Distributional Preferences for Market Design*

Federico Echenique† Teddy Mekonnen‡ M. Bumin Yenmez§

August 21, 2026

Abstract

Institutions selecting students, employees, or members value both who is selected and the resulting group's composition. We study "distributional" preferences over group composition, and identify an upper-bound property and two exchange properties. With the upper-bound property, the exchange properties are necessary and sufficient for two results: the greedy rule is the unique choice rule that is non-wasteful, distributionally maximal, and free of justified envy; it is also path independent. In matching markets, deferred acceptance is the unique mechanism satisfying the three axioms, individual rationality, and strategy-proofness. Our framework accommodates intersectional identities and subsumes models based on reserves and matroids.

Keywords: Distributional preferences, market design, deferred acceptance, path independence.

JEL Classifications: D47, C78, I24, I28.

*We are grateful to Junya Murayama for excellent research assistance, and to Battal Doğan and Koji Yokote for helpful discussions. We also wish to thank multiple seminar audiences for their feedback, and the attendees of EC'26. None of the authors have sources of research support, financial relationships, or other potential conflicts of interest to report.

†Department of Economics, UC Berkeley. Contact: fede@econ.berkeley.edu

‡Department of Economics, Brown University. Contact: mekonnen@brown.edu

§Department of Economics, Washington University in St. Louis, Durham University (UK), and Özyeğin University, Turkey. Contact: bumin@wustl.edu.

1. Introduction

Many institutions face multidimensional selection problems. They care about the quality of the individuals they select, but they also care about the composition of the selected group. A university may want to admit strong students while also forming a class that reflects socioeconomic, racial, geographic, or curricular diversity. A hospital may want highly qualified residents while also valuing a cohort with different backgrounds, specialties, and service interests. A firm may want talented employees while also caring about the diversity of skills and experiences in teams. Thus, a central question in such settings is: How should an institution select individuals according to priority or merit while also pursuing distributional objectives over the group as a whole?

In this paper, we answer this question in a canonical school-choice environment. While we focus on school choice, the same design problem arises in labor markets, public-sector hiring, and team formation. We use the school-choice framing because these multidimensional objectives are well documented and prevalent (Alves and Willie, 1987). Schools naturally wish to admit students with the highest priorities, which may be determined by standardized test performance, proximity of residence, and sibling attendance. Yet, they may also wish—and are often legally required—to guarantee access for various identity groups based on gender, race, veteran status, socioeconomic status, and other traits.

Formally, this paper develops a framework that takes as primitive a school's preference over sets of students, which we refer to as the school's *distributional preference*. Such a preference can capture the school's desire to attain a particular composition of identity groups or other relevant traits in the admitted cohort. Importantly, the framework allows us to articulate complex distributional objectives directly, including objectives involving intersectional or overlapping identity groups. Given this preference and a priority ranking over individual students, we ask when the school can choose a cohort that is distributionally desirable, respects priorities, behaves consistently across applicant pools, and can be embedded in a centralized matching mechanism.

This preference-based approach differs from much of the market-design literature on distributional objectives, which often starts from particular admissions policies or choice rules and asks when those policies are consistent with a school's distributional objectives. The focus is often on quotas, reserves, or floors-and-ceilings policies. In general, such policies specify a minimum, maximum, or reserved number of seats to be allocated to students from each identity category. Hence, these policies give a simple language for formalizing the school's distributional objective and often work well when each student belongs to a single identity category.

When a student belongs to multiple categories, however, these policies may fail to articulate the school's distributional objective over intersectional identities. Counting such a student toward only one of her identities or allowing her to count toward multiple identities both have shortcomings, as the following example illustrates.

Example 1. Consider a school with two seats. It evaluates the composition of its student body based on socioeconomic status (SES) and minority status. There are four applicants, $\mathcal{S} = \{s_1, s_2, s_3, s_4\}$, prioritized in descending order so that s_1 has the highest priority and s_4 has the lowest. Student s_1 is both a minority and a low-SES student; s_2 is neither a minority nor a low-SES student; s_3 is a low-SES student; and s_4 is a minority student.

