
Improving Policy-Constrained Kidney Exchange via Pre-Screening

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In barter exchanges, participants swap goods with one another without exchanging
2 money; these exchanges are often facilitated by a central clearinghouse, with the
3 goal of maximizing the aggregate quality (or number) of swaps. Barter exchanges
4 are subject to many forms of uncertainty—in participant preferences, the feasibility
5 and quality of various swaps, and so on. Our work is motivated by kidney exchange,
6 a real-world barter market in which patients in need of a kidney transplant swap
7 their willing living donors, in order to find a better match. Modern exchanges
8 include 2- and 3-way swaps, making the kidney exchange *clearing problem* NP-
9 hard. Planned transplants often *fail* for a variety of reasons—if the donor organ is
10 rejected by the recipient’s medical team, or if the donor and recipient are found
11 to be medically incompatible. Due to 2- and 3-way swaps, failed transplants can
12 “cascade” through an exchange; one US-based exchange estimated that about 85%
13 of planned transplants failed in 2019. Many optimization-based approaches have
14 been designed to avoid these failures; however most exchanges cannot implement
15 these methods, due to legal and policy constraints. Instead, we consider a setting
16 where exchanges can *query* the preferences of certain donors and recipients—asking
17 whether they would accept a particular transplant. We characterize this as a two-
18 stage decision problem, in which the exchange program (a) queries a small number
19 of transplants before committing to a matching, and (b) constructs a matching
20 according to fixed policy. We show that selecting these edges is a challenging
21 combinatorial problem, which is non-monotonic and non-submodular, in addition to
22 being NP-hard. We propose both a greedy heuristic and a Monte Carlo tree search,
23 which outperforms previous approaches, using experiments on both synthetic data
24 and real kidney exchange data from the United Network for Organ Sharing.

25 1 Introduction

26 We consider a multi-stage decision problem in which a decision-maker uses a fixed *policy* to solve
27 a hard (stochastic) problem. Before using the policy, the decision-maker can first *measure* some of
28 the uncertain problem parameters—in a sense, guiding the policy toward a better solution. Our primary
29 motivation is kidney exchange, a process where patients in need of a kidney transplant swap their
30 (willing) living donors, in order to find a better match. Many government-run kidney exchanges match
31 patients and donors using a *matching algorithm* that follows strict policy guidelines [9]; this matching
32 algorithm is often written into law or policy, and is not easily modified. Modern kidney exchanges
33 use both cyclical swaps and chain-like structures (initiated by an unpaired altruistic donor) [25], and
34 identifying the max-size or max-weight set of transplants is both NP- and APX-hard [1, 7].

35 In kidney exchange—as in many resource allocation settings—information used by the decision-
36 maker is subject to various forms of uncertainty. Here we are primarily concerned with uncertainty
37 in the *feasibility* of potential transplants: if a donor is matched with a potential recipient, will the
38 transplant actually occur? Planned transplants may *fail* for a variety of reasons: for example, medical

39 testing may reveal that the donor and recipient are incompatible (a *positive crossmatch*); the recipient
40 or their medical team may reject a donor organ in order to wait for a better match; or the donor
41 may decide to donate elsewhere before the exchange is planned. Failed transplants are especially
42 troublesome in kidney exchange, due to the cycle and chain structures used: for example, suppose
43 that a cyclical swap is planned between three patient/donor pairs; if any one of the planned transplants
44 fails, then none of the other transplants in that cycle can occur. Unfortunately, it is quite common for
45 planned transplants to fail. For example, the United Network for Organ Sharing (UNOS¹) estimates
46 that in FY2019, about 85% of their planned kidney transplants failed [18].

47 Various matching algorithms have been proposed that aim to mitigate transplant failures (for exam-
48 ple, using stochastic optimization [15, 3], robust optimization [22], or conditional value at risk [6]).
49 However, implementing these strategies would require modifying fielded matching algorithms—which
50 in many cases would require changing law or policy. One way to avoid failures without modifying
51 the matching algorithm is to *pre-screen* potential transplants [18, 10, 11], by communicating with the
52 recipients’ medical team and possibly using additional medical tests. Pre-screening transplants is
53 costly, as it requires scarce time and resources. Furthermore, there are often many thousand potential
54 transplants in any given exchange; selecting which transplants to screen is not easy.

55 In this paper we investigate methods for selecting a limited number of transplants to pre-screen, in
56 order to “guide” the fixed matching algorithm to a better outcome. We formalize this as a multistage
57 stochastic optimization problem, and we consider both an *offline* setting (where screenings are
58 selected all at once), and an *online* setting (where screenings are selected sequentially).

59 **Related Work.** While kidney exchange is known to be a hard packing problem, several algorithms
60 exist that are scalable in practice, and are used by fielded exchanges [14, 3, 20]. Prior work has
61 addressed potential transplant failures; our model is inspired by Dickerson et al. [15]. Pre-screening
62 potential transplants has also been addressed in prior work ([11, 23], and § 5.1 of [13]), and our model
63 is similar to stochastic matching and stochastic k -set packing [5]. However there are substantial
64 differences between these models and ours: (a) many prior approaches assume that a large number
65 of transplants may be pre-screened [11, 23]—on the order of one for each patient in the exchange;
66 we assume far fewer screenings are possible; (b) prior work often assumes a *query-commit* setting—
67 where successfully pre-screened transplants *must* be matched. Instead we assume that non-screened
68 transplants may also be matched—which more-accurately represents the way that modern exchanges
69 operate; (c) most prior work assumes that transplants that pass pre-screening are guaranteed to result
70 in a transplant. In reality, transplants often fail after pre-screening, a fact reflected in our model.

71 One of our approaches is based on *Monte Carlo Tree Search* (MCTS), which allows efficient explo-
72 ration of intractably large decision trees. While MCTS is primarily associated with Markov decision
73 processes and game-playing [12], it has been used successfully for combinatorial optimization [16].
74 We use a version of MCTS, Upper Confidence Bounds for Trees (UCT), which balances exploration
75 and exploitation by treating each tree node as a multi-armed bandit problem [4, 17].

76 Our Contributions

- 77 1. (§ 2) We formalize the *policy-constrained edge query problem*: where a decision-maker
78 (such as a kidney exchange program) selects a set of potential edges (potential transplants) to
79 pre-screen, prior to constructing a final packing (a set of transplants) using a fixed algorithm.
80 This model generalizes existing models in the literature, as edge failure probabilities depend
81 on whether or not the edge is pre-screened. Further, we allow for context-specific constraints,
82 such as those imposed by public policy or the particular hospital or exchange.
- 83 2. (§ 3) We prove that when the decision-maker uses a max-weight packing policy (the
84 most common choice among fielded exchanges), the edge query problem is both non-
85 monotonic and non-submodular in the set of queried edges. Despite these worst-case
86 findings we show that this problem is nearly monotonic for real and synthetic data, and
87 simple algorithms perform quite well. On the other hand, when the decision-maker uses
88 a *failure-aware* (stochastic) packing policy, the edge query problem becomes monotonic
89 under mild assumptions.
- 90 3. (§ 4) We conduct numerical experiments on both simulated and real exchange data from the
91 United Network for Organ Sharing (UNOS). We demonstrate that our methods substantially
92 outperform prior approaches and a randomized baseline.

¹UNOS is the organization tasked with overseeing organ transplantation in the US: <https://unos.org/>.

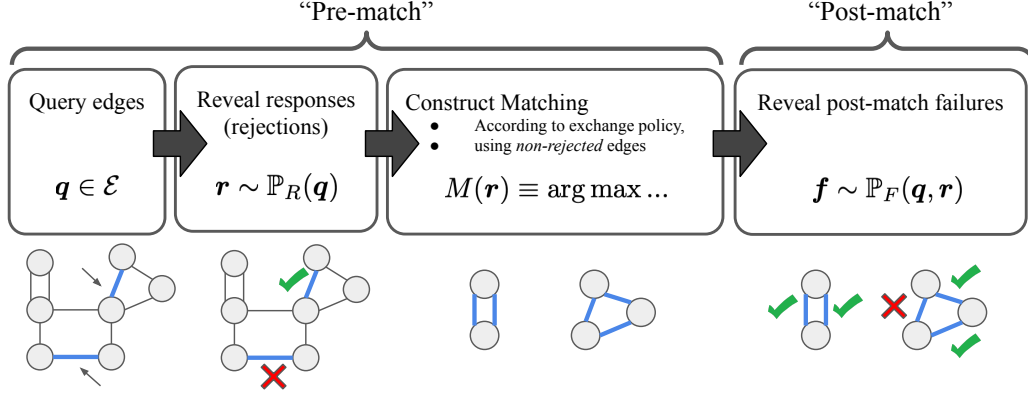


Figure 1: Single-stage edge selection: First, edges are selected to be queried, and responses revealed. Then, a final matching is constructed according to the exchange’s matching policy. Finally, the post-match edge failures are revealed.

93 2 The Policy-Constrained Edge Query Problem

94 Kidney exchanges are represented by a graph $G = (E, V)$ where vertices V represent (incom-
 95 compatible) patient-donor pairs, and non-directed donors (NDDs) who are willing to donate without
 96 receiving a kidney in return. Directed edges $e \in E$ between vertices represent potential transplants
 97 from the donor of one vertex to the patient of another. Edge weights represent the “utility” of an
 98 edge, and are typically set by exchange policy. Solutions to a kidney exchange problem (henceforth,
 99 *matchings*) consist of both directed *cycles* on G containing only patient-donor pairs, and directed
 100 *chains* beginning with an NDD and passing through one or more pairs; see Appendix A for an
 101 example exchange graph. Each vertex may participate in only one edge in a matching—as each vertex
 102 can donate and receive at most one kidney.

103 Vectors are denoted in bold, and are indexed by either cycles or edges: \mathbf{y}_e indicates the element of
 104 \mathbf{y} corresponding to edge e , and \mathbf{x}_c is the element of \mathbf{x} corresponding to cycle c . Our notation uses a
 105 *cycle-chain* representation for matchings²: let \mathcal{C} represent cycles and chains in G , where each cycle
 106 and chain corresponds to a list of edges; as is standard in modern exchanges, we assume that cycles
 107 and chains are limited in length. Matchings are expressed as a binary vector $\mathbf{x} \in \{0, 1\}^{|\mathcal{C}|}$, where
 108 $\mathbf{x}_c = 1$ if cycle/chain c is in the matching, and 0 otherwise. Let w_c be the weight of cycle/chain c
 109 (the sum of c ’s edge weights). Let \mathcal{M} denote the set of *legal matchings*—that is, the set of vertex-
 110 disjoint cycles and chains on G , with chains up to length L and cycles up to length C . Cycle length
 111 cap C and chain length cap L are set by the each exchange, typically $C = 3$ and $L = 4$. These
 112 length limits serve two purposes: (a) longer cycles and chains are risky, in that they are likely to
 113 be impacted by edge failure, and (b) policy often requires that all transplants in a cycle or chain
 114 are completed *simultaneously*, and most transplant centers can only accommodate a handful of
 115 simultaneous transplants. The total weight of a matching is simply the summed weights of all its
 116 constituent cycles and chains: $\sum_{c \in \mathcal{C}} \mathbf{x}_c w_c$. We denote *sets* of edges using binary vectors, where
 117 $\mathbf{q} \in \{0, 1\}^{|E|}$ represents the set of all edges with $\mathbf{q}_e = 1$.

118 In the remainder of this paper we refer to pre-screening a transplant as *querying an edge*, in order
 119 to be consistent with the literature.

120 **Selecting Edge Queries.** Our setting consists of two phases (see Figure 1): during *pre-match*, the
 121 decision-maker selects edges to query, and each queried edge is either accepted or rejected; then
 122 the decision-maker constructs a matching using a fixed policy. During *post-match*, each match edge
 123 either fails (no transplant) or succeeds (the transplant proceeds). We consider two version of the
 124 pre-match phase: in the *single-stage* version, the decision-maker selects all queries before observing
 125 edge responses (accept/reject); in the *multi-stage* version, one edge is selected at a time and responses
 126 are observed immediately.

127 Unlike most prior work, edges in our model may fail during both the pre- and post-match phase.
 128 For example, suppose the decision-maker queries an edge from a 60-year-old non-directed donor, to

²Our experiments use the position-indexed formulation, which is more compact and equivalent [14].

