and Roth [27] and Dickerson et al. [7]. Moreover, these assumptions are consistent with real U.S. blood type frequencies reported by the American Red Cross (2021), approximately 45% type O, 40% type A, 11% type B, and 4% type AB, and with observed sensitization patterns in national KPD data, where the majority of patients exhibit moderate or low PRA levels.

As an application of Proposition 3, we present two illustrative examples showing that the asymptotic upper bounds on the price of fairness are very small under empirically supported distributions. In the first example, suppose there are two ethnicity groups, 80% White American and 20% African American, with $\mu_{O|1} = 0.45$, $\mu_{A|1} = 0.40$, $\mu_{B|1} = 0.11$, $\mu_{AB|1} = 0.04$ and $\mu_{O|0} = 0.51$, $\mu_{A|0} = 0.26$, $\mu_{B|0} = 0.19$, $\mu_{AB|0} = 0.04$. Moreover, there are three sensitization levels $\{0.05, 0.45, 0.9\}$. Then, with high probability, the price of fairness converges to zero for any $0.05 < \bar{\gamma} < 0.4$.

In the second example, suppose there are two ethnicity groups, 90% White American and 10% Asian American, with $\mu_{O|1} = 0.45$, $\mu_{A|1} = 0.40$, $\mu_{B|1} = 0.11$, $\mu_{AB|1} = 0.04$ and $\mu_{O|0} = 0.40$, $\mu_{A|0} = 0.275$, $\mu_{B|0} = 0.255$, $\mu_{AB|0} = 0.07$. Moreover, there are five sensitization levels $\{0.05, 0.25, 0.45, 0.65, 0.9\}$. Then, with high probability, the price of fairness is less than 0.01 for any $0.05 < \bar{\gamma} < 0.09$ and converges to zero for any $0.09 \leq \bar{\gamma} < 0.40$.

These two examples demonstrate that, under empirically grounded population characteristics, the price of fairness in large kidney exchange graphs is negligible, typically below 1%, confirming that fairness constraints can be imposed with minimal efficiency loss in real-world settings.

Under a similar random graph model, [7] claim that the price of group fairness is no greater than $2/33$, as $n \rightarrow \infty$; [27] claim that the relative efficiency loss for individual rationality is only about 1% in multi-hospital kidney exchange. All these

results give very low efficiency loss because, in large random graph models, there is a rich set of edges from each vertex such that we can easily adjust for the optimal solution such that the fairness constraint is satisfied. Beyond the kidney exchange setting, [21] provide an upper bound for the price of proportional fairness and the price of max-min fairness, which are close to 1, in general allocation problems.

5 Numerical Studies

In this section, we present simulation studies based on the random graph models described in Sect. 4 and real data from the UNOS dataset. All the optimization problems are solved by Gurobi (version 11.0) in R.

5.1 Simulation Under the Random Graph Models

We first conduct a simulation study under the random graph models to evaluate the numerical performance of different fairness criteria. We consider a binary protected feature that indicates whether one is white or non-white. Specifically, we fix 80 white incompatible pairs (80%) and 20 non-white incompatible pairs (20%). Among the 80 white pairs, 56 patients (70%) are lowly sensitized, 16 patients (20%) are moderately sensitized, and 8 patients (10%) are highly sensitized; and among the 20 non-white pairs, 14 patients (70%) are lowly sensitized, 4 patients (20%) are moderately sensitized, and 2 patients (10%) are highly sensitized. The PRA scores are set to be 0.9, 0.45, and 0.05 for highly sensitized, moderately sensitized, and lowly sensitized patients, respectively. The distribution follows the analysis in [25].

For non-white donors and patients, we simulate the blood type from a multinomial distribution (O: 51%, A: 26%, B: 19%, AB: 4%). For white donors and patients, we simulate the blood type from another multinomial distribution (O: 45%, A: 40%, B: 11%, AB: 4%). If a donor and a patient are blood-type compatible, they are incompatible with the probability of the patient's PRA score; otherwise, they are incompatible with probability 1. Given the patients, we repeat sampling the donors and dropping compatible pairs until the numbers of incompatible pairs are reached. Similarly, based on the blood types and PRA scores, we simulate the edges of the graph. That is, we simulate the compatibility for the donor and patient of every two vertices in the graph. We only allow cycles of length at most 3, and we do not consider any recourse strategies in this section.

We compare three fairness criteria, group fairness, individual fairness, and our newly proposed fairness criterion, using the outcome with no fairness constraints as a baseline. For our new fairness criterion, we evaluate the two parameter choices introduced in Sect. 3.1. For group fairness, we consider two parameter-selection methods. The first selects the largest feasible value of α for which (2) remains solvable, following [7]. This corresponds to imposing the strongest possible constraint, requiring us to maximize the number of matched highly sensitized patients. The second method selects the largest α that preserves the overall utility achieved in (1), as in [8]. In other words, among all exchange plans that maximize the objective

in (2), we choose the one that yields the greatest number of matched highly sensitized patients. For individual fairness, we adopt two analogous parameter-selection approaches. Specifically, we minimize the L_1 norm of the difference vector $\pi - \bar{\pi}$ while retaining at least 80% and 100% of the maximum achievable utility, respectively. These correspond to a stronger and a weaker fairness constraint.

Figure 2 presents the average selection rates of different fairness criteria within each subgroup. From Fig. 2, group fairness works to increase the selection rates of highly sensitized patients, and individual fairness works to balance the selection rates of the six subgroups. Differently, our fairness balances the selection rates of the white and non-white patients within each sensitization stratum, instead of the selection rates of all six subgroups. Our stronger fairness constraint can almost eliminate the absolute difference in selection rates, while our weaker fairness constraint can also significantly reduce the absolute difference in selection rates compared to




Fig. 2 Simulation results under random graph models. The average selection rates within each subgroup are calculated over 100 data replications. The error bars represent the mean ± 1 standard deviation of the absolute differences in selection rates

other methods. For the price of fairness, the relative utility loss is around 6.3% for our stronger fairness constraint and 0.7% for our weaker fairness constraint.

5.2 UNOS Data Analysis

We next conduct a simulation study based on the National UNOS STAR dataset. The National UNOS STAR dataset provides comprehensive transplant records collected by the UNOS, covering donor and recipient characteristics, allocation details, and transplant outcomes. After removing the missing values, the dataset comprises 77,073 records of transplant information for donors and patients, which include details such as blood types, HLA antigens (A1, A2, B1, B2, DR1, DR2), PRA scores, and racial background. The protected feature is set to be a binary variable: white (64.8%) and other racial backgrounds (35.2%). Patients are categorized based on their PRA scores as follows: those with scores above 0.8 are labeled as highly sensitized; those with scores ranging from 0.1 to 0.8 are considered moderately sensitized; and patients with scores below 0.1 are labeled as lowly sensitized.