The school has two reserve categories: one for low-SES students and one for minority students. A reserve category is filled by the highest-priority eligible applicant who has not already been admitted. Once all reserve categories are filled, any remaining seats are assigned according to the priority ranking.

Suppose first that an intersectional student can count toward at most one reserve category. If the SES reserve is processed first, then s_1 fills the SES reserve; the minority reserve then goes to s_4, yielding $\{s_1, s_4\}$ as the admitted pair. If instead the minority reserve is processed first, then s_1 fills the minority reserve; the SES reserve then goes to s_3, yielding $\{s_1, s_3\}$. Thus, the outcome depends on the order in which the reserve categories are processed. Yet the canonical description of reserve policies does not determine which reserve category should be processed first.

Next, suppose intersectional students count toward both reserve categories simultaneously. In this case, regardless of which reserve category is processed first, s_1 fulfills both reserves. The remaining seat goes to the highest-priority remaining student, s_2, yielding the set $\{s_1, s_2\}$.

However, this violates path independence. If s_1 is removed from the applicant pool, then, in the reduced pool $\{s_2, s_3, s_4\}$, the highest-priority students who qualify for the SES and minority reserves are s_3 and s_4, respectively. Thus, the school selects $\{s_3, s_4\}$ and rejects s_2, who qualifies for neither reserve category. Hence, s_2 is admitted from the larger applicant pool but rejected after another applicant is removed. This inconsistency is normatively undesirable in a single-school choice problem and can undermine implementation in centralized matching markets.

The source of these shortcomings is not simply that the school values more than one category. Rather, the reserve policy does not articulate how different categories should be jointly evaluated when a student belongs to more than one of them. The basic underlying issue is that the focus on admission policies and choice rules leaves the school's underlying

distributional objective implicit. In contrast, our paper starts from an explicit model of the school's preferences over cohorts and studies the properties of admission policies induced by those preferences.

To that end, we first consider choice by a single school. Given an applicant pool, the school's distributional preference identifies a *frontier*: the collection of non-wasteful cohorts that are undominated according to the distributional preference. We then define the *distributional choice rule*, which is a greedy algorithm that considers students in priority order and admits a student whenever doing so remains compatible with eventually reaching some set on the frontier. Thus, the distributional preference determines which final cohorts lie on the frontier, while the priority ranking selects among those cohorts.

By construction, the distributional choice rule is non-wasteful, distributionally maximal, and free of justified envy. These three normative axioms capture the basic requirements that the school fill its available seats whenever possible, select an undominated cohort according to its distributional preference, and depart from the priority ranking only when doing so is justified by that preference. However, without imposing further structure on the distributional preference, these axioms neither uniquely identify the distributional choice rule nor ensure that the rule behaves consistently across applicant pools.

We therefore identify three structural properties of the distributional preference. The *upper bound property* imposes a "lattice-like" structure on the preference by requiring incomparable, non-wasteful cohorts to have a distributional upper bound. As a result, every cohort on the frontier is not merely maximal but optimal among the non-wasteful cohorts. The *maximizer property* then requires optimal cohorts within a fixed applicant pool to be connected by value-preserving one-for-one exchanges. Specifically, given two optimal cohorts, there must exist one student belonging to each cohort but not the other, whose exchange leaves both cohorts distributionally indifferent to their corresponding original cohorts. Finally, the *coherence property* imposes a corresponding exchange condition across applicant pools when students are added or removed.

The upper-bound and maximizer properties deliver our main single-school results. Our first set of results concerns choice from a fixed applicant pool. Given the upper-bound property, the maximizer property is necessary and sufficient for the distributional choice rule to be the unique choice rule that is non-wasteful, distributionally maximal, and free of justified envy. Moreover, any departure from the distributional choice rule must either leave some seats empty, select a distributionally dominated cohort, or select a cohort that is inferior according to the priority ranking. We also provide a revealed-priority foundation for the rule. When the priority ranking is not taken as primitive but rather is inferred from the school's choices, non-wastefulness, distributional maximality, and acyclicity of revealed priority com-

parisons imply that the choice rule must be the distributional choice rule for some priority ranking.