129 a 35-year-old recipient; if the recipient or their medical team rejects the elderly donor and decides
 130 to wait for a younger donor, this is a pre-match rejection. Instead suppose the edge is not queried,
 131 and it is included in the final matching; if medical screening reveals that the patient and donor are
 132 incompatible, this is a post-match failure. We refer to pre-match failures as *rejections* and post-match
 133 failures as *failures*; however we make no assumption about their cause. We represent potential failures
 134 and rejections using binary random variables: $\mathbf{r} \in \{0, 1\}^{|E|}$ denotes pre-match rejections, where
 135 $r_e = 1$ if e is queried and rejected, and 0 otherwise ($r_e = 0$ for all non-queried edges). Similarly
 136 $\mathbf{f} \in \{0, 1\}^{|E|}$ denotes post-match failures, where $f_e = 1$ if edge e fails post-match, and 0 otherwise.
 137 We assume that the distribution of rejections $\mathbf{r} \sim \mathbb{P}_R(\mathbf{q})$ is known, and depends on \mathbf{q} ; we assume the
 138 distribution of failures $\mathbf{f} \sim \mathbb{P}_F(\mathbf{q}, \mathbf{r})$ is known, and depends on both \mathbf{q} and \mathbf{r} .

Rejections and failures impact the matching through the *weight* of each cycle and chain. If any cycle edge fails, then *no* transplants in the cycle can proceed; if a chain edge fails, than all edges *following* it cannot proceed.³ Suppose we observe failures \mathbf{f} ; the *final matching weight* of c is

$$F(c, \mathbf{y}) \equiv \begin{cases} \sum_{e \in c} w_e & \text{if } \sum_{e \in c} \mathbf{y}_e = 0 \\ 0 & \text{if } c \text{ is a cycle and } \sum_{e \in c} \mathbf{y}_e > 0 \\ \sum_{e \in c'} w_e & \text{if } c \text{ is a chain, where } c' \text{ includes all edges up to the first failed edge.} \end{cases}$$

139 Thus the *post-match expected weight* of matching \mathbf{x} , due to both rejections \mathbf{r} and failures \mathbf{f} , is

$$W(\mathbf{x}; \mathbf{q}, \mathbf{r}) \equiv \mathbb{E}_{\mathbf{f} \sim \mathbb{P}_F(\mathbf{q}, \mathbf{r})} \left[\sum_{c \in \mathcal{C}} \mathbf{x}_c F(c, \mathbf{r} + \mathbf{f}) \right].$$

140 **Matching Policy** In this paper we assume that the final matching is constructed using a fixed matching
 141 policy, which uses only *non-rejected* edges; we denote this policy by $M(\mathbf{r})$. We focus primarily
 142 on the *max-weight* policy $M^{\text{MAX}}(\cdot)$, which is used by most fielded exchanges, and the *failure-aware*
 143 policy $M^{\text{FA}}(\cdot)$, which maximizes the expected post-match weight [15]:

$$M^{\text{MAX}}(\mathbf{r}) \in \arg \max_{\mathbf{x} \in \mathcal{M}} \sum_{c \in \mathcal{C}} \mathbf{x}_c F(c, \mathbf{r}), \quad M^{\text{FA}}(\mathbf{r}) \in \arg \max_{\mathbf{x} \in \mathcal{M}(\mathbf{r})} \mathbb{E}_{\mathbf{f} \sim \mathbb{P}_F(\mathbf{q}, \mathbf{r})} \left[\sum_{c \in \mathcal{C}} \mathbf{x}_c F(c, \mathbf{r} + \mathbf{f}) \right].$$

144 Evaluating this policy requires solving a kidney exchange clearing problem, which is NP-hard [1].
 145 However, state-of-the-art method can solve realistic kidney exchange clearing problems in fractions
 146 of a second (e.g., our experiments use the PICEF method of Dickerson et al. [14]); thus, throughout
 147 this paper we treat this policy as a low- or no-cost oracle.

148 Next we formalize the *edge selection problem*—the main focus of this paper. We denote by \mathcal{E} the
 149 set of “legal” edge subsets, subject to exchange-specific constraints; we assume that \mathcal{E} is a matroid
 150 with ground set E . For example, the decision-maker may limit the number of queries issued to any
 151 one medical team (vertex in G) or transplant center (group of vertices). We aim to select an edge set
 152 $\mathbf{q} \in \mathcal{E}$ which maximizes the *expected weight* of the final matching. These edges are selected using
 153 only the distribution of future rejections and failures; we take a *stochastic optimization* approach,
 154 maximizing the expected outcome over this uncertainty.

155 **Single-Stage Setting.** The single-stage policy-constrained edge selection problem (henceforth, the
 156 *edge selection problem*) is expressed as

$$\max_{\mathbf{q} \in \mathcal{E}} V^S(\mathbf{q}), \quad \text{with} \quad V^S(\mathbf{q}) \equiv \mathbb{E}_{\mathbf{r} \sim \mathbb{P}_R(\mathbf{q})} [W(M(\mathbf{r}); \mathbf{q}, \mathbf{r})], \quad (1)$$

157 where, $M(\mathbf{r})$ denotes the matching policy after observing rejections \mathbf{r} , and $W(\mathbf{x}; \mathbf{q}, \mathbf{r})$ denotes the
 158 post-match expected weight of matching \mathbf{x} . Exact evaluation of $V^S(\mathbf{q})$ is often intractable, as the
 159 support of $\mathbb{P}_R(\mathbf{q})$ grows exponentially in $|\mathbf{q}|$. In experiments we approximate $V^S(\mathbf{q})$ using sampling,
 160 and these approximations converge for a moderate number of samples (see Appendix B).

161 **Multistage Setting.** In the multi-stage setting, edge rejections are observed immediately after each
 162 edge is queried. The multi-stage problem is expressed as

$$\max_{\mathbf{q}^1 \in \mathcal{E}_1} \mathbb{E}_{\mathbf{r}^1 \sim \mathbb{P}_R(\mathbf{q}^1)} \left[\max_{\mathbf{q}^2 \in \mathcal{E}_1} \mathbb{E}_{\mathbf{r}^2 \sim \mathbb{P}_R(\mathbf{q}^2)} \left[\dots \max_{\mathbf{q}^K \in \mathcal{E}_1} \mathbb{E}_{\mathbf{r}^K \sim \mathbb{P}_R(\mathbf{q}^K)} [W(M(\mathbf{r}); \mathbf{q}, \mathbf{r})] \dots \right], \quad (2)$$

³This assumes that chains can be *partially* executed: for example, suppose that the 4th edge in a 10-edge chain fails; the first three edges can still be matched, and the post-failure chain weight sums only these three edges. Not all fielded exchanges use this policy: some exchanges cancel the entire chain if one of its edges fails.

163 where $\mathbf{q} \equiv \sum_{i=1}^K \mathbf{q}^i$ denotes all queried edges, $\mathbf{r} \equiv \sum_{i=1}^K \mathbf{r}^i$ denotes all rejections, and $\mathcal{E}_1 \subseteq \mathcal{E}$ be
 164 denotes the legal edge subsets containing only one edge. First, we observe that Problems 1 and 2
 165 require evaluating a matching policy $M(\mathbf{r})$. In the case of kidney exchange, evaluating both the
 166 max-weight policy $M^{\text{MAX}}(\cdot)$ and the failure-aware policy $M^{\text{FA}}(\cdot)$ require solving NP-hard problems;
 167 thus Problems 1 and 2 are at least NP-hard as well.

168 However, regardless how difficult the matching policy is, the question remains whether *edge*
 169 *selection* is hard. We observe that while these problems are difficult in principle, experiments (§ 4)
 170 show that they are easy in practice. Proofs of the following propositions can be found in Appendix D.

171 **Proposition 2.1.** *With matching policy $M^{\text{MAX}}(\cdot)$, the objective of Problem 1 is non-monotonic in the*
 172 *number of queried edges, even with independent edge distributions.*

173 In other words, querying additional edges can sometimes lead to a *worse* outcome. This is
 174 somewhat counter-intuitive; one might think that providing additional information to the matching
 175 policy would strictly improve the outcome. This is a worst-case result—and in fact our experiments
 176 demonstrate that querying edges almost always leads to a better final matching weight.

177 **Proposition 2.2.** *With matching policy $M^{\text{MAX}}(\cdot)$, the objective of Problem 1 is non-submodular in*
 178 *the set of queried edges.*

179 In other words, certain edges are *complementary* to each other—and querying complementary edges
 180 simultaneously can yield a greater improvement than querying them separately. Taken together, these
 181 propositions indicate that single-stage edge selection with matching policy $M^{\text{MAX}}(\cdot)$ is a challenging
 182 combinatorial optimization problem. On the other hand, using the failure-aware matching policy
 183 $M^{\text{FA}}(\cdot)$ allows us to avoid some of these issues under mild assumptions.

184 **Assumption 2.3.** *Let $\mathbf{q}, \mathbf{r} \in \{0, 1\}^{|E|}$ denote initial edge queries and responses. Let \mathbf{q}' be additional*
 185 *edges, such that $\mathbf{q} + \mathbf{q}' \in \{0, 1\}^{|E|}$ denotes an augmented edge set; let $\mathbf{r}' \in \{0, 1\}^{|E|}$ denote the*
 186 *responses to edges \mathbf{q}' only. We assume that for any such \mathbf{q}, \mathbf{r} , and \mathbf{q}' ,*

$$\mathbb{E}[\mathbf{r} + \mathbf{f} \mid \mathbf{q}, \mathbf{r}] \geq \mathbb{E}[\mathbf{r} + \mathbf{r}' + \mathbf{f} \mid \mathbf{q} + \mathbf{q}', \mathbf{r}].$$

187 Intuitively, Assumption 2.3 excludes distributions where queries arbitrarily increase edge failure or
 188 rejection. For example, Assumption 2.3 disallows the following distribution: suppose all edges are
 189 independent; all queried edges are accepted ($P(\mathbf{r}_e = 1 \mid \mathbf{q}) = 0$ for all \mathbf{q}), all accepted edges have
 190 failure probability 0.5 ($P(\mathbf{f}_e = 1 \mid \mathbf{q}_e = 1, \mathbf{r}_e = 0) = 0.5$), and all non-queried edges have failure
 191 probability 0.1 ($P(\mathbf{f}_e = 1 \mid \mathbf{q}_e = \mathbf{r}_e = 0) = 0.1$). In this case, if an edge is not queried, then it has
 192 overall rejection or failure probability 0.1 (i.e., $\mathbb{E}[\mathbf{r}_e + \mathbf{f}_e \mid \mathbf{q}, \mathbf{r}] = 0.1$ with $\mathbf{q}_e = 0$); if this edge is
 193 queried, then it has rejection or failure probability 0.5 (i.e., $\mathbb{E}[\mathbf{r}_e + \mathbf{r}'_e + \mathbf{f}_e \mid \mathbf{q} + \mathbf{q}', \mathbf{r}] = 0.5$ with
 194 $\mathbf{q}'_e = 1$).

195 We also assume that edge failures are *independent*.

Definition 2.4 (Edge Independence). *Two edges $e, e' \in E$ are independent if (a) their rejection*
distributions are conditionally independent, given whether or not they were queried:

$$\mathbf{r}_e \perp\!\!\!\perp \mathbf{r}_{e'} \mid \mathbf{q}_e \quad \text{and} \quad \mathbf{r}_e \perp\!\!\!\perp \mathbf{r}_{e'} \mid \mathbf{q}_{e'}$$

and (b) their failure distributions are conditionally independent, given whether or not they were
queried and rejected:

$$\mathbf{f}_e \perp\!\!\!\perp \mathbf{f}_{e'} \mid \mathbf{q}_e, \mathbf{r}_e \quad \text{and} \quad \mathbf{f}_e \perp\!\!\!\perp \mathbf{f}_{e'} \mid \mathbf{q}_{e'}, \mathbf{r}_{e'}.$$

196 **Proposition 2.5.** *If edges are independent and Assumption 2.3 holds, then with a failure-aware*
 197 *matching policy the objective of Problem 1 is monotonic in the set of queried edges.*

198 While Propositions 2.1 and 2.2 state that single-stage edge selection is challenging in the worst
 199 case, our computational results suggest that these problems are often easier on realistic exchanges.