We fix the number of incompatible donor-patient pairs to be 100. Donors and patients are randomly sampled from the dataset to form incompatible pairs independently. The overall compatibility is determined by both blood type and HLA compatibility. HLA compatibility is assessed based on the number of mismatches in the A, B, and DR alleles. Specifically, a donor and patient are considered HLA compatible if their level of HLA mismatch is less than 3. The HLA mismatch level is based on UK kidney matching policies, and it can be calculated using the R package `transplantr`. Again, we only allow cycles of length at most 3.

Figure 3 reports the average selection rates within each subgroup. From Fig. 3, group fairness, individual fairness, and our new fairness criterion all enhance equitable access to transplant opportunities according to their specific fairness criteria.

5.3 Experiment on Selection Probability Prediction

We next conduct a numerical experiment to evaluate the prediction accuracy of the selection probability. We assume there is an underlying population consisting of 80% white and 20% non-white incompatible pairs, following the same blood type distributions and sensitization distributions in Sect. 5.1. We assume the historical data $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$, current data $\{v_1, \dots, v_{N_1}\}$, and future data are independently sampled from the population. We fix $N_0 = 200$, $N = 100$, $B = 1000$, and vary N_1 in $\{20, 40, 60, 80\}$. For illustration purposes, we focus on our new fairness criterion.

Table 1 presents the numerical results of the selection probability prediction. As L increases, the prediction accuracy improves, evidenced by a lower mean squared error (MSE) and a coverage rate of the prediction interval approaching 95%. This improvement occurs because the size of the unobserved future data, $N - N_1$, decreases, which makes the prediction easier. Although these results indicate the




Fig. 3 Simulation results based on UNOS data. The average selection rates within each subgroup are calculated over 100 data replications. The error bars represent the mean $\pm$ 1 standard deviation of the absolute differences in selection rates

Table 1 Accuracy of the predicted selection probability of $v_1, \dots, v_{N_1}$. The results are averaged over 100 data replications

N_1	20	40	60	80
MSE	0.079	0.056	0.045	0.038
Coverage	0.977	0.965	0.962	0.958
Width	0.445	0.332	0.265	0.227

method's validity, there is a bias in the prediction interval due to the finite size of N_0 . When N_0 is large, the set $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$ closely approximates the underlying population distribution; while when N_0 is small, the distribution of $\{\tilde{v}_1, \dots, \tilde{v}_{N_0}\}$ may differ from the underlying population distribution.




Fig. 4 Simulation results with subset-recourse strategy under random graph models. The average selection rates within each subgroup are calculated over 100 data replications. The error bars represent the mean $\pm$ 1 standard deviation of the absolute differences in selection rates

6 Discussion

In this paper, we propose a new fairness criterion that balances selection probabilities within protected groups across each sensitization level. Based on the calibration principle in machine learning, this fairness criterion offers a meaningful and innovative approach in the context of kidney exchange. We propose an efficient solution to implement this criterion and conduct both theoretical and empirical evaluations to analyze the associated price of fairness.




Fig. 5 Simulation results with subset-recourse strategy based on UNOS data. The average selection rates within each subgroup are calculated over 100 data replications. The error bars represent the mean $\pm$ 1 standard deviation of the absolute differences in selection rates

7 Supplementary Information

The replication files for this paper are available at <https://github.com/mingrui229/Fair-KPD>.

Appendix A Proofs

The proof of Proposition 1 depends on the following lemma.

Lemma 1 Fix $k \geq 2$ and $p \in (0, 1)$. For each n , let $(A_1, \dots, A_k)$ be a k -partite undirected random graph with $|A_i| = n$ for all i , where every potential edge between distinct parts is present independently with probability at least p .

Then, with probability $1 - o(1)$ as $n \rightarrow \infty$, the graph contains n vertex-disjoint copies of the complete k -partite graph with one vertex in each part; that is, there exist n disjoint k -tuples, each consisting of one vertex from every part A_i , such that all $\binom{k}{2}$ inter-part edges among the vertices in each k -tuple are present.

Proof of Lemma 1 We build the statement in four steps.

Step 1: Bipartite building block. Let X, Y be disjoint sets of size n , and consider a bipartite graph in which each edge (x, y) is present independently with probability at least $q > 0$. By monotone coupling with $G(n, n, q)$, this random graph stochastically dominates an Erdős–Rényi bipartite graph. It is classical that $G(n, n, q)$ has a perfect matching with probability $1 - o(1)$ as $n \rightarrow \infty$, and ‘having a perfect matching’ is monotone in the edge set. Hence the original bipartite graph also has a perfect matching with probability $1 - o(1)$.

Step 2: Uniform representation and filtration. For each unordered pair (u, v) with $u \in A_i, v \in A_j, i \neq j$, let $U_{(u,v)}$ be independent $\text{Unif}[0, 1]$. If the designated probability of the edge (u, v) is $p_{(u,v)} \in [p, 1]$, we declare (u, v) present iff $U_{(u,v)} \leq p_{(u,v)}$. This yields exactly the model in the lemma.

For $t = 1, \dots, k - 1$, define

$$\mathcal{F}_t := \sigma(U_{(u,v)} : u \in A_i, v \in A_j, i \neq j, \{u, v\} \subseteq A_1 \cup \dots \cup A_{t+1}),$$

which reveals all edges whose endpoints lie in $A_1, \dots, A_{t+1}$.

Step 3: Inductive construction of $(t + 1)$ -tuples. We build the desired k -tuples in $k - 1$ rounds. At the end of round $t \geq 2$ we will have n disjoint sets $B_1^{(t)}, \dots, B_n^{(t)}$, each containing one vertex from each of $A_1, \dots, A_t$ and inducing a complete t -partite subgraph.

Round $t = 2$. Consider the bipartite graph between A_1 and A_2 . All edges are independent and present with probability at least p , so by Step 1 it has a perfect matching with probability $1 - o(1)$. Fix such a matching and write its edges as (x_j, y_j) , $j = 1, \dots, n$, with $x_j \in A_1, y_j \in A_2$, and define

$$B_j^{(2)} := \{x_j, y_j\}, \quad B_2 := \{B_1^{(2)}, \dots, B_n^{(2)}\}.$$

Inductive step. Assume for some $t \in \{2, \dots, k - 1\}$ that we have disjoint sets

$$B_t = \{B_1^{(t)}, \dots, B_n^{(t)}\}$$

such that each $B_j^{(t)}$ contains exactly one vertex from each of $A_1, \dots, A_t$ and induces a complete t -partite subgraph. Then B_t is $\mathcal{F}_{t-1}$ -measurable.

