

A Systematic Approach to Selection Problems

We discuss a systematic approach to Prioritized Selection Problems, in which an organization is presented with a set of individuals, and must choose which subset to accept. We assume that the organization has a complete priority ranking of the individuals. In addition, the selection rule may be constrained to select from a pre-defined set of feasible subsets, which may be specified explicitly or implicitly (through quotas or reserves). We identify a natural family of rules—sequential dictator rules—and characterize it using some natural axiomatic properties. In cases where the feasible subsets are implicitly specified through upper and lower quotas or through reserves, we identify important special cases for which there exists a selection that unambiguously selects higher priority individuals than any other feasible selection. Furthermore, we identify conditions under which this selection can be found in polynomial time, and conditions in which doing so is computationally hard.

Our general approach is motivated by the observation that it is easier for the general public and policymakers to reason about outcomes rather than algorithms. Our sequential dictator rules do not require policymakers to reason about the effect of tweaking some step of the algorithm. Rather, they rely only on a description of which final selections are acceptable. We believe that this feature makes these rules more “designer friendly,” promotes transparency, and makes them less likely to produce unintended consequences.

1 INTRODUCTION

We study *Prioritized Selection Problems*, in which an organization is presented with a set of individuals, and must choose which subset to accept. We assume that the organization has a complete priority ranking of individuals. This ranking may be determined by various factors, including date of application, test scores or other measures of merit, a lottery, or some combination of these. In many applications, the organization also has objectives or constraints that prevent it from simply selecting the highest-priority individuals. For two examples, consider the following:

- **US Diversity Immigrant Visa program.** Priority is determined by lottery. However, to ensure a diverse set of immigrants, there are upper quotas on the number of individuals selected from each country and region.
- **COVID-19 Vaccine Scheduling.** Priority is determined by factors including age, occupation, and health status. However, doses are spread across multiple locations, and individuals must be assigned to a location which they can reach.

Given a set of individuals $\mathcal{I}$, a *selection* is a subset of these individuals, and a *selection rule* φ is a function from priority orders over $\mathcal{I}$ to selections. In other words, $\varphi(>)$ gives the set of selected individuals when the priority order is $>$.

1.1 Feasible Selections

In some applications, including the two examples mentioned above, the selection rule is constrained to select from a pre-defined set of *feasible* selections $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{I}}$. In the case of the Diversity Visa Lottery, a selection is feasible (in $\mathcal{F}$) if it abides by all country and regional upper quotas. In the case of vaccine scheduling, a selection is feasible if it is possible to assign each selected individual to a location such that (i) every individual can reach their assigned location, and (ii) no location is assigned more individuals than it can serve.

In the presence of such constraints, it is natural to choose a feasible selection that “respects the priority ranking as much as possible.” Although there are several possible interpretations of this phrase, we say that a feasible selection *respects priorities* if, for every individual who is not selected, it is impossible to select this individual without displacing some higher-priority individual who is currently selected.

Proposition 1 establishes that there is always a unique feasible selection that respects priorities, and it can be found by a simple “serial dictatorship” (Algorithm 1). This algorithm processes individuals in priority order, and selects each individual if there is a feasible selection that includes that individual and all others who have been selected so far. Proposition 1 further implies that we can define a family of rules—“sequential dictator rules”—parametrized by the collection of feasible selections $\mathcal{F}$, and denoted $SD_{\mathcal{F}}$: Given any priority order $>$, $SD_{\mathcal{F}}(>)$ is the unique feasible selection that respects priorities.

1.2 Key Questions and Results

The key questions we will address in this work are:

- (1) What are some desirable features of the family of sequential dictator selection rules?
- (2) For which choices of $\mathcal{F}$ does using a “serial dictatorship” align with natural notions of efficiency such as maximizing the number of selected applicants?

Section 3 addresses question (1). For any collection of feasible selections $\mathcal{F}$, we show that the rule $SD_{\mathcal{F}}$ is

- **Monotonic:** If a selected individual’s priority increases (fixing the priority ranking of others), then that individual continues to be selected.
- **Non-bossy:** If a change in an individual’s priority ranking does not affect whether that individual is selected, then it does not affect any other individual.
- **Lower invariant:** Whether an individual is selected depends only on the ranking of higher priority individuals, and not on how lower priority individuals are ranked.

Furthermore, Theorem 1 establishes that these properties in fact *characterize* the class of sequential dictator selection rules. That is, if a selection rule satisfies these three properties, then it must be equivalent to $SD_{\mathcal{F}}$ for some suitably chosen $\mathcal{F}$. Finally, we show that these three properties are independent by giving examples of selection rules that satisfy any two of them, but not the third.

Section 4 addresses question (2) by introducing a model in which each individual i has a “value” v_i of being selected, and the value of a selection is the sum of the values of selected individuals. Given feasible subsets $\mathcal{F}$, a feasible selection is *efficient* if no feasible selection has a higher value.

Clearly, if the individuals with highest priority according to $>$ have low value, the goals of efficiency and respecting priority are incompatible. Therefore, we assume that the values “agree with” the priority order, meaning that higher-priority individuals are (weakly) more valuable to select. Even then, simple examples demonstrate that for some choices of $\mathcal{F}$ and $\mathbf{v}$, the value of the selection chosen by $SD_{\mathcal{F}}$ can be an arbitrarily small fraction of the value of an efficient selection. As a result, we consider two special cases for the feasible sets $\mathcal{F}$ which arise in many applications. We refer to these as *quotas* and *reserves*.

When using quotas, we consider models in which (i) there is a finite set of traits T ; (ii) each applicant has a subset of these traits; and (iii) the feasible selections are defined by upper and lower bounds (quotas) on the number of selected individuals with each trait. For this family of feasible subsets $\mathcal{F}$, we show that finding an efficient selection is NP-hard (Proposition ??). However, we identify special cases in which (i) the sequential dictator rule $SD_{\mathcal{F}}$ finds an efficient selection, and (ii) the sequential dictator rule may not find an efficient selection, but an efficient selection can be found by solving a linear program. Specifically, (i) holds if the traits are *nested* (Proposition 5), and (ii) holds if the traits form a *bihierarchy* (Proposition ??).

When using reserves, the feasible selections are defined by a finite set of positions, and a bipartite graph $\mathcal{G}$ of individuals and positions, indicating which individuals may be matched to which positions. For example, in the case of the vaccine allocation problem described above, there could be an edge from an individual to a particular dose if the individual can reach the location where that

dose is stored. In this case, it is natural to treat the edges of the compatibility graph as defining hard constraints: a selection is feasible if and only if it is possible to match every selected individual to a compatible position. For this family of feasible subsets $\mathcal{F}$, we show that $SD_{\mathcal{F}}$ can be computed in polynomial time, and $SD_{\mathcal{F}}$ selects a subset of individuals that *priority dominates* any other feasible selection (and therefore is efficient).

In some applications, the compatibility graph expresses priorities (“soft” constraints), rather than hard constraints. For example, universities in India often reserve some seats for students from “backwards classes” which have historically been marginalized. However, if not enough students from backwards classes apply, then it may be natural to make these seats available to others. Inspired by such settings, we formulate two definitions of *soft reserves*. Under both definitions, the number of selected individuals can be at most the number of positions. In the case of *maximal reserves*, feasible selections must also be a superset of the individuals matched in some maximal matching of $\mathcal{G}$. In the case of *maximum reserves*, feasible selections must be a superset of the individuals matched in some maximum matching of $\mathcal{G}$. We find that computing an efficient selection is NP-hard for the case of maximal reserves but can be done efficiently for the case of maximum reserves.

1.3 Motivation for Our Work

Although there are many contexts in which our results could be applied, we are primarily motivated by “affirmative action” style policies, which are intended to boost representation of specific groups. Such policies are widespread, and often implemented on an ad-hoc basis. Specifically, policies often describe an algorithm for selecting individuals. There are numerous cases in which the algorithms that are deployed do not have the desired effect, and even cases where they have the *opposite* of the intended effect. For example,

- Boston used to have a policy that reserved half of the seats at each school for students who lived within the school’s “walk zone”. Over a decade after this policy was enacted, Dur et al. [2018] reported that it led to outcomes which were nearly identical to what would have happened with no walk zone priority at all!
- In India, government positions are allocated based on applicant exam scores, but some positions are reserved for applicants from historically marginalized groups, including women and “backwards classes.” The complexities of dealing with overlapping reservations resulted in the use of an algorithm which sometimes selected a woman from a majority group over another woman from a backwards class *with a higher exam score*! This clearly violates the intention of the program. Such instances resulted in numerous lawsuits, and several instances in which the Supreme Court of India tried to specify the exact algorithm that should be used to allocate positions [Sönmez and Yenmez, 2022].
- In New York City, specialized high schools admit students on the basis of a test score. The “Discovery Program” reserves seats at these schools for students from disadvantaged backgrounds. However, high-scoring students who are eligible for these seats are often assigned to an “open” seat at their second or third choice rather than a reserved seat at their first choice [Faenza et al., 2022].
- Each year, up to 20,000 H1B visas are reserved for applicants with advanced degrees from US institutions. In 2009, swelling demand caused a change from first-come first-served to a lottery. Seemingly unintentionally, this change was made in a way that significantly increased the number of visas awarded to advanced degree holders [Pathak et al., 2020].

