

Efficient and Fair Healthcare Rationing

Haris Aziz

UNSW Sydney, Australia

HARIS.AZIZ@UNSW.EDU.AU

Florian Brandl

University of Bonn, Germany

FLORIAN.BRANDL@UNI-BONN.DE

Abstract

The rationing of healthcare resources has emerged as an important issue, which has been discussed by medical experts, policy-makers, and the general public. We consider a rationing problem where medical units are to be allocated to patients. Each unit is reserved for one of several categories, and each category has a priority ranking over the patients. We present a class of allocation rules that respect the priorities, comply with the eligibility requirements, allocate the largest feasible number of units, and do not penalize agents for rising in the priority ranking of a category. The rules characterize all possible allocations that satisfy the first three properties and are polynomial-time computable.

1. Introduction

The COVID-19 pandemic emerged as one of the major challenges the world faced. It resulted in a frantic race to produce the most effective and safe vaccine to stem the devastating effects of the pandemic. While the creation, testing, and approval of vaccines proceeded at an unprecedented pace, the initial scarcity posed the question of how to distribute them efficiently and fairly. The same problem arises for other scarce or costly healthcare resources such as ventilators and antiviral treatments. The issue of fair allocation of scarce societal resources is also essential in artificial intelligence (see, e.g., Das (2022), Aziz (2020)). One approach to allocating resources is to assign patients to priority groups, which capture how much a patient needs treatment. For example, medical practitioners and policy-makers highlighted three priority groups: healthcare workers, other essential workers and people in high-transmission settings, and people with medical vulnerabilities associated with poorer COVID-19 outcomes (Persad, Peek, & Emanuel, 2020; Truog, Mitchell, & Daley, 2020). Other concerns that have been discussed include racial equity (Bruce & Tallman, 2021).

It is, however, not enough to identify priority groups. There is also a need to algorithmically and transparently make allocation decisions based on the priorities (Emanuel, Persad, Upshur, Thome, Parker, Glickman, Zhang, Boyle, Smith, & Phillips, 2020; WHO, 2020). In a New York Times article, the issue has been referred to as one of the hardest decisions health organizations must make (Fink, 2020). Since the decisions need to be justified to the public, they must be aligned with ethical guidelines, such as respecting the patients' priorities. These decisions are not straightforward, especially when patients are eligible for multiple categories. In that case, the decision of which category is used to serve a patient can have compounding effects on what categories other agents can use. A competing objective is allocating the resources efficiently to maximize their social benefit. For example, no medical unit should be left unallocated. Thus, the following fundamental question arises.

How should we allocate scarce medical resources fairly and efficiently while accounting for various ethical principles and priority groups?

The problem of healthcare rationing has recently been formally studied by market designers. Pathak, Sönmez, Ünver, and Yenmez (2020, 2023) were among the first to frame the problem as a two-sided matching problem in which patients are on one side, and the resource units are on the other side. Thus, they linked the healthcare rationing problem with the rich field of two-sided matching (Roth & Sotomayor, 1990). Pathak et al. (2020, 2023) suggested dividing the units into different reserve categories, each with a subset of agents eligible for that category and its own priority ranking over the patients. The categories and the category-specific priorities represent the ethical principles and guidelines that a policy-maker may wish to implement.¹ For example, a category for senior people may prioritize the eldest citizens, or a category for healthcare workers may rank ICU personnel ahead of medical workers not directly exposed to the disease. Having a holistic framework that considers different types of priorities has been termed important in healthcare rationing.² The approach of Pathak et al. (2020, 2023) has been recommended or adopted by various organizations, including the NASEM (National Academies of Sciences, Engineering, and Medicine) Framework for Equitable Vaccine Allocation (NASEM-National Academies of Sciences, Engineering, and Medicine, 2020) and has been endorsed in medical circles (Persad et al., 2020; Sönmez, Pathak., Ünver, Persad, Truog, & White, 2020). The approach has been covered widely in the media, including the New York Times and the Washington Post.³

Pathak et al. (2020, 2023) proposed the Smart Reserves algorithm for category priorities consistent with a baseline priority ordering.⁴ It computes a maximum size matching satisfying the basic axioms (eligibility compliance, respect of priorities, and non-wastefulness). The Smart Reserves algorithm carefully decides which category each patient should use to achieve the maximum size property. However, the literature has not addressed the problem with possibly *heterogeneous* priorities. Allowing heterogeneous priorities for categories is very much in the spirit of incorporating different ethical values. For example, it seems desirable that the priority ordering in the category of older people can be different from the priority ordering in the category of front-line workers. The former may rank individuals by age, and the latter may favor energetic medical professionals.⁵ In this paper, we set out to address this issue and answer the following research problem.

-
1. See, for example, the book by Bognar and Hirose (2014) on the ethics of healthcare rationing that discusses many of these principles.
 2. In a report issued by the Deeble Institute for Health Policy Research, Martin (2015) writes that “To establish robust healthcare rationing in Australia, decision-makers need to acknowledge the various implicit and explicit priorities that influence the process and develop a decision-making tool that incorporates them.”
 3. <https://www.covid19reservesystem.org/media>
 4. The priority ranking of a category is consistent with a baseline priority ordering if it ranks the agents eligible for it according to the baseline priority.
 5. Even in other contexts such as immigration, where rationing policies are applied, respecting heterogeneous priorities can be important. For example, if a country has a quota for admitting engineers, the top engineering applicant who satisfies basic eligibility requirements may have a good case to be issued a visa.

For healthcare rationing with heterogeneous priorities, how should we allocate resources in a fair, economically efficient, and computationally tractable way?

Contribution We consider healthcare rationing with heterogeneous priorities. Our first contribution is introducing a class of allocation rules called Reverse Rejecting rules, which

- (i) comply with the eligibility requirements,
- (ii) respect the priorities of the categories (for each category, patients of higher priority are served first),
- (iii) yield a maximum size matching (one that allocates the largest feasible number of units to eligible patients),
- (iv) are non-wasteful (no unit is unused but could be used by some eligible patient),
- (v) respect improvements (do not penalize patients for rising in the priority ranking of a category), and
- (vi) are strongly polynomial-time computable.

We prove that Reverse Rejecting rules characterize all outcomes that satisfy the first three properties. We show how Reverse Rejecting rules can be used to obtain the *Smart Reverse Rejecting* rule, which processes a given number of unreserved units first and last and which incorporates the additional goal of allocating the largest feasible number of units from a designated subset of categories called preferential categories. The Smart Reverse Rejecting rule satisfies a new axiom we call order preservation, which is parameterized by how many unreserved units are processed first and last. Moreover, it generalizes two well-known reserves rules—over-and-above and minimum-guarantees (Galanter, 1961, 1984)—that are understood in the context of preferential categories having consistent priorities. Finally, we discuss how our algorithms and their properties extend to soft reserves.

Reverse Rejecting rules also apply to the school choice problem when students are only interested in being matched to one of their acceptable schools and to hiring settings when applicants are interested in one of the positions and each department has its own priorities. Lastly, they apply to other rationing scenarios, such as allocating limited slots at public events or visas to immigration applicants.

2. Related Work

The paper is related to an active area of research on matching with distributional constraints (see, e.g., Kojima (2019), Aziz, Biró, and Yokoo (2022)). One general class of distributional constraints examined in matching market design pertains to common quotas over institutions such as hospitals (Kamada & Kojima, 2015, 2017; Biró, Fleiner, Irving, & Manlove, 2010a; Goto, Iwasaki, Kawasaki, Kurata, Yasuda, & Yokoo, 2016).