200 2.1 Using the Max-Weight Matching Policy as a Baseline

201 It might seem as though our edge pre-screening procedure is simply compensating for a flawed
 202 matching policy ($M^{\text{MAX}}(\cdot)$). If matching policy $M^{\text{MAX}}(\cdot)$ performs poorly in practice, then why not
 203 use $M^{\text{FA}}(\cdot)$? Indeed, $M^{\text{FA}}(\cdot)$ can directly account for edge failure, and Proposition 2.5 states that the
 204 edge selection problem is in fact “easy” when using this policy. However this assumes that the edge
 205 failure distribution is accurately known. The failure-aware matching policy $M^{\text{FA}}(\cdot)$ is very sensitive
 206 to the specified edge failure distribution, and if the assumed distribution is *incorrect* then this policy
 207 can perform very poorly. This is in fact a reason that $M^{\text{FA}}(\cdot)$ is not used in practice: edge failure

208 distributions are not accurately known, and exchange programs are hesitant—with good reason—to
 209 guide their matching policy with a noisy estimation of edge failure.

210 In our edge pre-screening setting, the assumed edge distribution may also be noisy, however
 211 this does not directly affect the matching algorithm. In the worst case, an incorrect edge failure
 212 distribution will lead us to pre-screen edges that do not provide useful information to the exchange
 213 (this is the status quo). This is in stark contrast to $M^{FA}(\cdot)$: in the worst case, a bad estimation of the
 214 edge failure distribution will cause $M^{FA}(\cdot)$ to match risky cycles and chains, potentially decreasing
 215 the number of transplants.

216 2.2 A Note on Match Run Frequency

217 In the real world kidney exchange is a *dynamic* process: patients are constantly entering and
 218 leaving the pool, and participants are matched every few weeks. One way to deal with edge failure
 219 uncertainty is to match *more frequently*. For example if we match patients and donors once every
 220 month, then each failed edge adds one month to the waiting time for every patient who relied on the
 221 failed transplant. If we match participants every few hours, this increase in waiting time is far less
 222 severe.

223 However there are good reasons to match *infrequently*. First, exchanges benefit with the addition
 224 of more patients and donors—new edges “thicken” the compatibility graph, enabling more cycles
 225 and chains. Second, each transplant center participating in exchange needs a transplant coordinator to
 226 manage individual patients and donors. Many small hospitals do not have a full-time staff for organ
 227 exchange, and they cannot deal with frequent match offers. For these reasons, most exchanges match
 228 patients and donors very infrequently: the UK national kidney exchange matches patients and donors
 229 every quarter⁴ (once every three months); the Canadian national exchange matches participants once
 230 every four months⁵; UNOS currently matches patients once every week.⁶

231 3 Solving the Policy-Constrained Edge Query Problem

232 First we propose an exhaustive tree search which returns an optimal solution to Problem 1 given
 233 enough time. Building on this, we propose a Monte Carlo Tree Search algorithm and a simple greedy
 234 algorithm. Our multi-stage approaches are very similar to these, and can be found in Appendix E.

235 Our optimal exhaustive search uses a *search tree* where each tree node corresponds to an edge
 236 subset in \mathcal{E} . The children of node \mathbf{q} correspond to any $\mathbf{q}' \in \mathcal{E}$ which are equivalent to the parent
 237 \mathbf{q} , but include one additional edge: $C(\mathbf{q}) \equiv \{(\mathbf{q} + \mathbf{q}') \mid \forall \mathbf{q}' \in \mathcal{E} : |\mathbf{q}'| = 1 \mid (\mathbf{q} + \mathbf{q}') \in \mathcal{E}\}$. We say
 238 that edge sets (or tree nodes) containing L edges are on the L^{th} level of the tree. We refer to nodes
 239 with no children as *leaf nodes*. Unlike other tree search settings, the optimal solution to Problem 1
 240 may be at *any* node of the tree, not only leaf nodes; this is a consequence of non-monotonicity (see
 241 Proposition 2.1). The tree defined by root node $\mathbf{q} = \mathbf{0}$ and child function $C(\mathbf{q})$ contains all legal edge
 242 subsets in \mathcal{E} , when \mathcal{E} is a matroid. Thus, *any* exhaustive tree search algorithm (such as depth-first
 243 search) will identify an optimal solution, given enough time and memory.

244 Of course exhaustive search is only tractable if \mathcal{E} is small. Consider the class of *budgeted* edge sets
 245 $\mathcal{E}(\Gamma)$ used in our experiments: $\mathcal{E}(\Gamma) \equiv \{\mathbf{q} \in \{0, 1\}^{|E|} \mid |\mathbf{q}| \leq \Gamma\}$ (edge sets containing at most Γ
 246 edges). The number of edge sets in $\mathcal{E}(\Gamma)$ grows roughly exponentially in Γ and $|E|$, and is impossible
 247 to enumerate even for small graphs. Suppose a graph has 50 edges and we have an edge budget of
 248 five: there are over two million edge sets in $\mathcal{E}(5)$. Even small exchange graphs can have thousands of
 249 edges, and thus $\mathcal{E}(\Gamma)$ cannot be enumerated. Therefore, we propose search-based approach.

Monte Carlo Tree Search for Edge Selection (MCTS): We propose a tree-search algorithm for
 single-stage edge selection, MCTS, based on Monte Carlo Tree Search (MCTS), with the Upper
 Confidence for Trees (UCT) algorithm [17]. Our approach keeps track of a *value* (the objective value
 of Problem 1) and a UCB value estimate for each node, and these values are updated during sampling.
 The formula used to estimate a node’s UCB value is

$$\frac{U - V^{\min}}{V^{\max} - V^{\min}} + \sqrt{N^P/N}$$

⁴<https://nhsbtbde.blob.core.windows.net/umbraco-assets-corp/24443/po1274-4.pdf>

⁵<https://profedu.blood.ca/en/organs-and-tissues/programs-and-services/kidney-paired-donation-kpd-program>

⁶https://unos.org/wp-content/uploads/unos/KPD_emanual_how-matching-works.pdf

250 where U is the “UCB value estimate” calculated by MCTS, N is the number of visits to the node, N^P
 251 is the number of visits to the node’s parent, and V^{max} and V^{min} are the largest and smallest node
 252 values encountered during search.

253 When the set of tree nodes is too large to enumerate UCT can use a huge amount of memory—by
 254 storing values for each visited node. To limit both memory use and runtime, we incrementally
 255 search the tree from a temporary root node. Beginning from the root (the empty edge set), we use
 256 UCB sampling on the next L levels of nodes—where L is a small fixed integer. After a fixed time
 257 limit, sampling stops and we set the *new* root node to the current root’s best child according to its
 258 UCB estimate—using the method of [17]. This process repeats until we reach the final level of the
 259 search tree. Algorithm 1 gives a pseudocode description of MCTS, which uses Algorithm 2 as a
 260 submethod. While often successful, MCTS requires extensive training and parameter tuning. As a
 261 simpler alternative, we propose a greedy algorithm.

262 **Single-Stage Greedy Algorithm:** Greedy. Like MCTS, our greedy algorithm (Greedy) begins with
 263 the empty edge set as the root node, and iteratively searches deeper levels of the tree. However unlike
 264 MCTS, Greedy simply selects the child node with the greatest objective value in Problem 1—that is,
 265 *greedily* improving the objective value; see Appendix E for a pseudocode description.

266

ALGORITHM 1: MCTS: Tree Search for Single-Stage Edge Selection	ALGORITHM 2: Sample: Sampling function used by MCTS
(input) K : maximum size of any legal edge set	(input) q, M
(input) T : time limit per level	
(input) L : number of look-ahead levels	
$q^R \leftarrow \mathbf{0}$ root node (no edges)	$N[q] \leftarrow N[q] + 1$
$q^* \leftarrow \mathbf{0}$ the best visited node	$V[q] \leftarrow$ objective of edge set q in Problem 1
$V^* \leftarrow$ objective value of q^*	if $V[q] > V^*$ then
for $N = 1, \dots, K$ do	$q^* \leftarrow q, V^* \leftarrow V[q]$
$M \leftarrow \min\{N + L, K\}$	if q has no children then
$Q \leftarrow$ all nodes in levels N to M	return $V[q]$
$U[q] \leftarrow 0 \forall q \in Q$ UCB value estimate	if q has children then
$V[q] \leftarrow 0 \forall q \in Q$ objective value	if $ q < M$ then
$N[q] \leftarrow 0 \forall q \in Q$ number of visits	$q' \leftarrow \arg \max_{q \in C(q^R)} U[q] + \text{UCB}[q]$
while less than time T has passed do	$U[q] \leftarrow U[q] + \text{Sample}(q', M)$
$\text{Sample}(q^R, M)$	else
$q^R \leftarrow \arg \max_{q \in C(q^R)} U[q]$	$q' \leftarrow$ a random descendent of q at any level
Delete $U[\cdot], V[\cdot]$, and $N[\cdot]$	$V' \leftarrow$ objective value of q' in Problem 1
return q^*	if $V' > V^*$ then
	$q^* \leftarrow q', V^* \leftarrow V'$
	$U[q] \leftarrow U[q] + V'$

267 **Runtime.** Our methods rely on an “oracle” to solve the NP-hard kidney exchange matching problem;
 268 while state-of-the-art methods solve real-sized instances of these problems in fractions of a second,
 269 there is no guaranteed bound for absolute runtime. Instead, we can report the *number of calls* to
 270 this oracle for each method as a measure of complexity. Both benchmark methods (max-weight
 271 matching and failure-aware [15]) as well as IIAB [11] use exactly one oracle call; i.e., they are
 272 $O(1)$. Both Greedy and MCTS use a fixed number of samples (M) to evaluate the objective of an
 273 edge set. Greedy evaluates the objective of an edge set exactly Γ times; thus, Greedy is $O(M \cdot \Gamma)$.
 274 Finally, MCTS can in theory visit all potential edge sets of size at most Γ (i.e., an exhaustive search),
 275 which is $O(M \cdot \sum_{\gamma=1}^{\Gamma} \binom{|E|}{\gamma})$. Since this version of MCTS is intractable in both runtime and memory,
 276 Algorithm 1 imposes reasonable limits on our implementation.

277 4 Computational Experiments

278 We conduct a series of computational experiments using both synthetic data, and real kidney
 279 exchange data from UNOS; all code for these experiments is available online.⁷ In these experiments,
 280 “legal” edge sets are the budgeted edge sets defined as $\mathcal{E}(\Gamma) \equiv \{q \in \{0, 1\}^{|E|} \mid |q| \leq \Gamma\}$.
 281 In Sections 4.2 and 4.3 we present results in the single- and multi-stage edge selection settings,
 282 respectively. We use both real data and synthetic data for our experiments.

⁷Link removed for review.

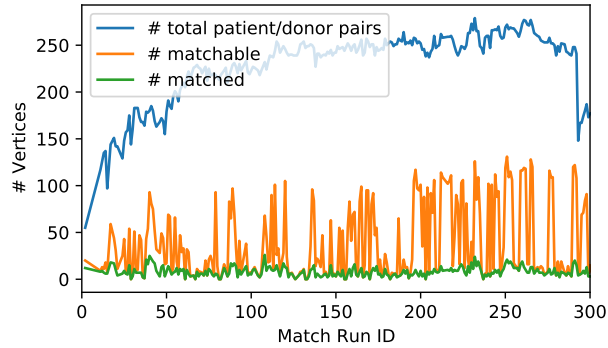


Figure 2: Number of patient donor pairs in each exchange, the number of matchable vertices (who can participate in a legal cycle or chain), and the number of vertices matched by $M^{\text{MAX}}(\cdot)$ in simulation.




Figure 3: Histogram of the number of *compatible* donors (edges) for each recipient, over all UNOS graphs.




Figure 4: Histogram of the number of *matchable* edges for each recipient, over all UNOS graphs.