Individually, each of these examples provides an interesting research opportunity. Collectively, they illustrate the challenge of reasoning about the behavior of complex algorithms. Recent work by Pathak et al. [2022] supports this argument, as it uses a survey of over 1,000 people to demonstrate

that the effects of changing the order in which positions reserved for specific groups are filled are poorly understood.

One of the motivations for our work is that we believe it is much easier for the public and policymakers to form opinions about what *outcomes* are acceptable, rather than the *algorithms* that should be used to find these outcomes. Our sequential dictator rules do not require policymakers to reason about the effect of tweaking some step of the algorithm. Rather, they rely only on a description of which final selections are acceptable. We believe that this feature makes these rules more “designer friendly” and less likely to produce unintended consequences. The special cases considered in Section 4 provide additional assistance by establishing practical choices for the collection of feasible sets under which sequential dictator rules are well-behaved.

2 RELATED WORK

Our paper contributes to the rich literature on assignment and matching problems with distributional constraints. Such considerations arise in the design of rules for assigning students to schools and colleges, for rationing health-care resources, etc.

Abdulkadiroğlu and Sönmez [2003] propose the use of deferred acceptance algorithms for school choice, and mention that one advantage of these algorithms is that they can easily be modified to accommodate racial quotas. Kojima [2012] observes that such a solution could have negative consequences for certain types of students, sometimes the very types of students the affirmative action policy is meant to favor. Specifically, Kojima [2012] considers a school choice model with two types of students—majority and minority—and finds that imposing an upper bound on the number of majority students at each school may end up making the minority students worse off: a strict upper bound for majority students at one school (even if the school has empty seats) creates increased competition from the majority students at other schools, and could have an adverse effect on minority students, especially if they are not interested in schools that have unfilled seats because of the upper quota constraint. This observation has led market designers to think more carefully about the consequences and unintended side-effects of such quota-based affirmative action policies.

These concerns motivate Hafalir et al. [2013] to consider effective affirmative action policies which rely on *minority reserves*: working with a model with two types of individuals, they propose a mechanism in which schools give higher priority to minority students up to the point that the minorities fill the reserves, and show that the resulting allocation are more consistent with the goals of the designer. A related model of Ehlers et al. [2014] attempts to achieve diversity goals by imposing both upper and lower bounds on the number of students of each type admitted at any school, while still giving students the freedom to rank schools. Not surprisingly, if these bounds are treated as hard constraints, a good assignment of students to schools may not exist; so these authors interpret these upper and lower bounds as “soft” constraints that should be satisfied “as much as possible” and design an algorithm to find a fair and non-wasteful assignment of students to schools. While the literature discussed so far is on two-sided matching problems with strict preferences and do not have any direct relevance to the models we discuss, they set the stage for the types of models that have been explored in the subsequent literature on one-sided assignment models with distributional constraints: models with lower and upper quotas, reserves, interpreting constraints as hard (if there is a feasible solution) or soft (when faced with infeasibility or unallocated objects) are all features of the models we study and were first explored in these two-sided matching contexts.

An important shortcoming of these early models is the assumption that agents are *partitioned* into types. Richer models arise if we relax this assumption and assume that each agent may have multiple types (“traits,” in our terminology). Motivated by college admissions in India [Baswana et al., 2019] [Aygun and Turhan, 2021], school choice in Chile [Correa et al., 2021], and other matching problems in practice, more recent papers consider diversity goals when agents have

multiple traits. In this case, there are many possible accounting conventions that could be used. One possibility is to allow this agent to count against the “quota” for each of their traits. An alternative is for the agent to be assigned to a position “reserved” for one of their traits. Yet another option is to distribute one unit of weight across the traits of the chosen individual [Delacretaz, 2021]. In this paper, we consider the first two possibilities.

Past work on reserves has considered settings where some seats are reserved for protected categories, while remaining seats are “open” to anybody. One key question facing a designer is this: if an applicant in a protected category is processed when both an open seat and a reserved seat are available, which seat shall they receive? Pathak et al. [2020] note that assigning these applicants to reserved seats increases the merit ranking of selected individuals, but reduces the number of reserve-eligible candidates who are selected, while assigning them to open seats has the opposite effect. Sönmez and Yenmez [2022] note that both policies are in place simultaneously for different types of reservations in India, and refers to the former as a “minimum guarantee” and the latter as an “over and above” reservation. We note that while both policies are well-defined selection rules, sequential dictator rules are best thought of as implementing minimum guarantees. They cannot replicate over-and-above policies, because these policies can be shown to be bossy.

3 RESPECTING PRIORITIES

3.1 Model

In a *selection problem* there is a set of applicants $\mathcal{I}$, from which a decision maker has to select a subset $S \subseteq \mathcal{I}$. There is a strict priority order $>$ over $\mathcal{I}$. We denote by $\mathbf{P}$ the set of all priority orders over $\mathcal{I}$. A *selection rule* is a mapping $\varphi : \mathbf{P} \rightarrow 2^{\mathcal{I}}$ that determines a selection for every priority order. We let $\varphi(>)$ denote the selection of rule φ for priority order $>$. In addition, there is a collection of feasible selections $\mathcal{F} \subseteq 2^{\mathcal{I}}$ which may capture capacity constraints, diversity goals or other restrictions imposed by law or by the designer.

3.2 The Sequential Dictator Selection Rule

In this section, our goal is to define a selection rule that (i) always returns a feasible selection and (ii) adheres to the priority order “as much as possible.”

The correct interpretation of the latter criteria is not obvious. Consider an example with four individuals, with $1 > 2 > 3 > 4$. Suppose that the feasible selections are $\mathcal{F} = \{\{1, 4\}, \{2, 3\}\}$. Then it is not clear which feasible selection to choose. In contrast, if $\mathcal{F} = \{\{1, 4\}, \{2, 3\}, \{1, 3\}\}$, then we claim that $\{1, 3\}$ is an obvious choice. The following definition formalizes this idea.

DEFINITION 1. *Selection S' **priority dominates** selection S if (and only if) $S' \neq S$ and there exists an injective function ψ from S to S' , such that for each $i \in S$, $\psi(i) \geq i$.*

Note in particular that if S' priority dominates S , then $|S'| \geq |S|$. One way to think of priority domination is as follows. For $k \leq |S|$, let i_k^S be the k^{th} highest priority individual in S , and $i_k^{S'}$ be the k^{th} highest priority individual in S' . If S' priority dominates S , then for each $k \leq |S|$, we must have $i_k^{S'} \geq i_k^S$. In our example above, $\{1, 3\}$ priority dominates both $\{1, 4\}$ and $\{2, 3\}$. However, neither of the selections $\{1, 4\}$ and $\{2, 3\}$ priority dominate the other.

If our goal is to respect the priority order as much as possible, then at a minimum, we should choose a feasible selection that is not priority dominated. However, this does not address the question of which selection to choose when there are multiple undominated feasible selections. To resolve this question, we propose the following criteria.

ALGORITHM 1: Sequential Dictator

Input: Set of applicants $\mathcal{I}$, feasible selections $\mathcal{F}$ and priority order $>$, with $i_1 > i_2 > \dots > i_n$.

$S := \emptyset$;

foreach $k = 1, \dots, n$ **do**

if $\exists S' \in \mathcal{F} : S \cup \{i_k\} \subseteq S'$ **then**

$S = S \cup \{i_k\}$

end

end

return S

DEFINITION 2. A feasible selection $S \in \mathcal{F}$ **respects priorities** if, for each unselected applicant $i \notin S$, it is impossible to select them without displacing a higher-priority applicant. Formally, there is no feasible selection $S' \in \mathcal{F}$ such that $\{i\} \cup \{j \in S : j > i\} \subseteq S'$.

