Assortment Optimization for Matching Patients and Providers

Abstract

Strong patient-provider relationships are critical for effective healthcare delivery. However, high provider turnover rates lead to situations where patients lack providers, which poses a significant logistical challenge. We address this by proposing automated patient-provider matching algorithms. We formulate patient-provider matching as an instance of assortment optimization, where patients are offered a set of provider options and respond sequentially. We then develop solutions built upon bipartite matching and demonstrate that our algorithms have approximation guarantees and improve match quality compared to baselines.

Keywords: Matching, Patient, Provider, Healthcare Operations, Assortment Planning

Data and Code Availability We use synthetically generated data, we plan to share the code/dataset, and attach an anonymized repository.

Institutional Review Board (IRB) Our work does not require an IRB.

1. Introduction

While providers play an essential role in the healthcare system (Pearson and Raeke, 2000; Wu et al., 2022), high provider turnover rates frequently leave patients without a provider and disrupt patient care (Reddy et al., 2015) This problem is especially pressing in primary care due to the need for care continuity (Kajaria-Montag et al., 2024). In these situations, healthcare administrators can manually rematch patients, but doing so is costly and inefficient.

We address these logistical burdens by studying algorithms for patient-provider matching. To allow for patients to have agency, we frame the problem using an assortment optimization framework so each patient receives a "menu" of potential providers from which to choose (Shi, 2016; Rios and Torrico, 2023; Davis et al., 2013). We release menus for all patients upfront (e.g. through a patient portal) then let patients respond and select providers. Because patient response order is random, prior work in assortment optimization fails to solve this problem, so we develop new matching algorithms that achieve good theoretical and empirical performance.

Our contributions are: (1) we formalize the patient-provider matching problem using assortment optimization, (2) we develop algorithms for patient-provider matching using bipartite matching, and (3) we validate our algorithms theoretically, through approximation bounds, and empirically, through improved match quality on a synthetic dataset.

2. Problem Setup

An instance of the patient-provider matching problem consists of N patients and M providers. A match between patient i and provider j has match quality $\theta_{i,j}$. Match quality encompasses factors that impact patient-provider relationships such as insurance compatibility, physical distance, and language concordance (Manson, 1988). We can learn match quality from data, such as patient surveys and clinical records, and use this to predict patient-provider compatibility. We note that match quality must be learned carefully to avoid perpetuating biases (Rogo-Gupta et al., 2018).

We offer menus $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N$ upfront to patients. Each menu, $\mathbf{x}_i \in \{0,1\}^M$, details which providers are offered to patient i, where $x_{i,j} = 1$ indicates provider j is offered to patient i. Patients then respond sequentially in a random order $\pi = \pi_1, \pi_2, \ldots, \pi_N$, where π_t is the t^{th} patient. Menus are offered upfront to reduce logistical burden, while patients respond in random order because each makes their selection at a random time. Patients select providers based on a choice function $f(\mathbf{x}_{\pi_t}, \mathbf{y}_t) \in \{0,1\}^M$, a 0-1 vector denoting which provider (if any) is selected. Here, $\mathbf{y}_t \in \{0,1\}^M$ indicates which providers are available when patient π_t is making a decision. Initially, all providers are available, $y_{1,j} = 1$, and providers transition from available to unavailable upon selection: $y_{t,j} = y_{t-1,j}(1 - f(\mathbf{x}_{\pi_{t-1}}, \mathbf{y}_{t-1})_j)$.

We select menus to optimize for match rate, $\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j$, and match quality, $\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}$, selecting these due to the needs of our healthcare partners. Each objective is optimized over all patient response orderings:

$$\max_{\mathbf{x}} \mathbb{E}_{\pi}[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}] \quad (1)$$

We focus on f as a *uniform choice model*. We define a uniform choice model as follows: with probability p, a patient selects their most preferred (i.e. highest match quality) available provider, and with probability $1-p$, selects no provider, for fixed p. We selected this due to its simplicity and flexibility; we leave investigation into alternative models for future work.

