Efficient, Fair, and Incentive-Compatible Healthcare Rationing

Haris Aziz[†] Florian Brandl[‡]

Fair and efficient rationing of healthcare resources has emerged as an important issue that has been discussed by medical experts, policy-makers, and the general public. We consider a healthcare rationing problem where medical units are to be allocated to patients. Each unit is reserved for one of several categories and the categories may have different priorities for the patients. We present a flexible allocation rule that respect the priorities, comply with the eligibility requirements, allocate the largest feasible number of units, and do not incentivize agents to hide that they qualify through a category. To the best of our knowledge, this is the first known rule with the aforementioned properties. Our rule also characterizes all possible outcomes that satisfy the first three properties. Moreover, it is polynomial-time computable.

Keywords: Allocation under priorities, healthcare rationing, assignment maximization.

1 Introduction

The COVID-19 pandemic has emerged as a major challenge that the world has faced. It has resulted in a frantic scientific race to produce the most effective and safe vaccine to stem the devastating effects of the pandemic. Whereas there is encouraging initial news on the creation of vaccines, there are still several scientific challenges on how to distribute, allocate, and administer them in an efficient and fair manner.

Since healthcare resources such as ventilators, antiviral treatments, and vaccines can be scarce or costly, a fundamental question that arises is who to prioritize when making allocation decisions. For example, three important priority groups that are highlighted by medical practitioners and policy-makers are (1) health care workers; (2) other essential workers and people in high-transmission settings; (3) and people with medical vulnerabilities associated with poorer

[†]UNSW Sydney and Data61 CSIRO, Australia, haris.aziz@unsw.edu.au

[‡]Princeton University, USA, brandl.ffx@gmail.com

COVID-19 outcomes (Persad et al., 2020; Truog et al., 2020). Other concerns that have been discussed include racial equity (Bruce and Tallman, 2021).

When healthcare resources need to be allocated, it is not enough to identify priority groups. There is also a need to algorithmically and transparently make these prioritized allocations decisions (Emanuel et al., 2020; WHO, 2020). In a New York Times article, the issue has been referred to as one of the hardest decisions health organizations need to make (Fink, 2020). Since the decisions need to be justified to the public, they must be aligned with the ethical guidelines, such as respecting priorities of various categories. These decisions are not straightforward, especially when a patient is eligible for more than one category. When patients are eligible for multiple categories, the decision on which category is used can have compounding effects on what categories other agents can use. A fundamental question that arises is the following one:

> *How do we allocate scarce medical resources fairly and efficiently while taking into account various ethical principles and priority groups?*

The question is not just fundamental but the solution to the problem is time–critical as various states, city councils, and municipalities start to roll out vaccines using their particular ethical guidelines. The problem of health care rationing has recently been formally studied by market designers. Pathak et al. (2020) were among the first to frame the problem as a two-sided matching problem in which patients are on one side and the resources units are on the other side. By doing so, they linked the healthcare rationing problem with the rich field of two-sided matching (Roth and Sotomayor, 1990).

Pathak et al. suggested dividing the units into different reserve categories, each with its own priority ranking of 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 have an age-specific priority ranking that puts the eldest citizens first. Having a holistic framework that considers different types of priorities has been termed important in healthcare rationing.[2] The approach of Pathak et al. 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 et al., 2020). The approach has been covered widely in the media, including the New York Times and the Washington Post.[3]

For their two-sided matching formulation, Pathak et al. (2020) proposed a solution for the problem. One of their key insights was that running the Deferred Acceptance algorithm (Gale

[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

and Shapley, 1962; Roth, 2008) on the underlying problem satisfies basic relevant axioms (eligibility compliance, respect of priorities, and non-wastefulness). They also showed when all the category priorities are consistent with a global baseline priority, then there is a smart reserves algorithm that computes a maximum size matching satisfying the basic axioms. The smart reserves approach makes careful decisions about which category should be availed by which patient to achieve the maximum size property. However, the problem of such a smart reserves approach for the general problem with general *heterogeneous* priorities has not been addressed in the literature. Allowing heterogeneous priorities for categories seems to be very much in the spirit of incorporating different ethical values. For example, one would expect the priority ordering for old people to be very different from a priority ordering for front-line workers which would favor energetic medical professionals.[4] In this paper, we set out to address this issue and answer the following research problem.

> *For the general healthcare rationing problem with heterogeneous priorities, how do we allocate resources in a fair, economically efficient, strategyproof, and computationally tractable way?*

Contribution For the general healthcare rationing problem, we first highlight that naively ascribing strict preferences over the categories to the agents can have adverse effects on the efficiency of the outcome when patients are eligible for multiple categories. If the eligibility requirements are treated as hard constraints, it leads to inefficient allocation of resources. If the eligibility requirements are treated as soft constraints, then the outcome does not allocate the resources optimally to the highest priority patients, thereby undermining important healthcare guidelines and ethical principles.

Our first contribution is to show that there exists a rule (Reverse Rejecting (REV) rule) that

(i) complies with the eligibility requirements

(ii) respects the priorities of the categories (for each category, patients of higher priority are served first)

(iii) yields a maximum size matching (one that allocates the largest feasible number of units to eligible patients)

(iv) is non-wasteful (there is no unit that is unused but could be used by some eligible patient)

(v) is strategyproof (does not incentivize agents to under-report the categories they qualify for)

(vi) is strongly polynomial-time computable.

[4]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.

We prove that REV characterizes all outcomes that satisfy the first three properties.[5] We show how REV can be extended to a more general class of rules called *Smart REV (S–REV)* in which a given number of unreserved units are to be processed earlier and later and in which an additional goal is to allocate the largest feasible number of units from a designated subset of categories called preferential categories. S–REV satisfies order preservation that is a new axiom that we propose and which is parametrized with respect to how many unreserved units are processes first and last. Our general class of rules 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. We also provide axiomatic characterizations of over-and-above and minimum-guarantees. Finally, we discuss how our algorithms and their properties extend to the case where the reservations are treated as soft reservations.

Our algorithm immediately applies to the school choice problem in which students are only interested in being matched to one of their acceptable schools. It also applies to hiring settings in which applicants are interested in one of the positions, and each of the departments has its own priorities. Finally, it applies to many other rationing scenarios, such as allocation of 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 et al., 2021). One general class of distributional constraints that have been examined in matching market design pertains to common quotas over institutions such as hospitals (Kamada and Kojima, 2015, 2017; Biró et al., 2010; Goto et al., 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 et al., 2013; Ehlers et al., 2014; Echenique and Yenmez, 2015; Kurata et al., 2017; Aygün and Turhan, 2020; Aygün and Bó, 2020; Aziz et al., 2020; Gonczarowski et al., 2019; Dur et al., 2018, 2020). Categories in healthcare rationing correspond to affirmative action types in school choice. For a brief survey, we suggest the book chapter by Heo (2019). Except for the special case in which students have exactly 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. Ahmed et al. (2017), Dickerson et al. (2019), and Ahmadi et al. (2020) consider optimisation approaches for diverse matchings, but their objective and models are different.