The motivation for this definition is that it offers a “simple” explanation for each individual’s outcome. Suppose that an individual i complains that they were not selected, while some lower priority individuals were. If we have chosen a selection that respects priorities, then we can tell i that these lower priority individuals were not the reason that i was not selected. Rather, i was not selected because any feasible selection including i would require the displacement of some individual with a *higher* priority.

Applying this concept to our previous example with $1 > 2 > 3 > 4$ and $\mathcal{F} = \{\{1, 4\}, \{2, 3\}\}$, the selection $\{1, 4\}$ respects priorities, while the selection $\{2, 3\}$ does not. Informally, if we choose $\{2, 3\}$, there is no “simple” explanation for why individual 1 was not selected. The following result establishes that there is always a unique selection that respects priorities. This selection is never priority dominated, and if there is a feasible selection that priority dominates all others, then this must be the selection that respects priorities.

PROPOSITION 1. For any feasible set $\mathcal{F}$ and any priority order $>$:

- (i) There is a unique feasible selection that respects priorities, which we denote $SD_{\mathcal{F}}(>)$.
- (ii) The selection $SD_{\mathcal{F}}(>)$ is not priority dominated by any other feasible selection.
- (iii) If there exists $S^* \in \mathcal{F}$ that priority dominates every other feasible selection, then $SD_{\mathcal{F}}(>) = S^*$.

PROOF. For part (i), let $SD_{\mathcal{F}}$ be the selection chosen by the ‘sequential dictator’ algorithm (Algorithm 1). Let S be any other feasible selection. Let i be the highest priority individual for which the selections $SD_{\mathcal{F}}$ and S differ. Because S is feasible and the two selections make identical choices for individuals with higher priority than i , it must be that $i \in SD_{\mathcal{F}}(>)$, and therefore that $i \notin S$. Therefore, S does not respect priorities, as $SD_{\mathcal{F}}(>)$ is feasible and selects i along with all higher-priority individuals selected by S .

For part (ii), suppose that there is a feasible selection S that priority dominates $SD_{\mathcal{F}}(>)$. Let i be the highest-ranking individual that is in S but not $SD_{\mathcal{F}}(>)$. By definition, S and $SD_{\mathcal{F}}(>)$ make identical decisions on individuals with higher priority than i : $\{j : j > i\} \cap S = \{j : j > i\} \cap SD_{\mathcal{F}}(>)$. Therefore, when $SD_{\mathcal{F}}(>)$ considered individual i , S was a feasible selection containing i and all higher-priority individuals that had already been selected. Therefore, $SD_{\mathcal{F}}(>)$ should select i , contradicting our assumption that S priority dominates $SD_{\mathcal{F}}(>)$.

From this, (iii) follows immediately. If there exists $S^* \in \mathcal{F}$ that priority dominates every other feasible selection, then if $SD_{\mathcal{F}}(>)$ is not equal to this selection, it is priority dominated by S^* , contradicting (ii). $\square$

3.3 An Axiomatic Characterization of Sequential Dictator Rules

The previous subsection defined the sequential dictator rule $SD_{\mathcal{F}}$ and showed that for any fixed $\mathcal{F}$, it is the only rule that is feasible and respects priorities. Often, policymakers get to choose the feasible selections $\mathcal{F}$. Different choices of $\mathcal{F}$ correspond to different selection rules: we refer to $(SD_{\mathcal{F}})_{\mathcal{F} \subseteq 2^I}$ as the family of respecting priorities selection rules. In this section, we give an axiomatic characterization of this family of selection rules.

In practice, selection rules are often defined not through a set of feasible outcomes, but rather by a description of an algorithm for considering and selecting applicants. In such cases, the feasible outcomes are often defined only implicitly. As one example, we consider the case of “over and above” reserves .

EXAMPLE 1. *There are four individuals, 1, 2, 3, 4. Individuals 1 and 2 belong to a minority group. We can select at most two individuals, and wish to reserve one seat for minorities. We decide to use an “over and above” selection rule φ^{OA} . This rule first selects the highest-priority individual, and then selects the highest priority minority that remains.*

Note that the selection rule φ^{OA} is defined algorithmically, and does not explicitly identify which selections are feasible. However, for any selection rule φ , we can define

$$\text{Image}(\varphi) = \{S \subseteq 2^I : S = \varphi(>) \text{ for some } > \in \mathbf{P}\}$$

to be the set of selections that φ might choose. In the case of φ^{OA} , it is easy to see that this consists of the five selections $\{1, 2\}$, $\{1, 3\}$, $\{1, 4\}$, $\{2, 3\}$ and $\{2, 4\}$. Note that $\{3, 4\}$ is never selected, as neither 3 nor 4 is a minority.

We claim that for any choice of $\mathcal{F}$ containing $\text{Image}(\varphi^{OA})$, the rule φ^{OA} does not respect priorities. To see this, note that when the priority order is $1 > 3 > 2 > 4$, φ^{OA} selects $\{1, 2\}$, even though $\{1, 3\}$ is feasible (meaning that it would have been chosen for a different priority order) and priority dominates $\{1, 2\}$.

Thus, while choosing $\mathcal{F}$ gives the designer substantial flexibility (as we explore in Section 4), restricting to the family of sequential dictator rules does rule out certain selection rules commonly used in practice. One might therefore ask, why should a designer use a sequential dictator rule? What advantages do these rules offer?

We have already argued informally that defining these rules using acceptable *outputs* (rather than *algorithms*) makes them more transparent. Now, we show that sequential dictator rules satisfy three desirable properties, and are in fact the *only* selection rules to satisfy all three properties. To define each of these properties, we introduce new notation. Given a priority order $>$ and a subset of individuals $S \subset I$, we let $>_{-S}$ be the restriction of $>$ to $I \setminus S$. In other words, for $i, j \in I \setminus S$,

$$i >_{-S} j \iff i > j.$$

When $S = \{i\}$ is a singleton, we write $>_{-i}$ in place of $>_{-\{i\}}$.

Our first property is *monotonicity*, which states that improvements to an individual’s priority should weakly improve their outcome .

DEFINITION 3. *A selection rule φ is **monotone** if for any applicant $i \in I$, and for every $>, >' \in \mathbf{P}$ such that (i) $>_{-i} = >'_{-i}$ and (ii) the priority of applicant i is higher in $>'$ than in $>$, then*

$$i \in \varphi(>) \implies i \in \varphi(>').$$

The second property is *non-bossiness*, which says that a change to an individual’s priority which does not affect whether that individual is selected should also not affect who else is selected.

DEFINITION 4. A selection rule φ is **non-bossy** if for any applicant $i \in \mathcal{I}$, and for every $\succ, \succ' \in \mathbf{P}$ such that $\succ_{-i} = \succ'_{-i}$,

$$\{i\} \cap \varphi(\succ) = \{i\} \cap \varphi(\succ') \Rightarrow \varphi(\succ) = \varphi(\succ').$$

Our final property is *lower-invariance*, which says that whether an individual is selected should depend only on the priority of *higher*-priority individuals, and not on the exact priority of lower-priority individuals.

DEFINITION 5. A selection rule φ is **lower invariant** if for any applicant $i \in \mathcal{I}$, and for every $\succ, \succ' \in \mathbf{P}$ such that $\succ_{-\{j:i>j\}} = \succ'_{-\{j:i>j\}}$,

$$\{i\} \cap \varphi(\succ) = \{i\} \cap \varphi(\succ').$$

We are now ready to present our characterization of the family of sequential dictator rules.

THEOREM 1.

- For any collection $\mathcal{F} \subseteq 2^{\mathcal{I}}$, the selection rule $SD_{\mathcal{F}}$ is monotone, non-bossy and lower invariant.
- If φ is a selection rule that is monotone, non-bossy and lower invariant, then there exists $\mathcal{F} \subseteq 2^{\mathcal{I}}$ such that $\varphi(\succ) = SD_{\mathcal{F}}(\succ)$ for every $\succ \in \mathbf{P}$.

4 SPECIAL CASES OF INTEREST: QUOTAS AND RESERVES

In general, choosing the selection that respects priorities may come at a significant cost. For example, suppose that there are only two feasible selections: $\{1\}$, and $\mathcal{I} \setminus \{1\}$. Then if individual 1 is at the top of the priority order, respecting priorities requires that we select this individual, and reject all others. However, this example is somewhat contrived and unlikely to arise in practice.