3. Algorithms for Matching

3.1. Greedy Algorithms

Greedy solutions to the patient-provider matching problem offer all providers to all patients:

Definition 1 ***Greedy Menu*** - *We define the greedy menu,* $\mathbf{x}^G$*, as* $x_{i,j}^G = 1$ *for all* i *and* j*.*

While greedy approaches perform well in other assortment optimization tasks (Aouad and Saban, 2023), here greedy approaches result in poor matches:

Lemma 2 *Consider an instance of the assortment optimization problem with* N *patients,* M *providers, and match quality* θ*. Let* $x_{i,j}^G$ *be the greedy menu. Let* $f(\mathbf{x}_{\pi_t}, \mathbf{y}_t)$ *be the uniform choice model with parameter* p*. Then let the match quality of the greedy algorithm be* $\text{ALG} = \mathbb{E}_{\pi}[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}^G, \mathbf{y}_i)_j \theta_{i,j}]$ *and let the optimal solution be* $\text{OPT} = \max_{\mathbf{x}} \mathbb{E}_{\pi}[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}]$ *Then, for any* p *and* ϵ*, there exists* θ *and* N*, so* $\text{ALG} \leq \epsilon\text{OPT}$*.*

We prove this by constructing scenarios where greedy menus achieve $O(\epsilon)$ reward (proofs in Appendix C).

3.2. Bipartite Matching Algorithm

To improve upon greedy solutions, we propose an algorithm based on bipartite matching. We first solve a bipartite matching problem between patients and providers, with edge weights $\theta_{i,j}$, and offer each patient their corresponding bipartite match.

Definition 3 ***Bipartite Matching Menu*** - *Let* $z_{i,j}$ *be the optimal solution to the bipartite matching problem between* N *patients and* M *providers with coefficients* $\theta_{i,j}$ *and 1-1 matching constraints. Then the bipartite matching menu,* $\mathbf{x}^B$*, is* $x_{i,j}^B = z_{i,j}$*.*

Our bipartite matching algorithm avoids the pitfalls of greedy solutions because it considers matches globally, which improves performance guarantees:

Theorem 4 *Consider an instance of the assortment optimization problem with* N*,* M*,* θ*. Let* $f(\mathbf{x}_{\pi_t}, \mathbf{y}_t)$ *be the uniform choice model with parameter* p*. Let* $\mathbf{x}_i^B$ *be the bipartite matching menu. Let* $\text{ALG} = \mathbb{E}_{\pi}[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}^B, \mathbf{y}_i)_j \theta_{i,j}]$ *and let* $\text{OPT} = \max_{\mathbf{x}} \mathbb{E}_{\pi}[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{\pi_i,j}]$*. Then* $\text{ALG} \geq p\text{OPT}$*.*

We prove this by upper bounding the optimal match rate with a bipartite matching problem, then showing that our algorithm achieves a p-fraction of this value.

3.3. Grouping Algorithm

To improve the performance of our algorithm for small p, we augment the bipartite matching menu with more options for patients. We do so by selectively grouping patients and aggregating their menus. For example, we might group patients 1, 2, and 3 so that each has a menu of providers X, Y, and Z.

Our procedure starts by considering all subsets of exactly B patients. For each subset, we compute the change in expected match quality from aggregating their menus, which we call α. We then sort all α values in descending order and form groups greedily: subsets with higher α become groups, and subsets only become groups if all members are still "ungrouped." We repeat this for subsets of size $B-1$ to 2 and any "ungrouped" patient keeps their single bipartite match. We present details in Algorithm 1.

We select such an approach because it improves match quality while preserving match rate, the only such augmenting method that does so:

Lemma 5 *Consider an instance of the patient-provider matching problem where* $N = M$ *with match quality* θ*. Let* $z_{i,j}$ *be the 1-1 bipartite matching solution, yielding an assortment* $\mathbf{x}^B$*. Let* $v_i = j$ *if* $z_{i,j} = 1$*. Next, consider a set of augmentations denoted through a graph* $G = (V, E)$*, where nodes correspond to patients, and an edge from* i *to* i' *means that* $x_{i',v_i} = 1$*. Then each patient is offered at least one available provider if and only if the graph* G *consists of connected components that are each complete.*

We prove this by constructing patient orderings so that non-complete graphs result in patients with empty menus, demonstrating that only group-based algorithms guarantee non-empty menus.