Pathak et al. (2020) were the first to frame a rationing problem with category priorities as a two-sided matching problem in which agents are simply interested in a unit of resource and the

[5] In a preliminary version of the paper, we proposed a different rule that satisfies the above properties. It requires solving as maximum weight matching of a corresponding graph. However, it does not characterize all outcomes that satisfy the first three properties.

resources are reserved for different categories. They show that artificially enforcing strict preferences of the agents 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. They then proposed to use the smart reserves approach of Sönmez and Yenmez (2020) for the restricted problem when all the preferential categories have consistent priorities. Our results can be viewed as simultaneously achieving the key axioms of the two approaches of Pathak et al. (2020). Firstly, we propose a new algorithm that achieves the same key axiomatic properties for heterogeneous priorities as the algorithms of Sönmez and Yenmez (2020) and Pathak et al. (2020) for homogeneous priorities. Secondly, our algorithm has an important advantage over the Deferred Acceptance formulation of Pathak et al. (2020) for the case of heterogeneous priorities as our approach additionally achieves the important property of maximality of size. It additionally satisfies a property called maximality in beneficiary assignment, which requires that the maximum number of units from the set of 'preferential' categories are used. Pathak et al. (2020) design a flexible feature of their smart reserves rule that gives agents a designated number of unreserved units before the other units are processed. By doing so, they elegantly capture two extreme approaches within their class that have wide-spread appeal. The first approach is based on *minimum-guarantees* that specifies the minimum number of units that are 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] Our $S{-}REV$ rule achieves these features even for the case of heterogeneous priorities. In follow-up work, Grigoryan (2020) considers optimisation approaches for variants of the problem but does not present any polynomial-time algorithm or consider incentive issues. In contrast to the papers on healthcare rationing discussed above, we also consider strategyproofness and monotonicity aspects and show that our rule complies with them.

In this paper, we attempt to compute what are essentially maximum size stable matchings. The problem of computing such matchings is NP-hard if both sides have strict preferences/priorities (Biró et al., 2010). In our problem, the agents essentially have dichotomous preferences (categories they are/are not eligible for) and, hence, we are able to obtain a polynomial-time algorithm for the problem.

Furthermore, our rules are strategyproof. In contrast, for other two-sided matching settings, it is known that maximizing the number of matched individuals results in incentive and fairness impossibilities (see, e.g., Afacan et al., 2020; Krysta et al., 2014). 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, 2016; Abraham et al., 2007; Bogomolnaia and Moulin, 2015).

[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)).

3 Model

We adopt the essential features of the healthcare rationing model of Pathak et al. (2020) with one generalization: we allow the categories' priorities over agents to be weak rather than strict. There are q identical and indivisible units of some resource, which are to be allocated to the agents in a set N 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 preorder 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 $I = (N, C, (\succsim_c), (q_c))$ is an instance (of the rationing problem). For convenience, we will write $(\succsim_c)$ and (q_c) for the profile of priorities and quotas in the sequel. We will also consider a baseline ordering $\succ_\pi$, which can be an arbitrary permutation of the agents. It could be interpreted as a global scale measuring the need for treatment.

A matching $\mu\colon N \to C \cup \{\emptyset\}$ 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 $\{\{i, \mu(i)\}\colon \mu(i) \neq \emptyset\}$.[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).

The first axiom we consider requires that matchings comply with 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 elderly people.

Definition 1 (Compliance with eligibility requirements)**.** A matching μ *complies with eligibility requirements* 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. 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 2 (Respect of priorities)**.** A matching μ *respects priorities* if for any $i, j \in N$ and $c \in C$, $\mu(i) = c$ and $\mu(j) = \emptyset \implies j \not\succ_c i$. 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 realise that the axiom "respect of priorities" is equivalent to *justified envy-freeness* in the context of school-choice matchings (Abdulkadiroğlu and 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.

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

Definition 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 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, page 69). 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 et al., 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, $E = \{\{i, c\} \colon i \succ_c \emptyset\}$. If G is any graph, we denote by $mw(G)$ the number of edges in a maximum size matching of G.

The following example illustrates the definitions above.

Example 1. 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.

$$\mu_1 = \emptyset \qquad \mu_2 = \{\{2, c_1\}\} \qquad \mu_3 = \{\{2, c_2\}\}$$
$$\mu_4 = \{\{3, c_1\}\} \qquad \mu_5 = \{\{2, c_2\}, \{3, c_1\}\}$$

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 5 (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 1 to 4 if $f(I)$ satisfies the axiom for all instances I. Moreover, we define a notion of strategyproofness for allocation rules. Note that all units are identical and agents have no preferences over the category of the

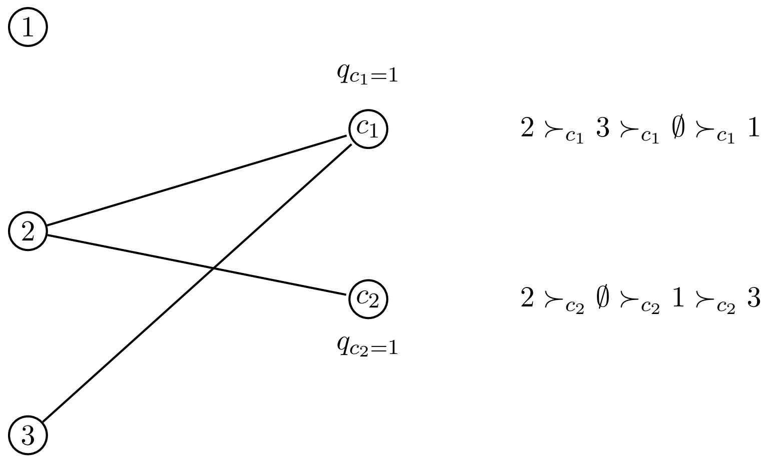

Figure 1: The instance I described in Example 1. The reservation graph B_I on the left has an edge from i to c if i is eligible for c. The priority rankings of the categories are depicted on the right.

unit they receive. However, they may have an incentive to hide being eligible for a category, or, more generally, to aim for a lower priority for some category.[8]

Formalizing strategyproofness requires the following definition. Let $(\succsim_c)$ and $(\succsim'_c)$ be priority profiles and $i \in N$. We say agent i's priority decreases from $(\succsim_c)$ to $(\succsim'_c)$ if for all $j, k \neq i$ and $c \in C$,

$$j \succsim_c k \longleftrightarrow j \succsim'_c k$$
$$j \succsim_c i \longrightarrow j \succsim'_c i \text{ and } j \succ_c i \longrightarrow j \succ'_c i$$

That is, the priority rankings over agents other than i are the same in both profiles and i can only move down in the priority rankings from $(\succsim_c)$ to $(\succsim'_c)$. We also say that i's priority decreases from $I = (N, C, (\succsim_c), (q_c))$ to $I' = (N, C, (\succsim'_c), (q_c))$. Strategyproofness requires that if i is unmatched for I, then i is also unmatched for I'.