In this section, we impose additional structure on the feasible sets $\mathcal{F}$, in ways intended to capture common real-world constraints. Our over-arching question is, in which cases can $SD_{\mathcal{F}}$ be shown to priority dominate all other feasible selections? In these cases, the normative case for using this selection rule is clear.

In essence, $SD_{\mathcal{F}}$ greedily selects individuals, so long as doing so would not violate feasibility. Greedy algorithms are well-known to solve optimization problems when the underlying constraints define a matroid.¹ We now reproduce this result in our setting.

DEFINITION 6. The feasible selections $\mathcal{F}$ **induce a matroid** if for any $I_1, I_2 \subseteq \mathcal{I}$ such that $I_1 \subseteq S_1 \in \mathcal{F}$, $I_2 \subseteq S_2 \in \mathcal{F}$, and $|I_1| > |I_2|$, there exists $i \in I_1 \setminus I_2$ and $S_3 \in \mathcal{F}$ such that $I_2 \cup \{i\} \subseteq S_3$.

PROPOSITION 2. If $\mathcal{F}$ induces a matroid, then for any priority order $\succ$, $SD_{\mathcal{F}}(\succ)$ priority dominates any other feasible assignment.

While Proposition 2 gives a complete answer to when $SD_{\mathcal{F}}$ will produce a priority dominant selection, verifying whether $\mathcal{F}$ induces a matroid can be non-trivial. In addition, in some cases, $\mathcal{F}$ may induce a matroid, but $SD_{\mathcal{F}}$ may nevertheless be NP-hard to compute. Therefore, below, we consider two common structures for $\mathcal{F}$ from practice: Quotas and Reserves.

4.1 Reserved Positions

We now study a particular family of feasible selections which we refer to as *Reserved Positions*. The representation of this family is based on a bipartite graph $G = (\mathcal{I}, \mathcal{P}, E)$, where each node of the left side correspond to an applicant $i \in \mathcal{I}$, each node on the right side correspond to a position $p \in \mathcal{P}$, and there is an edge $(i, p) \in E \subseteq \mathcal{I} \times \mathcal{P}$ only if position p is acceptable for applicant i . We refer to $G = (\mathcal{I}, \mathcal{P}, E)$ as the *compatibility graph*. A *matching* M of G is a subset of edges such that

¹See https://en.wikipedia.org/wiki/Matroid#Greedy_algorithm.




Fig. 1. Several examples comparing hard, maximum, and maximal reserves, with individuals in blue and positions in green. In all three cases, lower-numbered individuals have higher priority. On the left, hard reserves, maximum reserves, and maximal reserves coincide: under all three, any pair of agents is feasible. The unique selection that respects priorities is $\{1, 2\}$, but an “over and above” reserve with A as an open seat and B as a reserved seat selects $\{1, 3\}$. In the middle, under hard reserves and maximum reserves, it is infeasible to match 1. On the other hand, under maximal reserves $\{1, 2\}$ is feasible (but not $\{1, 3\}$). On the right, under hard reserves or maximum reserves, the selection that respects priorities is $\{2, 3, 4, 5\}$. By Theorem 2, this priority dominates any other feasible selection under these constraints. However, under maximal reserves, the selection $\{1, 2, 3, 4\}$ is feasible (justified by assigning 2 to B, 3 to A, 4 to C, and assigning 1 to D, as none of 1, 5, 6 are eligible for D and 1 has highest priority).

every node is incident to at most one edge in M . For any matching M of G , we denote by I_M the set of applicants that are incident to an edge in M , that is,

$$I_M = \{i \in I : \text{exists a position } p \in \mathcal{P} \text{ such that } (i, p) \in M\}. \quad (1)$$

In some environments, compatibility constraints must be strictly enforced. We model these situations using what we call *hard reserves*. Formally, given a compatibility graph $G = (I, \mathcal{P}, E)$, the feasible selections when using hard reserves are

$$\mathcal{F}^{\text{hard}}(G) = \{S \subseteq I : \text{exists a matching } M \text{ of } G \text{ such that } I_M = S\}. \quad (2)$$

Under hard reserves, a position may go unfilled if there is no unselected individual who is eligible for it. In many applications, it may be preferable to reduce waste by allowing someone who would otherwise not be eligible for this position to be assigned to it ([?] refer to this as “de-reserving” this position). To model such cases, we introduce two types of “soft” reserves: *maximum reserves* and *maximal reserves*. We define

$$\mathcal{F}^{\text{maximum}}(G) = \{S \subseteq I : |S| \leq |\mathcal{P}| \text{ and } \exists \text{ a maximum matching } M \text{ of } G \text{ such that } I_M \subseteq S\}. \quad (3)$$

$$\mathcal{F}^{\text{maximal}}(G) = \{S \subseteq I : |S| \leq |\mathcal{P}| \text{ and } \exists \text{ a maximal matching } M \text{ of } G \text{ such that } I_M \subseteq S\}. \quad (4)$$

In both cases, applicants can be matched to positions for which they are not eligible, so long as everyone who is eligible for that position is selected. Observe that for any compatibility graph G ,

$$\mathcal{F}^{\text{hard}}(G) \subseteq \mathcal{F}^{\text{maximum}}(G) \subseteq \mathcal{F}^{\text{maximal}}(G).$$

Intuitively, hard constraints interpret compatibility identified by the edges of G very strictly, while maximum and maximal reserves reduce waste and can increase the priority of selected individuals. Depending on the relative importance of compatibility and the priority order, the designer can choose any of these three descriptions of feasible sets. Figure 1 illustrates the differences between these feasible sets and the corresponding sequential dictator rules.

The main result of this section (Theorem 2) establishes that, for any compatibility graph G , the families of feasible selections $\mathcal{F}^{\text{hard}}(G)$, $\mathcal{F}^{\text{maximum}}(G)$ and $\mathcal{F}^{\text{maximal}}(G)$ induce a matroid. From

Proposition 2, it follows that in these cases the selection that respects priorities will also priority dominate all other feasible selections. Furthermore, the selection when using maximal reserves will priority dominate the selection when using maximum reserves, which will priority dominate the selection when using hard reserves.

THEOREM 2. *For any $G = (\mathcal{I}, \mathcal{P}, E)$, the feasible selections $\mathcal{F}^{\text{hard}}(G)$, $\mathcal{F}^{\text{maximum}}(G)$ and $\mathcal{F}^{\text{maximal}}(G)$ defined by (2), (3) and (4) each induce a matroid.*

We now comment on connections between this result and others in the matching literature. Hard constraints induce the well-known “matching matroid” or “transversal matroid”²

The horizontal reserves in Sönmez and Yenmez [2022] correspond to what we call maximum reserves: their non-wastefulness condition implies that reserves are soft, and their condition of “maximally complying with reservations” further specifies maximum (rather than maximal) reserves. They provide a “meritorious horizontal choice rule”, and their Theorem 1 establishes that this choice rule priority dominates all other feasible selections. This result parallels (and can be seen as a consequence of) our conclusion that maximum reserves induce a matroid. While their meritorious choice rule and our sequential dictator rule $SD_{\mathcal{F}^{\text{maximum}}(G)}$ always find the same selection, the descriptions of these algorithms differ substantially. In particular, their algorithm takes two passes: one to fill as many reserved seats as possible, and another to fill any remaining positions with the highest-priority remaining individuals. By contrast, our algorithm is described in a single pass, which selects each individual so long as doing so is feasible. We see advantages to both descriptions. Our description offers participants a simpler and more transparent description of why they were or were not selected. Meanwhile, their algorithm offers a description that can be more easily converted into functional software.

We are not aware of any paper that explicitly defines a feasible set corresponding to our maximal reserves. However, several papers study algorithms which may select matchings that are feasible under maximal reserves but not maximum reserves. For example, “over-and-above” choice rules which preferentially allocate reserve-eligible students to “open” seats always find a selection that is feasible under maximal reserves, but do *not* necessarily find the selection that respects priorities within this feasible set (see Example 1). Similarly, school choice algorithms in Chile have an exogenously specified order in which students are considered for various seats [Correa et al., 2021]. This always produces a feasible selection under maximal reserves, but not necessarily one that priority dominates all others.

To our knowledge, the definition of maximal reserves and the result that maximal reserves induce a matroid are new. In light of this result, it seems natural that instead of using an exogenously specified order to fill positions, one could endeavor to choose the selection that respects priorities. Perhaps one reason that this has not been done is that the problem of finding this selection is NP hard for maximal reserves, while it can be solved in polynomial time for maximum and hard reserves.