Our second set of results concerns the consistency of choice across applicant pools. Given the upper-bound property, the maximizer and coherence properties are jointly necessary and sufficient for the distributional choice rule to satisfy path independence for every priority ranking.

By making explicit the relevant comparisons over cohorts, we may address the choices in Example 1. Suppose that $\{s_1, s_3\}$ belongs to the frontier and weakly dominates both $\{s_1, s_4\}$ and $\{s_3, s_4\}$. Such comparisons are natural if the school places relatively greater weight on low-SES representation. The decisive remaining comparison is between $\{s_1, s_3\}$ and $\{s_1, s_2\}$. If $\{s_1, s_3\} \succ \{s_1, s_2\}$, then $\{s_1, s_2\}$ does not belong to the frontier. The distributional choice rule therefore admits s_1, rejects s_2, and admits s_3, yielding $\{s_1, s_3\}$. If instead $\{s_1, s_3\} \sim \{s_1, s_2\}$, then both cohorts belong to the frontier. Because s_2 has higher priority than s_3, the rule selects $\{s_1, s_2\}$. In this way, the preference-based approach replaces an exogenous convention about how to count an intersectional student with explicit comparisons over the resulting cohorts.

We next extend the analysis to centralized matching markets. We show that when every school's distributional preference satisfies the three structural properties, deferred acceptance based on the induced distributional choice rules is the unique mechanism that is individually rational, non-wasteful, distributionally maximal, strategy-proof, and free of justified envy. Thus, the same structural properties that ensure a well-behaved choice rule for a single school also support deferred-acceptance implementation in a market with multiple schools.

The framework both rationalizes familiar policies and extends beyond them. Standard and overlapping reserve rules arise from distributional preferences based on partition and transversal matroids, respectively. The framework also includes a graded formulation of floors and ceilings, represented by distributional preferences that rank violations of lower and upper bounds on the representation of each identity group. It further accommodates objectives that are not naturally expressed through reserves, including score-based preferences that assign different values to different combinations of intersectional identities, preferences that value linearly non-redundant combinations of traits, and preferences that value linkages across identity groups. Thus, the paper provides a unified preference foundation for distributional policies whose induced choice rules respect priorities, are non-wasteful, and support deferred-acceptance implementation.

1.1. Related Literature

The market-design literature on distributional policies has largely developed through reserves, quotas, and related constraints. Standard reserve rules were introduced and studied by Hafalir et al. (2013), floors-and-ceilings rules by Ehlers et al. (2014), and related axiomatic foundations were provided by Echenique and Yenmez (2015). Reserve policies have been applied in Brazil (Aygün and Bó, 2021), Chile (Correa et al., 2019), and India (Aygün and Turhan, 2017; Sönmez and Yenmez, 2022). Further work studies regional constraints (Kamada and Kojima, 2015, 2017, 2018), reserve processing order (Dur et al., 2018, 2020), and alternative systems of quotas, transfers, and reserve adjustment (Westkamp, 2013; Fragiadakis and Troyan, 2017; Aygün and Turhan, 2020; Kumano and Kurino, 2022; Doğan and Erdil, 2025). Imamura (2025) provides a recent axiomatic analysis of the trade-off between meritocracy and diversity.

Much of this literature starts with reserve categories, eligibility requirements, distributional targets, or quota constraints. Related preference-based work studies the rationalizability of affirmative-action choice rules and the implications of multidimensional identities (Çelebi, 2024). Our analysis takes as primitive a possibly incomplete distributional preference over cohorts and keeps it separate from the priority ranking over individual students. This formulation allows us to characterize when the two can be combined through a uniquely characterized and path-independent priority-based choice rule.