Definition 6 (Strategyproofness)**.** An allocation rule f is strategyproof if $f(I)(i) = \emptyset$ implies $f(I')(i) = \emptyset$ whenever i's priority decreases from I to I'.

In particular, with a strategyproof allocation rule, agents cannot benefit from hiding that they are eligible for a category.[9]

An allocation rule is non-bossy if no unmatched agent can decrease her priority to change the set of matched agents. Combined with the other axioms, this property turns out to be fairly demanding. Thus, we weaken it by requiring only that no unmatched agent can decrease her priority and thereby change which of the agents lower in the baseline ordering are matched.

[8]In the context of school choice, lowering oneself in the priority ranking of a school is akin to students deliberately under-performing in an entrance exam.

[9]This restricted version of strategyproofness under which agents do not have an incentive to hide their eligible categories, has been referred to as incentive-compatibility by Aygün and Bó (2020).

Definition 7 (Weak non-bossiness)**.** An allocation rule f is weakly non-bossy if $f(I)(i) = \emptyset$ implies that any j with $i \succ_\pi j$ is matched in $f(I)$ if and only if j is matched in $f(I')$ whenever i's priority decreases from I to I'.

4 Issues with Breaking Ties and Applying Deferred Acceptance

The approach of Pathak et al. (2020) is to frame the healthcare rationing problem as a two-sided matching problem. They showed that if one artificially introduces (strict) preferences for the agents over the categories they are eligible for and applies the Deferred Acceptance algorithm, the resulting matching complies with the eligibility requirements, respects priorities, and is non-wasteful (Pathak et al., 2020, Theorem 2). They state the algorithmic implications of their result.

> *"Not only is this result a second characterization of matchings that satisfy our three basic axioms, it also provides a concrete procedure to calculate all such matchings."*

Although considering all possible artificial preferences and running Deferred Acceptance gives us all the matchings satisfying the three axioms, it is computationally expensive to consider $|C|!^{|N|}$ different preference profiles.

Not every preference profile leads to a compelling outcome even if the categories have strict priorities. For example, many preference profiles lead to matchings that are not maximum size. The next example highlights this issue.[10]

Example 2. Consider the instance in Example 1. Suppose we run the Deferred Acceptance algorithm assuming all agents prefer c_1 to c_2 to c_3. Assuming agents can only be matched to categories they are eligible for (compliance with eligibility requirements), the resulting matching is $\mu_2 = \{\{2, c_1\}\}$. This matching is however, not the most efficient use of the resources because it is possible to allocate all units while still satisfying the axioms in Definitions 1 to 3: $\mu_5 = \{\{3, c_1\}, \{2, c_2\}\}$.

Hence, artificially inducing preferences of agents and running Deferred Acceptance can lead to inefficient allocations. Even if we ignore computational concerns and can assign preferences to agents so that the matching selected by Deferred Acceptance is of maximum size and respects priorities, it is not clear whether such a rule satisfies strategyproofness and monotonicity properties like the ones we introduced above. We propose a rule that circumvents both issues.

5 The Reverse Rejecting Rule

The Reverse Rejecting Rule (REV) depends on the baseline ordering $\succ_\pi$. It iteratively goes over the agents in the reverse order of the baseline ordering. When considering agent i, it

[10] The issue is also evident from the discussion by Pathak et al. (2020) where they point out that sequential treatment of categories may not give a maximum size matching.

decides if i is rejected, that is, placed in the set of rejected agents R, or not. Agent i is rejected if and only if the agents $N \setminus (R \cup \{i\})$ can form a matching μ such that (i) μ has the same size as a maximum size matching of the reservation graph B_I, and (ii) no agent in $R \cup \{i\}$ has justified envy for any agent matched by μ. The second constraint is captured as follows. If an agent $i \in R$ is rejected, we do not allow any agent j to be matched to a category c if $i \succ_c j$. This can be ensured by removing the edge $\{j, c\}$ if $i \in R$ and $i \succ_c j$. In particular, we will use the graph $B_I^{-R} = ((N \setminus R) \cup C, E)$ where $E = \{\{j, c\} : j \succ_c \emptyset \text{ and } \nexists i \in R \text{ such that } i \succ_c j\}$. Note that B_I^{-R} is a subgraph of B_I. More generally, if $R_2 \supset R_1$, then $B_I^{-R_2}$ is a subgraph of $B_I^{-R_1}$. The rule is formalized as Algorithm 1. For an instance I, we will denote by R_I as the set of agents who are not matched by Algorithm 1.

The methodology of *REV* is different from the horizontal envelop rule of Sönmez and Yenmez (2020) and the smart reserves rule of Pathak et al. (2020). In contrast to iteratively or instantly selecting agents that will be matched, the *REV* rule goes in the reverse order of the baseline ordering and decides which agents to reject. More importantly, *REV* works for heterogeneous priorities.

Algorithm 1 Reverse Rejecting (*REV*) Rule

Input: $I = (N, C, (\succsim_c), (q_c))$; a baseline ordering $\succ_\pi$ over agents in N

Output: A matching of B_I.

```
Construct the reservation graph B_I = (N ∪ C, E) where {i, c} ∈ E if i ≻_c ∅.
Set of rejected agents R ⟵ ∅
for agent i in reverse order of ≻_π do
    Consider graph B_I^{-R} = ((N \ R) ∪ C, E) where
        E = {{i, c} : i ≻_c ∅ and ∄j ∈ R such that j ≻_c i}
    if mw(B_I^{-R∪{i}}) = mw(B_I) then
        Add i to R
    end if
end for
R_I ⟵ R
Compute a maximum size matching μ of B_I^{-R}.
return μ
```

Example 3 (Illustration of Algorithm 1)**.** 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 $1 \succ_{c_1} 4 \succ_{c_1} 2 \succ_{c_1} \emptyset$ and $1 \succ_{c_2} 3 \succ_{c_2} \emptyset$. For $1 \succ_\pi 2 \succ_\pi 3 \succ_\pi 4$, let us simulate *REV*. The reservation graph B_I is depicted in Figure 2a which 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 1. *The REV rule*

(i) terminates with a feasible matching,

(ii) complies with eligibility requirements,

(iii) returns a matching of maximum size among feasible matchings,

(iv) respects priorities,

(v) is weakly non-bossy,

(vi) is strategyproof, and

(vii) can be computed in strongly polynomial time.