Within the umbrella of work on matching with distributional constraints, particularly relevant to healthcare rationing is the literature on school choice with diversity constraints and reserve systems (Hafalir, Yenmez, & Yildirim, 2013; Ehlers, Hafalir, Yenmez, & Yildirim, 2014; Echenique & Yenmez, 2015; Kurata, Hamada, Iwasaki, & Yokoo, 2017;

Aygün & Turhan, 2020; Aygün & Bó, 2020; Aziz, Gaspers, & Sun, 2020; Gonczarowski, Nisan, Kovalio, & Romm, 2019; Dur, Kominers, Pathak, & Sönmez, 2018; Dur, Pathak, & Sönmez, 2020). Categories in healthcare rationing correspond to affirmative action types in school choice. We suggest the book chapter by Heo (2019) for a brief survey. Except for the special case in which students have precisely one type (see, e.g., Ehlers et al. (2014)), most of the approaches do not achieve diversity goals optimally, whereas, for the healthcare rationing problem we consider, we aim to find matchings that maximize the number of units allocated to eligible patients. This literature also considers optimization approaches for diverse matchings but for different models and objectives (see, e.g., Ahmed, Dickerson, and Fuge (2017), Dickerson, Sankararaman, Srinivasan, and Xu (2019), Ahmadi, Ahmed, Dickerson, Fuge, and Khuller (2020)).

Pathak et al. (2020, 2023) were the first to frame a rationing problem with category priorities as a two-sided matching problem in which agents are interested in a single unit of resource, and the resources are reserved for different categories. They show that artificially endowing agents with strict preferences over the categories and running the Deferred Acceptance algorithm results in desirable outcomes for the rationing problem. They note, however, that this approach may lead to matchings that are not Pareto optimal (cf. Remark 4.5). By contrast, Reverse Rejecting rules always allocate the largest feasible number of units and thus give a Pareto optimal allocation.

Pathak et al. (2020, 2023) also consider a setting where given numbers of “unreserved” units must be allocated before or after allocating the units reserved for categories. They propose to use the Smart Reserves approach of Sönmez and Yenmez (2020) for the restricted problem when all categories have consistent priorities. The thereby obtained class of Smart Reserve rules includes two extreme approaches with widespread appeal. The first approach is based on *minimum-guarantees* specifying the minimum number of units kept for a particular agent group. The second approach is *over-and-above*; it sets aside the specified number of units for an agent group and only uses them once all the unreserved units are allocated (for which the agent group is eligible as well).⁶ Pathak et al. (2020, 2023) assume that all the priorities of categories are consistent with a baseline ordering. In contrast, we allow the priorities of categories to be *heterogeneous*. Our Smart Reverse Rejecting rule achieves the features of the Smart Reserves rule for heterogeneous priorities. In follow-up work, Grigoryan (2020) considers optimization approaches for variants of the problem but does not present polynomial-time algorithms.

In this paper, we attempt to compute particular maximum size stable matchings. The problem of computing such matchings is NP-hard if both sides have strict preferences (priorities, respectively) (Biró, Manlove, & Mittal, 2010b). In our problem, the agents have dichotomous preferences (categories they are/are not eligible for), which allows us to obtain a polynomial-time algorithm for the problem. Computing outcomes that match as many agents as possible has also been examined in related but different contexts (see, e.g., Aziz (2018), Andersson and Ehlers (2020), Abraham, Blum, and Sandholm (2007), Bogomolnaia and Moulin (2015)).

Sonmez and Yenmez (2024) model an affirmative action program in India for positions in the public service. In this model, individuals are to be assigned to job types based on

6. Both the “minimum-guarantees” and “over-and-above” approaches have been discussed in the context of reserves systems in India (see, e.g., Galanter (1961, 1984)).

their preferences over job types. Each individual belongs to a single category and has a merit score for each job type (and thus, priorities are heterogeneous across job types). A number of positions in each job type are reserved for each category. Additionally, they consider traits of individuals, which we ignore in this discussion. One way of comparing our settings is to identify their job types with our categories, assume that all individuals are indifferent between all their acceptable job types, and assume that no individual belongs to any of their categories. Then, the mechanism they propose complies with the eligibility requirements, respects priorities, and is non-wasteful; however, Sonmez and Yenmez (2024) do not consider the maximum size property or respect of improvements. In their setting with non-indifferent individuals, their mechanism is strategyproof. Other ways of identifying the two settings result in homogeneous priorities and thus miss out on an essential component of our model. Similarly, Sonmez and Yenmez (2022) examined the particular model of an affirmative action program implemented in India. In the paper, the priorities are not heterogeneous.

Balinski and Sönmez (1999) introduced the *respecting improvements* property in the college admissions model of Gale and Shapley (1962). In that context, it states that students should not be penalized for higher priority among the colleges (e.g., because of better test scores). They show that the student proposing the Deferred Acceptance algorithm is the unique stable mechanism that respects improvements.

3. Model

We adopt the model of Pathak et al. (2020, 2023) with two generalizations: we do not assume that the categories' priorities are consistent with a baseline ordering and we allow the priorities to be weak rather than strict. There are q identical and indivisible units of some resource, a finite set C of categories, and a set N of agents with $|N| = n$. Each category c has a quota $q_c \in \mathbb{N}$ with $\sum_{c \in C} q_c = q$ and a priority ranking \succsim_c , which is a weak order on $N \cup \{\emptyset\}$. An agent i is *eligible* for category c if $i \succ_c \emptyset$. We denote by N_c the set of agents who are eligible for c . We say that $I = (N, C, (\succsim_c)_{c \in C}, (q_c)_{c \in C})$ is an instance (of the rationing problem). We will write (\succsim_c) and (q_c) for profiles of priorities and quotas in the sequel.

A matching $\mu: N \rightarrow C \cup \{\emptyset\}$ is a function that maps each agent to a category or to \emptyset and satisfies the capacity constraints: for each $c \in C$, $|\mu^{-1}(c)| \leq q_c$. For an agent $i \in N$, $\mu(i) = \emptyset$ means that i is unmatched (that is, does not receive any unit), and $\mu(i) = c$ means that i receives a unit reserved for category c . When convenient, we will identify a matching μ with the set of agent-category pairs $\{\{i, \mu(i)\}: \mu(i) \neq \emptyset\}$.⁷

We introduce four axioms in the context of allocating medical units that are well-grounded in practice. For further motivation of these axioms, we recommend the detailed discussions by Pathak et al. (2020, 2023).

The first axiom requires matchings to comply with the eligibility requirements. It specifies that a patient should only take a unit of a category for which the patient is eligible. For example, a young person should not take a unit from the units reserved for older people. We sometimes refer to a matching that complies with the eligibility requirements as a *feasible matching*.

7. In graph theoretic terms, μ is a b -matching because multiple edges in μ can be adjacent to a category c .

Definition 3.1 (Compliance with eligibility requirements). A matching μ *complies with the eligibility requirements* (or is *feasible*) if for any $i \in N$ and $c \in C$, $\mu(i) = c$ implies $i \succ_c \emptyset$.

The second axiom concerns the respect of priorities of categories, and it is the central fairness concept we consider. It rules out that a patient is matched to a unit of some category c while some other agent with a higher priority for c is unmatched.

Definition 3.2 (Respect of priorities). A matching μ *respects priorities* if for any $i, j \in N$ and $c \in C$, $\mu(i) = c$ and $\mu(j) = \emptyset$ implies $j \not\succeq_c i$. If there exist $i, j \in N$ and $c \in C$ with $\mu(i) = c$, $\mu(j) = \emptyset$, and $j \succ_c i$, we say that j has *justified envy* towards i for category c .

An astute reader who is familiar with the theory of stable matchings will immediately realize that respect of priorities is equivalent to *justified envy-freeness* in the context of school-choice matchings (Abdulkadiroğlu & Sönmez, 2003).

Next, non-wastefulness requires that if an agent is unmatched despite being eligible for a category, then all units reserved for that category are matched to other agents.

Definition 3.3 (Non-wastefulness). A matching μ is *non-wasteful* if for any $i \in N$ and $c \in C$, $i \succ_c \emptyset$ and $\mu(i) = \emptyset$ implies $|\mu^{-1}(c)| = q_c$.

Not all non-wasteful matchings allocate the same number of units. In particular, some may not allocate as many units as possible. A stronger efficiency notion prescribes that the number of allocated units is maximal subject to compliance with the eligibility requirements.

Definition 3.4 (Maximum size matching). A matching μ is a *maximum size matching* if it has maximum size among all matchings complying with the eligibility requirements.