PROPOSITION 3. *The problem of evaluating $SD_{\mathcal{F}^{\text{maximal}}(G)}$ is NP-complete.*

For any graph G , the rules $SD_{\mathcal{F}^{\text{maximum}}(G)}$ and $SD_{\mathcal{F}^{\text{hard}}(G)}$ can be evaluated in polynomial time.

4.2 Quotas

In the second family of constraints we study, there are multiple groups of applicants. For each group, there is an upper and lower limit on the number of its members that can be selected. Formally, there is a finite set of traits T . For each trait $t \in T$, there is a set of applicants with the trait $\mathcal{I}_t \subseteq \mathcal{I}$, and an upper and lower quota $u_t, \ell_t \in \mathbb{N}$ on the number of selected applicants with trait t . We use

²https://en.wikipedia.org/wiki/Matroid#Matroids_from_graph_theory.

$\mathcal{I}_T$ as shorthand for $\{I_t\}_{t \in T}$. Then, a selection is feasible if it satisfies both quotas for every trait. That is,

$$\mathcal{F}(\mathcal{I}_T, \mathbf{u}, \boldsymbol{\ell}) = \{S \subseteq \mathcal{I} : \text{for every } t \in T, \ell_t \leq |S \cap \mathcal{I}_t| \leq u_t\}. \quad (5)$$

In contrast with the reserved positions defined in Section 4.1, the problem of determining whether there exists a feasible selection is NP-complete. Furthermore, even when a feasible selection exists, there may be no feasible selection that priority dominates all others.

PROPOSITION 4. *Determining whether $\mathcal{F}(\mathcal{I}_T, \mathbf{u}, \boldsymbol{\ell})$ defined in (5) is empty is NP-complete.*

If $\mathcal{F}(\mathcal{I}_T, \mathbf{u}, \boldsymbol{\ell})$ is non-empty, there may be no feasible selection that priority dominates all others.

Our proof in the appendix establishes NP-completeness result even in special cases in which there is just a single global lower quota or a single global upper quota. These proofs reduce from the independent set problem and the set cover problem, respectively. Meanwhile, Example 2 shows that in general, there may be no feasible selection that priority dominates all others.

EXAMPLE 2. *Consider an instance with three applicants $\mathcal{I} = \{1, 2, 3\}$ and two traits $T = \{A, B\}$. Let $\mathcal{I}_A = \{1, 2\}$, $\mathcal{I}_B = \{1, 3\}$ and $u_A = u_B = 1 = \ell_A = \ell_B$. In this instance, there are only two feasible selections: $\mathcal{F}(\mathcal{I}_T, \mathbf{u}, \boldsymbol{\ell}) = \{\{1\}, \{2, 3\}\}$. If $1 > 2 > 3$, then neither feasible selection priority dominates the other.*

4.2.1 Hierarchies. As Example 2 illustrates, under quotas there might be tension between respecting priorities and other goals, such as maximizing the number of selected individuals. We now provide conditions under which the feasible selection that respects priorities priority dominates all others.

DEFINITION 7. *The traits $\mathcal{I}_T = \{I_t\}_{t \in T}$ are **nested** if for every $t, t' \in T$, either $I_t \subset I_{t'}$ or $I_{t'} \subset I_t$.*

*The traits $\mathcal{I}_T = \{I_t\}_{t \in T}$ form a **hierarchy** if for every $t, t' \in T$, either I_t and $I_{t'}$ are disjoint ($I_t \cap I_{t'} = \emptyset$) or one contains the other ($I_t \subset I_{t'}$ or $I_{t'} \subset I_t$).*

PROPOSITION 5. *If $\mathcal{I}_T$ are nested, then for any upper and lower quotas $\mathbf{u}, \boldsymbol{\ell} \in \mathbb{N}^T$ such that the set of feasible selections is non-empty, the feasible sets $\mathcal{F}(\mathcal{I}_T, \mathbf{u}, \boldsymbol{\ell})$ induce a matroid. Therefore, the selection that respects priorities also priority dominates all other feasible selections.*

If the collection of sets $\mathcal{I}_T = \{I_t\}_{t \in T}$ is not a hierarchy, then there exist upper and lower quotas $\mathbf{u}, \boldsymbol{\ell} \in \mathbb{N}^T$ and a priority order $>$ such that no feasible selection priority dominates all others.

We are confident that the positive result in Proposition 5 continues to hold when the traits form a hierarchy, so that this structure is in some sense “necessary and sufficient” for there to exist a selection that priority dominates all others. Furthermore, Algorithm 2 will always find this selection. A proof of this result will appear in future versions of this paper.

REFERENCES

- Atila Abdulkadiroğlu and Tayfun Sönmez. 2003. School Choice: A Mechanism Design Approach. *American Economic Review* 93, 3 (June 2003), 729–747. <https://doi.org/10.1257/000282803322157061>
- Orhan Aygun and Bertan Turhan. 2021. How to Deserve Reserves. *arXiv preprint arXiv:2103.05899* (2021).
- Surender Baswana, Partha Pratim Chakrabarti, Sharat Chandran, Yashodhan Kanoria, and Utkarsh Patange. 2019. Centralized admissions for engineering colleges in India. *INFORMS Journal on Applied Analytics* 49, 5 (2019), 338–354.
- Jose Correa, Natalie Epstein, Rafael Epstein, Juan Escobar, Ignacio Rios, Nicolás Aramayo, Bastian Bahamondes, Carlos Bonet, Martin Castillo, Andres Cristi, et al. 2021. School choice in Chile. *Operations Research* (2021).
- David Delacretaz. 2021. Processing Reserves Simultaneously. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*.
- Umut M. Dur, Scott Duke Kominers, Parag A Pathak, and Tayfun Sönmez. 2018. Reserve Design: Unintended Consequences and the Demise of Boston’s Walk Zones. *Journal of Political Economy* 126, 6 (December 2018).
- Lars Ehlers, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153, C (2014), 648–683.
- Yuri Faenza, Swati Gupta, and Zhang Xuan. 2022. Discovering Opportunities in New York City’s Discovery Program: an Analysis of Affirmative Action Mechanisms. (2022).
- Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. 2013. Effective affirmative action in school choice. *Theoretical Economics* 8, 2 (2013), 325–363. <https://doi.org/10.3982/TE1135>
- Joseph Douglas Horton and Kyriakos Kilakos. 1993. Minimum edge dominating sets. *SIAM Journal on Discrete Mathematics* 6, 3 (1993), 375–387.
- Fuhito Kojima. 2012. School choice: Impossibilities for affirmative action. *Games and Economic Behavior* 75, 2 (2012), 685–693.
- Parag A. Pathak, Alex Rees-Jones, and Tayfun Sonmez. 2020. Immigration Lottery Design: Engineered and Coincidental Consequences of H-1B Reforms. (2020).
- Parag A. Pathak, Alex Rees-Jones, and Tayfun Sonmez. 2022. Reversing Reserves. *Management Science* (2022).
- Tayfun Sönmez and M. Bumin Yenmez. 2022. Affirmative Action in India via Vertical, Horizontal, and Overlapping Reservations. *Econometrica* (2022).

ALGORITHM 2:

Input: Individuals and traits $\mathcal{I}_T$, upper and lower quotas $\mathbf{u}, \ell \in \mathbb{N}_0^T$ and priority order $\succ$.

$S := \emptyset$;

foreach $k = 1, \dots, |T|$ **do**

while $|S \cap \mathcal{I}_{t_k}| < \ell_{t_k}$ **do**

 Let i be the highest-priority individual in $\mathcal{I}_{t_k}$ such that $S \cup \{i\}$ does not violate any upper quotas.

$S = S \cup \{i\}$

end

end

foreach $j = 1, \dots, |I|$ **do**

if $i_j \notin S$ and $S \cup \{i_j\}$ does not violate any upper quotas **then**

$S = S \cup \{i_j\}$

end

end

return S

A PROOFS FROM SECTION 3

A.1 Proof of Theorem 1

We divide the proof into two parts: Lemma 1 establishes that for any collection of feasible selections $\mathcal{F} \subseteq 2^I$, the rule $SD_{\mathcal{F}}$ is monotonic, non-bossy and lower invariant. Lemma 2 shows that if a selection rule satisfies these three properties, then it must be equivalent to $SD_{\mathcal{F}}$ for some suitably chosen $\mathcal{F}$.