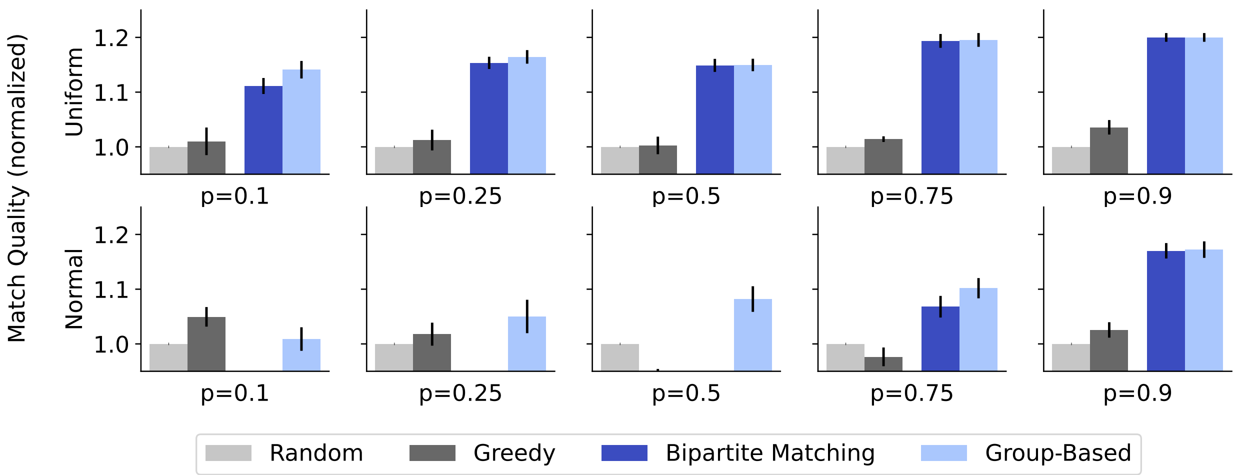

Figure 1: Bipartite matching algorithms outperform random and greedy baselines for uniformly distributed θ Group-based approaches build upon bipartite matching, and its improvement is most pronounced when p is small and θ is normally distributed, due to the need to reoffer providers.

Algorithm 1 Grouping algorithm

Input: Bipartite Menu, $\mathbf{x}^B$, and match quality θ
Output: Grouping menu, $\mathbf{x}^R$
Let $v_i = j$ if $x_i^B = j$
Let $c_i = 1$ for $1 \leq i \leq N$
Initialize $\mathbf{x}^R = \mathbf{x}^B$
for $k = B$ **to** 2 **do**
  for all $S \subseteq \{1, \dots, N\}$, $|S| = k$ **do**
    Let $\mathbf{x}' = \mathbf{x}^R$ and $x'_{i,v_j} = 1$ for all $i, j \in S$
    Let $\alpha_S = \mathbb{E}_\pi[\sum_{i \in S} \sum_{j=1}^{M} f(\mathbf{x}'_{\pi_i}, \mathbf{y}_i)_j \theta_{\pi_i,j}] - \mathbb{E}_\pi[\sum_{i \in S} \sum_{j=1}^{M} f(\mathbf{x}^B_{\pi_i}, \mathbf{y}_i)_j \theta_{\pi_i,j}]$
  end for
  Let $\alpha'_1, \alpha'_2, \dots, \alpha'_l$ be α sorted in descending order, with corresponding subsets $S_1, S_2, \dots, S_l$
  for $i = 1$ **to** l **do**
    if $c_j = 1$ for all $j \in S_i$ and $\alpha'_i > 0$ **then**
      Let $x^R_{j,v_{j'}} = 1$ for all $j, j' \in S_i$
      Let $c_j = 0$ for all $j \in S_i$
    end if
  end for
end for

4. Experiments

4.1. Experimental Setup + Datasets

We compare our algorithms against random and greedy baselines on a synthetic dataset. We construct our synthetic dataset by randomly generating θ according to one of two distributions: (i) uniform: $\theta_{i,j} \sim U(0,1)$ and (ii) normal: $\theta_{i,j} \sim \mathcal{N}(\mu_j, \sigma)$, $\mu_j \sim U(0,1)$. The former corresponds to situations where all match qualities are independent, while the latter corresponds to more and less popular providers. Because our algorithms maximize match rate (see Lemma 5), we compare algorithms according to the normalized match quality, which is the match quality divided by that of the random algorithm. We include further details in Appendix B.