Define a bipartite graph Γ_t whose left part is B_t and whose right part is A_{t+1} . For $B \in B_t$ and $a \in A_{t+1}$, put an edge between B and a in Γ_t if and only if a is adjacent to every vertex in B in the original graph, i.e.,

$$(B, a) \in E(\Gamma_t) \iff U_{(v,a)} \leq p_{(v,a)} \text{ for all } v \in B.$$

Conditional on $\mathcal{F}_{t-1}$, B_t is fixed and the variables $\{U_{(v,a)} : v \in B\}$ are independent, with $|B| = t$ and $p_{(v,a)} \geq p$. Hence

$$\Pr((B, a) \in E(\Gamma_t) \mid \mathcal{F}_{t-1}) = \prod_{v \in B} p_{(v,a)} \geq p^t =: q_t > 0.$$

For distinct pairs $(B, a) \neq (B', a')$, the sets of uniforms $\{U_{(v,a)} : v \in B\}$ and $\{U_{(v',a')} : v' \in B'\}$ are disjoint, so the edge indicators in Γ_t are independent, each with conditional probability at least q_t .

Thus, conditional on $\mathcal{F}_{t-1}$, Γ_t satisfies the assumptions of Step 1 with parameter q_t , and since $|B_t| = |A_{t+1}| = n$, it has a perfect matching with probability $1 - o(1)$. On that event, write the matching as $(B_j^{(t)}, a_j)$, $j = 1, \dots, n$, and set

$$B_j^{(t+1)} := B_j^{(t)} \cup \{a_j\}, \quad B_{t+1} := \{B_1^{(t+1)}, \dots, B_n^{(t+1)}\}.$$

Each $B_j^{(t+1)}$ is then a complete $(t+1)$ -partite subgraph with one vertex in each of $A_1, \dots, A_{t+1}$.

Step 4: Union bound over rounds. There are $k-1$ rounds. For each fixed t , the conditional failure probability is $o(1)$ as $n \rightarrow \infty$, hence the unconditional failure probability is $o(1)$. Since k is fixed, a union bound shows that with probability $1 - o(1)$ all rounds succeed, yielding n disjoint k -tuples as in the statement. $\square$

Proof of Proposition 1 We organize the proof into six steps:

- Step 1 defines coarse and fine types and their frequencies.
- Step 2 introduces template types and a coarse-type linear program.
- Step 3 shows that the unconstrained optimal value is asymptotic to the LP optimum.
- Step 4 splits coarse usage into fine types in a way that encodes priorities.
- Step 5 realizes these fine-type counts in G_N using Lemma 1.
- Step 6 combines everything.

Throughout, we work under the assumptions of Sect. 4.

Step 1: Coarse and fine types and their frequencies. Each vertex v (an incompatible donor–patient pair) has

$$b_1(v), b_2(v) \in \{O, A, B, AB\}, \quad r(v) \in \{r_1, \dots, r_M\}, \quad A(v) \in \{0, 1\}.$$

Define the coarse type

$$\tilde{\tau}(v) := (b_1(v), b_2(v), r(v)) \in \tilde{\mathbb{T}} := \{O, A, B, AB\}^2 \times \{r_1, \dots, r_M\},$$

and the fine type

$$\tau(v) := (\tilde{\tau}(v), A(v)) \in \mathbb{T} := \tilde{\mathbb{T}} \times \{0, 1\}.$$

For each N , let V_N denote the vertex set of G_N , with $|V_N| = N$. For $t \in \mathbb{T}$ and $\tilde{t} \in \tilde{\mathbb{T}}$, define

$$n_t(N) := |\{v \in V_N : \tau(v) = t\}|, \quad n_{\tilde{t}}(N) := \sum_{a=0}^1 n_{(\tilde{t},a)}(N),$$

and let $\pi_t > 0, \pi_{\tilde{t}} > 0$ be the corresponding type probabilities under the i.i.d. type distribution.

By the strong law of large numbers, almost surely

$$\frac{n_t(N)}{N} \rightarrow \pi_t, \quad \frac{n_{\tilde{t}}(N)}{N} \rightarrow \pi_{\tilde{t}} \quad \text{as } N \rightarrow \infty \text{ for all } t, \tilde{t}. \quad (\text{A1})$$

We henceforth work on any realization for which (A1) holds.

Step 2: Template types and a coarse-type linear program. Ignoring the protected feature, we classify cycles (or relevant subsets) into finitely many “templates”. Two sets $c_1, c_2 \in \mathcal{C}(G_N)$ have the same *template type* if there is a bijection $\phi : V(c_1) \rightarrow V(c_2)$ such that:

- $\tilde{\tau}(v) = \tilde{\tau}(\phi(v))$ for all $v \in V(c_1)$; and
- $(v_1, v_2) \in E_N$ iff $(\phi(v_1), \phi(v_2)) \in E_N$.

Because the size $|c|$ is uniformly bounded and $\tilde{\mathbb{T}}$ is finite, there are only finitely many template types. Index them by $i = 1, \dots, K$.

For each template i , fix a representative directed pattern H_i with vertex set $V(H_i)$ and $|V(H_i)| = s_i$. For each coarse type $\tilde{t} \in \tilde{\mathbb{T}}$, let

$$\tilde{a}_{\tilde{t},i} := |\{w \in V(H_i) : \tilde{\tau}(w) = \tilde{t}\}|.$$

Since $u(c)$ depends only on the induced subgraph and the coarse types (and is invariant under relabelings preserving this information), $u_i := u(c)$ is well-defined for each template i .

We consider the coarse-type linear program

$$\begin{aligned} \max_{z \geq 0} \quad & U^T z := \sum_{i=1}^K u_i z_i, \\ \text{s.t.} \quad & \sum_{i=1}^K \tilde{a}_{\tilde{t},i} z_i \leq \pi_{\tilde{t}}, \quad \forall \tilde{t} \in \tilde{\mathbb{T}}. \end{aligned} \quad (\text{A2})$$

Feasibility is clear and boundedness follows from $u(c)/|c| \leq C$ and bounded template sizes. Let ρ be the optimal value and z^* an optimal solution.