Proof. (i) Termination with a feasible matching. As we iterate through the for loop and move in the reverse direction of the baseline ordering, we keep the invariant that $mw(B_I^{-R}) = mw(B_I)$. We stop when we cannot add any further agent into R such that $mw(B_I^{-R}) = mw(B_I)$. At this point, the algorithm returns a maximum size matching of B_I^{-R}.

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

(iii) Maximum size matching. The returned matching is by construction a matching of size $mw(B_I)$ and hence of maximum size among all matching satisfying compliance with eligibility.

(iv) Respect of priorities. Suppose the returned matching μ does not satisfy respect of priorities. First, we claim that for a feasible matching of B_I^{-R} provides no justified envy for any agent in R. For contradiction, suppose there exists some $j \in R$ that is not matched in the outcome matching μ, $i \in \mu(c)$ for some $c \in C$, and $j \succ_c i$. But in that case $\{i, c\}$ is not an edge in B_I^{-R}, a contradiction.

We have proved that no agent in R can have justified envy in a matching μ of B_I^{-R}. If $|R| = n - mw(B_I)$, then we are done. Suppose that $|R| < n - mw(B_I)$. Then, during the course of the algorithm, we arrive at a point such that $|N \setminus R| > mw(B_I)$ but we cannot add any agent i to R because there exists no matching of $B_I^{-(R\cup\{i\})}$ of size $mw(B_I)$ that respects priorities with respect to rejected agents $R \cup \{i\}$. In that case, consider an instance $I' = (N \setminus R, C, (\succsim_c)', (q_c))$ where $\succsim_c' = \succsim_c$ except that if $\emptyset \succ_c j$ if $\{c, j\}$ not in the edge set of B_I^{-R}. Note that the maximum size matching of $B_{I'}$ has size at least $mw(B_I)$

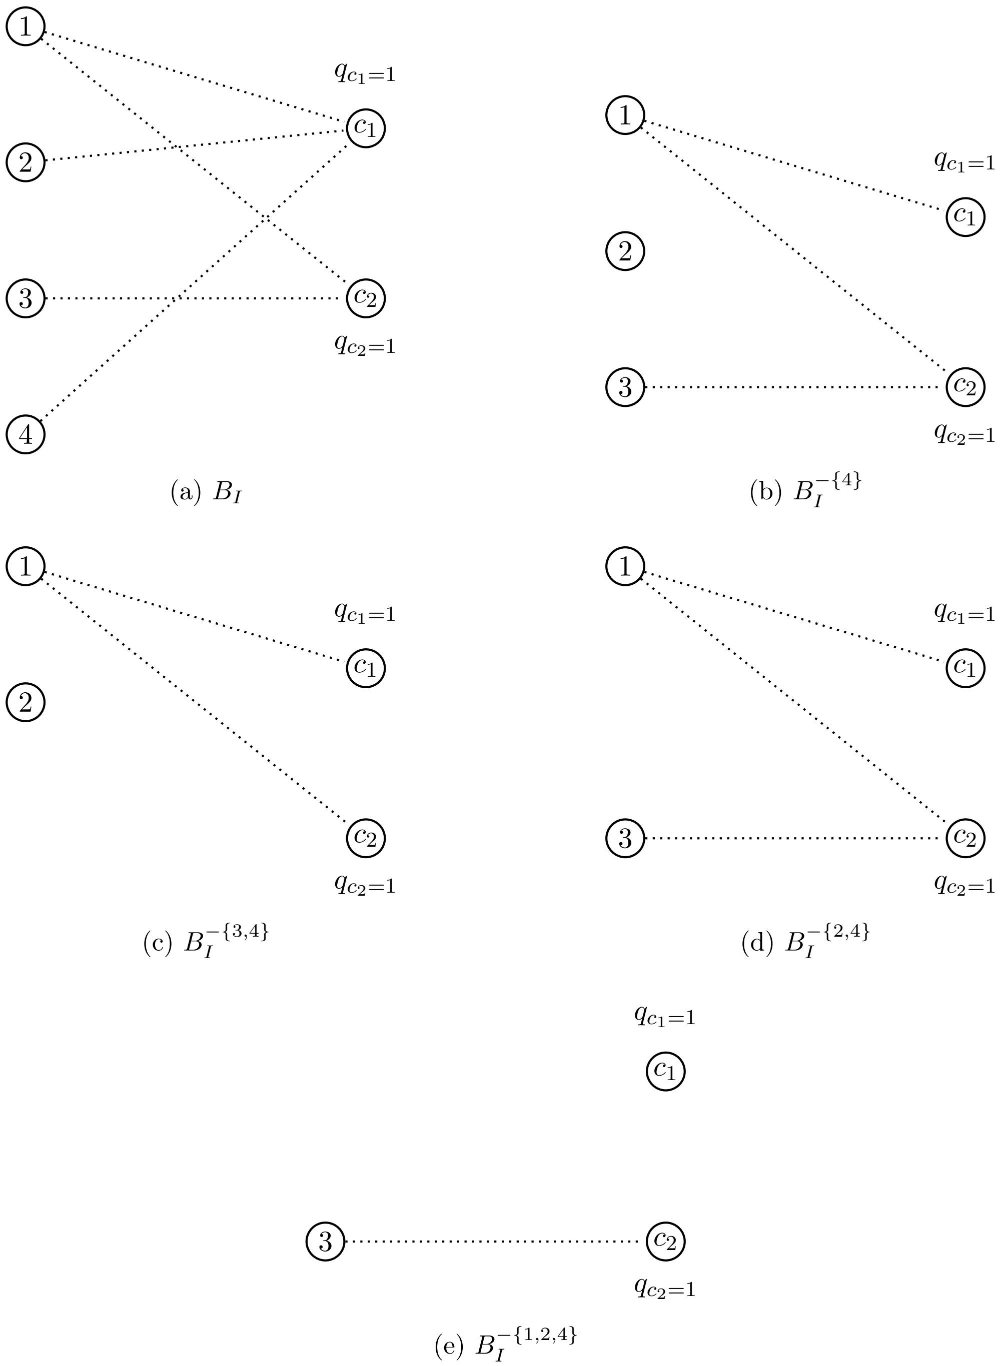

(a) B_I

(b) $B_I^{-\{4\}}$

(c) $B_I^{-\{3,4\}}$

(d) $B_I^{-\{2,4\}}$

(e) $B_I^{-\{1,2,4\}}$

Figure 2: Graphs for the instance I in Example 3.

as $B_{I'} = B_I^{-R}$. For instance I' we can run REV to compute a maximum size matching μ of $B_{I'}$ that also satisfies respect of priorities with respect to agents in $N \setminus R$. No agent in $j \in R$ can have justified envy for any agent k matched in μ because if $j \succ_{\mu(k)} k$, then $\emptyset \succ_{\mu(k)} k$ in instance I'. Hence μ is a matching that respects priorities under instance I, $|\mu| = mw(B_I)$, the set of agents matched by μ is a strict subset of $N \setminus R$. Hence, we have arrived at a contradiction.