4.2. Algorithm Comparison

We compare our algorithm to greedy and random baselines while varying θ and p. We fix $N = M = 25$, vary $p \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$, and vary $\theta_{i,j}$ to be either uniformly or normally distributed.

In Figure 1, we show that when $p \geq 0.25$ our group-based algorithm outperforms baselines ($p < 0.005$). When θ is uniformly distributed, group-based algorithms outperform baselines by at least 13%, while for normally distributed θ with $p \geq 0.25$, group-based algorithms outperform all baselines by at least 4%. Group-based algorithms perform poorly for $p = 0.1$ because the low match rate encourages larger menus; one solution is to increase menu size B.

Our bipartite matching and group-based algorithms perform similarly to each other for uniformly

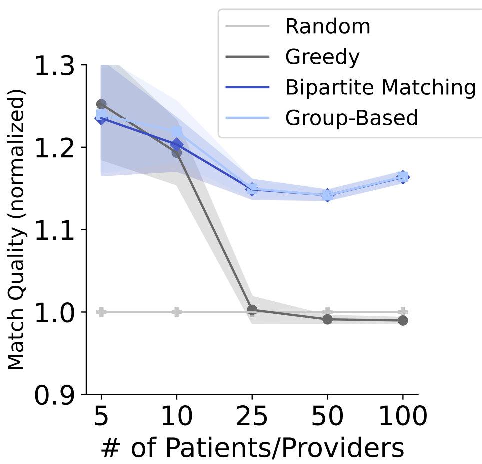

Figure 2: Our bipartite matching and group-based algorithms offer the biggest improvement over baselines when N is large, as larger N increases problem complexity.

distributed θ (within 3%), while for normally distributed θ with $p \leq 0.75$, group-based algorithms are better ($p < 0.0001$). For normally distributed θ it is advantageous to reoffer a provider to multiple patients, which is why group-based outperforms bipartite matching; this is because popular providers should be offered to various patients.

4.3. Varying Patients and Providers

To understand the impact of patient and provider numbers on algorithm performance we vary N and M. We vary $N = M \in \{5, 10, 25, 50, 100\}$ while letting $p = 0.5$ and θ be uniformly distributed. In Appendix A, we experiment with settings where $N \neq M$.

In Figure 2, we find that larger N or M increases the gap between baselines and our algorithms. When $N = 5$, we see that greedy and bipartite matching algorithms perform similarly (within 3%). However, for $N \geq 10$, greedy algorithms perform worse than both of our methods ($p < 0.001$), which occurs due to increased problem complexity with large N.

5. Related Works

To construct patient-provider matches, prior work has investigated algorithms using techniques including genetic programming (Zhu et al., 2023), clustering (Chen et al., 2019), and deferred acceptance (Chen et al., 2020). These works consider matching in a batch setting, and we extend these ideas into a sequential setting with patient deferrals.

We frame patient-provider matching using assortment optimization, a technique that has been applied to domains including retail (Aouad and Saban, 2023), school choice (Shi, 2016), and matching markets (Rios and Torrico, 2023). Within assortment optimization, different response settings have been studied, including online response, where agents make choices sequentially (Aouad and Saban, 2023), and offline response, where agents make choices in-batch (Davis et al., 2013). Our work can be seen as an intermediate between these two extremes.

Our work can leverage clinical information to compute match qualities, θ. For example, prior work has discussed factors that impact patient-provider relationships, including gender (Greenwood et al., 2018), race (Greenwood et al., 2020), and language (Manson, 1988). We focus on patient-provider match quality because it can impact downstream health outcomes, such as medication intake (Nguyen et al., 2020), and mortality rate (Alsan et al., 2019).