283 4.1 Data

284 **Real Data.** We use exchange graphs from the United Network for Organ Sharing (UNOS), represent-
 285 ing UNOS match runs between 2010 and 2019. Some of these exchange graphs only have the trivial
 286 matching (no cycles or chains), or they have only one non-trivial matching. We ignore these graphs
 287 because the matching policy is a “constant” function (to return the one feasible matching) and edge
 288 queries cannot change the outcome. Removing these, we are left with 324 UNOS exchange graphs.

289 These exchange graphs are relatively small and sparse: most graphs have fewer than 250 vertices,
 290 and fewer of half of these can be matched via a legal cycle or chain (with cycle cap $C = 3$ and chain
 291 cap $L = 4$); we refer to these vertices as *matchable*. The number of NDDs (who can initiate chains)
 292 is extremely small: most exchanges have zero or 1 NDD. Figure 2 shows the number of vertices, the
 293 number of matchable vertices, and the number matched by $M^{\text{MAX}}(\cdot)$ (in simulation). These graphs are
 294 also sparse: Figure 3 shows the number of edges (compatible donors) for each patient over all UNOS
 295 graphs. The median number of edges per patients is 29, and 14% have only one edge. Furthermore,
 296 while many patients have multiple *compatible* edges, very few of these edges are *matchable* (see
 297 Figure 4): 72% of all patients cannot be matched via a legal cycle or chain; of the *matchable* patients
 298 (with one or more matchable donor), the median number of matchable donors is 3.

299 **Synthetic Data.** We generate random kidney exchange graphs based on directed Erdős-Rényi graphs
 300 defined using parameters N and p : let V be a fixed set of N vertices; for each pair of vertices (V_1, V_2)
 301 there is an edge from V_1 to V_2 with probability p , and an edge from V_2 to V_1 with probability p
 302 (independent of the edge from V_1 to V_2). Any vertices with no incoming edges are considered NDDs.
 303 In simulations we generate graphs which are smaller and more-sparse than the average UNOS graph
 304 ($N \leq 100$ and $p = 0.01$). These are meant to represent a toy-model of kidney exchange, rather than
 305 a realistic imitation of exchange graphs.

Table 1: Left: Optimality gap for Greedy, over 100 random graphs with $p = 0.01$ and various N , with edge budget $\Gamma = 3$; bottom row shows the maximum value of %OPT over all graphs. Right: Single-stage results on UNOS graphs using the variable IIAB edge budget (top rows), and the failure-aware method (bottom row). Columns P_X indicates the X^{th} percentile of Δ^{MAX} over all UNOS graphs.

%OPT	Num. Graphs (out of 100)			Method	Simple edge dist.			KPD edge dist.		
	$N = 50$	$N = 75$	$N = 100$		P_{10}	P_{50}	P_{90}	P_{10}	P_{50}	P_{90}
[0, 0.1]	93	93	90	MCTS	0.40	0.67	1.11	0.05	0.45	3.44
(0.1, 1]	5	4	9	Greedy	0.47	0.64	1.00	0.02	0.47	3.44
(1, 2]	1	3	1	Random	0.00	0.10	0.46	-0.11	0.00	0.63
(2, 100]	1	0	0	IIAB	0.21	0.45	0.89	-0.27	0.12	2.24
Max %OPT	2.8	1.5	1.0	Fail-Aware	0.00	0.09	0.23	-0.27 [†]	0.00 [†]	2.17 [†]

306 In these experiments edge rejections and failures are independently distributed for each edge e ; let
307 P_R be the rejection probability, P_Q is the post-match success probability if e is queried/accepted, and
308 P_N is the success probability if e is not queried. To simulate edge rejections and failures we use two
309 synthetic edge distributions: *Simple* and *KPD*. In the *Simple* distribution, $P_R = 0.5$, $P_Q = 1$, and
310 $P_N = 0.5$ for all edges. The *KPD* distribution is inspired by the fielded exchange setting from which
311 we draw our real underlying compatibility graphs. According to UNOS, about 34% of all edges are
312 rejected by a donor or recipient pre-match [18]; we draw P_R uniformly from $U(0.25, 0.43)$ for each
313 edge. Edges ending in highly-sensitized patients (who are often less healthy and more likely to be
314 incompatible) are considered high-risk; for these edges we draw P_Q from $U(0.2, 0.5)$ and P_N from
315 $U(0.0, 0.2)$. For other edges we draw P_Q from $U(0.9, 1.0)$ and P_N from $U(0.8, 0.9)$.

316 4.2 Single-Stage Edge Selection Experiments

317 In this section we compare against the baseline of a max-weight matching *without* edge queries
318 (using policy $M^{\text{MAX}}(\cdot)$). Many fielded kidney exchanges use a variant of this matching policy, so by
319 comparing against this baseline we are illustrating the impact of edge queries on the state-of-the-art
320 matching policies used in many real exchanges. Let V_X be the objective⁸ of Problem 1 achieved by
321 method X , we calculate Δ^{MAX} (the relative difference from baseline) as $\Delta^{\text{MAX}} \equiv (V_X - V^S(\mathbf{0})) / V^S(\mathbf{0})$.
322 A value of $\Delta^{\text{MAX}} = 0$ means that method X did not improve over the baseline, a value of $\Delta^{\text{MAX}} = 1$
323 means that X achieved an objective 100% greater than the baseline, and so on. Furthermore a value
324 of $\Delta^{\text{MAX}} > 0$ means that method X *increases* the objective by querying edges, while $\Delta^{\text{MAX}} < 0$ means
325 that method X *decreases* the objective by querying edges.

326 **Result: Greedy is essentially Optimal with small random graphs.** First we investigate the *dif-*
327 *ficulty* of edge selection. Using random graphs, we compare Greedy to the *optimal* solution to
328 Problem 1, found by exhaustive search (OPT). We generate three sets of 100 random graphs with
329 $N = 50, 75$, and 100 vertices, and each with $p = 0.01$. For all graphs we run both OPT and Greedy
330 with edge budget 3; we calculate the *optimality gap* of Greedy as %OPT $\equiv 100 \times (V_{\text{OPT}} - V_{\text{Greedy}}) / V_{\text{OPT}}$,
331 where V_X denotes the objective achieved by method X . ($V_{\text{OPT}} > 0$ in all graphs used in these exper-
332 iments.) If %OPT = 0 then Greedy returns an optimal solution, and %OPT > 0 means that Greedy
333 is not optimal. Table 1 (left) shows the number of random graphs binned by %OPT, as well as the
334 maximum %OPT over all graphs. For each N , Greedy returns an optimal solution for at least 90 of
335 the 100 graphs; the *maximum* %OPT over all graphs is 2.8.

336 In other words, Greedy always returns an *optimal* or nearly-optimal set of edges to query for small
337 random graphs. This is somewhat unexpected, since the edge selection problem is both non-monotone
338 and non-submodular (see Section 2).

339 **Result: Greedy is essentially monotonic with UNOS graphs.** We test Greedy on real UNOS
340 graphs, using maximum budget 100. Figure 5a shows the median Δ^{MAX} over all UNOS graphs,
341 with shading between the 10th and 90th percentiles. Larger edge budgets almost never decrease
342 the objective achieved by Greedy, and Greedy *never* produces a worse outcome than the baseline.

⁸All objective values are estimated using up to 1000 sampled rejection scenarios (see Appendix B), as it is intractable to evaluate the exact objective of large edge sets.

[†]We use an approximation of Fail-Aware for the *KPD* dist.; *true* Fail-Aware should always have $\Delta^{\text{MAX}} > 0$.

343 Thus—in our setting—single-stage edge selection is effectively monotonic in our setting, and Greedy
344 is an effective method.

345 **Result: MCTS and Greedy are nearly equivalent with UNOS graphs.** We compare all methods on
346 UNOS graphs, using smaller, more-realistic edge budgets from 1 to 10. For MCTS we use a 1-hour
347 time limit per edge (Γ hours total). Figures 5b and 5d compare Δ^{MAX} for MCTS, Greedy, and random
348 edge selection, for the *Simple* and *KPD* edge distributions, respectively. We draw two conclusions
349 from these results: (1) MCTS and Greedy produce almost identical results, further suggesting that
350 Greedy is nearly optimal in our setting; (2) in our setting, edge selection is *effectively* monotonic, as
351 Δ^{MAX} almost never decreases. However Figure 5d gives an example of non-monotonicity for both
352 Greedy and Random: in some cases, querying edges can lead to a *worse* outcome than querying no
353 edges.

354 **Result: Both MCTS and Greedy outperform benchmarks from the literature.** We also compare
355 against two state-of-the-art approaches: the edge selection approach of [11] (IIAB), which uses a
356 *variable* edge budget that depends on the graph structure; and the failure-aware matching policy
357 of [15] (Fail-Aware⁹), which does not query edges. To our knowledge, IIAB is the only edge selection
358 method in the literature. We compare against the Fail-Aware method because it is a state-of-the-art
359 kidney exchange matching policy which aims to maximize the expected matching weight, under a
360 similar edge failure model to ours; we compare against this approach to further illustrate the utility of
361 querying edges.

362 Table 1 (right) shows a comparison of all edge-selection methods—each using the variable edge
363 budget of IIAB; the bottom row shows results for Fail-Aware. Both MCTS and Greedy achieve greater
364 Δ^{MAX} (in distribution) than both benchmark methods. This is expected in both cases: IIAB uses a
365 heuristic to select edges to query, which does not consider the final matching weight—the objective of
366 our edge selection problem; on the other hand, both MCTS and Greedy are designed to maximize this
367 objective. We do not expect Fail-Aware to out-perform any edge selection methods, since Fail-Aware
368 does not have access to information revealed after edge queries.

369 It is notable that Greedy performs better than MCTS (in distribution). This likely means that MCTS
370 is *under-trained*—that the time and memory limits used in our implementation are too restrictive;
371 alternatively, this indicates that Greedy is simply very effective in our setting.

372 4.3 Multi-Stage Edge Selection Experiments on UNOS Graphs

373 We run initial multi-stage edge selection experiments on all UNOS graphs with the *Simple* edge
374 distribution. For each graph we test our multi-stage variants of MCTS and Greedy, and compare
375 with a baseline of random edge selection; as before, MCTS uses a 1-hour training time per level. It
376 is substantially harder to evaluate the multi-stage objective, as each edge edge-selection method
377 changes depending on rejections observed in prior stages. Similarly, the MCTS search tree is orders
378 of magnitude larger in the multi-stage setting: each node in tree corresponds to both an edge set *and*
379 a rejection scenario (see Appendix E).

380 In these initial experiments we evaluate each method on 10 edge rejections *realizations* (only a
381 small subset). We estimate Δ^{MAX} for each method and each graph by averaging the final matching
382 weight over all realizations. Figure 5c shows the results of these experiments.

383 These initial multi-stage results are quite similar to our single-stage results. However it is notable
384 that the objective value in the multi-stage setting is somewhat higher than in the single-stage setting—
385 even using the simple method Greedy. Further, this suggests that more can be gained by developing
386 a more sophisticated multi-stage edge selection policy. We leave this for future work.

387 5 Conclusions and Future Research Directions

388 Many planned kidney exchange transplants *fail* for a variety of reasons; these failures greatly
389 reduce the number of transplants that an exchange can facilitate, and increase the waiting time for
390 many patients in need of a kidney. Avoiding transplant failures is a challenge, as exchanges are often
391 constrained by policy and law in how they match patients and donors. We consider a setting where
392 exchanges can *pre-screen* certain transplants, while still matching patients and donors using a fixed
393 policy. We formalize a multi-stage optimization problem based on realistic assumptions about how
394 transplants fail, and how exchanges match patients and donors; we emphasize that these important
395 assumptions are not included in prior work. While this problem is challenging in theory, we show

⁹For the *KPD* distribution we use an approximation of Fail-Aware, which assumes a uniform edge failure probability.