(v) Weak non-bossiness. For simplicity, suppose the baseline ordering $\succ_\pi$ is $1 \succ_\pi 2 \succ_\pi \cdots \succ_\pi n$. Let $i \in N$, I, I' be instances so that i's priority decreases from I to I'. Assume that i is unmatched in $REV(I)$. Note that $mw(B_I) = mw(B_{I'})$ since the edge set of $B_{I'}$ is a subset of the edge set of B_I and $REV(I)$ does not match i. Let $R_I^i = \{j \in R_I \colon i \succ_\pi j\}$ be the agents who are unmatched in $REV(I)$ and are lower in the baseline ordering than i. Define $R_{I'}^i$ similarly. We need to 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 $R = R_I^i \cap \{j+1, \ldots, n\} = R_I \cap \{j+1, \ldots, n\}$. If $j \in R_I^i$, this is because there exists a matching μ of $B_I^{-(R \cup \{j\})}$ with $|\mu| = mw(B_I)$. Since $R_I \cap \{j, \ldots, n\} \subset R \cup \{j\}$, μ does not match any agent in $R_I \cap \{j, \ldots, n\}$. Since i is unmatched in $REV(I)$, 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 $|\mu| = mw(B_{I'}) = mw(B_I)$. 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 $B_{I'}^{-(R \cap \{j\})}$, 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 $R_I^i = R_{I'}^i$. We remark that the argument above breaks down for agents j with $j \succ_\pi i$ since if $i \in S$, not all edges of $B_{I'}^{-S}$ are also edges of B_I^{-S}.

(vi) Strategyproofness. Suppose agent i's priority decreases from I to I' and i is unmatched in $REV(I)$. Hence, $i \in R_I$. By the proof of weak non-bossiness above, $R_I^i = R_{I'}^i$. We prove that i is rejected under I'. Since $R_I^i = R_{I'}^i$, and since $i \in R_I$, it follows that $mw(B_I^{-(R_I^i \cup \{i\})}) = mw(B_I)$. Consider a maximum size matching of $B_I^{-(R_I^i \cup \{i\})}$. Such a matching does not match i. Note that for each $j \in N \setminus (R_I^i \cup \{i\})$ and $c \in C$, $\{j, c\}$ is in the edge set of $B_I^{-(R_I^i \cup \{i\})}$) if and only if $\{j, c\}$ is in the edge set of $B_{I'}^{-(R_{I'}^i \cup \{i\})}$. It follows that μ is also a feasible matching of $B_{I'}^{-(R_{I'}^i \cup \{i\})}$. Therefore $i \in R_{I'}$.

(vii) Polynomial time computability. The rule makes at most n calls to computing maximum cardinality matching of the underlying reservation graphs so runs in strongly polynomial time.

□

Remark 1. *REV* is *not* non-bossy. 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 $1 \succ_{c_1} 4 \succ_{c_1} 2 \succ_{c_1} \emptyset$ and $1 \succ_{c_2} 3 \succ_{c_2} \emptyset$. For $1 \succ_\pi 2 \succ_\pi 3 \succ_\pi 4$, *REV* yields the matching $\mu = \{\{1, c_1\}, \{3, c_2\}\}$. If agent 4 reports that they are not eligible for c_1, agent 2 moves to the second equivalence class in the priority order of c_1 and *REV* yields the matching $\mu' = \{\{1, c_2\}, \{2, c_1\}\}$. Since $\mu \neq \mu'$, *REV* violates non-bossiness.

Note that among the maximum size matchings respecting the eligibility requirements and the priorities, *REV* returns one that matches the set of agents that is maximal according to the upward lexicographic ordering on subsets of agents induced by the baseline ordering. Next, we present a characterization of possible outcomes of *REV*.

Theorem 2 (Characterization of *REV* outcomes)**.** *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 for some baseline ordering.*

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 a baseline ordering $\succ_\pi$ such that $i \succ_\pi j$ for all $i \in S$ and $j \notin S$. We prove that each agent $j \notin S$ is in R_I. Consider the graph $B_I^{-(R_I^j \cup \{j\})}$, where $R_I^j = \{i \in R_I : j \succ_\pi i\}$. Its vertex set contains S by definition of $\succ_\pi$. The graph $B_I^{-(R_I^j \cup \{j\})}$ is obtained by restricting B_I to $N \setminus (R_I^j \cup \{j\})$ and removing the edge $\{i, c\}$ if there is $k \in R_I^j \cup \{j\}$ with $k \succ_c i$. Since μ is a matching of $B_I^{-(N \setminus S)}$, it is also a matching of $B_I^{-(R_I^j \cup \{j\})}$. Hence, each agent $j \notin S$ is placed in R_I. Once all the agents in $N \setminus S$ are rejected, no further agents can be rejected. For otherwise $N \setminus R_I$ would be smaller than $mw(B_I) = |\mu|$. Since μ is a maximum size matching of $B_I^{-(N \setminus S)}$, it follows that μ is a possible outcome of *REV* under the baseline ordering $\succ_\pi$.

The converse follows from Theorem 1. □

6 Treating Unreserved Units Asymmetrically

We have thus far treated all categories symmetrically. Now we designate one category $c_u \in C$ as the *unreserved category*. All agents are eligible for the unreserved category and the priority ranking for the unreserved category equals the baseline ordering $\succ_\pi$. We refer to the units reserved for the unreserved category as *unreserved units* and call the remaining categories *preferential categories*. The set of preferential categories is denoted by C_p. There are two reasons for introducing the unreserved category. Firstly, we want to consider maximum beneficiary assignments, which maximize the number of agents matched to preferential categories.[11]

Definition 8 (Maximum beneficiary assignment)**.** A matching μ is a *maximum beneficiary assignment* if it maximizes the number of edges that share a node with C_p.

Secondly, in various reserves problems, unreserved units are treated in special ways such as being allocated later or earlier. We discuss both of these issues.

We observe that applying REV 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.

6.1 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" reserve 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" reserve 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 (Galanter, 1984, Chapter 13, Part B) 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.[12]

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.[13]

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

[12] Galanter (1961, 1984) was one of the first researchers to explore the differences between the minimum-guarantees and over-and-above reservation rules in depth.

[13] 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 ("merit") and preferential reservation is only enforced if the reserves are not maximally used (Galanter, 1984).

Over-and-Above Consider the agents in the order of the baseline ordering. For each agent i, match her to the unreserved category if (i) an unreserved unit remains and (ii) if agent i is eligible for some preferential category, say, c, and there are still at least $\min\{q_c, |N_c|\}$ agents from N_c 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.