6. Conclusion and Real-World Impact

Strong patient-provider relationships are key to preventive care, but patients are frequently left without any provider. To address this, we propose algorithms to automatically match patients and providers through an assortment optimization. We demonstrate approximation guarantees for our algorithms and show that our algorithms improve upon baselines on a synthetic dataset. We provide three research directions to help bring such algorithms to practice:

1. **Provider workload balance** - Our algorithms currently optimize for match quality, but provider-side objectives such as provider workload balance should impact matches.

2. **Varied Choices Models** - Alternate choice models might better capture patient decision-making and lead to more realistic models.

3. **Real data** - We are currently working with healthcare partners to obtain real-world data that allows us to better estimate N, M, and θ.

References

Marcella Alsan, Owen Garrick, and Grant Graziani. Does diversity matter for health? experimental evidence from oakland. *American Economic Review*, 109(12):4071–4111, 2019.

Ali Aouad and Daniela Saban. Online assortment optimization for two-sided matching platforms. *Management Science*, 69(4):2069–2087, 2023.

Ruimin Chen, Mutong Chen, and Hui Yang. Dynamic physician-patient matching in the healthcare system. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 5868–5871. IEEE, 2020.

Xi Chen, Liu Zhao, Haiming Liang, and Kin Keung Lai. Matching patients and healthcare service providers: a novel two-stage method based on knowledge rules and owa-nsga-ii algorithm. *Journal of Combinatorial Optimization*, 37:221–247, 2019.

James Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment planning under the multinomial logit model with totally unimodular constraint structures. *Work in Progress*, 2013.

Brad N Greenwood, Seth Carnahan, and Laura Huang. Patient–physician gender concordance and increased mortality among female heart attack patients. *Proceedings of the National Academy of Sciences*, 115(34):8569–8574, 2018.

Brad N Greenwood, Rachel R Hardeman, Laura Huang, and Aaron Sojourner. Physician–patient racial concordance and disparities in birthing mortality for newborns. *Proceedings of the National Academy of Sciences*, 117(35):21194–21200, 2020.

Harshita Kajaria-Montag, Michael Freeman, and Stefan Scholtes. Continuity of Care Increases Physician Productivity in Primary Care. *Management Science*, 2024.

Aaron Manson. Language concordance as a determinant of patient compliance and emergency room use in patients with asthma. *Medical care*, 26(12):1119–1128, 1988.

AM Nguyen, N Siman, M Barry, CM Cleland, H Pham-Singer, O Ogedegbe, C Berry, and D Shelley. Patient-physician race/ethnicity concordance improves adherence to cardiovascular disease guidelines. *Health Serv Res*, 55(Suppl 1):51, 2020.

Steven D Pearson and Lisa H Raeke. Patients' trust in physicians: many theories, few measures, and little data. *Journal of general internal medicine*, 15:509–513, 2000.

Ashok Reddy, Craig E Pollack, David A Asch, Anne Canamucio, and Rachel M Werner. The effect of primary care provider turnover on patient experience of care and ambulatory quality of care. *JAMA internal medicine*, 175(7):1157–1162, 2015.

Ignacio Rios and Alfredo Torrico. Platform design in matching markets: a two-sided assortment optimization approach. *arXiv preprint arXiv:2308.02584*, 2023.

Lisa J Rogo-Gupta, Carolyn Haunschild, Jonathan Altamirano, Yvonne A Maldonado, and Magali Fassiotto. Physician gender is associated with press ganey patient satisfaction scores in outpatient gynecology. *Women's Health Issues*, 28(3):281–285, 2018.

Peng Shi. Assortment planning in school choice. Technical report, mimeo, 2016.

Dezhi Wu, Paul Benjamin Lowry, Dongsong Zhang, and Youyou Tao. Patient trust in physicians matters—understanding the role of a mobile patient education system and patient-physician communication in improving patient adherence behavior: Field study. *Journal of Medical Internet Research*, 24(12):e42941, 2022.

Wangqi Zhu, Jitao He, and Hangxin Guo. Doctor-patient bilateral matching considering diagnosis and treatment perception in the absence of public health resources. *Frontiers in Public Health*, 10:1094523, 2023.