These four axioms capture the first guideline put forth in the report by the National Academies of Sciences, Engineering, and Medicine: “ensure that allocation maximizes benefit to patients, mitigates inequities and disparities, and adheres to ethical principles” (NASEM-National Academies of Sciences, Engineering, and Medicine, 2020). Requiring matchings to be of maximum size is aligned with the principle to “gain the best value we possibly can from the expenditure of that resource” (Dawson, Isaacs, Jansen, Jordens, Kerridge, Kihlbom, Kilham, Preisz, Sheahan, & Skowronski, 2020).

It will be useful to associate a graph B_I , called a *reservation graph*, with an instance I . $B_I = (N \cup C, E)$ is a bipartite graph with an edge from i to c if i is eligible for c . That is, $E = \{\{i, c\} : i \succ_c \emptyset\}$. If B_I is any reservation graph, we denote by $ms(B_I)$ the number of edges in a maximum size matching of B_I subject to given quotas (q_c). A maximum size matching of a bipartite graph can be computed in polynomial time (more precisely, cubic time in the number of vertices of the graph) by any of several well-established methods such as Kuhn’s algorithm (Kuhn, 1955) or the Hopcroft-Karp-Karzanov algorithm (Hopcroft & Karp, 1973; Karzanov, 1973).⁸

8. For example, Kuhn’s algorithm proceeds as follows. An *alternating path* with respect to a matching μ is a path in G in which edges alternate between those in μ and those not in μ . An *augmenting path* is an alternating path that starts and ends with an edge not in μ . Kuhn’s algorithm repeatedly applies Berge’s lemma, which states that a matching has maximum size if and only if there is no augmenting path with respect to that matching. Hence, starting with any matching and repeatedly alternating it along an augmenting path as long as an augmenting path exists results in a maximum size matching. The polynomial runtime follows from the fact that augmenting paths can be found in polynomial time and each alteration increases the size of the matching by one edge.




Figure 1: The instance I described in Example 3.5. The reservation graph B_I on the left has an edge from i to c if i is eligible for c . The priority rankings of the categories are depicted on the right.

The following example illustrates the definitions above.

Example 3.5. Suppose there are three agents and two categories with one reserved unit each.

$$N = \{1, 2, 3\}, \quad C = \{c_1, c_2\}, \quad q_{c_1} = 1, q_{c_2} = 1.$$

The priority ranking of c_1 is $2 \succ_{c_1} 3 \succ_{c_1} \emptyset \succ_{c_1} 1$ and the priority ranking of c_2 is $2 \succ_{c_2} \emptyset \succ_{c_2} 1 \succ_{c_2} 3$. Figure 1 illustrates this instance I of the rationing problem.

Note that agent 1 is not eligible for any category, agent 2 is eligible for c_1 and c_2 , and agent 3 is eligible only for c_1 . Thus, the following matchings comply with the eligibility requirements.

$$\begin{aligned} \mu_1 &= \emptyset & \mu_2 &= \{\{2, c_1\}\} & \mu_3 &= \{\{2, c_2\}\} \\ \mu_4 &= \{\{3, c_1\}\} & \mu_5 &= \{\{2, c_2\}, \{3, c_1\}\} \end{aligned}$$

All of these matchings except for μ_4 respect priorities. Only μ_2 and μ_5 are non-wasteful. The only maximum size matching is μ_5 .

We are interested in allocation rules, which, for each instance, return a matching.

Definition 3.6 (Allocation rule). An allocation rule maps every instance I to a matching for I .

We say that an allocation rule f satisfies one of the axioms in Definitions 3.1 to 3.4 if $f(I)$ satisfies the axiom for all instances I .

An allocation rule respects improvements if it never hurts an agent to rise in the categories' priorities. More precisely, suppose an agent is assigned a unit for a given priority profile. Then this agent should still be assigned a unit if her priority for each category increases or stays the same and the priorities among all other agents stay the same. This property has been introduced by Balinski and Sönmez (1999) for two-sided matching problems. More formally, let (\succ_c) and (\succ'_c) be priority profiles and $i \in N$. We say agent i 's priority increases from (\succ_c) to (\succ'_c) if for all $j, k \neq i$ and $c \in C$,

$$\begin{aligned} j \succ_c k &\text{ if and only if } j \succ'_c k \\ i \succ_c j &\text{ implies } i \succ'_c j \text{ and } i \succ_c j \text{ implies } i \succ'_c j \end{aligned}$$

That is, the priority rankings over agents other than i are the same in both profiles and i can only move up in the priority rankings from (\succ_c) to (\succ'_c) . We also say that i 's priority increases from $I = (N, C, (\succ_c), (q_c))$ to $I' = (N, C, (\succ'_c), (q_c))$. Respecting improvements requires that if i is matched under I , then i is also matched under I' .

Definition 3.7 (Respecting improvements). An allocation rule f respects improvements if $f(I)(i) \neq \emptyset$ implies $f(I')(i) \neq \emptyset$ whenever i 's priority increases from I to I' .

Remark 3.8 (Strategyproofness). In healthcare rationing, priorities are based on verifiable characteristics (such as age, occupation, or pre-existing health conditions). However, agents may be able to hide some of these characteristics (for example, a pre-existing health condition) and thereby lower their priority for some categories. Respecting improvements can thus be interpreted as a notion of strategyproofness: assuming that agents prefer receiving a unit to not receiving a unit, each agent's dominant strategy is to declare all characteristics that increase her priority. The restricted version of respecting improvements requiring only that it is never beneficial for an agent to become ineligible for a category has been referred to as incentive compatibility by Aygün and Bó (2020).

4. Reverse Rejecting Rules


We introduce the class of Reverse Rejecting rules. Each Reverse Rejecting rule is based on a linear order \succ_π , over the agents. The Reverse Rejecting rule based on \succ_π (REV_π) considers the agents in ascending order of \succ_π . It rejects an agent if the agents who have not been rejected thus far can form a maximum size matching (among matchings in the reservation graph) for which none of the rejected agents has justified envy towards a matched agent. The second constraint is implemented by removing those edges from the reservation graph B_I that would cause justified envy by rejected agents. That is, if agent i is rejected, we cannot allow any agent j to be matched to a category c if $i \succ_c j$, and so remove the edge $\{j, c\}$ from B_I . More generally, if R is the set of rejected agents, $B_I^{-R} = ((N \setminus R) \cup C, E)$ with

$$E = \{\{j, c\} : j \succ_c \emptyset \text{ and there is no } i \in R \text{ such that } i \succ_c j\}$$

is the corresponding reduced reservation graph. Note that if $R \supset R'$, then B_I^{-R} is a subgraph of $B_I^{-R'}$. In particular, B_I^{-R} is a subgraph of B_I . The REV_π rule then works as follows:

Let the set of rejected agents R be empty at the start and consider the agents in ascending order of \succ_π . When considering agent i , add i to the rejected agents R if and only if $ms(B_I^{-R \cup \{i\}}) = ms(B_I)$. After the last agent has been considered, let R_I be the final set of rejected agents and choose a maximum size matching of the reduced reservation graph $B_I^{-R_I}$.

The methodology of Reverse Rejecting rules is different from the Horizontal Envelope rule of Sönmez and Yenmez (2020) and the Smart Reserves rule of Pathak et al. (2020, 2023). In contrast to iteratively or instantly selecting agents that will be matched, Reverse Rejecting rules proceed in ascending order of order over agents and decide which agents to reject. More importantly, Reverse Rejecting rules work for heterogeneous and weak priorities.


 Figure 2: Graphs for the instance I in Example 4.1.

Example 4.1 (Illustration of Reverse Rejecting rules). Consider an instance with

$$N = \{1, 2, 3, 4\}, \quad C = \{c_1, c_2\}, \quad q_{c_1} = q_{c_2} = 1$$

The priorities are $1 \sim_{c_1} 4 \succ_{c_1} 2 \succ_{c_1} \emptyset$ and $1 \sim_{c_2} 3 \succ_{c_2} \emptyset$. For $1 \succ_{\pi} 2 \succ_{\pi} 3 \succ_{\pi} 4$, let us simulate the REV_{π} rule. The reservation graph B_I is depicted in Figure 2a. It has a maximum size matching of size 2. First agent 4 is considered. Since the graph $B_I^{-\{4\}}$ depicted in Figure 2b admits a matching of size 2, agent 4 is placed in R . Next, agent 3 is considered and not placed in R since the graph $B_I^{-\{3,4\}}$ depicted in Figure 2c does not admit a matching of size 2. The graph $B_I^{-\{2,4\}}$ depicted in Figure 2d admits a matching of size 2. Hence, agent 2 is placed in R . Lastly, since $B_I^{-\{1,2,4\}}$ depicted in Figure 2e does not admit a matching of size 2, agent 1 is not placed in R . The final outcome is a maximum size matching of the graph $B_I^{-\{2,4\}}$ as depicted in Figure 2d.