LEMMA 1. *For any collection $\mathcal{F} \subseteq 2^I$, the selection rule $SD_{\mathcal{F}}$ is monotone, non-bossy and lower invariant.*

PROOF OF LEMMA 1. In what follows, we fix an arbitrary collection of feasible selections $\mathcal{F} \subseteq 2^I$ and an arbitrary applicant $i \in I$.

We begin by showing that $SD_{\mathcal{F}}$ is lower invariant. Let $>, >' \in \mathbf{P}$ be such that $>_{\{j:j>i\}} = >'_{\{j:j>i\}}$. We must prove that

$$\{i\} \cap SD_{\mathcal{F}}(>) = \{i\} \cap SD_{\mathcal{F}}(>'). \quad (6)$$

In order to generate $SD_{\mathcal{F}}(>)$ and $SD_{\mathcal{F}}(>')$, we use Algorithm 1 taking as input $(\mathcal{F}, >)$ and $(\mathcal{F}, >')$, respectively. Since (i) Algorithm 1 iterates over each applicant in decreasing order of priorities, and (ii) before applicant i orders $>$ and $>'$ are identical, then both executions of Algorithm 1 coincide in every iteration up to applicant i . In particular, in both executions the decision made over applicant i will be the same and (6) holds.

We will now prove that $SD_{\mathcal{F}}$ is monotone. For any priority order $> \in \mathbf{P}$ and applicant $i \in I$, we denote the ranking of applicant i in $>$ as $rank_{>}(i)$. It suffices to show that for any $>, >' \in \mathbf{P}$ such that $>_{-i} = >'_{-i}$ and $rank_{>}(i) = rank_{>'}(i) + 1$, if $i \in SD_{\mathcal{F}}(>)$ then $i \in SD_{\mathcal{F}}(>')$. We claim that

$$\{i\} \cup \{j \in SD_{\mathcal{F}}(>') : j >' i\} \subseteq SD_{\mathcal{F}}(>). \quad (7)$$

From (7) and the fact that $SD_{\mathcal{F}}(>')$ respects priorities, it follows that $i \in SD_{\mathcal{F}}(>')$. We now prove our claim (7). Let $j^* \in I$ be the individual that precedes i in the order $>$, that is, $rank_{>}(i) = rank_{>}(j^*) + 1$. Because $>_{-i} = >'_{-i}$ and $rank_{>'}(i) = rank_{>}(j^*)$, it follows that

$$>_{\{j:j>j^*\}} = >'_{\{j:j>i\}}.$$

Hence, from the fact that $SD_{\mathcal{F}}$ is lower invariant, we have that

$$\{j \in SD_{\mathcal{F}}(>') : j >' i\} = \{j \in SD_{\mathcal{F}}(>) : j > j^*\}. \quad (8)$$

Moreover, since $j^* > i$ it follows that

$$\{j \in SD_{\mathcal{F}}(>) : j > j^*\} \subseteq \{j \in SD_{\mathcal{F}}(>) : j > i\} \subseteq SD_{\mathcal{F}}(>). \quad (9)$$

Combining (8), (9) and the assumption that $i \in SD_{\mathcal{F}}(>)$ yields (7).

Finally, we show that $SD_{\mathcal{F}}$ is non-bossy. The proof is by induction on the number of agents.

For single-applicants problems, the rule $SD_{\mathcal{F}}$ is trivially non-bossy.

Assuming we have proved non-bossiness of $SD_{\mathcal{F}}$ for instances with fewer than n agents. Consider an instance with n agents and suppose we are given orderings $>$ and $>'$ that agree on the relative ordering of every pair of applicants not involving applicant i , that is, $>_{-i} = >'_{-i}$. Our goal is to establish that if applicant i 's outcome is the same for the two orderings $>$ and $>'$, then the rule $SD_{\mathcal{F}}$ selects the same element of $\mathcal{F}$ for the two orderings as well, that is, if $\{i\} \cap SD_{\mathcal{F}}(>) = \{i\} \cap SD_{\mathcal{F}}(>')$ then $SD_{\mathcal{F}}(>) = SD_{\mathcal{F}}(>')$.

Suppose there is an applicant k who is not part of any feasible selection. We can omit applicant k from the problem and obtain a smaller instance for which the result follows by induction. Thus we may assume that each applicant is part of some member of $\mathcal{F}$. We consider the following cases:

- (1) Applicant $j \neq i$ is in the top position in both $>$ and $>'$: Let $\mathcal{G}$ be the elements of $\mathcal{F}$ that includes applicant j ; remove applicant j from the problem to get an instance involving $(n-1)$ applicants; by the induction hypothesis, the rule $SD_{\mathcal{G}}$ is non-bossy on the reduced problem, and the result follows.
- (2) Applicant j is in the top position in $>$ and applicant i is in the top position in $>'$: We start with two observations: (i) the top two applicants in $>'$ are i and j ; and (ii) applicant i must be selected in $>'$ and applicant j must be selected in $>$. For a violation of non-bossiness to occur, applicant i must be selected in $>$ as well. This means there is a member of $\mathcal{F}$ that contains both i and j . We claim that the selection for the ordering $>''$, obtained from $>'$ by swapping the top two applicants, is the same as for the ordering $>'$. Now, note that applicant j is the top agent in both $>$ and $>''$; by (1) above, we can infer non-bossiness in this case.

All that remains is to prove our claim. Consider two priority orders $>, >' \in \mathbf{P}$ in which the top two applicants are swapped: in the ordering $>$, the first applicant is i and the second is j , whereas in the ordering $>'$, the first applicant is j and the second is i . If there is a member of $\mathcal{F}$ that includes both i and j , then both applicants are selected in $SD_{\mathcal{F}}(>)$ and $SD_{\mathcal{F}}(>')$ as both selections respects priorities. Then it's easy to see that $SD_{\mathcal{F}}(>) = SD_{\mathcal{F}}(>')$. $\square$

LEMMA 2. *Let φ be a selection rule that is monotone, non-bossy and lower invariant. Then there exists $\mathcal{F} \subseteq 2^I$ such that*

$$\varphi(>) = SD_{\mathcal{F}}(>) \text{ for any } > \in \mathbf{P}. \quad (10)$$

PROOF. Let

$$\mathcal{F} = \{\varphi(>) : \text{for every } > \in \mathbf{P}\}. \quad (11)$$

We will show that $SD_{\mathcal{F}}$ replicates the behavior of φ , that is, (10) holds. For the sake of contradiction, suppose that there exists an order $>$ that induces different selections under φ and $SD_{\mathcal{F}}$. Formally,

$$\varphi(>) \neq SD_{\mathcal{F}}(>).$$

Let $i \in I$ be the highest priority applicant for which $\varphi(>)$ and $SD_{\mathcal{F}}(>)$ differ. Clearly, applicant i must be selected under $SD_{\mathcal{F}}(>)$, otherwise, $SD_{\mathcal{F}}(>)$ would not respect priorities. From (11), there must exist an order $>'$ such that

$$\{i\} \cup \{j \in \varphi(>) : j > i\} \subseteq \varphi(>'). \quad (12)$$

We claim that there exists an order $>''$ that satisfies the following conditions:

- 1) Identical to $>$ in the initial positions up to applicant i , and
- 2) Induces the same selection as $>'$, that is, $\varphi(>') = \varphi(>')$.

From (12) and condition 2), it follows that individual i is selected in $\varphi(>')$. This contradicts the assumption that φ is lower invariant, notice that orders $>$ and $>''$ coincide up to individual i , but individual i is only selected under $>''$.

In order to prove the existence of order $>''$, we provide an algorithm that starts with order $>'$, and sequentially improves the priority of a single applicant, until conditions 1) and 2) are satisfied:

- Start from order $>'$. Let r be the position of applicant i in $>'$.
- For $k = 1, \dots, r$,
 - Let j be the applicant in the k -th position in $>'$.
 - Move applicant j to the k -th position, while maintaining the relative order of all the remaining applicants.

Observe that by construction, at the end of step k , the first k positions of the order generated and $>$ coincide. Hence, condition 1) immediately holds.