Step 3: Asymptotic behavior of the unconstrained optimum. We show that

$$\text{Opt}(G_N) = \rho N + o(N) \quad \text{almost surely.} \quad (\text{A3})$$

Step 3A: Upper bound. Let $(x_c)_{c \in \mathcal{C}(G_N)}$ be an optimal solution of the unconstrained problem (5) on G_N . For each template type i , define

$$Z_i^{(N)} := \sum_{c \in \mathcal{C}(G_N)} x_c \mathbf{1}_{\{c \text{ has template type } i\}}, \quad y^{(N)} := \frac{1}{N} Z^{(N)}.$$

Vertex-disjointness implies that, for each coarse type $\tilde{i}$,

$$\sum_{i=1}^K \tilde{a}_{\tilde{i},i} Z_i^{(N)} \leq n_{\tilde{i}}(N).$$

Dividing by N and letting $N \rightarrow \infty$, any subsequential limit y of $\{y^{(N)}\}$ satisfies

$$\sum_{i=1}^K \tilde{a}_{\tilde{i},i} y_i \leq \pi_{\tilde{i}}, \quad \forall \tilde{i} \in \tilde{\mathbb{T}},$$

so y is feasible for (A2) and $U^T y \leq \rho$.

The objective value on G_N is

$$\text{Opt}(G_N) = \sum_{i=1}^K u_i Z_i^{(N)} = N U^T y^{(N)},$$

so

$$\limsup_{N \rightarrow \infty} \frac{\text{Opt}(G_N)}{N} = \limsup_{N \rightarrow \infty} U^T y^{(N)} \leq \rho.$$

Step 3B: Lower bound. Define integer vectors

$$Z^{(N)} := (\lfloor N z_1^* \rfloor, \dots, \lfloor N z_K^* \rfloor), \quad y^{(N)} := \frac{1}{N} Z^{(N)}.$$

Then $y^{(N)} \rightarrow z^*$, and for each coarse type $\tilde{i}$,

$$\sum_{i=1}^K \tilde{a}_{\tilde{i},i} y_i^{(N)} \leq \sum_{i=1}^K \tilde{a}_{\tilde{i},i} z_i^* \leq \pi_{\tilde{i}},$$

so $\sum_i \tilde{a}_{\tilde{i},i} Z_i^{(N)} \leq N \pi_{\tilde{i}} + O(1)$. Together with (A1), this implies $\sum_i \tilde{a}_{\tilde{i},i} Z_i^{(N)} \leq n_{\tilde{i}}(N)$ for all large N , so the coarse-type supply of vertices suffices to realize $Z^{(N)}$ templates.

To realize these templates, fix i . Label the vertices of H_i as $w_1, \dots, w_{s_i}$ with coarse types $\tilde{i}_1, \dots, \tilde{i}_{s_i}$. For each (i, ℓ) choose a subset

$$A_\ell^{(i)}(N) \subseteq \{v \in V_N : \tilde{\tau}(v) = \tilde{i}_\ell\}$$

of size $Z_i^{(N)}$, such that all sets $A_\ell^{(i)}(N)$ are disjoint across all (i, ℓ) ; this is possible by the previous capacity inequality.

For fixed i , consider the induced s_i -partite undirected graph on parts $A_1^{(i)}(N), \dots, A_{s_i}^{(i)}(N)$. Let L_i be the number of directed edges in H_i . If $u \in A_\ell^{(i)}(N)$ and $v \in A_m^{(i)}(N)$ with $\ell \neq m$, we connect u and v by an undirected edge if all directed edges between w_ℓ and w_m required by H_i (and medically admissible) are present

between u and v in G_N . Each such undirected edge appears with probability at least $p_0^{L_i} =: p_i > 0$, and different pairs (u, v) use disjoint sets of underlying directed edges, so these edges are independent.

Thus, for each i , the s_i -partite graph on $A_1^{(i)}(N), \dots, A_{s_i}^{(i)}(N)$ satisfies the hypotheses of Lemma 1 with $k = s_i$, $n = Z_i^{(N)}$, and parameter $p_i > 0$. If $z_i^* > 0$, then $Z_i^{(N)} \rightarrow \infty$ and Lemma 1 implies that this graph contains $Z_i^{(N)}$ vertex-disjoint complete s_i -partite subgraphs with probability $1 - o(1)$. Each such subgraph corresponds to one copy of H_i . If $z_i^* = 0$, then $Z_i^{(N)} = O(1)$ and failure to realize these copies changes the total utility by at most $O(1) = o(N)$.

Since there are finitely many templates, a union bound shows that with probability $1 - o(1)$ we can realize all templates H_i with counts $Z_i^{(N)}$ simultaneously. On this event, we obtain an exchange plan with template counts $Z^{(N)}$ and utility

$$\sum_{i=1}^K Z_i^{(N)} u_i = N U^T z^* + O(1) = \rho N + o(N).$$

Hence $\liminf_{N \rightarrow \infty} \frac{\text{Opt}(G_N)}{N} \geq \rho$, proving (A3).

Step 4: Splitting coarse usage into fine types and encoding priority. We now refine coarse-type usage to fine types (including the protected feature) in a way that captures the desired subgroup prioritization.

Step 4A: Coarse usage per type. For each coarse type $\tilde{t}$, define

$$\lambda_{\tilde{t}} := \sum_{i=1}^K \tilde{a}_{\tilde{t},i} z_i^* \leq \pi_{\tilde{t}}.$$

Step 4B: Splitting into fine-type usages. Choose nonnegative numbers λ_t for $t = (\tilde{t}, a) \in \mathbb{T}$ such that

$$\lambda_{(\tilde{t},0)} + \lambda_{(\tilde{t},1)} = \lambda_{\tilde{t}}, \quad \lambda_t \leq \pi_t \quad \forall t.$$

For each prioritized quadruple $(b_1, b_2, i, j) \in \mathcal{P}$ with coarse type $\tilde{t} = (b_1, b_2, r_j)$ and fine types

$$t^{(\text{prio})} = (\tilde{t}, i), \quad t^{(\text{unprio})} = (\tilde{t}, 1 - i),$$

we further require that either $\lambda_{t^{(\text{unprio})}} = 0$ or $\lambda_{t^{(\text{prio})}} = \pi_{t^{(\text{prio})}}$. This is achieved by the rule

$$(\lambda_{t^{(\text{prio})}}, \lambda_{t^{(\text{unprio})}}) = \begin{cases} (\lambda_{\tilde{t}}, 0), & \lambda_{\tilde{t}} \leq \pi_{t^{(\text{prio})}}, \\ (\pi_{t^{(\text{prio})}}, \lambda_{\tilde{t}} - \pi_{t^{(\text{prio})}}), & \lambda_{\tilde{t}} > \pi_{t^{(\text{prio})}}. \end{cases}$$

For coarse types not in any prioritized pair, we pick any feasible split.