Theorem 4.2 (Properties of Reverse Rejecting rules). *For each linear order \succ_{π} over agents, the REV_{π} rule*

- (i) *complies with eligibility requirements,*
- (ii) *respects priorities,*
- (iii) *returns a matching of maximum size among feasible matchings,*
- (iv) *respects improvements, and*
- (v) *can be computed in strongly polynomial time.*

Proof. (i) Compliance with eligibility requirements. The outcome is a matching of $B_I^{-R_I}$, which is a subgraph of B_I . Any edge $\{i, c\}$ of B_I is such that $i \succ_c \emptyset$. Therefore the outcome satisfies eligibility requirements.

(ii) Respect of priorities. Let μ be the matching returned by the REV_π rule for the instance I . First, we claim that μ provides no justified envy for any agent in R_I . If $j \in R_I$, $i \in N$, and $c \in C$ with $\mu(i) = c$ and $j \succ_c i$, then $\{i, c\}$ was deleted from the reservation graph when rejecting j and is thus not an edge of $B_I^{-R_I}$. Since μ is a matching of B_I^{-R} , it does not match i to c .

Second, we show that μ matches all agents in $N \setminus R_I$. In particular, no agent in $N \setminus R_I$ can have (justified) envy. Assume for contradiction that $|\mu| < n - |R_I|$. This means that $ms(B_I^{-(R_I \cup \{i\})}) < ms(B_I) = ms(B_I^{-R_I}) = |\mu|$ for all $i \in N \setminus R_I$. Among all maximum size matchings of $B_I^{-R_I}$, let μ' be one that is Pareto optimal with respect to the priorities. If there were $i, j \in N \setminus R_I$ so that j has justified envy towards i in μ' ($\mu'(i) = c$, $\mu'(j) = \emptyset$, and $j \succ_c i$), then the matching $\mu' \setminus \{\{i, c\}\} \cup \{\{j, c\}\}$ would Pareto dominate μ' with respect to the priorities, which cannot be. Thus, μ' respect priorities, so that for $j \in N \setminus R_I$ with $\mu'(j) = \emptyset$, μ' is a matching of $B_I^{-(R_I \cup \{j\})}$. It follows that $ms(B_I^{-(R_I \cup \{j\})}) = ms(B_I^{-R_I})$, which contradicts that the REV_π rule rejects none of the agents in $N \setminus R_I$.

(iii) Maximum size matching. The returned matching is by construction a matching of size $ms(B_I)$ and hence of maximum size among all matching complying with the eligibility requirements.

(iv) Respecting improvements.

For simplicity, suppose \succ_π is $1 \succ_\pi 2 \succ_\pi \dots \succ_\pi n$. Let $i \in N$ and I, I' be instances so that i 's priority increases from I to I' . Assume that i is unmatched in $REV_\pi(I')$. We need to show that i is unmatched in $REV_\pi(I)$.

Note that $ms(B_I) = ms(B_{I'})$ since the edge set of B_I is a subset of the edge set of $B_{I'}$ and $REV_\pi(I')$ does not match i . Let $R_I^i = \{j \in R_I : i \succ_\pi j\}$ be the agents who are unmatched in $REV_\pi(I)$ and are lower in the order \succ_π than i . Define $R_{I'}^i$, similarly. We first show that $R_I^i = R_{I'}^i$. To this end, we prove by induction that $j \in R_I^i$ if and only if $j \in R_{I'}^i$ for $j = n$ to $i + 1$.

The base case is trivial since the set of rejected agents is the empty set at the start of the algorithm under both instances I and I' .

For the induction step, let $j > i$ and suppose $j' \in R_I^i$ if and only if $j' \in R_{I'}^i$ for $j' = n$ to $j + 1$. We prove that $j \in R_I^i$ if and only if $j \in R_{I'}^i$. Denote by R the set of rejected agents at the beginning of the round where agent j is considered. Note that $R = R_{I'}^i \cap \{j + 1, \dots, n\} = R_{I'} \cap \{j + 1, \dots, n\}$. If $j \in R_{I'}^i$, this is because there exists a matching μ of $B_{I'}^{-(R \cup \{j\})}$ with $|\mu| = ms(B_{I'})$. Since $R_{I'} \cap \{j, \dots, n\} \subset R \cup \{j\}$, μ does not match any agent in $R_{I'} \cap \{j, \dots, n\}$. Since i is unmatched in $REV_\pi(I')$, it follows that $i \in R_{I'}$, and we can assume that μ does not match i . Now consider $B_I^{-(R \cup \{j\})}$. For each $k \in N \setminus (R \cup \{j\} \cup \{i\})$ and $c \in C$, $\{k, c\}$ is an edge of $B_I^{-(R \cup \{j\})}$ if and only if $\{k, c\}$ is an edge of $B_{I'}^{-(R \cup \{j\})}$ since the priorities for all agents other than

i are the same in I and I' . It follows that μ is also a feasible matching of $B_I^{-(R \cup \{j\})}$. Therefore $j \in R_I^i$.

For the other direction, suppose $j \in R_{I'}^i$. Hence, there exists a matching μ of $B_I^{-(R \cup \{j\})}$ with $|\mu| = ms(B_I) = ms(B_{I'})$. As above, for each $k \in N \setminus (R \cup \{j\} \cup \{i\})$ and $c \in C$, $\{k, c\}$ is an edge of $B_I^{-(R \cup \{j\})}$ if and only if $\{k, c\}$ is an edge of $B_{I'}^{-(R \cup \{j\})}$. For i and any $c \in C$, note that if $\{i, c\}$ is an edge of $B_I^{-(R \cup \{j\})}$, then $\{i, c\}$ is an edge of $B_{I'}^{-(R \cup \{j\})}$. Hence μ is also a feasible matching of $B_{I'}^{-(R \cup \{j\})}$. Thus, $j \in R_{I'}^i$.

This completes the proof that $R_I^i = R_{I'}^i$.

Now we prove that i is unmatched in I , that is, $i \in R_I$. Since $i \in R_{I'}$, it follows that $ms(B_{I'}^{-(R_{I'}^i \cup \{i\})}) = ms(B_{I'})$. Moreover, $B_{I'}^{-(R_{I'}^i \cup \{i\})}$ is a subgraph of $B_I^{-(R_I^i \cup \{i\})}$ since $R_I^i = R_{I'}^i$ and i 's priority increases from I to I' . Hence, $ms(B_I^{-(R_I^i \cup \{i\})}) = ms(B_{I'}^{-(R_{I'}^i \cup \{i\})}) = ms(B_{I'})$. Therefore $i \in R_I$.

- (v) Polynomial time computability. The rule makes at most n calls to computing a maximum size matching of the underlying reservation graphs. Thus, it runs in strongly polynomial time. □

Remark 4.3 (Relation to school choice). Theorem 4.2 can be rephrased in the context of school choice (see e.g., Abdulkadiroğlu and Sönmez (2003)). Agents correspond to students and categories correspond to schools. Suppose each student finds a subset of schools acceptable and is indifferent between all acceptable schools (which is a restrictive assumption). The schools have priorities over the students. Then, Theorem 4.2 reads as follows.

Theorem. *Consider the school choice problem where the students partition the schools into acceptable and unacceptable schools. Then, there is an allocation rule that only matches students to acceptable schools, admits no justified envy, matches the maximal feasible number of students, and does not penalize students for rising in the schools' priorities or incentivize students to hide that they find a school acceptable.*

Note that among the matchings that comply with the eligibility requirements, respect priorities, and have maximum size among feasible matchings, REV_π returns one that matches the set of agents that is maximal according to the upward lexicographic ordering on subsets of agents induced by \succ_π . Hence, the outcome of REV_π depends heavily on the order \succ_π . We show that in fact, every matching satisfying these three properties is an outcome of REV_π for some order \succ_π . Thus, these properties characterize all possible outcomes of Reverse Rejecting rules.