Appendix A. Other Choices of Patients and Providers

We evaluate the impact of varying $N \in \{5, 10, 25, 50, 100\}$ while fixing $M = 25$, and varying $N \in \{5, 10, 25, 50, 100\}$ while fixing $M = 25$, and plot this in Figure 3. We find that when $N \leq M$, our algorithm performs better than baselines. However, when N is much larger than M, we find that baselines can perform better. This is due to provider scarcity, making it important to show any provider, even with low match quality.

Appendix B. Experimental Details

We run experiments for 6 seeds and 10 trials per seed. We resample θ for different seeds and fix θ but vary π for different trials. We let $B = 3$ for all experiments, and restrict menus to be of size at most $R = 5$ to model real-world scenarios (where patient menus cannot be of arbitrary size). For menus larger than R, we randomly sample a subset of the menu.

Appendix C. Proofs

Lemma 2:
Consider an instance of the assortment optimization problem with N patients, M providers, and match quality $\theta_{i,j}$. Let $x^G_{i,j}$ be the greedy menu. Let $f(\mathbf{x}_{\pi_t}, \mathbf{y}_t)$ be the uniform choice model with parameter p. Then let the match quality of the greedy algorithm be Let ALG $= \mathbb{E}_\pi[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}^G_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}]$ and let the optimal solution be

$$\text{OPT} = \max_{\mathbf{x}} \mathbb{E}_\pi[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}] \quad (2)$$

Then, for any p and ϵ, there exists some $\theta_{i,j}$ and N, so that ALG $\leq \epsilon$OPT

Proof We construct a problem instance where the greedy algorithm performs an ϵ fraction of the optimal algorithm; that is ALG $\leq \epsilon$OPT. To do so, we consider an instance of the assortment optimization problem with $N = M$ patients and providers. Let $\theta_{1,1} = 1$, while $\theta_{i,1} = 2\delta$ for $i \neq 1$, where $\delta \leq \frac{1}{2}$. Let $\theta_{i,j} = \delta$ for all i and for $j \neq 1$. In other words, provider 1 has a match quality of 1 with patient 1, and 2δ for all other patients. All other provider-patient pairs have a match quality of δ.

In this scenario, the optimal selection is to let $\mathbf{x}_i = \mathbf{e}_i$, so that each patient gets a menu of size one, with only provider i being available. This results in an expected total match quality of OPT $= p(\delta(N-1)+1)$. Note that this corresponds to the bipartite matching algorithm.

Next, consider the greedy algorithm, where $\mathbf{x}_i = \mathbf{1}$. Note that all patients prefer provider $j = 1$. By symmetry, each patient has an equal chance of receiving provider $j = 1$. Therefore, with probability at most $\frac{1}{N}$, patient i is first (that is $\pi_1 = i$). therefore, with probability $\frac{1}{N}$, $\pi_1 = i$, and we receive a match quality of 1, while with probability $1 - \frac{1}{N}$, we receive a match quality of 2δ. For all the $M - 1 = N - 1$ other providers, we receive a utility of δ, with a probability p of accepting each. Combining gives that our expected total utility is $\frac{1}{N} + \frac{2\delta(N-1)}{N} + p(N-1)\delta$.

Therefore, we get the following:

$$\begin{aligned}
&\frac{\text{ALG}}{\text{OPT}} \\
&\quad = \frac{\frac{1}{N} + \frac{2\delta(N-1)}{N} + p(N-1)\delta}{p(\delta(N-1)+1)} \\
&\quad \leq \frac{\frac{1}{N} + 2\delta + p(N-1)\delta}{p(\delta(N-1)+1)} \leq \frac{\frac{1}{N} + 2\delta + p(N-1)\delta}{p} \\
&\quad \leq \frac{1}{Np} + 2\frac{\delta}{p} + N\delta
\end{aligned}$$

We let $\frac{1}{Np} \leq \frac{\epsilon}{3}$, so we let $N = \frac{3}{\epsilon p}$. Additionally, we let $N\delta \leq \frac{\epsilon}{3}$, so $\delta \leq \frac{\epsilon}{3N}$, and $2\frac{\delta}{p} \leq \frac{\epsilon}{3} = \delta \leq \frac{2p\epsilon}{3}$. Letting $\delta = \min(\frac{2p\epsilon}{3}, \frac{\epsilon}{3N})$, shows that

$$\frac{\text{ALG}}{\text{OPT}} \leq \frac{1}{Np} + 2\frac{\delta}{p} + N\delta \leq \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon \quad (3)$$