Step 4C: Fine-type template counts. We refine the coarse counts $\tilde{a}_{\tilde{t},i}$ into $b_{t,i} \geq 0$ so that

$$\sum_{a=0}^1 b_{(\tilde{t},a),i} = \tilde{a}_{\tilde{t},i} \quad \forall \tilde{t}, i, \quad \sum_{i=1}^K b_{t,i} z_i^* = \lambda_t \quad \forall t \in \mathbb{T}. \quad (\text{A4})$$

If $\lambda_{\bar{t}} > 0$, we set

$$b_{(\bar{t},a),i} := \tilde{a}_{\bar{t},i} \frac{\lambda_{(\bar{t},a)}}{\lambda_{\bar{t}}}, \quad a \in \{0, 1\},$$

and if $\lambda_{\bar{t}} = 0$, we set $b_{(\bar{t},a),i} = 0$. A direct check shows that (A4) holds.

Step 5: Constructing matching fine-type counts in G_N . We now pass from the limiting fractions λ_t and $b_{t,i}$ to integer counts in G_N .

Step 5A: Target number of matched vertices per fine type. For each fine type t and N , let

$$M_t^{(N)} := \lfloor N\lambda_t \rfloor.$$

By (A1) and $\lambda_t \leq \pi_t$, we have $M_t^{(N)} \leq n_t(N)$ for all large N , so we can choose

$$V_t^{\text{use}}(N) \subseteq \{v \in V_N : \tau(v) = t\}, \quad |V_t^{\text{use}}(N)| = M_t^{(N)}.$$

For each prioritized quadruple $(b_1, b_2, i, j) \in \mathcal{P}$ with fine types $t^{(\text{prio})}, t^{(\text{unprio})}$, our construction of λ_t implies that, up to $o(N)$ vertices, either only the prioritized type is used or the prioritized type is essentially saturated. This exactly matches the asymptotic priority requirement in the definition of $\text{Opt}_{\varepsilon_N}^{\text{prio}}(G_N)$ for some sequence $\varepsilon_N \downarrow 0$.

Step 5B: Matching fine-type/template counts and realization. Let $\gamma_{t,i} := b_{t,i}z_i^*$. We seek integers $m_{t,i}^{(N)}$ such that

$$\sum_{i=1}^K m_{t,i}^{(N)} = M_t^{(N)} \quad \forall t \in \mathbb{T}, \quad \sum_{t \in \mathbb{T}} m_{t,i}^{(N)} = s_i Z_i^{(N)} \quad \forall i,$$

and $m_{t,i}^{(N)} = N\gamma_{t,i} + O(1)$ as $N \rightarrow \infty$. This can be obtained by a standard rounding argument on the finite transportation polytope: start from $m_{t,i}^{(N)} := \lfloor N\gamma_{t,i} \rfloor$ and adjust finitely many entries by $O(1)$ to match the row and column sums exactly. The total adjustment per (t, i) is $O(1)$, so this changes the utility by at most $O(1) = o(N)$.

For each pair (t, i) , partition $V_t^{\text{use}}(N)$ into disjoint subsets of size $m_{t,i}^{(N)}$ and assign each subset to the appropriate position in the $Z_i^{(N)}$ copies of H_i . This yields, for each template i , parts $A_1^{(i)}(N), \dots, A_{s_i}^{(i)}(N)$ of size $Z_i^{(N)}$ with prescribed fine (and hence coarse) types.

We now apply exactly the same Lemma 1-based construction as in Step 3B, restricted to these parts. The independence and lower bounds on edge probabilities are unchanged, so with probability $1 - o(1)$ we can realize the desired $Z_i^{(N)}$ disjoint copies of H_i using precisely these vertices, for all i with $z_i^* > 0$. Templates with $z_i^* = 0$ contribute $O(1)$ vertices and can be handled arbitrarily, with $o(N)$ effect.

Thus, with probability $1 - o(1)$, for all large N there exists an exchange plan that:

- uses exactly the vertices in $\bigcup_t V_t^{\text{use}}(N)$ up to $O(1)$,
- has template counts $Z^{(N)}$, and
- satisfies the asymptotic priority requirement with sequence ε_N .

The corresponding utility is

$$\sum_{i=1}^K Z_i^{(N)} u_i + O(1) = \rho N + o(N),$$

so almost surely

$$\text{Opt}_{\varepsilon_N}^{\text{prio}}(G_N) \geq \rho N + o(N).$$

Step 6: Conclusion. From Step 3 we know $\text{Opt}(G_N) = \rho N + o(N)$ almost surely, and from Step 5 we know $\text{Opt}_{\varepsilon_N}^{\text{prio}}(G_N) \geq \rho N + o(N)$ almost surely. Hence

$$\text{Opt}(G_N) - \text{Opt}_{\varepsilon_N}^{\text{prio}}(G_N) = o(N) \quad \text{almost surely,}$$

proving the proposition. $\square$

Proof of Proposition 2 We bound the price of fairness in terms of the number of matched vertices and show it is at most the explicit quantity L in the statement. The argument is purely type-level and ignores graph structure until the final step.

Step 1: Local loss for a fixed (b_1, b_2, r) . Fix (b_1, b_2) and r , and consider the two fine types

$$(b_1, b_2, r, 1) \quad \text{and} \quad (b_1, b_2, r, 0),$$

with proportions $\mu_{b_1, b_2, r, 1}$ and $\mu_{b_1, b_2, r, 0}$. Let

$$\bar{\mu}_{r,a} := \sum_{b_1, b_2} \mu_{b_1, b_2, r, a}, \quad a \in \{0, 1\},$$

be the total mass at level r in group $A = a$.

In the unconstrained optimum, the total fraction of matched vertices of these two types can be written as

$$p(\mu_{b_1, b_2, r, 1} + \mu_{b_1, b_2, r, 0})$$

for some $p \in [0, 1]$, which we can think of as the local selection probability at (b_1, b_2, r) . We now compare this with an allocation that enforces the stronger requirement that the selection probabilities at level r for $A = 0$ and $A = 1$ are exactly equal:

$$\frac{\text{matches in } (A = 1, R = r)}{\bar{\mu}_{r,1}} = \frac{\text{matches in } (A = 0, R = r)}{\bar{\mu}_{r,0}}.$$

This is stricter than the original fairness constraint, so any loss bound obtained here is conservative.