Theorem 4.4 (Characterization of outcomes of Reverse Rejecting rules). *A matching complies with the eligibility requirements, respects priorities, and has maximum size among feasible matchings if and only if it is a possible outcome of REV_π for some order \succ_π .*

Proof. Consider a matching μ with the three properties. Suppose it matches the set of agents $S \subseteq N$. Our first claim is that μ is feasible matching of $B_I^{-(N \setminus S)}$. Since μ satisfies the eligibility requirements, it is a matching of the graph B_I constrained to the vertex set

$S \cup C$. Since μ respects priorities, there exists no edge $\{i, c\} \in \mu$ such that $j \succ_c i$ for some $j \in N \setminus S$. Therefore, it follows that μ is a matching of $B_I^{-(N \setminus S)}$.

Now consider an order \succ_π such that $i \succ_\pi j$ for all $i \in S$ and $j \in N \setminus S$. Denote by R_I the set of unmatched agents in $REV_\pi(I)$. We claim that $R_I = N \setminus S$. Since $|\mu| = ms(B_I)$ by assumption and μ is a matching of $B_I^{-(N \setminus S)}$ as shown above, it follows that $N \setminus S \subset R_I$. The inclusion cannot be strict since then $ms(B_I^{-R_I}) < |\mu| = ms(B_I)$. Thus, $R_I = N \setminus S$. Since μ is a maximum size matching of $B_I^{-(N \setminus S)}$, $REV_\pi(I) = \mu$ is possible.

The converse follows from Theorem 4.2. □

Remark 4.5 (Characterization of outcomes of the Deferred Acceptance algorithm). Pathak et al. (2020, 2023) obtain a result analogous to Theorem 4.4 for the Deferred Acceptance algorithm. More precisely, they show that a matching complies with the eligibility requirements, respects priorities, and is non-wasteful if and only if there is a profile of strict preferences for the agents over the categories so that the matching is the outcome of the Deferred Acceptance algorithm applied to the resulting two-sided matching instance. Note that not every matching with these properties has maximum size.⁹

5. Treating Reserved and Unreserved Units Asymmetrically

Various reserves problems share two features. First, there is a designated category $c_u \in C$ interpreted as the *unreserved category* to which every agent has access, and the remaining categories C_p are *preferential categories* to which only some agents have access. For example, only some minority groups may qualify for the preferential categories in affirmative action programs. Second, there is a baseline ordering \succ_π over agents. For job selection under diversity concerns, the baseline ordering can be interpreted as an order of merit, or for healthcare rationing and school choice, the baseline ordering could be the need for treatment or the score in a central entrance exam, respectively. We assume that the priority ranking for the unreserved category equals the baseline ordering and refer to the units reserved for the unreserved category as *unreserved units*.

In this section, we explain why the asymmetry between the unreserved category and the preferential categories is helpful to capture common policy goals, which we formalize by a new axiom called *order preservation*. We then adapt Reverse Rejecting rules to incorporate this asymmetry, resulting in Smart Reverse Rejecting rules, which are shown to satisfy order preservation in addition to the properties they inherit from Reverse Rejecting rules.

As in the previous section, the emphasis lies on finding matchings that maximize the benefit to the agents. This is captured by maximum beneficiary assignments, which maximize the number of agents matched to preferential categories.¹⁰

Definition 5.1 (Maximum beneficiary assignment). A matching μ is a *maximum beneficiary assignment* if it maximizes the number of agents matched to a preferential category.

9. For the instance in Example 3.5, assume that all agents prefer c_1 to c_2 . Then applying the Deferred Acceptance algorithm results in the matching $\mu_2 = \{\{2, c_1\}\}$, which does not have maximum size.

10. In our context, the combination of maximum beneficiary assignment and non-wastefulness implies maximum size.

We observe that applying the *REV* rule to the preferential categories C_p and then allocating the unreserved units among the unmatched agents, say, according to the baseline ordering, yields a maximum beneficiary assignment.

5.1 Order Preservation

Certain policy goals may require allocating a designated number of unreserved units before allocating the units reserved for preferential categories. For example, the rationale for the “over-and-above” reserves approach is that agents first get a bite at the designated unreserved units before they utilize the preferential category units for which they are eligible. By contrast, the “minimum-guarantees” reserves approach first assigns agents to preferential categories and then matches the remaining agents to the unreserved units. We first define minimum-guarantees and over-and-above reservation rules (Chapter 13, Part B, Galanter (1984)) when the agents are eligible for at most one preferential category and all categories have priorities that are consistent with the baseline ordering in the sense that the agents that are eligible for a category are ranked according to the baseline ordering.¹¹

Minimum-guarantee Consider the agents in the order of the baseline ordering. For each agent, match her to a preferential category if (i) the agent is eligible for a preferential category and (ii) not all units reserved for this category have been allocated. Otherwise, match the agent to the unreserved category if an unreserved unit remains.¹²

Over-and-above Consider the agents in the order of the baseline ordering. For each agent i , match her to the unreserved category only if (i) an unreserved unit remains, and (ii) agent i is eligible for some preferential category, say, c , and there are still at least q_c agents from N_c other than i who are unmatched. Then, fill up the preferential categories as follows: for each $c \in C_p$, the $\min\{q_c, |N_c|\}$ highest priority unmatched agents are given one unit each from c .¹³

We present an example adapted from the book of Galanter (1984) to illustrate the difference between the minimum-guarantees and the over-and-above approach.¹⁴

Example 5.2. Consider the case where $N = \{1, 2, 3, 4\}$, $C = \{c, c_u\}$, $q_c = q_{c_u} = 1$, $N_c = \{1, 4\}$, and $4 \succ_\pi 3 \succ_\pi 2 \succ_\pi 1$. The outcome of the minimum-guarantees rule is that agent 3 and 4 are selected and the matching is $\{\{3, c_u\}, \{4, c\}\}$. The outcome of the over-and-above rule is that agents 1 and 4 are selected and the matching is $\{\{1, c\}, \{4, c_u\}\}$. In the

11. Galanter (1961, 1984) was one of the first to study the differences between the minimum-guarantees and over-and-above reservation rules in depth.

12. There is another version of the minimum-guarantees rule called the *Partha method* that gives an equivalent outcome but operates differently as an algorithm. In the Partha method, the units are allocated according to the baseline ordering and preferential reservation is only enforced if the reserves are not maximally used (Galanter, 1984).

13. Note that in the first stage, agents may be matched to the unreserved category but not to a preferential category. Hence, in the second stage, for any preferential category c , there are q_c units left.

14. Galanter (1984) studied minimum-guarantees and over-and-above in the context of job allocation in India where the baseline ordering represents the merit ranking and the preferential categories are historically disadvantaged groups. He observes that the minimum-guarantees rule “insures that the amount of effective reservation is somehow commensurate with the backwardness that inspired it.” On the other hand, he observes that the minimum-guarantees rule may “overstate the effective amount of reservation” especially if the disadvantaged groups are doing well enough on merit (page 461).

example, the minimum-guarantees rule coincides with the rule that solely uses the baseline ordering. On the other hand, the over-and-above rule provides additional representation of agents from the preferential category c .

We note that the minimum-guarantees approach allocates the unreserved units at the end whereas the over-and-above approach allocates the unreserved units first. We now explicitly distinguish between unreserved units that are processed earlier and later. To be precise, let $C = C_p \cup \{c_u^1, c_u^2\}$, where c_u^1 represents the unreserved units to be treated first and c_u^2 the unreserved units to be treated at the end. We view c_u^1 and c_u^2 as subcategories of the unreserved category c_u and the q_{c_u} slots for c_u as being partitioned into $q_{c_u^1}$ slots for c_u^1 and $q_{c_u^2}$ slots for c_u^2 . Hence, $q_{c_u^1} + q_{c_u^2} = q_{c_u}$ and $\succ_{c_u^1} = \succ_{c_u^2} = \succ_\pi$. If an agent is matched to c_u^1 or c_u^2 , we say that she receives an unreserved unit.