We now prove that, at each step, the overall selection remains unchanged. This implies condition 2) above. Fix an arbitrary step k , and let j be the individual whose priority is improved. We claim that the outcome of individual j remains unchanged. Because φ is non-bossy and the relative order of all remaining individuals is preserved, our claim implies that the overall selection must remain unchanged. All that remains is to prove our claim. From the observation above and the assumption that φ is lower-invariant, the outcome of j will be the same as in $\varphi(>)$. If j is selected in $\varphi(>)$, from (12) it must also be selected in $\varphi(>')$ so its outcome remains unchanged. If j is not selected in $\varphi(>)$, then in $\varphi(>')$ it must not be selected either. This follows as we are improving the priority of individual j and φ is monotone, hence, individual j must be weakly better. $\square$

A.2 Examples Demonstrating Independence of Our Three Axioms

EXAMPLE 3 (MONOTONE, LOWER-INVARIANT BUT BOSSY). Consider a setting where there are multiple minority groups that are disjoint, that is, each applicant can be a member of at most one minority. There are two types of seats: (i) unreserved seats that can be allocated to any applicant, and (ii) for each minority, there are reserved seats that can be allocated only to members of the minority. Let φ be the selection rule that, for any priority order $>$, generates a selection according to the following procedure:

- (1) Unreserved seats are allocated based on the priority order $>$.
- (2) For each of the (mutually exclusive) minorities, positions that are set aside are allocated to the remaining members of the minority again based on the priority order $>$.

We start by showing that no selection rule $SD_{\mathcal{F}}$ can replicate the behavior of φ . Consider an instance with three applicants $I = \{1, 2, 3\}$, one open seat and one seat reserved for members of the minority group $M = \{2, 3\}$. Let

$$1 > 2 > 3 \text{ and } 2 >' 1 >' 3.$$

From the definition of φ , we must have that $SD_{\mathcal{F}}(>) = \{1, 2\}$ so $\{1, 2\} \in \mathcal{F}$. In addition, it should be that $SD_{\mathcal{F}}(>') = \{2, 3\}$. However, selection $\{2, 3\}$ doesn't respect priorities $>'$ as $1 >' 3$ and $\{1, 2\} \in \mathcal{F}$.

From the instance above, we can verify that φ is bossy. Notice that under both orders (i) applicant 2 is selected, and (ii) applicant 1 is preferred to applicant 3. However, the overall selections differ.

EXAMPLE 4 (MONOTONE, NON-BOSSY BUT NOT LOWER-INVARIANT). Consider an instance with three applicants $I = \{1, 2, 3\}$. We define the selection rule φ as follows:

- (1) If 3 is not the worst priority applicant, then select applicant 3.
- (2) Otherwise, select applicants 1 and 2.

We begin by showing that no selection rule $SD_{\mathcal{F}}$ can replicate the behavior of φ . Let

$$1 > 2 > 3 \text{ and } 1 >' 3 >' 2.$$

By (2) have that $SD_{\mathcal{F}}(>) = \{1, 2\}$ so $\{1, 2\} \in \mathcal{F}$. In addition, by (1) we have that $SD_{\mathcal{F}}(>') = \{3\}$. However, selection $\{3\}$ doesn't respect priorities $>'$ as $1 >' 3$ and $\{1, 2\} \in \mathcal{F}$.

It is easy to see that φ is non-bossy. Notice that the two feasible selections are disjoint. Hence, if the outcome of an applicant remains unchanged, then the overall selection must be the same.

We now show that φ is monotone. On one hand, if the priority of applicant 3 is improved, then she will be selected and weakly better. On the other hand, applicants $\{1, 2\}$ are only selected when they are in the first two positions. In this situation, the only improvement feasible is to increase the priority of the applicant in the second position. Of course, this maintains the overall selection and this agent is weakly better.

Finally, we see that φ is not lower-invariant. Notice that applicant 1 has the top priority in $>$ and $>'$. However, φ only selects applicant 1 under $>$.

EXAMPLE 5 (NON-BOSSY, LOWER-INVARIANT BUT NOT MONOTONE). Consider an instance with two applicants $I = \{1, 2\}$. Let the selection rule φ be defined as follows:

- (1) If applicant 1 has the top priority, select all applicants.
- (2) Otherwise, select nobody.

We start by showing that no selection rule $SD_{\mathcal{F}}$ can replicate the behavior of φ . Let

$$1 > 2 \text{ and } 2 >' 1.$$

From (1) above, we must have that $SD_{\mathcal{F}}(>) = \{1, 2\}$ so $\{1, 2\} \in \mathcal{F}$. In addition, it should be that $SD_{\mathcal{F}}(>) = \emptyset$. Of course, selecting no applicants does not respects priorities $>'$ as applicant 2 has the first priority and $\{1, 2\} \in \mathcal{F}$.

From the instance above, we see that φ is not monotone. Notice that the priority of applicant 2 is better in $>'$ than in $>$ but it is only selected under $>$.

Clearly, φ is non-bossy. Note that φ either selects all applicants or nobody. Hence, if the status of an applicant remains unchanged so does the overall selection.

Finally, we show that φ is lower-invariant. This follows from the fact that the selection is determined by the applicant who has the top priority.

B PROOFS FROM SECTION 4

PROOF OF PROPOSITION 2. For the sake of contradiction, suppose that there exists a priority order $>$ and a feasible selection S , such that S is not priority dominated by $SD_{\mathcal{F}}(>)$. Since $SD_{\mathcal{F}}(>)$ respects priorities and doesn't dominates S , we must have that $|SD_{\mathcal{F}}(>)| = |S|$. For each selection S and integer $k \in \{1, \dots, |S|\}$, we let $S^k \subseteq S$ be the subset containing the k highest priority individuals in S , that is,

$$S_k = \{i_1^S, \dots, i_k^S\}.$$

Let $k^* \in \{1, \dots, |S|\}$ be the highest index for which

$$i_{k^*}^S > i_{k^*}^{SD_{\mathcal{F}}(>)}.$$

The existence of k^* is guaranteed as S is not dominated by $SD_{\mathcal{F}}(>)$. Consider $SD_{\mathcal{F}}(>)^{k^*-1}$ and S^{k^*} . Because $\mathcal{F}$ induces a matroid, there must exists $j \in S^{k^*} \setminus SD_{\mathcal{F}}(>)^{k^*-1}$ such that $SD_{\mathcal{F}}(>)^{k^*-1} \cup \{j\} \in \mathcal{F}$. Then $SD_{\mathcal{F}}(>)$ doesn't respects priorities as $j \notin SD_{\mathcal{F}}(>)$ but can be selected without displacing any higher priority individual in $SD_{\mathcal{F}}(>)$. □

B.1 Reserved Positions

PROOF OF THEOREM 2. We start by considering the case of hard reserves. Here, our result follows immediately by noticing that $\mathcal{F}^{hard}$ induces the matching matroid.

We now turn to maximum and maximal reserves. Let S and S' be two feasible selections of maximum cardinality, that is, $|S| = |\mathcal{P}| = |S'|$. It suffices to show that for any $i \in S \setminus S'$ there is a $j \in S' \setminus S$ such that $S \setminus \{i\} \cup \{j\}$ is a feasible selection.

Suppose there is a maximal matching M of G such that $I_M \subseteq S$ and such that $i \notin I_M$. In this case, agent i is a "free agent" in the sense that i is given a position that they are not eligible for; observe that i can be replaced by *any* agent $j \in S' \setminus S$ to obtain another feasible selection.

Thus we may assume that $i \in I_M$ for *every* maximal matching M of G with $I_M \subseteq S$.

Fix any such M . Let p be the position assigned to i under matching M . Starting with node i , build the following alternating tree:

- (1) Start with the edge (i, p) .

- (2) Add all edges of the graph connecting the agents in S' to p ; Let W be the set of agents added in this step. Note that W is non-empty as otherwise S' cannot be a superset of a maximal matching (some neighbor of p can be added).
- (3) If for some $j \in W$, $j \notin S$, then replacing i with j results in a feasible selection.
- (4) Every member of $W \cap S$ should be matched to a position in M : otherwise, we can find a maximal matching in which i is not matched to p . So the tree can be extended with the corresponding matching edges.

Repeating steps (2) – (4) and noting that the process has to end, and it can only end in step (3), observe that we eventually find a j such that replacing i with j results in a feasible selection. $\square$

B.1.1 Soft Reserves.