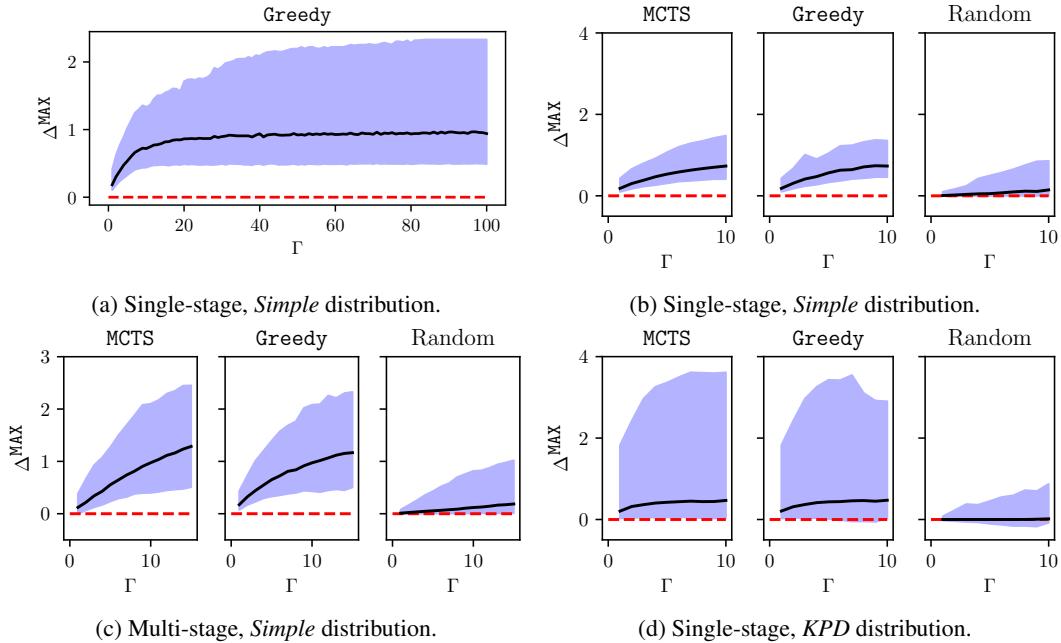


Figure 5: Results for UNOS graphs. Right: edge budget up to 10 for the *Simple* distribution (top) and the *KPD* distribution (bottom). Top-left: Greedy with edge budget up to 100, for the simple distribution. Bottom-left: multi-stage methods using the *Simple* distribution. In all plots, a solid line indicates median Δ^{MAX} over all UNOS graphs, and shading is between the 10th and 90th percentiles; a dotted line indicates the baseline.

396 that it is much easier in practice—with computational experiments using both synthetic data and real
 397 data from the United Network for Organ Sharing. In experiments, we find that pre-screening even a
 398 small number of potential transplants (around 10) significantly increases the overall quality of the
 399 final match—by more than 100% of the original match weight.

400 Our initial study of the pre-screening problem suggests several areas for future work. First we
 401 assume that the distribution of transplant failures is known, when in reality only rough approximations
 402 of these distributions are available. Second, we assume that exchange participants (donors, recipients,
 403 hospitals) are not strategic. In reality, strategic behavior plays a substantial role in real exchanges [2];
 404 we expect that participants might behave strategically when responding to pre-screening requests.
 405 Third, our model does not account for equitable treatment of different patients [21]. For example,
 406 it may be the case that pre-screening a transplant decreases the likelihood of the transplant being
 407 matched. That might disproportionately impact highly-sensitized patients, which are both sicker and
 408 more difficult to match than other patients.

409 **Broader Impact**

410 This work lives within the broader context of kidney exchange research. For clarity, we separate
411 our broader impacts into two sections: first we discuss the impact of kidney exchange in general; then
412 we discuss our work in particular, within the context of kidney exchange research and practice.

413 **Impacts of Kidney Exchange** Patients with end-stage renal disease have only two options: receive
414 a transplant, or undergo dialysis once every few days, for the rest of their lives. In many countries
415 (including the US), these patients register for a deceased donor waiting list—and it can be months
416 or years before they receive a transplant. Many of these patients have a friend or relative willing to
417 donate a kidney, however many patients are incompatible with their corresponding donor. Kidney
418 exchange allows patients to “swap” their incompatible donor, in order to find a *higher-quality* match,
419 *more quickly* than a waiting list. Transplants allow patients a higher quality of life, and cost far less,
420 than lifelong dialysis. About 10% of kidney transplants in the US are facilitated by an exchange.

421 Finding the “most efficient” matching of kidney donors to patients is a (computationally) hard
422 problem, which cannot be solved by hand in most cases. For this reason many fielded exchanges
423 use algorithms to quickly find an efficient matching of patients and donors. Many researchers study
424 kidney exchange from an algorithmic perspective, often with the goal of improving the number or
425 quality of transplants facilitated by exchanges. Indeed, this is the purpose of our paper.

426 **Impacts of Our Work** In this paper we investigate the impact of pre-screening certain potential
427 transplants (edge) in an exchange, prior to constructing the final patient-donor matching. To our
428 knowledge, some modern fielded exchanges pre-screen potential transplants in an ad-hoc manner;
429 meaning they do not consider the impacts of pre-screening on the final matching. We propose methods
430 to estimate the importance of pre-screening each edge, as measured by the change in the overall
431 number and quality of matched transplants.¹⁰ Importantly, our methods do not require a change in
432 matching policy; instead, they indicate to policymakers which potential transplants are important to
433 pre-screen, and which are not. The impacts of our contributions are summarized below:

434 **Some potential transplants cannot be matched**, because they cannot participate in a “legal” cyclical
435 or chain-like swap (according to the exchange matching policy). Accordingly, there is no “value”
436 gained by pre-screening these transplants; our methods will identify these potential transplants, and
437 will recommend that they not be pre-screened. Pre-screening requires doctors to spend valuable time
438 reviewing potential donors; removing these unmatchable transplants from pre-screening will allow
439 doctors to focus only on transplants that are relevant to the current exchange pool.

440 **Some transplants are more important to pre-screen than others**, and our methods help identify
441 which are most important for the final matching. We estimate the value pre-screening of each
442 transplant by *simulating* the exchange matching policy in the case that the pre-screened edge is
443 pre-accepted, and in the case that it is pre-refused.

444 **To estimate the value of pre-screening each transplant, we need to know (a) the likelihood**
445 **that each transplant is pre-accepted and pre-refused, and (b) the likelihood that each planned**
446 **transplant fails for any reason, after being matched.** These likelihoods are used as input to our
447 methods, and they can influence the estimated value of pre-screening different transplants. Importantly,
448 it may not be desirable to calculate these likelihoods for each potential transplant (e.g., using data from
449 the past). For example if a patient is especially sick, we may estimate that any potential transplant
450 involving this patient is very likely to fail prior to transplantation (e.g., because the patient is to ill to
451 undergo an operation). In this case, our methods may estimate that all potential transplants involving
452 this patient have very low “value”, and therefore recommend that these transplants should not be
453 pre-screened. One way to avoid this issue is to use the same likelihood estimates for all transplants.

454 **To estimate the impact of our methods (and how they depend on the assumed likelihoods, see**
455 **above), we recommend using extensive modeling of different pre-screening scenarios before**
456 **deploying our methods in a fielded exchange.** This is important for several reasons: first, exchange
457 programs cannot always *require* that doctors pre-screen potential transplants prior to matching. Since
458 we cannot be sure which transplants will be pre-screened and which will not, simulations should
459 be run to evaluate each possible scenario. Second, theoretical analysis shows that pre-screening
460 transplants can—in the worst case—negatively impact the final outcome. While this worst-case

¹⁰Quality and quantity of transplants is measured by transplant weight, a numerical representation of transplant quality (e.g., see UNOS/OPTN Policy 13 regarding KPD prioritization points https://optn.transplant.hrsa.gov/media/1200/optn_policies.pdf).

461 outcome is possible, our computational experiments show that it is very unlikely; this can be addressed
462 further with more experiments tailored to a particular exchange program.

463 References

- 464 [1] D. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets:
465 Enabling nationwide kidney exchanges. In *Proceedings of the ACM Conference on Electronic*
466 *Commerce (EC)*, pages 295–304, 2007.
- 467 [2] N. Agarwal, I. Ashlagi, E. Azevedo, C. R. Featherstone, and Ö. Karaduman. Market failure in
468 kidney exchange. *American Economic Review*, 109(11):4026–70, 2019.
- 469 [3] R. Anderson, I. Ashlagi, D. Gamarnik, and A. E. Roth. Finding long chains in kidney exchange
470 using the traveling salesman problem. *Proceedings of the National Academy of Sciences*, 112
471 (3):663–668, 2015.
- 472 [4] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem.
473 *Machine learning*, 47(2-3):235–256, 2002.
- 474 [5] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the cure for
475 your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762,
476 2012.
- 477 [6] H. Bidkhorji, J. P. Dickerson, K. Ren, and D. C. McElfresh. Kidney exchange with inhomoge-
478 neous edge existence uncertainty. In *Proceedings of the Conference on Uncertainty in Artificial*
479 *Intelligence (UAI)*, forthcoming, 2020.
- 480 [7] P. Biró, D. F. Manlove, and R. Rizzi. Maximum weight cycle packing in directed graphs, with
481 application to kidney exchange programs. *Discrete Mathematics, Algorithms and Applications*,
482 1(04):499–517, 2009.
- 483 [8] P. Biró, B. Haase-Kromwijk, T. Andersson, E. I. Ásgeirsson, T. Baltsová, I. Boletis, C. Bolot-
484 inha, G. Bond, G. Böhmig, L. Burnapp, et al. Building kidney exchange programmes in
485 europe—an overview of exchange practice and activities. *Transplantation*, 103(7):1514, 2019.
- 486 [9] P. Biró, J. van de Klundert, D. Manlove, W. Pettersson, T. Andersson, L. Burnapp, P. Chromy,
487 P. Delgado, P. Dworzak, B. Haase, et al. Modelling and optimisation in european kidney
488 exchange programmes. *European Journal of Operational Research*, 2019.
- 489 [10] A. Blum, A. Gupta, A. D. Procaccia, and A. Sharma. Harnessing the power of two crossmatches.
490 In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 123–140, 2013.
- 491 [11] A. Blum, J. P. Dickerson, N. Haghtalab, A. D. Procaccia, T. Sandholm, and A. Sharma. Ignor-
492 ance is almost bliss: Near-optimal stochastic matching with few queries. *Operations Research*,
493 2020. Earlier version appeared in the ACM Conference on Economics and Computation (EC),
494 2015.
- 495 [12] B. Bouzy. Associating shallow and selective global tree search with Monte Carlo for 9×9 Go.
496 In *International Conference on Computers and Games*, pages 67–80. Springer, 2004.
- 497 [13] J. P. Dickerson. A unified approach to dynamic matching and barter exchange. Technical report,
498 Doctoral Dissertation, Carnegie Mellon University, Pittsburgh, PA, 2016.
- 499 [14] J. P. Dickerson, D. Manlove, B. Plaut, T. Sandholm, and J. Trimble. Position-indexed for-
500 mulations for kidney exchange. In *Proceedings of the ACM Conference on Economics and*
501 *Computation (EC)*, 2016.
- 502 [15] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Failure-aware kidney exchange. *Manage-*
503 *ment Science*, 65(4):1768–1791, 2019. Earlier version appeared in the ACM Conference on
504 Economics and Computation (EC), 2013.
- 505 [16] B. Kartal, E. Nunes, J. Godoy, and M. Gini. Monte Carlo tree search with branch and bound for
506 multi-robot task allocation. In *The IJCAI-16 workshop on autonomous mobile service robots*,
507 volume 33, 2016.