Pathak et al. (2020, 2023) formulated a family of rules, called *Smart Reserves rules*, that allow agents to be eligible for multiple preferential categories and generalize the minimum-guarantee and over-and-above rules to this case. In this section, we capture these approaches via an axiom for matchings called *order preservation* and then propose a new rule that also works for heterogeneous priorities. Order preservation is parametrized by the number of unreserved units that are placed in categories c_u^1 and c_u^2 . It captures the idea that there is an order of the categories (c_u^1 , C_p , and then c_u^2), and no two agents should be able to swap their matches so that eligibility requirements are not violated, and an earlier category gets a higher priority agent after the swap.

Definition 5.3 (Order preservation). Consider a matching μ of agents to categories in $C_p \cup \{c_u^1, c_u^2\}$. We say that μ is *order preserving* (with respect to c_u^1 and c_u^2) for the baseline ordering \succ_π if for any two agents $i, j \in N$,

- (i) $\mu(i) \in C_p \cup \{c_u^2\}$, $\mu(j) = c_u^1$, and j is eligible for category $\mu(i)$ implies $j \succ_\pi i$, and
- (ii) $\mu(j) \in C_p \cup \{c_u^1\}$, $\mu(i) = c_u^2$, and i is eligible for category $\mu(j)$ implies $j \succ_{\mu(j)} i$.¹⁵

There are two extreme ways unreserved units can be treated under order preservation. The first is if $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$. The other is if $q_{c_u^1} = q_{c_u}$ and $q_{c_u^2} = 0$. The conceptual contribution of Definition 5.3 is that instead of describing over-and-above and minimum-guarantees rules as consequences of certain sequential allocation methods, order preservation captures a key property of their resulting *matchings*. It is formulated so that it allows for heterogeneous priorities or agents being eligible for multiple categories.

We collect some of the properties of the minimum-guarantees rule and the over-and-above rule when each agent is eligible for at most one preferential category and categories have consistent priorities. These properties characterize the minimum-guarantees rule for the corresponding notion of order-preservation.

Proposition 5.4 (Properties of minimum-guarantees and over-and-above). *Assume each agent is eligible for at most one preferential category and all categories have consistent priorities. Then, the outcome of the minimum-guarantees (over-and-above) rule*

¹⁵ The qualification that i is eligible for category $\mu(j)$ could be omitted without altering the content of the definition since if i is not eligible for $\mu(j)$, $j \succ_{\mu(j)} i$ holds trivially. We keep it to maintain the symmetry with (i).

- (i) complies with the eligibility requirements,
- (ii) is a maximum beneficiary assignment,
- (iii) respects priorities,
- (iv) is non-wasteful, and
- (v) satisfies order preservation for $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$ ($q_{c_u^1} = q_{c_u}$ and $q_{c_u^2} = 0$).

Moreover, the outcome of the minimum-guarantees rule is the unique matching satisfying (i) to (v) with $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$.

Proof. First consider the minimum-guarantees rule. It complies with the eligibility requirements and is non-wasteful. It also yields a maximum beneficiary assignment because for each preferential category, the maximum possible number of agents is matched. The unreserved units are matched later to the unmatched agents. Therefore, the matching respects priorities and satisfies order preservation for $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$.

We prove that there is exactly one matching satisfying (i) to (v) with $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$. Hence, the outcome of the minimum-guarantees rule is the unique such matching. Suppose for contradiction there are two such matchings μ' and μ'' . Since each agent is eligible for at most one category in C_p and μ' and μ'' satisfy maximum beneficiary assignment, it follows that $|\mu'(c)| = |\mu''(c)|$ for all $c \in C_p$. We prove that for either of μ' and μ'' , the agents matched to $c \in C_p$ are the $\min\{q_c, |N_c|\}$ highest priority eligible agents for c . Suppose this is not the case for μ' . Then, there exist $i, j \in N_c$ such that $\mu'(j) = c$, $\mu'(i) \neq c$, and $i \succ_c j$. If $\mu'(i) = \emptyset$, then μ' does not respect priorities. If $\mu'(i) = c_u^2$, then μ' does not satisfy order preservation as i and j can swap their matches without violating eligibility requirements. We have established that for each $c \in C_p$, $\mu'(c) = \mu''(c)$. Respect of priorities, non-wastefulness, and the fact that every agent is eligible for c_u^2 imply that the agents matched to c_u^2 for either of μ' and μ'' are the highest priority agents who are not matched to categories in C_p . Hence, $\mu'(c_u^2) = \mu''(c_u^2)$.

Now consider the over-and-above rule. It complies with the eligibility requirements and is non-wasteful. It also yields a maximum beneficiary assignment because for each preferential category, the maximum possible number of agents are matched. The unreserved units are matched to the highest priority agents possible subject to enabling a maximum beneficiary assignment. Therefore, the matching respects priorities and satisfies order preservation for $q_{c_u^1} = q_{c_u}$ and $q_{c_u^2} = 0$. \square

Remark 5.5 (Non-uniqueness of over-and-above). The outcome of the over-and-above rule is not necessarily the unique matching satisfying (i) to (v) with $q_{c_u^1} = q_{c_u}$ and $q_{c_u^2} = 0$. Let $N = \{1, 2, 3, 4\}$ and $C = \{c_u^1, c_1, c_2\}$ with each category having capacity 1. Suppose the priorities are $4 \succ_{c_u^1} 3 \succ_{c_u^1} 2 \succ_{c_u^1} 1 \succ_{c_u^1} \emptyset$, $4 \succ_{c_1} 2 \succ_{c_1} \emptyset$, and $3 \succ_{c_2} 1 \succ_{c_2} \emptyset$. Then, the outcome of the over and above rule is $\{\{4, c_u^1\}, \{2, c_1\}, \{3, c_2\}\}$. Another matching that satisfies the properties is $\{\{3, c_u^1\}, \{4, c_1\}, \{1, c_2\}\}$.

	$S\text{-}REV_\pi$	Smart Reserves	DA
compliance with eligibility requirements	✓	✓	✓
maximum beneficiary assignment	✓	✓	–
respect of priorities	✓	✓*	✓
respects improvements	✓	n/a	✓
order preservation	✓	✓*	–
polynomial-time computability	✓	✓	✓

Table 1: Properties satisfied by $S\text{-}REV_\pi$, the Smart Reserves rule of Pathak et al. (2020, 2023), and the Deferred Acceptance algorithm described in Remark 4.5. An asterisk indicates that the property holds if priorities are strict and consistent with the baseline ordering. N/a indicates that the rule assumes homogenous priorities but the property allows for changes in the priorities that may result in inhomogeneous priorities.

5.2 The Smart Reverse Rejecting Rule

The Smart Reserves rule of Pathak et al. (2020, 2023) gives agents the unreserved units from c_u^1 as long as the set of remaining agents can be matched to get a maximum beneficiary assignment. Whereas Reverse Rejecting rules are not equipped to handle unreserved categories, the Smart Reserves rule is not designed to handle heterogeneous priorities. The ideas of both approaches can be combined to obtain the *Smart Reverse Rejecting* ($S\text{-}REV_\pi$) rule. $S\text{-}REV_\pi$ first determines which agents get an unreserved unit from c_u^1 , then allocates the units reserved for preferential categories to the remaining agents using the REV_π rule, and gives the c_u^2 units to the remaining agents according to the baseline ordering. More precisely, $S\text{-}REV_\pi$ works as follows.

Let the set of agents to be given unreserved units from c_u^1 , N_1 , be empty at the start and consider the agents in order of the baseline ordering in descending order. When agent i is considered, add i to N_1 if N_1 contains fewer than $q_{c_u^1}$ agents and the agents in $N \setminus (N_1 \cup \{i\})$ can form a maximum beneficiary assignment. After the last agent has been considered, give each agent in N_1 an unreserved unit from c_u^1 . Use REV_π to allocate the units reserved for the preferential categories C_p to the remaining agents. Lastly, give the unreserved units from c_u^2 to the remaining agents in descending order of the baseline ordering.