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

Example 4. Consider the case where $N = \{1, 2, 3, 4\}$, $C = \{c, c_u\}$, $q_c = q_{c_u} = 1$, $N_c = \{1, 4\}$, and $1 \succ_\pi 2 \succ_\pi 3 \succ_\pi 4$. The outcome of the minimum-guarantees rule is that agent 1 and 2 are selected and the matching is $\{\{1, c\}, \{2, c_u\}\}$. 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 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 will assume that $q_{c_u^1} + q_{c_u^2} = q_{c_u}$ and $\succ_{c_u^1} = \succ_{c_u^1} = \succ_\pi$. One of the contributions of Pathak et al. (2020) was to formulate a family of rules, called *smart reserves rules*, that not only allows agents to be eligible for multiple preferential categories, but generalizes the minimum-guarantees and over-and-above rule when agents may be eligible for multiple preferential categories. In this section, we capture these approaches via an axiom for matchings called *order preserving* 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 category 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 9 (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 baseline ordering $\succ_\pi$ if for any two agents $i, j \in N$

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

[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).

(ii) $\mu(j) \in C_p \cup \{c_u^1\}$, $i \succ_{\mu(j)} j$, and i is eligible for category $\mu(j)$ implies $\mu(i) \neq c_u^2$.[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 9 is that instead of explaining or describing over-and-above and minimum-guarantees rules as consequences of certain sequential allocation methods, Definition 9 captures the key property of their resulting *matching*. It is formulated so that it allows for heterogeneous priorities or agents being eligible for multiple categories. The following propositions provide characterizations of the two well-known allocation rules. Both characterizations use order preservation.

Proposition 1 (Characterization 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, a matching is the outcome of the minimum-guarantees (over-and-above) rule if and only if it*

(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 $q_{c_u^2} = 0$).

Proof. Minimum-guarantees. First, the minimum-guarantees rule 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 later to the unmatched agents. Therefore, the matching respects priorities and satisfies order preservation for $q_{c_u^1} = 0$ and $q_{c_u^2} = q_{c_u}$.

Next, we prove that there is exactly one matching satisfying the five properties. Suppose for contradiction there are two such outcomes μ' and μ''. Then either there must be some category c that has different matches N' and N'' in the two outcomes μ' and μ'', respectively or the set of agents matched to u_c are different. First assume the former. Among them the agents in $N' \cup N''$, consider the least priority agent i in $(N' \setminus N'') \cup (N'' \setminus N')$. Suppose $i \in N' \setminus N''$. Then it means that there exists some $j \in N'' \setminus N'$ such that $j \succ_\pi i$. Hence, μ' does not satisfy order preservation or respect of priorities. The other case is that all the categories have the same agent matches but u_c has different sets of matched agents in μ' and μ''. But then, one of μ' and μ'' does not satisfy respect of priorities or non-wastefulness.

Over-and-above. First, the over-and-above rule complies with the eligibility requirements and is non-wasteful. It also yields a maximum beneficiary assignment because for each preferential

[15] It follows from $i \succ_{\mu(j)} j$ that i is eligible for category $\mu(j)$. We state it explicitly in (ii) to maintain the symmetry with (i).

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 $q_{c_u^2} = 0$.

Next, we prove that there is exactly one matching satisfying the five properties. Suppose for contradiction there is an outcome μ'' different than μ' the outcome of the over-and-above rule that also satisfies the four properties. Since μ'' is also a maximum beneficiary assignment, $|\mu''(c)| = |\mu'(c)|$ for each $c \in C_p$. Suppose $\mu'(c_u) \neq \mu''(c_u)$. Then consider the highest priority $i \in \mu'(c_u)$ such that $i \notin \mu''(c_u)$. Then, since μ'' satisfies order preservation for $q_{c_u^1} = q_{c_u}$ and $q_{c_u^2} = 0$ as well as respect of priorities, there is no $j \in \mu''(c_u)$ such that $i \succ_\pi j$. Since $|\mu''(c)| = |\mu'(c)|$, this implies that the number of agents from N_c who are unmatched in μ'' is higher than in μ' which implies that μ'' violates non-wastefulness. We have now concluded that $\mu''(c_u) = \mu'(c_u)$. Non-wastefulness and respect of priorities imply that $\mu''(c) = \mu'(c)$ for all $c \in C$. □

Pathak et al. (2020) take a flexible approach towards treating unreserved units. Their smart reserves rule 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 with respect to the preferential categories. Whereas the *REV* rule does not provide this flexible feature of processing units from c_u^1 earlier, the smart reserves rule is not equipped to handle heterogeneous priorities. We note that the ideas from our *REV* rule and the smart reserve rule complement each other and can be combined to obtain a *Smart Reverse Rejecting (S–REV)* rule.

6.2 Smart Reverse Rejecting Rule

Our Smart Reverse Rejecting (*S–REV*) rule tries to give unreserved units from c_u^1 to agents according to the baseline ordering as long as the remaining agents (who do not yet have any unit) can be matched to the preferential categories to get a maximum beneficiary assignment. While allocating unreserved units from c_u^1 to agents, if giving an agent i a unit from c_u^1 leads to a situation that empty-handed agents cannot form a maximum beneficiary assignment, then i is not given a unit from c_u^1. Once the agents who get a unit from c_u^1 are finalized, we can then use the *REV* rule to match the remaining agents to C_p. Finally, *S–REV* matches the remaining agents to c_u^2 according to the baseline ordering.

Next, we show that while *S–REV* gives a maximum beneficiary assignment and gives additional flexibility, it preserves the key properties satisfied by *REV*.

Lemma 1. *S–REV complies with the eligibility requirements.*

Proof. By construction, each agent is matched to a category they are eligible for so *S–REV* complies with the eligibility requirements. □

Lemma 2. *S–REV yields a maximum beneficiary assignment.*

Algorithm 2 Smart REV ($S\text{–}REV$) Rule

Input: $I = (N, C, (\succsim_c), (q_c))$; a baseline ordering $\succ_\pi$ over N.

Output: A matching.

$N_1 \longleftarrow \emptyset$ {Agents to be given unreserved units from c_u^1}
for agent i down the ordering in $\succ_\pi$ **do**
  if $|N_1| < q_{c_u^1}$ and there exists a maximum beneficiary assignment of agents in $N \setminus (N_1 \cup \{i\})$ in graph $B_{I'}$ where $I' = (N \setminus (N_1 \cup \{i\}), C_p, (\succsim_c), (q_c))$ **then**
    $N_1 \longleftarrow N_1 \cup \{i\}$
  end if
end for
Allocate the units reserved for C_p to agents in $N \setminus N_1$.

$$\mu' \longleftarrow REV((N \setminus N_1, C_p, (\succsim_c), (q_c)), \pi)$$

Allocate the units reserved for c_u^2 to agents in $N \setminus N_1$ who are unmatched in μ' in order of the baseline ordering. Add the corresponding edges to μ'.
Let μ be the matching obtained by adding to μ' the edges from agents in N_1 to c_u^1.
return μ

Proof. By construction, $S\text{–}REV$ yields a maximum beneficiary assignment. During the course of the algorithm, we only put an agent in N_1 if the agents in $N \setminus N_1$ can be matched to categories in C_p to obtain a maximum beneficiary assignment. □

One consequence is that $S\text{–}REV$ also gives a maximum size matching for the matching to all categories in $C = C_p \cup \{c_u^1, c_u^2\}$.