- 508 [17] L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In *European conference on*
509 *machine learning*, pages 282–293. Springer, 2006.
- 510 [18] R. Leishman. Challenges in match offer acceptance in the optn kidney paired donation pilot
511 program. In *INFORMS Annual Meeting*, 2019. Presentation in Session TB94 - Kidney
512 Allocation & Exchange.
- 513 [19] Z. Li, N. Gupta, S. Das, and J. P. Dickerson. Equilibrium behavior in competing dynamic
514 matching markets. In *Proceedings of the International Joint Conference on Artificial Intelligence*
515 *(IJCAI)*, pages 389–395, 2018.
- 516 [20] D. Manlove and G. O’Malley. Paired and altruistic kidney donation in the UK: Algorithms and
517 experimentation. *ACM Journal of Experimental Algorithmics*, 19(1), 2015.
- 518 [21] D. C. McElfresh and J. P. Dickerson. Balancing lexicographic fairness and a utilitarian objective
519 with application to kidney exchange. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- 520 [22] D. C. McElfresh, H. Bidkhori, and J. P. Dickerson. Scalable robust kidney exchange. In
521 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1077–1084,
522 2019.
- 523 [23] M. Molinaro and R. Ravi. Kidney exchanges and the query-commit problem. Manuscript, 2013.
- 524 [24] F. T. Rapaport. The case for a living emotionally related international kidney donor exchange
525 registry. *Transplantation Proceedings*, 18:5–9, 1986.
- 526 [25] M. Rees, J. Kopke, R. Pelletier, D. Segev, M. Rutter, A. Fabrega, J. Rogers, O. Pankewycz,
527 J. Hiller, A. Roth, T. Sandholm, U. Ünver, and R. Montgomery. A nonsimultaneous, extended,
528 altruistic-donor chain. *New England Journal of Medicine*, 360(11):1096–1101, 2009.
- 529 [26] A. Roth, T. Sönmez, and U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):
530 457–488, 2004.
- 531 [27] A. Roth, T. Sönmez, and U. Ünver. A kidney exchange clearinghouse in New England. *American*
532 *Economic Review*, 95(2):376–380, 2005.
- 533 [28] A. Roth, T. Sönmez, and U. Ünver. Pairwise kidney exchange. *Journal of Economic Theory*,
534 125(2):151–188, 2005.
- 535 [29] UNOS. United Network for Organ Sharing (UNOS). <http://www.unos.org/>.

536 **A Kidney Exchange and Edge Failures**

537 **Brief history.** Rapaport [24] proposed the initial idea for kidney exchange, while the first organized
 538 kidney exchange, the New England Paired Kidney Exchange (NEPKE), started in 2003–04 [26, 27,
 539 28]. NEPKE has since ceased to operate; at the point of cessation, its pool of patients and donors was
 540 merged into the United Network for Organ Sharing (UNOS) exchange in late 2010. That exchange
 541 now contains over 60% of transplant centers in the US, and performs matching runs via a purely
 542 algorithmic approach (as we discuss in Sections 1 and 2, and in much greater depth by UNOS [29],
 543 which is mandated to transparently and publicly reveal its matching process).

544 There are also two large private kidney exchanges in the US, the National Kidney Registry (NKR)
 545 and the Alliance for Paired Donation (APD). They typically only work with large transplant centers.
 546 NKR makes their matching decisions manually and greatly prefers matching incrementally through
 547 chains. APD makes their decisions through a combination of algorithmic and manual decision
 548 making. There are also several smaller private kidney exchanges in the US. They typically only
 549 involve one or a couple of transplant centers. These include an exchange at Johns Hopkins University,
 550 a single-center exchange at the Methodist Specialty and Transplant Hospital in San Antonio, and a
 551 single-center exchange at Barnes-Jewish Hospital affiliated with the Washington University in St.
 552 Louis. Largely, these exchanges also make their matching decisions via a combined algorithmic and
 553 manual process. These exchanges compete in a variety of ways (e.g., by allowing patient-donor pairs
 554 to register in multiple exchange programs); this competition can lead to loss in efficiency [2] as well
 555 as sub-optimal changes to individual exchanges’ matching policies [19].

556 There are now established kidney exchanges in the UK [20], Italy, Germany, Netherlands, Canada,
 557 England, Portugal, Israel, and many other countries. European countries are also explicitly exploring
 558 connecting their individual exchanges together in various ways [8].

559 **Edge failures.** The dilemma of edge failures is illustrated in the example exchange graph shown in
 560 Figure 6. This exchange consists of a 3-chain (dashed edges) and two 2-cycles (solid edges). Suppose
 561 the decision-maker queries edge e_A : if e_A is accepted, then the chain from the NDD (n) through pairs
 562 (d_1, p_1) , (d_2, p_2) , and (d_3, p_3) , i.e., the dashed edges, can be included in the matching. However
 563 if e_A is queried and rejected, then the NDD cannot initiate the chain, and only the cycles may be
 564 matched. In our model, if e_A is not queried then it may still be matched.




Figure 6: Sample exchange graph with a 3-chain (dashed edges) and two 2-cycles (solid edges). The NDD is denoted by n , and each patient (and associated donor) is denoted by p_i (d_i). If edge e_1 is not queried, or queried and *accepted*, then the chain may be included in the final matching. However if edge e_A is queried and *rejected*, then only the 2-cycles may be included in the final matching.

565 **B Estimating The Objective of Problem 1**

566 The objective of the single-stage edge selection problem requires evaluating all rejection scenarios
 567 $\mathbf{r} \sim \mathbb{P}_R(\mathbf{q})$, and the support of this distribution grows exponentially in the number of edges $|\mathbf{q}|$. In
 568 computational experiments, to estimate the objective of Problem 1, we sample up to 1000 scenarios
 569 from $\mathbb{P}_R(\mathbf{q})$. More explicitly: we *exactly* evaluate the objective of edge sets with fewer than 10 edges;
 570 for larger edge sets, we sample the objective using 1000 draws from $\mathbb{P}_R(\mathbf{q})$.

571 Using bootstrapping experiments we demonstrate that our sampling approach is sufficient to
 572 accurately estimate the true objective, even for large edge sets. For 152 UNOS graphs, we computed
 573 edge sets by running Greedy with edge budgets ranging from 1 to 100. For each edge set, we then
 574 sample a subset of $N \in \{10, 30, 50, 100, 1000\}$ rejection scenarios, with replacement, from the set
 575 of all sampled edge outcomes. For each edge set and choice of N we repeat 200 times and calculate
 576 the sample mean for each replication. We then compute the standard deviations of these bootstrap

Edge budgets	$N = 10$	$N = 30$	$N = 50$	$N = 100$	$N = 1000$
1-10	0.10	0.06	0.04	0.03	0.01
11-20	0.12	0.07	0.05	0.04	0.01
21-30	0.13	0.08	0.06	0.04	0.01
31-40	0.14	0.08	0.06	0.04	0.01
41-50	0.14	0.08	0.06	0.04	0.01
51-60	0.15	0.08	0.07	0.05	0.01
61-70	0.15	0.09	0.07	0.05	0.02
71-80	0.16	0.09	0.07	0.05	0.02
81-90	0.17	0.10	0.08	0.05	0.02
91-100	0.18	0.10	0.08	0.06	0.02

Table 2: Median normalized standard deviation of the bootstrap mean, over 200 bootstrap samples for each sample size N , binned by edge budget.

Table 3: Single-stage results on random graphs with the *Simple* edge distribution, using the variable IIAB edge budget (top rows), and the failure-aware method (bottom row). Columns P_X indicates the X^{th} percentile of Δ^{MAX} over all 30 random graphs, for graphs with $N = 50, 75,$ and 100 vertices.

Method	$N = 50$			$N = 75$			$N = 100$		
	P_{10}	P_{50}	P_{90}	P_{10}	P_{50}	P_{90}	P_{10}	P_{50}	P_{90}
MCTS	0.22	0.30	0.38	0.11	0.33	0.46	0.23	0.33	0.38
Greedy	0.21	0.30	0.38	0.12	0.32	0.48	0.27	0.39	0.43
Random	0.12	0.19	0.23	0.10	0.19	0.28	0.12	0.19	0.23
IIAB	0.07	0.24	0.34	0.11	0.22	0.41	0.07	0.24	0.34
Fail-Aware	0.00	0.02	0.10	0.00	0.06	0.18	0.00	0.02	0.10

577 sample means to estimate the variance due to sampling. For each N , we calculate the mean sample
578 standard deviation, normalized by the sample mean. Table 2 shows the median normalized standard
579 deviation for all experiments under each N , with edge budgets aggregated into 10 bins. We find that
580 with $N = 1000$ samples, the standard deviation was on average only about 2% of the overall mean
581 value, even for large edge budgets.

582 C Additional Computational Results

583 First we show results for both single-stage and multi-stage edge selection on random graphs (see
584 § 4 for a description of these graphs). For $N = 50, 75,$ and 100 , we generate 30 random graphs with
585 N vertices and $p = 0.01$. For each graph we run single-stage experiments with $\Gamma = 1, \dots, 10$ and
586 multi-stage experiments with $\Gamma = 1, \dots, 15$. Unlike experiments on UNOS graphs we use a time
587 limit of 20 minutes per edge; all other parameters are the same. Figure 7a and 7b show single-stage
588 and multi-stage results for all random graphs, respectively. Table 3 shows comparisons to IIAB and
589 Fail-Aware for random graphs with $N = 50, 75,$ and 100 .

590 As with UNOS graphs, results for MCTS and Greedy are quite similar, and both methods achieve
591 larger Δ^{MAX} than Random, IIAB, and Fail-Aware. We make two observations: (1) Greedy appears to
592 achieve larger Δ^{MAX} than MCTS in the single-stage setting, likely because of insufficient training time
593 for MCTS; (2) in the multi-stage setting, MCTS performs *at least* as well as Greedy, and often better.
594 Observation (2) is consistent with our experiments on UNOS graphs, and is somewhat surprising given
595 that MCTS used less training time in these experiments. This suggests that MCTS may substantially
596 improve over Greedy in the multi-stage setting; we leave further investigation to future work.

597 D Proofs for Section 2

598 In the proofs of Proposition 2.1 and Proposition 2.2 we consider a setting where all edges' pre-
599 match rejections and post-match failures are i.i.d., where $P_R = 0.5$ is the pre-match rejection
600 probability, $P_Q = 1.0$ is the post-match success probability if the edge is queried-and-accepted,




Figure 7: Results for 30 random graphs with edge probability $p = 0.01$ and $N = 50$ vertices (top row), $N = 75$ (middle row), and $N = 100$ (bottom row). All experiments use the *Simple* edge distribution. In all plots, a solid line indicates median Δ^{MAX} over all 30 random graphs, and shading is between the 10th and 90th percentiles; a dotted line indicates the baseline.




Figure 8: Exchange graph for Propositions 2.1 and 2.2. All edges have weight 1 except for edge (E, B) , which has weight 1.5.

601 and $P_N = 0.5$ is the success probability if e is not queried. That is, queried edges have rejection
 602 probability 0.5, accepted edges have zero failure probability, and non-queried edges have failure
 603 probability 0.5.

604 **D.1 Proof of Proposition 2.1**

605 (Proof by counterexample.) We provide an example where querying a single edge results in a *lower*
 606 objective value in Problem 1 (i.e., final expected matching weight) than querying no edges—when
 607 using the max-weight matching policy $M^{\text{MAX}}(\cdot)$.

608 Consider the exchange graph in Figure 8; edge (E, B) has weight 1.5, while all other edges have
 609 weight 1. First we consider the objective due to querying no edges, $V^S(\mathbf{0})$. In this case, no edges
 610 can be rejected pre-match, the max-weight matching includes cycle (C, D, F) (expected weight
 611 $3 \times (1/2)^3 = 3/8$) and cycle (A, B) (expected weight $2 \times (1/2)^2 = 1/2$), with total expected
 612 matching weight $7/8$. That is, $V^S(\mathbf{0}) = 7/8$.

613 Next consider the objective due to querying only edge $e_3 = (C, D)$, and let q' denote edge set
 614 $\{e_3\}$. With probability $1/2$, e_3 is rejected and cycle (B, C, E) is the max-weight matching – with

615 expected weight $3.5/8$. With probability $1/2$, e_3 is accepted and the max-weight matching includes
616 cycles (A, B) (with expected weight $1/2$) and (C, D, F) (with expected weight $3/4$); this matching
617 has total expected weight $5/4$. Thus, $V^S(\mathbf{q}) = 27/32 < 7/8 = V^S(\mathbf{0})$, which concludes the proof.

618 **D.2 Proof of Proposition 2.2**

619 (Proof by counterexample.) We provide an example where the objective value in Problem 1 (i.e.,
620 final expected matching weight) is non-submodular—when using the max-weight matching policy
621 $M^{\text{MAX}}(\cdot)$. We use the same rejection and failure distribution as in the proof of Proposition 2.1.