We show that $S\text{-}REV_\pi$ preserves the properties of REV_π while allowing to treat reserved and unreserved units asymmetrically. Note that part of $S\text{-}REV_\pi$ is applying the Reverse Rejecting rule for the baseline order \succ_π rather than some arbitrary order over agents. This is required for $S\text{-}REV_\pi$ to respect improvements. Table 1 summarizes the properties of $S\text{-}REV_\pi$, the Smart Reserves rule, and the approach based on the Deferred Acceptance algorithm described in Remark 4.5.

Theorem 5.6 (Properties of the Smart Reverse Rejecting rule). *The $S\text{-}REV_\pi$ rule*

- (i) *complies with eligibility requirements,*

- (ii) yields a maximum beneficiary assignment,
- (iii) respects priorities,
- (iv) respects improvements,
- (v) satisfies order preservation, and
- (vi) is polynomial-time computable.

Proof. (i) and (ii) are clear by construction. Note that a maximum beneficiary assignment is a maximum size matching of agent to all categories in $C = C_p \cup \{c_u^1, c_u^2\}$ since every agent is eligible for c_u^1 and c_u^2 .

(iii) We first prove that no unmatched agent can have justified envy for an agent matched to c_u^1 . Suppose an unmatched agent i has higher baseline priority than an agent j matched to c_u^1 . Then, i would have been considered before j when determining which agents get units from c_u^1 . The chosen matching shows that neither i nor the agents in N_1 (at the time i was considered) are needed for a maximum beneficiary assignment. Hence, i would have been added to N_1 and received a unit from c_u^1 , which is a contradiction.

Second, no unmatched agent can have justified envy towards an agent matched to c_u^2 because each unmatched agent comes later in the baseline ordering than each agent matched to c_u^2 . Finally, no unmatched agent can have justified envy towards an agent matched to a category in C_p as this would contradict the fact that REV_π respects priorities.

(iv) Suppose that i 's priority decreases from I to I' and i is unmatched under $S-REV_\pi(I)$. We show that i is unmatched under $S-REV_\pi(I')$. We first show that agent i is not matched to c_u^1 under $S-REV_\pi(I')$. Each time an agent j is added to N_1 , it is because the agents in $N \setminus (N_1 \cup \{j\})$ can be matched to C_p to obtain a maximum beneficiary assignment. Since i is not matched under $S-REV_\pi(I)$, i is not needed to obtain a maximum beneficiary assignment for I , which implies that i is not needed to obtain a maximum beneficiary assignment for I' . Hence, i cannot affect the selection of agents preceding her in the baseline ordering that are added to N_1 and hence matched to c_u^1 . Since i is not matched to c_u^1 under $S-REV_\pi(I)$, it means that when i was considered to be added to N_1 , the $q_{c_u^1}$ units of c_u^1 had already been used up. Therefore, i cannot get matched to c_u^1 under $S-REV_\pi(I')$.

We have shown that i cannot affect the set N_1 , that is, which agents are matched to c_u^1 . Since i is not matched to a category in C_p under $S-REV_\pi(I)$ and REV_π respects improvements, i is not matched to a category in C_p under $S-REV_\pi(I')$.

To show that i is not matched to c_u^2 under $S-REV_\pi(I')$, we observe the following.

- (i) Since agent i is not matched to a category in C_p under $S-REV_\pi(I)$ and REV_π yields a maximum size matching, agent i cannot affect the *number* of agents matched to categories in C_p .
- (ii) As observed in the proof that REV_π respects improvements, the set of agents *lower* in the baseline ordering (who compete with i to be matched to c_u^2) is the same under I and I' .¹⁶

16. This is the only point in the proof where we use that the Reverse Rejecting rule parameterized by the baseline ordering is used.

The above two facts imply that the number of agents *higher* up baseline ordering than i who are not matched to a category in $C_p \cup \{c_u^1\}$ and hence compete to be matched to c_u^2 is the same under I and I' . Therefore, i is not matched to c_u^2 under $S\text{-}REV_\pi(I')$.

(v) Consider a matching μ returned by $S\text{-}REV_\pi$. Suppose it does not satisfy order preservation. Then there exist two agents $i, j \in N$ such that one of the following holds:

- (i) $\mu(i) = c_u^1$, $\mu(j) \neq c_u^1$, $j \succ_{\mu(i)} i$, and i is eligible for category $\mu(j)$.
- (ii) $\mu(j) = c_u^2$, $\mu(i) \in C_p \cup \{c_u^1\}$, $j \succ_{\mu(i)} i$, and j is eligible for category $\mu(i)$.

We first consider the violation of the first type: $\mu(i) = c_u^1$, $\mu(j) \neq c_u^1$, $j \succ_\pi i$, and i is eligible for category $\mu(j)$. We examine the step at which i is considered when determining which agents get units from c_u^1 . Since $\mu(i) = c_u^1$, agent i is added to N_1 . Thus, a maximum beneficiary assignment of the agents in $N \setminus (N_1 \cup \{i\})$ exists. One such matching is μ . The matching μ' obtained from μ by swapping the matches of i and j is a maximum beneficiary assignment (since i is eligible for $\mu(j)$) for the agents in $N \setminus (N_1 \cup \{j\})$. Since $j \succ_\pi i$ and N_1 weakly increases in every step, N_1 cannot have been larger when the algorithm considered agent j . Hence, at this earlier step, μ' was also a maximum beneficiary assignment for agents in $N \setminus (N_1 \cup \{j\})$. But the existence of such a matching is the condition for adding j to N_1 , which contradicts that $\mu(j) \neq c_u^1$.

Next we consider a violation of the second type: $\mu(j) = c_u^2$, $\mu(i) \in C_p \cup \{c_u^1\}$, $j \succ_\pi i$, and j is eligible for category $\mu(i)$. Since a violation of the first type cannot happen, we may assume that $\mu(i) \in C_p$. But since j is not matched to a category in C_p , this implies that the REV rule does not respect priorities, a contradiction.

(vi) When agents are added iteratively to N_1 , the algorithm requires checking if there exists a maximum beneficiary assignment of agents in $N \setminus (N_1 \cup \{i\})$. This can be checked in polynomial time by algorithms for computing a maximum size b-matching. Once N_1 is fixed, we call the REV rule, which we have already shown to be polynomial-time computable. After that the remaining units can be allocated in linear-time by going down the baseline ordering. □

Remark 5.7 (Soft reserves). We assumed that only matchings that satisfy the eligibility requirements are feasible. The disadvantage of this approach is that some preferential category units may not be utilized even though some agents are unmatched. If we do not impose eligibility requirements as hard constraints, the setting is referred to as the case of “soft reserves”. In that case, we can first compute a matching that complies with the eligibility requirements using $S\text{-}REV_\pi$. If the resulting matching leaves some units from C_p unassigned, we allocate those to unmatched agents in order of the baseline ordering \succ_π . Assuming the preferential categories’ priorities over *ineligible* agents are consistent with the baseline ordering, the resulting rule satisfies all the properties in Theorem 5.6 except for “hard” compliance with the eligibility requirements. It also respects improvements by an argument similar to the proof of Theorem 5.6(iv). (Decreasing the priority of an agent does not affect the number of unmatched agents higher up in the baseline ordering.)

The natural analog of Theorem 4.4 to the setting with unreserved units—a matching complies with the eligibility requirements, respects priorities, is a maximum beneficiary

assignment, and satisfies order preservation if and only if it is a possible outcome of $S\text{-}REV_\pi$ for some baseline ordering \succ_π —fails. This is because order preservation is defined relative to a baseline ordering, leaving no freedom in choosing an ordering to be used for $S\text{-}REV_\pi$.

Acknowledgements

This material is based on work supported by the Deutsche Forschungsgemeinschaft under the Excellence Strategy EXC-2047 and the Grant BR 5969/1-1. An extended abstract of this paper appeared in the proceedings of the 22nd ACM Conference on Economics and Computation (2021). The authors thank Adi Ganguly, Tayfun Sönmez, and the participants of the following events for helpful comments: Online Social Choice and Welfare Seminar Series, Hong Kong Summer School on Game Theory and Social Choice, Bonn BGSE Workshop, and COMSOC Video Seminar. The authors are furthermore indebted to the editor and anonymous referees for their constructive feedback.