The paper most closely related to ours is Hafalir et al. (2026). They represent an institution's distributional objective by a real-valued function of type distributions and study a distribution-first, merit-second choice rule. Under ordinal concavity, they show that their distribution-conscious rule maximizes merit among distributionally optimal cohorts, can be computed in polynomial time, and satisfies path independence; their analysis also develops related matroid characterizations. Our analysis instead imposes non-wastefulness, allows the primitive distributional preference to be incomplete, and interprets the ranking over individuals as a system of priority rights. Conditional on the upper-bound property, we provide necessary and sufficient conditions for the greedy priority refinement to be uniquely characterized by non-wastefulness, distributional maximality, and freedom from justified envy, and to satisfy path independence under every priority ranking. When restricted to fixed-cardinality cohorts, the ordinal concavity condition studied by Hafalir et al. (2026) induces distributional preferences satisfying our structural properties.

A complementary rationalization result is provided by Yokote et al. (2025), who show that a choice rule is path independent if and only if it can be rationalized by an ordinally concave utility function. Their result asks whether a given choice rule admits an appropriate utility representation. We instead hold a primitive distributional preference fixed and characterize

when combining it with any priority ranking through the greedy rule yields path-independent choice. More broadly, Kojima et al. (2018) use $M^\natural$-concavity to design matching mechanisms under distributional constraints. Echenique et al. (2024) take a different but related approach, using a generalization of majorization to compare diversity relative to distributional targets and to axiomatize the associated choice rule.

Our treatment of overlapping identities is closely related to the literature on horizontal and overlapping reserves (Sönmez and Yenmez, 2020, 2022). These policies assign each candidate to at most one eligible reserve category while filling as many reserved positions as possible. We show that the resulting rule can be rationalized by a distributional preference based on the rank function of a transversal matroid; standard reserves arise analogously from partition matroids. More broadly, Fleiner (2003) shows how greedy choice under matroidal constraints can be combined with deferred acceptance, while Bonet et al. (2024) introduce *explainability* for affirmative-action policies and characterize outcome-based greedy rules. Whereas these approaches begin with a reserve system, feasible-outcome family, or matroidal constraint, we characterize when such a structure is generated by a primitive distributional preference. Conditional on the upper-bound property, the maximizer property holds if and only if each frontier forms the bases of a matroid, while the coherence property identifies when these frontiers are consistent across applicant pools.

We also contribute to the characterization of deferred-acceptance mechanisms. Alcalde and Barberà (1994) characterize deferred acceptance when priorities are exogenously given; alternative characterizations appear in Kojima and Manea (2010) and Ehlers and Klaus (2016). Abdulkadiroğlu and Grigoryan (2025) study deferred acceptance coupled with a reserves-and-quotas choice rule that minimizes priority violations subject to distributional targets and other axioms. We instead begin from primitive distributional preferences and use the meta-characterization of Doğan et al. (2025) to lift our choice-rule characterization to the mechanism level. This yields the unique deferred-acceptance mechanism satisfying the corresponding requirements of non-wastefulness, distributional maximality, and freedom from justified envy together with individual rationality and strategy-proofness. We also provide a revealed-priority foundation for the greedy distributional choice rule when the priority ranking is not taken as primitive.

Finally, the two-stage structure of the distributional choice rule is related to sequentially rationalizable choice. Manzini and Mariotti (2007) study procedures in which ordered rationales successively refine the set of alternatives. In our setting, the distributional preference first determines the frontier of distributionally maximal non-wasteful cohorts, and the priority ranking then selects among those cohorts. The two criteria represent an institutional objective and individual priority rights, respectively. Our results characterize when their or-

dered application yields a unique path-independent choice rule and supports implementation through deferred acceptance.

The rest of the paper is organized as follows: Section 2 introduces the model. Section 3 develops the structural properties and establishes the single-institution results. Section 4 extends the analysis to centralized markets. Section 5 presents applications, and Section 6 concludes. All proofs are in the Appendix.

2. Model

2.1. Preliminaries

We denote the set of non-negative integers by $\mathbb{Z}_+$ and the set of positive integers by $\mathbb{Z}_{++}$. For any finite set $\mathcal{E}$, let $2^{\mathcal{E}}$ denote its power set and $|\mathcal{E}|$ its cardinality. A *strict linear order* on $\mathcal{E}$ is a transitive, asymmetric, and total binary relation.