Lemma 3. *$S\text{–}REV$ returns a matching that satisfies order preservation.*

Proof. Consider a matching μ returned by $S\text{–}REV$. 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 in which Algorithm 2 considers agent i. 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: $\mu(j) = c_u^2$, $\mu(i) \in C_p \cup \{c_i^1\}$, $j \succ_\pi i$, 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 REV does not respect priorities, a contradiction. □

Lemma 4. *S–REV returns a matching that respects priorities.*

Proof. We first prove that no unmatched agent can have justified envy for an agent matched to c_u^1. Suppose an unmatched agent i comes earlier than an agent j matched to c_u^1. Then, i would have been selected in the set N_1 in the 'for' loop of Algorithm 2.

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 respects priorities. □

Lemma 5. *S–REV is strategyproof.*

Proof. We show that if an agent i is unmatched under S–REV, she cannot misreport to get matched. We first show that agent i cannot misreport to get matched to u_c^1. 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 truthful reporting, i is not needed to obtain a maximum beneficiary assignment even if she reports all her eligible categories, which implies that i is not needed to obtain a maximum beneficiary assignment if she reports a strict subset of her eligible categories. 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 was not matched to u_c^1, 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 could not have manipulated her priorities with respect to C_p to get one of them.

We have shown that i cannot affect the set N_1, that is, which agents are matched to c_u^1. So we suppose that agents matched to u_c^1 are already fixed. Since REV is strategyproof, agent i cannot get matched to a category in C_p by misreporting.

The remaining case is that, by misreporting, agent i affects the set of agents who are not matched to a category in $C_p \cup \{c_u^1\}$ and, hence, compete with i to be matched to c_u^2. We observe the following:

(i) Since agent i is not matched to a category in C_p and REV yields a maximum size matching, agent i cannot affect the *number* of agents matched to categories in C_p.

(ii) Since REV is weakly non-bossy, the set of agents *lower* in the baseline ordering who compete to be matched to c_u^2 is unchanged under a misreport by i.

The above two facts imply that under a misreport by i, the number of agents with a *higher* 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 does not change. Therefore, under any misreport, agent i is not matched to c_u^2. □

Lemma 6. *S–REV is weakly non-bossy.*

Proof. Consider an agent i who is unmatched under S–REV. By the proof of Lemma 5, it follows that i cannot affect

(i) the *set* of agents who are matched to c_u^1,

(ii) the *number* of agents with higher baseline ordering who are not matched to a category in $C_p \cup \{c_u^1\}$, and

(iii) the *set* of agents with lower baseline ordering who are not matched to a category in $C_p \cup \{c_u^1\}$.

Under a truthful report, agent i is unmatched and only agents with higher baseline ordering are matched to c_u^2. Hence, it follows that a misreport of agent i does not affect the set of agents with lower baseline ordering who are matched to c_u^2. □

Lemma 7. *S–REV is polynomial-time computable.*

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

Theorem 3. *The rule S–REV*

(i) complies with eligibility requirements,

(ii) yields a maximum size matching among feasible matchings,

(iii) yields a maximum beneficiary assignment,

(iv) respects priorities,

(v) is strategyproof,

(vi) is weakly non-bossy,

(vii) satisfies order preservation, and

(viii) is polynomial-time computable.

7 Discussion

We presented allocation rules that apply to resource allocation problems in which the resources are reserved for categories, each of which has a priority ranking over agents. The rules have several properties that are desirable in applications. Table 1 summarizes the properties satisfied by the main healthcare rationing algorithms discussed in this paper. Since $S\text{–}REV$ provided additional flexibility and still satisfies the main properties of REV, we do not include REV in the table.

	$S\text{–}REV$	Smart Reserves	DA
compliance with eligibility requirements	✓	✓	✓
non-wastefulness	✓	✓	✓
maximum size matching	✓	✓	–
respect of priorities	✓	✓*	✓
strategyproofness	✓	n/a	✓
weak non-bossiness	✓	n/a	✓
order preservation	✓	✓*	–
maximum beneficiary assignment	✓	✓	–
polynomial-time computability	✓	✓	✓

Table 1: Properties satisfied by $S\text{–}REV$, the Smart Reserves Rule of Pathak et al. (2020), and the Deferred Acceptance Algorithm described in Section 4. 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.

Our allocation rule involves a baseline ordering $\succ_\pi$ over the agents, which gives rise to a natural ordering in which patients are allocated units. We can go down the ordering $\succ_\pi$ and give a unit to the agent that was matched by the allocation rule.

Throughout the paper, we assumed that only matchings that satisfy the eligibility requirements are feasible. The disadvantage of such an approach is that it is possible that some preferential category units are not utilized even though some agents are unmatched. If we do not impose eligibility requirements as a hard constraint, the setting is referred to as the case of "soft reserves". In the case of soft reserves, we can do the following. We first use REV or $S\text{–}REV$ to compute a matching that complies with the eligibility requirements. If some units from C_p are not used and some agents are unmatched, we can use the baseline ordering to greedily allocate the agents to the units. Assuming each preferential category's priorities over *ineligible* agents are consistent with the baseline ordering, the resultant rule with the post step satisfies all the above properties except for "hard" compliance with the eligibility requirements. The argument for strategyproofness is similar to the proof of Lemma 5. By misreporting, an agent cannot affect the number of agents with higher baseline order who are unmatched.

We assumed that the categories and their capacities are primitives of the model. A separate research problem is to decide on the distribution of the units over the categories with the aim to reduce the pandemic for society. Finally, it will be useful to consider a more fine-grained model that allows quantifying how much a patient needs a unit.

We can reframe the main theorem in the context of school choice (Abdulkadiroğlu and Sönmez, 2003) by viewing agents as students and categories as schools. The students are indifferent between all schools that are acceptable for them. The schools have priorities over the students. Then Theorem 1 reads as follows.

Theorem 4. *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, has no justified envy, is non-wasteful, matches the maximal feasible number of students, and is strategyproof for students.*

Acknowledgements

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 5969/1-1. The authors thank Adi Ganguly for helpful comments and Tayfun Sönmez for additional pointers to the literature.

References

A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

A. Abdulkadiroğlu and T. Sönmez. Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory*, 112(1):157–172, 2003.

D. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce (ACM-EC)*, pages 295–304. ACM Press, 2007.