If the ratio $\mu_{b_1, b_2, r, 1} / \mu_{b_1, b_2, r, 0}$ exceeds the global ratio $\bar{\mu}_{r,1} / \bar{\mu}_{r,0}$, then this type is “too rich” in $A = 1$ individuals. A straightforward two-type calculation (keeping all possible $A = 0$ matches and enforcing the global ratio) shows that the fair allocation must lose at most

$$\frac{\mu_{b_1, b_2, r, 1} \bar{\mu}_{r, 0} - \mu_{b_1, b_2, r, 0} \bar{\mu}_{r, 1}}{(\mu_{b_1, b_2, r, 1} + \mu_{b_1, b_2, r, 0}) \bar{\mu}_{r, 0}}$$

of the local matched mass. When the inequality is reversed, the type is “too rich” in $A = 0$, and by symmetry (exchanging the roles of $A = 0$ and $A = 1$) the maximal local loss is

$$\frac{\mu_{b_1, b_2, r, 0} \bar{\mu}_{r, 1} - \mu_{b_1, b_2, r, 1} \bar{\mu}_{r, 0}}{(\mu_{b_1, b_2, r, 1} + \mu_{b_1, b_2, r, 0}) \bar{\mu}_{r, 1}}.$$

Taking the larger of these two expressions gives the loss bound for fixed (b_1, b_2, r) .

Step 2: Global bound across all (b_1, b_2, r) . At a fixed sensitization level r , the total matches in the unconstrained optimum are a sum over (b_1, b_2) of $M_{b_1, b_2, r}^*(N)$, and the number of matches in any fair allocation is the corresponding sum after trimming. Hence, at level r , the (asymptotic) fractional loss is a convex combination of the local losses $\text{Loss}_{b_1, b_2, r}$, and so is bounded by their maximum over (b_1, b_2) .

Taking also the maximum over r yields the loss bound in the statement.

Step 3: Implementability in G_N via subgroup prioritization. We now argue that the type-level bound L governs the asymptotic price of fairness in the random graphs G_N .

For each triple (b_1, b_2, r) , the trimming operation in Step 1 can be implemented by prioritizing between the fine types $(b_1, b_2, r, 1)$ and $(b_1, b_2, r, 0)$, as in Sect. 4.1: we decide, for each (b_1, b_2, r) , whether to prioritize $A = 1$ or $A = 0$, and collect these choices into a priority set $\mathcal{P}$.

Consider the constrained optimization problem with subgroup prioritization constraints (6)–(7) associated with $\mathcal{P}$. By Proposition 1, there exists a sequence $\varepsilon_N \downarrow 0$ such that we can impose the asymptotic priority requirements for all $(b_1, b_2, i, j) \in \mathcal{P}$ while losing only $o(N)$ matches relative to the unconstrained optimum.

Within this prioritized class, we can then discard matches exactly as in the type-level construction of Step 1 to enforce equality of selection probabilities between $A = 0$ and $A = 1$ at each r . The additional loss caused by this trimming is, by Step 2, at most a fraction L of the total number of matches, that is, at most $L \text{Opt}(G_N) + o(N)$ matches in G_N .

Thus, almost surely, there exists an allocation satisfying the stronger condition of equal selection probabilities at each r whose total number of matches is at least $(1 - L) \text{Opt}(G_N) - o(N)$. Since the fairness constraint (4) only requires the difference in selection probabilities at each r_j to be bounded by l_j and not necessarily equal to zero, any optimal solution under (4) is at least as efficient as this equalized allocation. Consequently, the asymptotic price of fairness due to (4) is no greater than L . $\square$

Proof of Proposition 3 We work in the random graph model of Sect. 4 specialized to the case where the objective is to maximize the number of transplants (no recourse strategies), and we adopt the additional structural assumptions stated before Proposition 3.

Notational convention. In the rest of the paper we index blood types of a pair as

$$(b_1, b_2) = (\text{donor blood type}, \text{patient blood type}).$$

By contrast, Ashlagi and Roth describe a pair of type X – Y when the *patient* has blood type X and the *donor* has blood type Y . Thus their type X – Y corresponds to our ordered pair

$$(b_1, b_2) = (Y, X).$$

Whenever we refer below to X – Y pair types or to their classification as self-/over-/under-demanded from [11], this is always in the *patient*–*donor* sense of Ashlagi–Roth; when we translate those statements into our formulas in terms of $\mu_{b_1|a}$ and $\mu_{b_2|a}$, we use the above correspondence.

The proof has four main steps.

- Step 1 recalls the blood-type structure of an asymptotically optimal allocation without fairness constraints, following [11], and translates it into our (b_1, b_2) notation.
- Step 2 quantifies how much matched “mass” can be reallocated between $A = 0$ and $A = 1$ within partially matched blood-type classes; this yields R_0 and R_1 .
- Step 3 uses this flexibility and, if necessary, discards some matches to enforce equality of selection probabilities across $A = 0$ and $A = 1$ at each sensitization level r , and computes the resulting loss.
- Step 4 uses Proposition 1 to argue that this type-level construction can be implemented in the random graphs G_N up to $o(N)$ loss and that the weaker fairness constraint (4) has no higher asymptotic price.

Step 1: Structure of the unconstrained optimal allocation. Fix N and let G_N be the random compatibility graph on N incompatible pairs. For each $a \in \{0, 1\}$ and $r \in \{r_1, \dots, r_M\}$, write

$$V_{a,r} := \{v \in V_N : A(v) = a, R(v) = r\}.$$

By the strong law and the multiplicative structure,

$$\frac{|V_{a,r}|}{N} \longrightarrow \sum_{b_1, b_2} \mu_{b_1, b_2, r, a} = c r \mu_a \quad \text{a.s. as } N \rightarrow \infty.$$

In [11], pair types are indexed as X – Y with X the *patient* blood type and Y the *donor* blood type. Translating into our (b_1, b_2) notation via $(b_1, b_2) = (Y, X)$, a type X – Y in [11] corresponds to *patient type* X and *donor type* Y , i.e., our $(b_1, b_2) = (Y, X)$.

Using that convention, [11] partition the 4×4 blood-type pairs X – Y as follows:

- self-demanded: X – X (patient and donor have the same blood type);
- over-demanded: X – Y with $Y \rightarrow X$ and $X \neq Y$ (donor offers a more flexible blood type than the patient needs);

- under-demanded: $X-Y$ with $X \rightarrow Y$ and $X \neq Y$ (patient seeks a more flexible type than the donor offers);
- reciprocally demanded: $A-B$ and $B-A$.