PROOF OF PROPOSITION 3. We start by proving the hardness result for maximal reserves. We give a reduction from the problem “minimum maximal matching” (MMM) of deciding whether a bipartite graph G admits a maximal matching of size at most k . This problem was shown to be NP-complete by [?]. Consider an instance of MMM, that is, a bipartite graph $G = (U, V, E)$ and a positive integer k . We generate an instance $I(G)$ of the selection problem under maximal reserves as follows. We set $\mathcal{P} = V$ and $\mathcal{I} = \mathcal{I}^1 \cup \mathcal{I}^2$, where $\mathcal{I}^1 = U$ and $\mathcal{I}^2$ consists of $|V|$ individuals who are not connected to any position. Let the priority order $>$ place all individuals in $\mathcal{I}^2$ ahead of all individuals in $\mathcal{I}^1$. Then G admits a maximal matching of size at most k if and only if there is a selection in $\mathcal{F}^{\text{maximal}}(G)$ containing the first $|V| - k$ individuals from $>$.

We now prove the positive result for maximum reserves. Consider the compatibility graph $G = G(\mathcal{I}, \mathcal{P}, E)$. By the Gallai-Edmonds (actually Dulmage-Mendelsohn) decomposition, the set of applicants $\mathcal{I}$ can be partitioned into three parts $\mathcal{I}_o, \mathcal{I}_e$ and $\mathcal{I}_u$ and similarly, the positions can be partitioned into three parts $\mathcal{P}_u, \mathcal{P}_e$, and $\mathcal{P}_o$ such that every maximum cardinality matching in G has the following structure:

- There is a perfect matching of the agents in $\mathcal{I}_e$ with the positions in $\mathcal{P}_e$.
- All the positions in $\mathcal{P}_o$ are assigned to agents in $\mathcal{I}_u$, and it is possible to find a matching that omits any specified agent in $\mathcal{I}_u$.
- All the agents in $\mathcal{I}_o$ are assigned to some position in $\mathcal{P}_u$, and it is possible to find a matching that omits any specified position in $\mathcal{P}_u$.

In particular, the agents in $\mathcal{I}_e$ and $\mathcal{I}_o$ are matched in *all* maximum cardinality matchings, and the agents do not consume any position in $\mathcal{P}_o$; and all positions in $\mathcal{P}_o$ are allocated to some subset of agents in $\mathcal{I}_u$.

Moreover, this decomposition can be found efficiently using, for example, an algorithm to find a maximum cardinality matching.

Given this result, an priority dominant selection under maximum reserves can be computed as follows: match all agents in $\mathcal{I}_e$ and $\mathcal{I}_o$; find a maximum-weight matching of the positions in $\mathcal{P}_o$ to the agents in $\mathcal{I}_u$; finally, pick, in decreasing order of priority, as many of the unassigned agents from $\mathcal{I}_u$ as the number of unassigned positions in $\mathcal{P}_u$. $\square$

B.2 Quotas

PROOF OF PROPOSITION 4. This follows from Lemmas 3 and 4, which each prove the result for a special case of the general quotas problem. $\square$

LEMMA 3. Suppose there is a trait $0 \in T$, such that $\mathcal{I}_0 = \mathcal{I}$ and $\ell_0 = q$. Additionally, for every other trait $t \in T \setminus \{0\}$, $\ell_t = 0$. If the collection of feasible selections is defined as in (5), then

- *determining whether there exists a feasible selection is NP-complete, and*
- *the selection that respects priorities might not priority dominate all other feasible selections.*

PROOF OF LEMMA 3. We first show that determining whether there exists a feasible selection is NP-complete, via a reduction from the Independent Set Problem. An instance of the Independent Set Problem is given by a graph $G = (V, E)$ and an integer k . An *independent set* of graph G is a set S of vertices such that no two vertices in S are adjacent. The goal is to determine whether there exists an independent set containing at least k elements. Given an instance of the Independent Set Problem, we construct an instance of the Selection Problem with upper quotas as follows: create one individual for each vertex, and one trait (with only two individuals) for each edge. Let each trait have an upper quota of 1, and let the overall lower quota be k . Then for any selection, the upper quotas are satisfied if and only if the set of vertices associated with the selected individuals is independent. Therefore, deciding whether there exists a feasible selection is equivalent to determining whether there is an independent set of size at least k .

We conclude by noting that Example 2 gives an instance satisfying the conditions of Lemma 3 in which the selection that respects priorities does not priority dominate all other feasible selections. $\square$

LEMMA 4. *Suppose there is a trait $0 \in T$, such that $I_0 = \mathcal{I}$ and $u_0 = q$. Additionally, for every other trait $t \in T \setminus \{0\}$, $u_t = |\mathcal{I}|$. If the collection of feasible selections is defined as in (5), then*

- *determining whether there exists a feasible selection is NP-complete, and*
- *the selection that respects priorities might not priority dominate all other feasible selections.*

PROOF OF LEMMA 4. First, we show that determining whether there exists a feasible selection is NP-complete, using a reduction from the Set Cover Problem. An instance of the Set Cover Problem is given by a ground set $\mathcal{U}$, a collection $\mathcal{S}$ of subsets of $\mathcal{U}$, and an integer k . A *cover* for $\mathcal{U}$ is a subcollection $C \subseteq \mathcal{S}$ of sets whose union is $\mathcal{U}$. The objective is to determine whether there exists a set cover $C \subseteq \mathcal{S}$ for $\mathcal{U}$ containing at most k elements. Given an instance of the Set Cover Problem, we construct an instance of the Selection Problem with lower quotas as follows: create one trait for each element in $\mathcal{U}$, and one individual for each set in $\mathcal{S}$ (having traits identified by the set). There is a lower quota of 1 for each trait, and an overall upper quota equal to k . Then a selection of individuals satisfies all minimum quotas if and only if the corresponding subsets cover $\mathcal{U}$. Therefore, determining whether there exists a feasible selection is equivalent to determining whether there is a set cover with at most k elements.

Secondly, we give an instance where the selection that respects priorities does not priority dominate all other feasible selections. Consider an instance with four individuals 1, 2, 3, 4, and three traits $T = \{0, A, B\}$. Let $\mathcal{I}_A = \{2, 4\}$, $\mathcal{I}_B = \{3, 4\}$, $\ell_A = 1 = \ell_B$ and $u_0 = 2$. Suppose that $1 > 2 > 3 > 4$. Then the feasible selection that respects priorities is $\{1, 4\}$. However, this selection doesn't priority dominate selection $\{2, 3\}$. $\square$

PROOF OF PROPOSITION 5. We can assume without loss of generality that there exists $\tau(1) \leq \dots \leq \tau(|T|)$ such that

$$\mathcal{I}_t = \{i_1, \dots, i_{\tau(t)}\}. \quad (13)$$

From Proposition 2, it suffices to show that $\mathbf{M} = (I, J(\mathcal{F}))$ is a matroid. Let $I, I' \in J(\mathcal{F})$ be such that $|I| < |I'|$. We must show that there exists $i \in I' \setminus I$ such that $I \cup \{i\} \in J(\mathcal{F})$. For the sake of contradiction, suppose that for every $i \in I' \setminus I$,

$$I \cup \{i\} \notin J(\mathcal{F}). \quad (14)$$

We claim that $|I| \geq |I'|$, a contradiction. Let i be the applicant in $I' \setminus I$ with the highest index. Since (14) holds, there must exist $\mathcal{I}_t$ such that $i \in \mathcal{I}_t$ and $|I \cap \mathcal{I}_t| = u_t$. Moreover, as $I' \in J(\mathcal{F})$ it must be

that

$$|I' \cap \mathcal{I}_t| \leq u_t. \quad (15)$$

Because i is the applicant in $I' \setminus I$ with the highest index, then $I' \setminus I \subseteq J(\mathcal{F})_t$. This implies that $|I| \geq |I'|$.

To show the necessity, suppose that $\mathcal{I}_T$ does not form a hierarchy. Then there exist traits t_1 and t_2 and individuals i_1, i_2, i_3 such that $i_1 \in \mathcal{I}_{t_1} \setminus \mathcal{I}_{t_2}$, $i_2 \in \mathcal{I}_{t_2} \setminus \mathcal{I}_{t_1}$, and $i_3 \in \mathcal{I}_{t_1} \cap \mathcal{I}_{t_2}$.

Set all lower quotas to be zero, set an upper quota of 1 for each of traits t_1 and t_2 , and a (vacuous) upper quota of $|\mathcal{I}|$ for all other traits. Let $3 > 2 > 1$ be the three highest-ranked individuals according to $>$. Then the sequential dictator rule selects individual 3, along with everyone who does not have traits t_1 or t_2 . However, it is also feasible to select individuals 1 and 2 along with everyone who does not have traits t_1 or t_2 . This selection is not priority dominated by the selection of the sequential dictator rule. By Proposition 1, it follows that there is no feasible selection that priority dominates all others.

□