M. O. Afacan, I. Bó, and B. Turhan. Assignment maximization, 2020.

S. Ahmadi, F. Ahmed, J. P. Dickerson, M. Fuge, and S. Khuller. An algorithm for multi-attribute diverse matching. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3–9, 2020.

F. Ahmed, J. P. Dickerson, and M. Fuge. Diverse weighted bipartite b-matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 35–41. AAAI Press, 2017.

T. Andersson and L. Ehlers. Assigning refugees to landlords in Sweden: Stable maximum matchings. Technical Report 2016-08, Université de Montréal Papyrus Institutional Repository, 2016.

O. Aygün and I. Bó. College admission with multidimensional privileges: The Brazilian affirmative action case. *American Economic Journal: Microeconomics*, 2020.

O. Aygün and B. Turhan. Dynamic reserves in matching markets: Theory and applications. *Journal of Economic Theory*, 188, 2020.

H. Aziz. Mechanisms for house allocation with existing tenants under dichotomous preferences. *Journal of Mechanism and Institution Design*, 2018.

H. Aziz, S. Gaspers, and Z. Sun. Mechanism design for school choice with soft diversity constraints. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 153–159, 2020.

H. Aziz, P. Biró, and M. Yokoo. Matching market design with constraints. 2021.

P. Biró, T. Fleiner, R. W. Irving, and D. F. Manlove. The college admissions problem with lower and common quotas. *Theoretical Computer Science*, 411(34):3136 – 3153, 2010.

P. Biró, D. F. Manlove, and S. Mittal. Size versus stability in the marriage problem. *Theoretical Computer Science*, 411(16-18):1828–1841, 2010.

G. Bognar and I. Hirose. *The Ethics of Health Care Rationing: An Introduction*. Routledge, 2014.

A. Bogomolnaia and H. Moulin. Size versus fairness in the assignment problem. *Games and Economic Behavior*, 90:119–127, 2015.

L. Bruce and R. Tallman. Promoting racial equity in covid-19 resource allocation. *Journal of Medical Ethics*, 2021.

A. Dawson, D. Isaacs, M. Jansen, C. Jordens, I. Kerridge, U. Kihlbom, H. Kilham, A. Preisz, L. Sheahan, and G. Skowronski. An ethics framework for making resource allocation decisions within clinical care: Responding to covid-19. *Journal of Bioethical Inquiry*, 17(4):749–755, 2020.

J. P. Dickerson, K. A. Sankararaman, A. Srinivasan, and P. Xu. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1877–1884, 2019.

U. Dur, S. D. Kominers, P. A. Pathak, and T. Sönmez. Reserve Design: Unintended Consequences and the Demise of Boston's Walk Zones. *Journal of Political Economy*, 126(6): 2457–2479, 2018.

U. Dur, P. A. Pathak, and T. Sönmez. Explicit vs. statistical targeting in affirmative action: Theory and evidence from Chicago's exam schools. *Journal of Economic Theory*, 187:104996, 2020.

F. Echenique and M. B. Yenmez. How to control controlled school choice. *American Economic Review*, 105(8):2679–94, August 2015.

L. Ehlers, I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory*, 153:648—683, 2014.

E. J. Emanuel, G. Persad, R. Upshur, B. Thome, M. Parker, A. Glickman, C. Zhang, C. Boyle, M. Smith, and J. P. Phillips. Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, 382(21):2049–2055, 2020.

S. Fink. The hardest questions doctors may face: Who will be saved? who won't? *The New York Times*, March 21, 2020.

M. Galanter. Equality and protective discrimination in India. *Rutgers Law Review*, 16(1):42–74, 1961.

M. Galanter. *Competing Equalities: Law and the Backward Classes in India.* Univ of California, 1984.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

Y. A. Gonczarowski, N. Nisan, L. Kovalio, and A. Romm. Matching for the Israeli "Mechinot" gap year: Handling rich diversity requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation*, pages 321–321, 2019.

M. Goto, A. Iwasaki, Y. Kawasaki, R. Kurata, Y. Yasuda, and M. Yokoo. Strategyproof matching with regional minimum and maximum quotas. *Artificial intelligence*, 235:40–57, 2016.

A. Grigoryan. Effective, fair and equitable pandemic rationing. Technical Report 3646539, SSRN, Nov 2020.

I. E. Hafalir, M. B. Yenmez, and M. Yildirim. Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363, 2013.

E. J. Heo. Equity and diversity in college admissions. In J. Laslier, H. Moulin, R. Sanver, and W. Zwicker, editors, *The Future of Economic Design*. Springer, 2019.

Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *The American Economic Review*, 105(1):67–99, 2015.

Y. Kamada and F. Kojima. Recent developments in matching with constraints. *The American Economic Review*, 107(5):200–204, 2017.

F. Kojima. New directions of study in matching with constraints. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design*. Springer-Verlag, 2019.

P. Krysta, D. F. Manlove, B. Rastegari, and J. Zhang. Size versus truthfulness in the house allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 453–470. ACM Press, 2014.

R. Kurata, N. Hamada, A. Iwasaki, and M. Yokoo. Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research*, 58:153–184, 2017.

E. Martin. Rationing in healthcare. Technical Report 8, Deeble Institute, 2015.

NASEM-National Academies of Sciences, Engineering, and Medicine. Framework for Equitable Allocation of COVID-19 Vaccine. Technical report, Washington, DC, 2020.

P. A. Pathak, T. Sönmez, M. U. Ünver, and M. B. Yenmez. 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, July 2020. URL https://ideas.repec.org/p/boc/bocoec/1015.html.

G. Persad, M. E. Peek, and E. J. Emanuel. Prioritizing groups for access to covid-19 vaccines. *JAMA*, 324(16):1601–1602., 2020.

A. E. Roth. Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory*, 36:537—569, 2008.

A. E. Roth and M. A. O. Sotomayor. *Two-Sided Matching: A Study in Game Theoretic Modelling and Analysis*. Cambridge University Press, 1990.

T. Sönmez and M. B. Yenmez. Affirmative action with overlapping reserves. Manuscript, 2020. URL http://fmwww.bc.edu/EC-P/wp990.pdf.

T. Sönmez, P. A. Pathak., M. U. Ünver, G. Persad, R. D. Truog, and D. B. White. Categorized priority systems: A new tool for fairly allocating scarce medical resources in the face of profound social inequities. *CHEST*, 2020.

R. D. Truog, C. Mitchell, and G. Q. Daley. The toughest triage — allocating ventilators in a pandemic. *New England Journal of Medicine*, 382(21):1973–1975, 2020.

WHO. A global framework to ensure equitable and fair allocation of COVID-19 products and potential implications for COVID-19 vaccines. Technical report, 2020.