Under the numerical assumptions on blood-type frequencies and sensitization, they show that in almost every large pool there exists an optimal allocation (with exchanges of length at most 3) whose limiting structure is deterministic in terms of these blood-type classes. After translating their $X-Y$ types into our (b_1, b_2) notation, the key consequences *within each group* ($A = a, R = r$) are:

- all self-demanded types are asymptotically fully matched;
- all $A-B$ and $B-A$ *patient-donor* types are asymptotically fully matched via 2-cycles and 3-cycles;
- all $AB-O$ *patient-donor* types are matched using 3-cycles involving overdemanded types;
- all remaining overdemanded types are matched to their corresponding under-demanded types in 2-cycles;
- only the genuinely underdemanded types beyond available overdemanded partners remain unmatched.

Counting these matched classes in our (b_1, b_2) notation, and using the product form

$$\mu_{b_1, b_2, r, a} = c r \mu_a \mu_{b_1|a} \mu_{b_2|a},$$

gives exactly the expressions $T_a(r)$ and $S_a(r)$ stated before the proposition. In particular, for each a and r the unconstrained selection probability in group ($A = a, R = r$) converges to

$$p_a(r) := \frac{T_a(r)}{S_a(r)}.$$

Aggregating across a gives the total matchable mass at level r , proportional to $Q(r)$ as defined in the statement.

Step 2: Reallocatable mass between groups and R_a . The Ashlagi–Roth construction leaves some blood-type classes (particularly those involving underdemanded types) only partially utilized. Within such classes we can reassign which fraction of matched pairs comes from $A = 0$ versus $A = 1$, provided we respect:

- for each blood-type class, the total matched mass in that class; and
- for each group $A = a$, the total number of available pairs of that class.

For each “bottleneck” class (e.g., patient–donor $B-AB$, $O-AB$, $A-AB$, $A-O$, $AB-O$, $O-B$ in the Ashlagi–Roth notation) this gives a small linear program in two variables (mass from $A = 0$ and from $A = 1$). Solving these programs and translating back to our (b_1, b_2) representation gives the quantities R_a as in the proposition:

$$R_a = \min \{ \phi_{B,AB}, 2\mu_a \mu_{B|a} \mu_{AB|a} - \phi_{B,AB} \} + \min \{ \phi_{O,AB} + \phi_{A,AB}, 2\mu_a \mu_{A|a} \mu_{AB|a} - \phi_{O,AB} - \phi_{A,AB} \} \\ + \min \{ \phi_{A,O} + \phi_{AB,O}, 2\mu_a \mu_{O|a} \mu_{AB|a} - \phi_{A,O} - \phi_{AB,O} \} + \min \{ \phi_{O,B}, 2\mu_a \mu_{O|a} \mu_{B|a} - \phi_{O,B} \}.$$

Informally, R_a is the maximal amount of matched mass that can be shifted in favour of group $A = a$ by reassigning responsibility for matches within these partially utilized blood-type classes.

Step 3: Enforcing equality of selection probabilities at each r . Fix a sensitization level r . We aim for the stronger condition

$$\frac{\# \text{matches in } (A = 1, R = r)}{|V_{1,r}|} = \frac{\# \text{matches in } (A = 0, R = r)}{|V_{0,r}|},$$

i.e., $p_1(r)$ and $p_0(r)$ are made equal. This is stricter than (4), so any loss bound obtained here is valid for our weaker fairness constraint.

Suppose first that $p_1(r) < p_0(r)$, so group $A = 1$ is disadvantaged.

(a) *Using R_1 to help $A = 1$.* By reallocating within partially matched classes we can increase the matched mass of $(A = 1, R = r)$ by at most R_1 , without changing the total mass at level r . This yields a new numerator $T_1(r) + R_1$ for $A = 1$ and reduces the numerator for $A = 0$ by the same R_1 .

If this already suffices for equality,

$$\frac{T_1(r) + R_1}{S_1(r)} \geq \frac{T_0(r)}{S_0(r)},$$

then equality can be achieved with no loss of total matches at level r .

(b) *If reallocation is insufficient, discarding from $A = 0$.* Otherwise we must additionally discard matches from $(A = 0, R = r)$. Let $x(r) \geq 0$ be the extra amount of matched mass we remove from group $A = 0$ at level r , beyond the R_1 already reallocated. Equality of selection probabilities requires

$$\frac{T_1(r) + R_1}{S_1(r)} = \frac{T_0(r) - x(r)}{S_0(r)},$$

hence

$$x(r) = \frac{S_0(r)T_1(r) - S_1(r)T_0(r) - S_0(r)R_1}{S_1(r)}.$$

(c) *Symmetric case $p_1(r) \geq p_0(r)$.* If $p_1(r) \geq p_0(r)$, the roles of the two groups are reversed: we can first reallocate up to R_0 in favour of $A = 0$ and, if needed, discard an additional amount $x'(r) \geq 0$ from group $A = 1$, where

$$x'(r) = \frac{S_1(r)T_0(r) - S_0(r)T_1(r) - S_1(r)R_0}{S_0(r)}.$$

(d) *Relative loss at level r .* At level r , the total matchable mass in the unconstrained optimum is proportional to $Q(r)$; hence the relative loss (price of fairness) at level r is at most

$$\max \left\{ \frac{S_1(r)T_0(r) - S_0(r)T_1(r) - S_0(r)R_1}{S_1(r)Q(r)}, \frac{S_0(r)T_1(r) - S_1(r)T_0(r) - S_1(r)R_0}{S_0(r)Q(r)}, 0 \right\},$$

which is exactly the quantity appearing in the statement.

Step 4: Implementability via subgroup prioritization and Proposition 1.

Up to now the construction has been at the level of limiting type proportions. We now argue that it can be implemented in the random graphs G_N (with exchanges of bounded length) up to $o(N)$ loss.

For each triple (b_1, b_2, r) and $a \in \{0, 1\}$, the reallocation and discarding operations in Step 3 determine whether, at that type, the fair allocation effectively *prioritizes* group $A = 1$ or $A = 0$. Collect these choices in a priority set $\mathcal{P}$ as in Sect. 4.1, and consider the optimization problem with subgroup-prioritization constraints (6)–(7) associated with $\mathcal{P}$.

By Proposition 1, there exists a sequence $\varepsilon_N \downarrow 0$ such that, almost surely, we can impose the asymptotic priority requirements (7) for all $(b_1, b_2, i, j) \in \mathcal{P}$ while losing only $o(N)$ matches relative to the unconstrained optimum. In other words, the prioritized solutions implement, up to $o(N)$, the type-level pattern constructed in Step 3.