Therefore, for any choice of ϵ and p, there exists a choice of N and θ (implicitly chosen through δ), so that ALG $\leq \epsilon$OPT. ■

Theorem 4:
Let π be unknown, so that we aim to maximize:

$$\max_{\mathbf{x}} \mathbb{E}_\pi[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}] \quad (4)$$

Let $\mathbf{x}^M_i$ be the bipartite matching menu. Then let resulting objective value be ALG $= \mathbb{E}_\pi[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}^M_{\pi_i}, \mathbf{y}_i)_j \theta_{i,j}]$ and let the optimal solution be

$$\text{OPT} = \max_{\mathbf{x}} \mathbb{E}_\pi[\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{M} f(\mathbf{x}_{\pi_i}, \mathbf{y}_i)_j \theta_{\pi_i,j}] \quad (5)$$

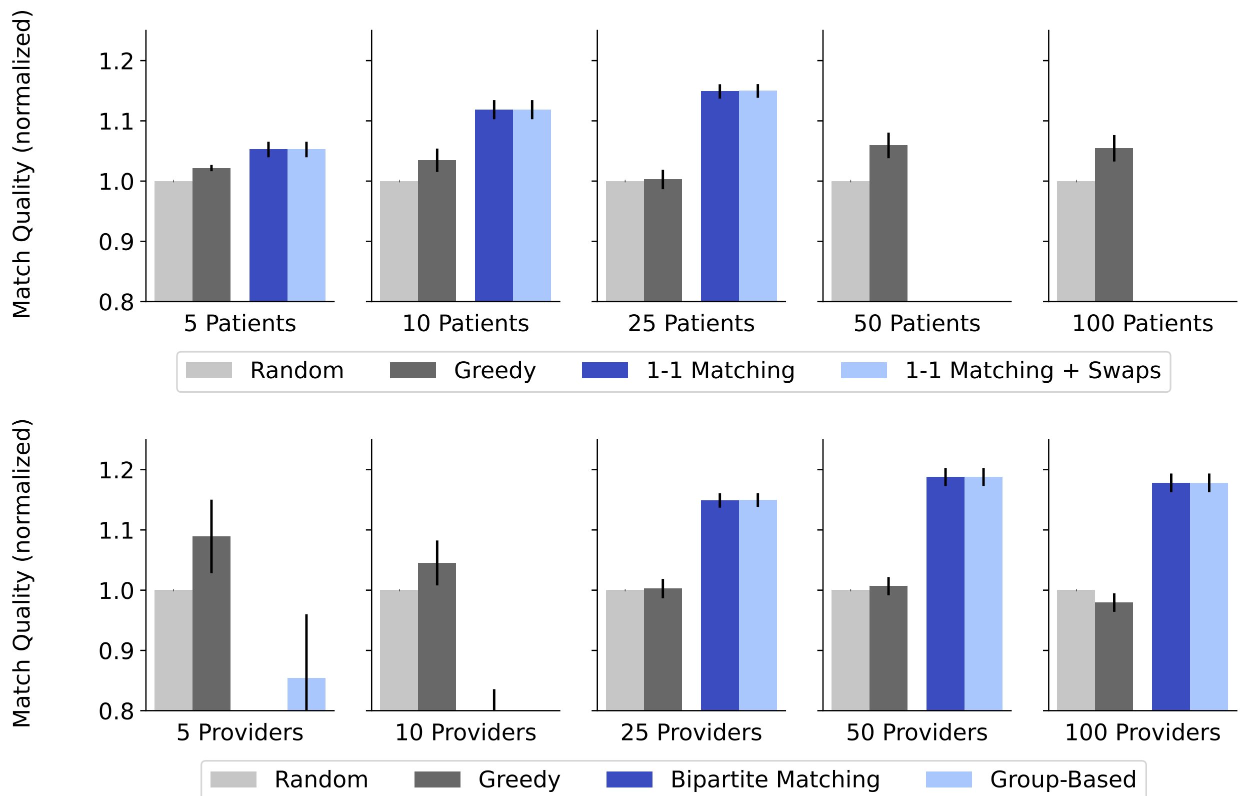

Figure 3: We evaluate scenarios with $M \neq N$, fixing $M = 25$ while varying N, and vice versa. In most situations, our algorithm outperforms baselines, while only in situations where $N > M$ do baselines outperform our method. This occurs due to provider scarcity, requiring us to match with every provider, which greedy algorithms do well.