References

- Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93(3), 729–747.
- Abdulkadiroğlu, A., & Sönmez, T. (2003). Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory*, 112(1), 157–172.
- Abraham, D., Blum, A., & Sandholm, T. (2007). Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce (ACM-EC)*, pp. 295–304. ACM Press.
- Ahmadi, S., Ahmed, F., Dickerson, J. P., Fuge, M., & Khuller, S. (2020). An algorithm for multi-attribute diverse matching. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 3–9.
- Ahmed, F., Dickerson, J. P., & Fuge, M. (2017). Diverse weighted bipartite b-matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 35–41. AAAI Press.
- Andersson, T., & Ehlers, L. (2020). Assigning refugees to landlords in sweden: Efficient, stable, and maximum matchings. *The Scandinavian Journal of Economics*, 122(3), 937–965.
- Aygün, O., & Bó, I. (2020). College admission with multidimensional privileges: The Brazilian affirmative action case. *American Economic Journal: Microeconomics*.
- Aygün, O., & Turhan, B. (2020). Dynamic reserves in matching markets: Theory and applications. *Journal of Economic Theory*, 188.
- Aziz, H. (2018). Mechanisms for house allocation with existing tenants under dichotomous preferences.. *Journal of Mechanism and Institution Design*.
- Aziz, H. (2020). Developments in multi-agent fair allocation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 13563–13568.
- Aziz, H., Biró, P., & Yokoo, M. (2022). Matching market design with constraints. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 12308–12316.

- Aziz, H., Gaspers, S., & Sun, Z. (2020). Mechanism design for school choice with soft diversity constraints. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 153–159.
- Balinski, M., & Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, *84*(1), 73–94.
- Biró, P., Fleiner, T., Irving, R. W., & Manlove, D. F. (2010a). The college admissions problem with lower and common quotas. *Theoretical Computer Science*, *411*(34), 3136 – 3153.
- Biró, P., Manlove, D. F., & Mittal, S. (2010b). Size versus stability in the marriage problem. *Theoretical Computer Science*, *411*(16-18), 1828–1841.
- Bognar, G., & Hirose, I. (2014). *The Ethics of Health Care Rationing: An Introduction*. Routledge.
- Bogomolnaia, A., & Moulin, H. (2015). Size versus fairness in the assignment problem. *Games and Economic Behavior*, *90*, 119–127.
- Bruce, L., & Tallman, R. (2021). Promoting racial equity in covid-19 resource allocation. *Journal of Medical Ethics*.
- Das, S. (2022). Local justice and the algorithmic allocation of scarce societal resources. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 12250–12255. AAAI Press.
- Dawson, A., Isaacs, D., Jansen, M., Jordens, C., Kerridge, I., Kihlbom, U., Kilham, H., Preisz, A., Sheahan, L., & Skowronski, G. (2020). An ethics framework for making resource allocation decisions within clinical care: Responding to covid-19. *Journal of Bioethical Inquiry*, *17*(4), 749–755.
- Dickerson, J. P., Sankararaman, K. A., Srinivasan, A., & Xu, P. (2019). Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, pp. 1877–1884.
- Dur, U., Kominers, S. D., Pathak, P. A., & Sönmez, T. (2018). Reserve Design: Unintended Consequences and the Demise of Boston’s Walk Zones. *Journal of Political Economy*, *126*(6), 2457–2479.
- Dur, U., Pathak, P. A., & Sönmez, T. (2020). Explicit vs. statistical targeting in affirmative action: Theory and evidence from Chicago’s exam schools. *Journal of Economic Theory*, *187*, 104996.
- Echenique, F., & Yenmez, M. B. (2015). How to control controlled school choice. *American Economic Review*, *105*(8), 2679–94.
- Ehlers, L., Hafalir, I. E., Yenmez, M. B., & Yildirim, M. A. (2014). School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, *153*, 648—683.
- Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, *382*(21), 2049–2055.

- Fink, S. (2020). The hardest questions doctors may face: Who will be saved? who won't?. *The New York Times*, March 21.
- Galanter, M. (1961). Equality and protective discrimination in India. *Rutgers Law Review*, 16(1), 42–74.
- Galanter, M. (1984). *Competing Equalities: Law and the Backward Classes in India*. Univ of California.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9–15.
- Gonczarowski, Y. A., Nisan, N., Kovalio, L., & Romm, A. (2019). Matching for the Israeli "Mechinot" gap year: Handling rich diversity requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pp. 321–321.
- Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., & Yokoo, M. (2016). Strategyproof matching with regional minimum and maximum quotas. *Artificial intelligence*, 235, 40–57.
- Grigoryan, A. (2020). Effective, fair and equitable pandemic rationing. Tech. rep. 3646539, SSRN.
- Hafalir, I. E., Yenmez, M. B., & Yildirim, M. (2013). Effective affirmative action in school choice. *Theoretical Economics*, 8(2), 325–363.
- Heo, E. J. (2019). Equity and diversity in college admissions. In Laslier, J., Moulin, H., Sanver, R., & Zwicker, W. (Eds.), *The Future of Economic Design*. Springer.
- Hopcroft, J. E., & Karp, R. M. (1973). An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.*, 2(4), 225–231.
- Kamada, Y., & Kojima, F. (2015). Efficient matching under distributional constraints: Theory and applications. *The American Economic Review*, 105(1), 67–99.
- Kamada, Y., & Kojima, F. (2017). Recent developments in matching with constraints. *The American Economic Review*, 107(5), 200–204.
- Karzanov, A. V. (1973). An exact estimate of an algorithm for finding a maximum flow, applied to the problem on representatives. *Problems in Cybernetics*, 5, 66–70.
- Kojima, F. (2019). New directions of study in matching with constraints. In Laslier, J.-F., Moulin, H., Sanver, R., & Zwicker, W. S. (Eds.), *The Future of Economic Design*. Springer-Verlag.
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2), 83–97.
- Kurata, R., Hamada, N., Iwasaki, A., & Yokoo, M. (2017). Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, 58, 153–184.
- Martin, E. (2015). Rationing in healthcare. Tech. rep. 8, Deeble Institute.
- NASEM-National Academies of Sciences, Engineering, and Medicine (2020). Framework for Equitable Allocation of COVID-19 Vaccine. Tech. rep., Washington, DC.

- Pathak, P. A., Sönmez, T., Ünver, M. U., & Yenmez, M. B. (2020). Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Health Care Rationing. Boston college working papers in economics 1015, Boston College Department of Economics.
- Pathak, P. A., Sönmez, T., Ünver, M. U., & Yenmez, M. B. (2023). Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Healthcare Rationing. *Management Science*.
- Persad, G., Peek, M. E., & Emanuel, E. J. (2020). Prioritizing groups for access to covid-19 vaccines. *JAMA*, *324*(16), 1601–1602.
- Roth, A. E., & Sotomayor, M. A. O. (1990). *Two-Sided Matching: A Study in Game Theoretic Modelling and Analysis*. Cambridge University Press.
- Sönmez, T., Pathak., P. A., Ünver, M. U., Persad, G., Truog, R. D., & White, D. B. (2020). Categorized priority systems: A new tool for fairly allocating scarce medical resources in the face of profound social inequities. *CHEST*.
- Sönmez, T., & Yenmez, M. B. (2020). Affirmative action with overlapping reserves. Manuscript.
- Sonmez, T., & Yenmez, M. B. (2022). Affirmative Action in India via Vertical, Horizontal, and Overlapping Reservations. *Econometrica*, 1143–1176.
- Sonmez, T., & Yenmez, M. B. (2024). Constitutional Implementation of Affirmative Action Policies in India. Tech. rep., arXiv.
- Truog, R. D., Mitchell, C., & Daley, G. Q. (2020). The toughest triage — allocating ventilators in a pandemic. *New England Journal of Medicine*, *382*(21), 1973–1975.
- WHO (2020). A global framework to ensure equitable and fair allocation of COVID-19 products and potential implications for COVID-19 vaccines.. Tech. rep..