Within this prioritized class, we can then carry out the same trimming as in Step 3 to enforce equality of selection probabilities at each r . The additional loss in the number of matches is, at level r , bounded by the relative loss above, and hence the overall asymptotic price of fairness is bounded by

$$\max_{r \in \{r_1, \dots, r_M\}} \max \left\{ \frac{S_1(r)T_0(r) - S_0(r)T_1(r) - S_0(r)R_1}{S_1(r)Q(r)}, \frac{S_0(r)T_1(r) - S_1(r)T_0(r) - S_1(r)R_0}{S_0(r)Q(r)}, 0 \right\}.$$

Finally, the original fairness constraint (4) requires only that the difference in selection probabilities at each r_j be bounded by l_j , not that they be exactly equal. Thus an optimal solution under (4) is at least as efficient as the equalized allocations we have constructed, so its asymptotic price of fairness is no larger than the bound stated in Proposition 3.

This completes the proof. $\square$

Appendix B More Simulations with Failure-Aware Strategies

We consider a similar data-generating process in Sect. 5.1 but with $N = 50$, and we repeat exactly the same data-generating process in Sect. 5.2. With additional vertex and edge uncertainties, let the failure probability p_v and p_{v_i, v_j} be independently sampled from a uniform distribution $U(0, 0.3)$. We consider the subset-recourse strategy and choose $\mathcal{S}$ to be the set of relevant subsets of size $(3, 1)$.

Figures 4 and 5 present the average selection rates of different fairness criteria within each subgroup. The numerical performance with failure-aware strategies is similar to that without failure-aware strategies.

Funding Li's research was partially supported by NSF grant CIF-2102227, NIH grants R01AG062542 and R01AG080043.

Dai was supported in part by NIH grant R01DK142026, a Merck Research Award, and a Hellman Fellowship Award.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

References

- Klimentova X, Pedroso JP, Viana A (2016) Maximising expectation of the number of transplants in kidney exchange programmes. *Comput Operat Res* 73:1–11
- Li Y, Song PX-K, Zhou Y, Leichtman AB, Rees MA, Kalbfleisch JD (2014) Optimal decisions for organ exchanges in a kidney paired donation program. *Stat Biosci* 6:85–104
- Pedroso JP (2014) Maximizing expectation on vertex-disjoint cycle packing. In: *Computational Science and Its Applications–ICCSA 2014: 14th International Conference*, Guimarães, Portugal, June 30–July 3, 2014, Proceedings, Part II 14, pp. 32–46. Springer
- Bray M, Wang W, Song PX-K, Kalbfleisch JD (2018) Valuing sets of potential transplants in a kidney paired donation network. *Stat Biosci* 10:255–279
- Wang W, Bray M, Song PX-K, Kalbfleisch JD (2017) A look-ahead strategy for non-directed donors in kidney paired donation. *Stat Biosci* 9:453–469
- McElfresh DC, Bidkhor H, Dickerson JP (2019) Scalable robust kidney exchange. *Proc AAAI Conf Artif Intel* 33(01):1077–1084
- Dickerson J, Procaccia AD, Sandholm T (2014) Price of fairness in kidney exchange. *13th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2014* 2, 1013–1020
- Freedman R, Borg JS, Sinnott-Armstrong W, Dickerson JP, Conitzer V (2020) Adapting a kidney exchange algorithm to align with human values. *Artif Intell* 283:103261
- Farnadi G, St-Arnaud W, Babaki B, Carvalho M (2021) Individual fairness in kidney exchange programs. *Proc AAAI Conf Artif Intel* 35(13):11496–11505
- St-Arnaud W, Carvalho M, Farnadi G (2022) Adaptation, comparison and practical implementation of fairness schemes in kidney exchange programs. *arXiv preprint [arXiv:2207.00241](https://arxiv.org/abs/2207.00241)*
- Ashlagi I, Roth AE (2014) Free riding and participation in large scale, multi-hospital kidney exchange. *Theor Econ* 9(3):817–863
- Klimentova X, Viana A, Pedroso JP, Santos N (2021) Fairness models for multi-agent kidney exchange programmes. *Omega* 102:102333
- Carvalho M, Lodi A (2023) A theoretical and computational equilibria analysis of a multi-player kidney exchange program. *Eur J Oper Res* 305(1):373–385
- Bromberger B, Spragan D, Hashmi S, Morrison A, Thomasson A, Nazarian S, Sawinski D, Porrett P (2017) Pregnancy-induced sensitization promotes sex disparity in living donor kidney transplantation. *J Am Soc Nephrol* 28(10):3025–3033
- García-Soriano D, Bonchi F (2020) Fair-by-design matching. *Data Min Knowl Disc* 34:1291–1335
- Karni G, Rothblum GN, Yona G (2022) On fairness and stability in two-sided matchings. In: Braverman M (ed.) *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 215, pp. 92–19217. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany. <https://doi.org/10.4230/LIPIcs.ITCS.2022.92>
- Hardt M, Price E, Srebro N (2016) Equality of opportunity in supervised learning. *Adv Neural Inf Proc Syst*
- Chouldechova A (2017) Fair prediction with disparate impact: a study of bias in recidivism prediction instruments. *Big data* 5(2):153–163

19. Kleinberg J, Mullainathan S, Raghavan M (2017) Inherent Trade-Offs in the Fair Determination of Risk Scores. In: 8th Innovations in Theoretical Computer Science Conference (ITCS 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 67, pp. 43–14323
20. Liebl D, Reimherr M (2023) Fast and fair simultaneous confidence bands for functional parameters. *J R Stat Soc Ser B Stat Methodol* 85(3):842–868
21. Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper Res* 59(1):17–31
22. Viviano D, Bradic J (2024) Fair policy targeting. *J Am Stat Assoc* 119(545):730–743
23. Wang W, Bray M, Song PX, Kalbfleisch JD (2019) An efficient algorithm to enumerate sets with fallbacks in a kidney paired donation program. *Operat Res Health Care* 20:45–55
24. Corbett-Davies S, Gaebler JD, Nilforoshan H, Shroff R, Goel S (2023) The measure and mismeasure of fairness. *J Mach Learn Res* 24(1):14730–14846
25. Saidman SL, Roth AE, Sönmez T, Ünver MU, Delmonico FL (2006) Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges. *Transplantation* 81(5):773–782
26. Erdős P, Rényi A (1968) On random matrices ii. *Studia Sci Math Hungar* 3:459–464
27. Ashlagi I, Roth A (2011) Individual rationality and participation in large scale, multi-hospital kidney exchange. In: *Proceedings of the 12th ACM Conference on Electronic Commerce*, pp. 321–322

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.