Consider the exchange graph in Figure 8; edge (E, B) has weight 1.5, while all other edges
have weight 1. With some abuse of notation, we will denote by $V^S(\{e_a, \dots, e_N\})$ the objective
of Problem 1 due to edge set $\{e_a, \dots, e_N\}$. Our counterexample for submodularity is that, for this
graph,

$$V^S(X \cup \{e_1, e_2\}) + V^S(X) > V^S(X \cup \{e_1\}) + V^S(X \cup \{e_2\}),$$

622 with set $X \equiv \{e_3\}$. That is, the objective increase due to querying *both* edges e_1 and e_3 is greater
623 than the combined increase due to querying both edges separately. Next we explicitly calculate each
624 of the above terms.

625 $V^S(X) = V^S(\{e_3\})$. There are two cases to consider:

- 626 • e_3 is accepted, with probability $1/2$. The max-weight matching is cycles (A, B) and
627 (C, D, F) , with expected weight $(1/2 + 3/4)$,
- 628 • e_3 is rejected, with probability $1/2$. The max-weight matching is cycle (B, C, E) , with
629 expected weight $3.5/8$.

630 Thus, $V^S(X) = (1/2)(1/2 + 3/4) + (1/2)(3.5/8) = 27/32$.

631 $V^S(X \cup \{e_1\}) = V^S(\{e_1, e_3\})$. There are four cases to consider:

- 632 • e_1 and e_3 are accepted, with probability $1/4$. The max-weight matching is cycles (A, B)
633 and (C, D, F) , with expected weight $(1 + 3/8)$,
- 634 • e_1 is rejected and e_3 is accepted, with probability $1/4$. The max-weight matching is cycle
635 (B, C, E) , with expected weight $3.5/8$.
- 636 • e_1 is accepted and e_3 is rejected, with probability $1/4$. The max-weight matching is cycle
637 (B, C, E) , with expected weight $3.5/8$.
- 638 • e_1 and e_3 are rejected, with probability $1/4$. The max-weight matching is cycle (B, C, E) ,
639 with expected weight $3.5/8$.

640 Thus the objective is $V^S(X \cup \{e_1\}) = (1/4)(1 + 3/8) + (3/4)(3.5/8) = 43/64$.

641 $V^S(X \cup \{e_2\}) = V^S(\{e_2, e_3\})$. There are three cases to consider

- 642 • e_3 is accepted, with probability $1/2$. The max-weight matching is cycles (A, B) and
643 (C, D, F) , with expected weight $(1/2 + 3/4)$,
- 644 • e_3 is rejected and e_3 is accepted, with probability $1/4$. The max-weight matching is cycle
645 (B, C, E) , with expected weight $3.5/4$,
- 646 • e_3 and e_2 are rejected, with probability $1/4$. The max-weight matching is cycle (A, B) ,
647 with expected weight $1/2$.

648 Thus the objective is $V^S(X \cup \{e_2\}) = (1/2)(1/2 + 3/4) + (1/4)(3.5/4) + (1/4)(1/2) = 31/32$.

649 $V^S(X \cup \{e_1, e_2\}) = V^S(\{e_1, e_2, e_3\})$. There are four cases to consider:

- 650 • e_1 and e_3 are accepted, with probability $1/4$. The max-weight matching is cycles (A, B)
651 and (C, D, F) , with expected weight $(1 + 3/4)$,
- 652 • e_1 is accepted and e_2 is rejected, with probability $1/4$ (the response from e_3 is irrelevant).
653 The max-weight matching is (A, B) and (C, D, F) , with expected weight $1 + 3/8$.
- 654 • e_1 is rejected and e_2 is accepted (the response from e_3 is irrelevant), with probability $1/4$.
655 The max-weight matching is cycle (B, C, E) , with expected weight $3.5/4$.
- 656 • e_1 and e_2 are rejected (the response from e_3 is irrelevant), with probability $1/4$. The
657 max-weight matching is cycle (C, D, F) , with expected weight $3/8$.

658 Thus the objective is $V^S(X \cup \{e_1, e_2\}) = (1/4)(1 + 3/4) + (1/4)(1 + 3/8) + (1/4)(3.5/4) +$
659 $(1/4)(3/8) = 35/32.$

660 Finally, we have:

$$\begin{aligned} V^S(X \cup \{e_1, e_2\}) + V^S(X) &= 35/32 + 27/32 \\ &= 1.9375 \end{aligned}$$

661 and

$$\begin{aligned} V^S(X \cup \{e_1\}) + V^S(X \cup \{e_2\}) &= 43/64 + 31/32 \\ &= 1.640625 \end{aligned}$$

662 Therefore, $V^S(X \cup \{e_1, e_2\}) + V^S(X) > V^S(X \cup \{e_1\}) + V^S(X \cup \{e_2\})$, which concludes the
663 proof.

664 D.3 Proof of Proposition 2.5

665 For the proof of Proposition 2.5 we make one assumption about the distribution of edge rejections
666 and failures: querying *additional* edges cannot increase the overall probability of rejection or failure
667 for any edge.

668 First we prove a handful of useful results.

669 **Lemma D.1.** *If all edges are independent and Assumption 2.3 holds, then additional edge queries*
670 *cannot decrease expected post-match cycle and chain weights. Formally,*

$$\mathbb{E}[F(c, \mathbf{r} + \mathbf{f}) \mid \mathbf{q}, \mathbf{r}] \leq \mathbb{E}[F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \mid \mathbf{q} + \mathbf{q}', \mathbf{r}]$$

671 for any $\mathbf{q}, \mathbf{q}' \in \{0, 1\}^{|E|}$ such that $\mathbf{q} + \mathbf{q}' \in \{0, 1\}^{|E|}$, for any $\mathbf{r} \in \{0, 1\}^{|E|}$, and for all $c \in \mathcal{C}$.

672 *Proof.* We address cycles and chains separately.

673 **Cycles.** Conditional on fixed \mathbf{q} and \mathbf{r} , the expected weight of cycle $c = (e_1, \dots, e_L)$ is expressed
674 as

$$\begin{aligned} \mathbb{E}[F(c, \mathbf{r} + \mathbf{f}) \mid \mathbf{q}, \mathbf{r}] &= \left(\sum_{e \in c} w_e \right) \mathbb{E} \left[\prod_{e \in c} (1 - \mathbf{r}_e - \mathbf{f}_e) \mid \mathbf{q}, \mathbf{r} \right] \\ &= \left(\sum_{e \in c} w_e \right) \prod_{e \in c} (1 - \mathbb{E}[\mathbf{r}_e + \mathbf{f}_e \mid \mathbf{q}, \mathbf{r}]) \end{aligned}$$

675 where the second step is due to the fact that all \mathbf{f}_e are independent. Similarly, for fixed \mathbf{q}' ,

$$\mathbb{E}[F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \mid \mathbf{q} + \mathbf{q}', \mathbf{r}] = \left(\sum_{e \in c} w_e \right) \prod_{e \in c} (1 - \mathbb{E}[\mathbf{r}_e + \mathbf{r}'_e + \mathbf{f}_e \mid \mathbf{q} + \mathbf{q}', \mathbf{r}]) .$$

Due to Assumption 2.3, the following inequality holds for all edges $e \in E$

$$\mathbb{E}[\mathbf{r}_e + \mathbf{f}_e \mid \mathbf{q}, \mathbf{r}] \geq \mathbb{E}[\mathbf{r}_e + \mathbf{r}'_e + \mathbf{f}_e \mid \mathbf{q} + \mathbf{q}', \mathbf{r}] ,$$

and it follows that

$$\mathbb{E}[F(c, \mathbf{r} + \mathbf{f}) \mid \mathbf{q}, \mathbf{r}] \leq \mathbb{E}[F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \mid \mathbf{q} + \mathbf{q}', \mathbf{r}] .$$

676 **Chains.** Similarly, the expected weight of chain $c = (e_1, \dots, e_L)$ is expressed as

$$\begin{aligned} \mathbb{E}[F(c, \mathbf{r} + \mathbf{f}) \mid \mathbf{q}, \mathbf{r}] &= \sum_{k=1}^L \left(\sum_{j=1}^k w_j \right) \mathbb{E} \left[\prod_{j=1}^k (1 - \mathbf{r}_{e_j} - \mathbf{f}_{e_j}) \mid \mathbf{q}, \mathbf{r} \right] \\ &= \sum_{k=1}^L \left(\sum_{j=1}^k w_j \right) \prod_{j=1}^k (1 - \mathbb{E}[\mathbf{r}_{e_j} + \mathbf{f}_{e_j} \mid \mathbf{q}, \mathbf{r}]) , \end{aligned}$$

677 where the second step is due to the fact that \mathbf{f}_e are independent. Similarly,

$$\mathbb{E}[F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \mid \mathbf{q} + \mathbf{q}', \mathbf{r}] = \sum_{k=1}^L \left(\sum_{j=1}^k w_j \right) \prod_{j=1}^k (1 - \mathbb{E}[\mathbf{r}_{e_j} + \mathbf{r}'_{e_j} + \mathbf{f}_{e_j} \mid \mathbf{q} + \mathbf{q}', \mathbf{r}]) .$$

as before, due to Assumption 2.3 it follows that

$$\mathbb{E}[F(c, \mathbf{r} + \mathbf{f}) \mid \mathbf{q}, \mathbf{r}] \leq \mathbb{E}[F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \mid \mathbf{q} + \mathbf{q}', \mathbf{r}].$$

678

□

Lemma D.2. *With a failure-aware matching policy, and if all edges are independent, adding a single edge to any edge query set weakly improves the objective of Problem 1. Formally, for any $\mathbf{q}, \mathbf{q}' \in \{0, 1\}^{|E|}$ with $\mathbf{q} + \mathbf{q}' \in \{0, 1\}^{|E|}$ and $|\mathbf{q}'| = 1$, and $M(\mathbf{r}) \equiv M^{\text{FA}}(\mathbf{r})$,*

$$V^S(\mathbf{q}) \leq V^S(\mathbf{q} + \mathbf{q}')$$

679 *Proof.* The objective of Problem 1 for edge set \mathbf{q} is expressed as

$$\begin{aligned} V^S(\mathbf{q}) &= \mathbb{E}_{\mathbf{r}|\mathbf{q}} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q}, \mathbf{r}} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r}) F(c, \mathbf{r} + \mathbf{f}) \right] \right] \\ &= \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q}}(\mathbf{r}) \mathbb{E}_{\mathbf{f}|\mathbf{q}, \mathbf{r}} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r}) F(c, \mathbf{r} + \mathbf{f}) \right] \\ &= \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q}}(\mathbf{r}) \sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r}) \mathbb{E}_{\mathbf{f}|\mathbf{q}, \mathbf{r}} [F(c, \mathbf{r} + \mathbf{f})] \end{aligned}$$

680 For edge set $\mathbf{q} + \mathbf{q}'$ we partition response variables into $\mathbf{r}, \mathbf{r}' \in \{0, 1\}^{|E|}$, where \mathbf{r}_e is the response
681 variable for all edges $e \in \mathbf{q}$, and $\mathbf{r}_e = 0$ for all other edges (including the edge in \mathbf{q}'). Similarly, \mathbf{r}'_e is
682 the response variable for edge \mathbf{q}' , and $\mathbf{r}'_e = 0$ for all other edges. The objective of $\mathbf{q} + \mathbf{q}'$ is expressed
683 as

$$\begin{aligned} V^S(\mathbf{q} + \mathbf{q}') &= \mathbb{E}_{\mathbf{r}, \mathbf{r}'|\mathbf{q} + \mathbf{q}'} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r} + \mathbf{r}') F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \right] \\ &= \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q} + \mathbf{q}'}(\mathbf{r}) \mathbb{E}_{\mathbf{r}'|\mathbf{q} + \mathbf{q}'} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r} + \mathbf{r}')^\top F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \right] \\ &= \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q}}(\mathbf{r}) \mathbb{E}_{\mathbf{r}'|\mathbf{q} + \mathbf{q}'} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r} + \mathbf{r}') F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \right], \end{aligned}$$

684 where in the final line we replace $P_{\mathbf{q} + \mathbf{q}'}(\mathbf{r})$ with $P_{\mathbf{q}}(\mathbf{r})$, because each \mathbf{r}_e is conditionally independent,
685 given \mathbf{q}_e .