Then ALG $\geq p$OPT.

Proof First, in our offering of the bipartite matching menu, we note that each patient is only offered one provider, and no provider is offered to more than one patient. Under this scenario, the results of each patient are independent Bernoulli variables, with probability of success p, scaled by the appropriate θ values. That is, the expected match quality is $\sum_{i=1}^{N} x_{i,j}^{B}\theta_{i,j}p$

Next, suppose that there exists some allocation of menus so that the match quality under such an assortment is higher than that of the corresponding linear program. Let the matches from such an allocation be denoted $u_{i,j}$. Then $\sum_{i=1}^{N} u_{i,j}\theta_{i,j} > \sum_{i=1}^{N} x_{i,j}^{B}\theta_{i,j}$. However, such a statement is a contradiction, as by the definition of z, it maximizes $\sum_{i=1}^{N} x_{i,j}^{B}\theta_{i,j}$. Therefore, no solution can improve upon the match quality of the bipartite match, which implies that OPT $\leq \sum_{i=1}^{N} x_{i,j}^{B}\theta_{i,j}$, while ALG $= p\sum_{i=1}^{N} x_{i,j}^{B}\theta_{i,j}$, so our algorithm achieves a reward of pOPT. ∎

Lemma 5:
Consider a situation where $N = M$. Let $z_{i,j}$ be the 1-1 bipartite matching solution to a problem with coefficients $\theta_{i,j}$. Let $v_i = j$ if $z_{i,j} = 1$. Consider the bipartite matching assortment, $\mathbf{x}_i^B$ Next, consider a set of augmentations denoted through a graph $G = (V, E)$, where nodes correspond to patients, and an edge from i to i' means that $x_{i',v_i} = 1$. Then each patient is offered at least one available provider if and only if the graph G consists of connected components that are each connected.

Proof We will first prove the forward direction; that if the set of assortments is complete, then the resulting menu is non-empty. Consider patient i in a complete graph of size k, so that the size of the menu for patient i is also k. Each time a member of the complete graph, i' is chosen before i in the ordering π, the set of available options in the menu decreases by 1. Because there are k members in the complete graph, at most $k-1$ things can come before i in the ordering, and so the menu size is at least $k-(k-1)=1$, which implies the menu is non-empty.

Next, we will prove that if the resulting menu is always non-empty, then the underlying graph must be a complete graph structure. We first consider some node u in the graph, corresponding to some patients. Suppose that u has a menu of size k, indicating that there exist k edges from u to some node. Suppose that there exists a node, v such that (v, u) is an edge, but (u, v) is not. Then consider the ordering that places u after its neighbors and v. In this ordering, we let patient v have u on its menu, and we suppose that v selects the provider assigned to u. Next, we let each of the neighbors for u, i, select themselves. This results in neither u nor its neighbors being available when u must select a patient, leaving an empty menu.

Therefore, for menus to be non-empty, it must be the case that any edge (i, u) must also have (u, i). Next, we consider the scenario where there exists a node w such that w is a two-hop neighbor of u but not an immediate neighbor of u. Suppose w is adjacent to v, so that v is on w's menu. Order the patients such that w comes first, then u's neighbors, then u. Let w select v, let v select u, and let all of u's other neighbors select themselves. This results in an empty menu for u; therefore, when all menus are nonempty, there must not exist any two-hop neighbors that are not also one-hop neighbors. Because this is undirected as shown before, for any node, all of its neighbors are a distance 1 away. Therefore, in any component, it is the case that all nodes are connected. This implies that the graph consists of complete graphs. ∎