Next, by definition

$$\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r} + \mathbf{r}') F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \geq \mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} \mathbf{x}_c F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \quad \forall \mathbf{x} \in \mathcal{M}.$$

686 That is, M^{FA} is guaranteed to maximize this expectation, and thus

$$V^S(\mathbf{q} + \mathbf{q}') \geq \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q}}(\mathbf{r}) \mathbb{E}_{\mathbf{r}'|\mathbf{q} + \mathbf{q}'} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} \left[\sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r}) F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f}) \right] \right] \quad (\text{B})$$

$$= \sum_{\mathbf{r} \in \{0, 1\}^{|\mathbf{q}|}} P_{\mathbf{q}}(\mathbf{r}) \sum_{c \in \mathcal{C}} M_c^{\text{FA}}(\mathbf{r}) \mathbb{E}_{\mathbf{r}'|\mathbf{q} + \mathbf{q}'} \left[\mathbb{E}_{\mathbf{f}|\mathbf{q} + \mathbf{q}', \mathbf{r}, \mathbf{r}'} [F(c, \mathbf{r} + \mathbf{r}' + \mathbf{f})] \right] \quad (\text{C})$$

Finally, combining (B) and (C) with Lemma D.1, the following inequality holds

$$V^S(\mathbf{q}) \leq V^S(\mathbf{q} + \mathbf{q}').$$

687

□

688 Using the above lemmas, the proof of Proposition 2.5 is straightforward:

689 **Proposition 2.5** *If edges are independent, and Assumption 2.3 holds, then with a failure-aware*
 690 *matching policy the objective of Problem 1 is monotonic in the set of queried edges.*

Proof. Let $\mathbf{q}', \mathbf{q}'' \in \mathcal{E}$ be two edge sets such that $\mathbf{q}' \subseteq \mathbf{q}''$. It remains to show that, with matching policy $M(\mathbf{r}) \equiv M^{\text{FA}}(\mathbf{r})$,

$$V^S(\mathbf{q}'') \leq V^S(\mathbf{q}').$$

691 First note that because \mathcal{E} is a matroid, there is a sequence of edges $(\mathbf{q}^{e_1}, \dots, \mathbf{q}^{e_L})$ (with each
 692 $|\mathbf{q}^{e_i}| = 1$) such that $\mathbf{q}'' + \mathbf{q}^{e_1} + \dots + \mathbf{q}^{e_L} = \mathbf{q}'$. Due to Lemma D.2, the following sequence of
 693 inequalities hold:

$$\begin{aligned} V(\mathbf{q}'') &\leq V(\mathbf{q}'' + \mathbf{q}^{e_1}) \\ &\leq V(\mathbf{q}'' + \mathbf{q}^{e_1} + \mathbf{q}^{e_2}) \\ &\dots \\ &\leq V(\mathbf{q}'' + \mathbf{q}^{e_1} + \dots + \mathbf{q}^{e_L}) \\ &= V(\mathbf{q}') \end{aligned}$$

694 which concludes the proof. □

695 E Algorithm Descriptions

696 Here we describe more explicitly the algorithms for Greedy and MCTS, for both the single-stage
 697 and multi-stage settings.

698 E.1 UCB Value Estimates for MCTS

Both the single- and multi-stage versions of MCTS use the method of [17] to select the next child node to explore. The formula used to estimate a node's UCB value is

$$\frac{U}{\frac{N}{N^P} - V^{\min}} + \sqrt{N^P/N}$$

699 where U is the ‘‘UCB value estimate’’ calculated by MCTS, N is the number of visits to the node,
 700 N^P is the number of visits to the node's parent, and V^{\max} and V^{\min} are the largest and smallest
 701 *node values* encountered during search. In single-stage MCTS, all nodes have both a *node value* (the
 702 objective value of Problem 1) and a UCB value estimate; as described below, in multi-stage MCTS only
 703 query nodes have a UCB value estimate, and only leaf nodes have a *node value* (expected matched
 704 weight, after observing responses from all queried edges).

705 E.2 Greedy Single-Stage Edge Selection

706 Algorithm 3 gives a pseudocode description of Greedy for the single-stage setting.

ALGORITHM 3: Greedy: Greedy Search Heuristic for Single-Stage Edge Selection

(input) \mathcal{E} : legal edge sets

```

 $\mathbf{q}^R \leftarrow \mathbf{0}$   the root node (no edges)
 $V^* \leftarrow$  objective value of  $\mathbf{q}^R$  Problem 1
while  $\mathbf{q}^R$  has children do
     $\mathbf{q}' \leftarrow$  child node of  $\mathbf{q}^R$  with maximal objective value in Problem 1
     $\mathbf{q}^R \leftarrow \mathbf{q}'$ 
return  $\mathbf{q}^R$ 
  
```

707 E.3 Multi-Stage Edge Selection

708 In the following sections we describe multi-stage versions of MCTS and Greedy. Unlike in the
 709 single-stage setting, these algorithms take as input a set of previously-queried edges $\mathbf{q} \in \{0, 1\}^{|E|}$
 710 and a corresponding set of observed rejections $\mathbf{r} \in \{0, 1\}^{|E|}$; they output the *next* edge to query.

Multi-Stage MCTS. The multi-stage search tree is somewhat more complicated than in the single-stage setting, as each node in the search tree corresponds to both a set of queried edges and a set of observed rejections. For this purpose we use two types of nodes: *outcome* nodes, and *query* nodes. Outcome nodes consist of previously-queried edges \mathbf{q} and previously-observed rejections \mathbf{r} , and are represented by tuple (\mathbf{q}, \mathbf{r}) . (The root of the search tree corresponds to *no* queries or observed rejections, $(\mathbf{0}, \mathbf{0})$.) The children of an outcome node are *query* nodes, represented by the next edge to query from the parent (outcome), represented by tuple $(\mathbf{q}, \mathbf{r}, e)$. Each outcome node has one child for every edge that has not yet been queried:

$$C^O(\mathbf{q}, \mathbf{r}) \equiv \{(\mathbf{q}, \mathbf{r}, e) \mid \forall e \in E : \mathbf{q} + \mathbf{u}^e \in \mathcal{E}\}$$

where \mathbf{u}^e is the unit vector for element e ($u_i^e = 0$ for all $i \neq e$, and $u_e^e = 1$). Each query node has exactly two children: one where the queried edge is accepted, and one where the queried edge is rejected,

$$C^Q(\mathbf{q}, \mathbf{r}, e) \equiv \{(\mathbf{q} + \mathbf{u}^e, \mathbf{r}), (\mathbf{q} + \mathbf{u}^e, \mathbf{r} + \mathbf{u}^e)\}$$

711 As before, the *level* of a node refers to the number of queried edges: $|\mathbf{q}|$ for outcome nodes, and
 712 $|\mathbf{q}| + 1$ for query nodes.

As before we refer to nodes with no children as leaf nodes; note that only outcome nodes are leaf nodes. Unlike the single-stage version of MCTS, in the multi-stage setting we only consider the value of leaf nodes¹¹. The value of a leaf (outcome) is

$$V^O(\mathbf{q}, \mathbf{r}) \equiv W(M(\mathbf{r}); \mathbf{q}, \mathbf{r}),$$

713 where as before $M(\mathbf{r})$ denotes the matching policy, and $W(\mathbf{x}; \mathbf{q}, \mathbf{r})$ denotes the expected matching
 714 weight of \mathbf{x} , subject to \mathbf{q} and \mathbf{r} . The value of leaf outcome nodes is used to by `QSample` and `OSample`
 715 to guide multi-stage MCTS.

¹¹This decision was made in part because initial results indicate that edge selection is essentially monotonic.

716 Algorithm 4 describes the multi-stage version of MCTS, taking previously-queried edges and
 717 observed responses as input. This algorithm initializes the value estimate $U[\cdot]$ and number of visits
 718 $N[\cdot]$ for query nodes in the next L levels—these quantities are used in the UCB calculation.

ALGORITHM 4: Multi-Stage MCTS

(input) \mathcal{E} : legal edge sets
 (input) K : maximum size of any legal edge set
 (input) T : time limit
 (input) L : number of look-ahead levels
 (input) \mathbf{q}^R : previously-queried edges
 (input) \mathbf{r}^R : previously-observed rejections

719

$M \leftarrow \min\{N + L, K\}$
 $Q \leftarrow$ all query nodes which are descendants of $(\mathbf{q}^R, \mathbf{r}^R)$, up to level M
 $U[(\mathbf{q}, \mathbf{r}, e)] \leftarrow 0 \forall (\mathbf{q}, \mathbf{r}, e) \in Q$ UCB value estimate
 $N[(\mathbf{q}, \mathbf{r}, e)] \leftarrow 0 \forall \mathbf{q} \in Q$ number of visits
while less than time T has passed **do**
 QSample($\mathbf{q}^R, \mathbf{r}^R, M$)
 $(\mathbf{q}^R, \mathbf{r}^R, e^*) \leftarrow$ child node of $(\mathbf{q}^R, \mathbf{r}^R)$ with the greatest UCB estimate
return e^*

ALGORITHM 5: QSample: Function for sampling query nodes in multi-stage MCTS

(input) (\mathbf{q}, \mathbf{r}) : outcome node
 (input) M : maximum level to sample from

720

if (\mathbf{q}, \mathbf{r}) has no children **then**
 return $V^O(\mathbf{q}, \mathbf{r})$ (return the value of this outcome node)
if (\mathbf{q}, \mathbf{r}) has children **then**
 if $|\mathbf{q}| < M - 1$ **then**
 $(\mathbf{q}, \mathbf{r}, e') \leftarrow$ child node of (\mathbf{q}, \mathbf{r}) with the greatest UCB estimate
 OSample($\mathbf{q}, \mathbf{r}, e$)
 else
 $(\mathbf{q}', \mathbf{r}') \leftarrow$ random leaf node, descendant from (\mathbf{q}, \mathbf{r})
 return $V^O(\mathbf{q}', \mathbf{r}')$

ALGORITHM 6: OSample: Function for sampling outcome nodes in multi-stage MCTS

(input) $(\mathbf{q}, \mathbf{r}, e)$: query node

721

$N[(\mathbf{q}, \mathbf{r}, e)] \leftarrow N[(\mathbf{q}, \mathbf{r}, e)] + 1$
 $\mathbf{q}' \leftarrow \mathbf{q} + \mathbf{u}^e$ (new query vector with edge e added)
 $Z \leftarrow$ randomly sample a response to edge e (0 if accept, 1 if reject)
 $\mathbf{r}' \leftarrow \mathbf{r} + Z\mathbf{u}^e$ (updated rejection vector)
 $U[(\mathbf{q}, \mathbf{r}, e)] \leftarrow U[(\mathbf{q}, \mathbf{r}, e)] + \text{QSample}(\mathbf{q}', \mathbf{r}')$

722

Algorithm 5 (QSample) samples query nodes from an outcome node, while Algorithm 6 (OSample)
 723 samples outcome nodes from a query node (and updates the query node’s UCB value estimate).

724

Multi-Stage Greedy. Algorithm 7 gives a pseudocode description of the multi-stage version of
 725 Greedy. This search heuristic returns the next edge to query with the highest expected final matching
 726 weight, *ignoring all future queries*. In other words, this approach treats every edge as the *last* edge;
 727 one might call this heuristic “myopic” as well as greedy.

ALGORITHM 7: Greedy Heuristic for Multi-Stage Edge Selection

(input) \mathcal{E} : legal edge sets

(input) q : previously-queried edges

(input) r : previously-observed rejections

$e^* \leftarrow \emptyset$ $V^* \leftarrow 0$

for all q' **in** q 's children **do**

$e' \leftarrow$ the new edge queried in child node q'

$r^A \leftarrow r$

$r^R \leftarrow r$

$r_{e'}^A \leftarrow 0$ (response scenario where e' is accepted, and $r_{e'} = 0$)

$r_{e'}^R \leftarrow 1$ (response scenario where e' is rejected, and $r_{e'} = 1$)

$p^A \leftarrow$ probability that e is accepted, conditional on previous responses

$p^R \leftarrow$ probability that e is rejected, conditional on previous responses

$V' \leftarrow p^A \cdot W(M(r^A); q', r^A) + p^R W(M(r^R); q', r^R)$ (value of querying edge e')

if $V' > V^*$ **then**

$e^* \leftarrow e'$

$V^* \leftarrow V'$

return e^*