A *preorder* $\succsim$ on $2^{\mathcal{E}}$ is a reflexive and transitive binary relation. For $E, E' \subseteq \mathcal{E}$, we write $E \succ E'$ if $E \succsim E'$ and $E' \not\succsim E$, and $E \sim E'$ if $E \succsim E'$ and $E' \succsim E$. Sets E and E' are *incomparable* under $\succsim$ if neither $E \succsim E'$ nor $E' \succsim E$.

A *matroid* is a pair $(\mathcal{E}, \mathcal{I})$, where $\mathcal{E}$ is a finite *ground set* and $\mathcal{I}$ is a collection of subsets of $\mathcal{E}$, called *independent sets*, satisfying

(I1) *Non-emptiness*: $\mathcal{I}$ is non-empty.
(I2) *Hereditary*: If $I \in \mathcal{I}$ and $I' \subseteq I$, then $I' \in \mathcal{I}$.
(I3) *Augmentation*: If $I, I' \in \mathcal{I}$ and $|I| > |I'|$, then there exists $e \in I \setminus I'$ such that $I' \cup \{e\} \in \mathcal{I}$.

An independent set that is maximal under set inclusion is called a *basis*, i.e., $B \in \mathcal{I}$ is a basis if no $I \in \mathcal{I}$ exists such that $I \supsetneq B$.

2.2. Setup

We begin with a single school and a finite, nonempty set of students $\mathcal{S}$. The school has capacity $q \in \mathbb{Z}_{++}$, a *priority ranking* π over students, and a *distributional preference* $\succsim$ over subsets of students. The priority ranking is a strict linear order on $\mathcal{S}$ and represents the school's ordering over individual students, often based on a student's test performance, proximity of residence, or sibling attendance. The distributional preference is a preorder on $2^{\mathcal{S}}$ and represents the school's preferences over the composition of student groups. The tuple $(\mathcal{S}, \pi, \succsim, q)$ thus constitutes the primitives of the single-school model. We extend this model to two-sided matching markets with multiple schools in Section 4.

The priority ranking π may not be related to the distributional preference $\succsim$ in any systematic way. One useful interpretation is that each student has a collection of demographic traits, and that $\succsim$ compares subsets of students according to their aggregate trait composition. Under such a preference, two sets $S, S' \subseteq \mathcal{S}$ may satisfy $S \sim S'$ whenever they have the same trait composition, even if they contain different students.

Because we do not require $\succsim$ to be complete (we only assume that it is a preorder), it may not compare all pairs of sets. For instance, a set consisting only of low-income students and a set consisting only of minority students may be incomparable if the school does not rank one identity category against the other. At the same time, $\succsim$ can compare sets involving overlapping identities, as in the motivating example in the Introduction. We illustrate the generality of the framework using the following examples (the examples are fleshed out as applications of our methods in Section 5).

Example 2. Consider the four students in Example 1. Suppose each student s receives a distributional score $f(s)$ given by

$$f(s) = \begin{cases} 2 & \text{if } s \text{ is both minority and low-SES,} \\ 1 & \text{if } s \text{ is minority but not low-SES,} \\ 0.5 & \text{if } s \text{ is low-SES but not minority,} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the school assigns different scores to different combinations of traits. Among students with exactly one protected trait, the school values minority students more than it does students with low-SES. At the same time, it also places a premium on intersectionality with $f(s_3)+f(s_4) < f(s_1)$, so that a student with both traits has a higher score than the combined value of the two traits when held by separate students.

We derive a distributional preference over $2^{\mathcal{S}}$ from the student-level distributional scores as follows: for any $S, S' \subseteq \mathcal{S}$, let $S \succsim S'$ if and only if $\sum_{s \in S} f(s) \geq \sum_{s \in S'} f(s)$. The induced distributional preference is complete, so any two subsets are comparable. For example, $\{s_1, s_2\} \sim \{s_1\} \succ \{s_3, s_4\}$.

Example 3. Once again, consider the four students in Example 1. Let $P_{SE} = \{s_1, s_3\}$ be the set of students eligible for a protected category based on socioeconomic status, and let $P_M = \{s_1, s_4\}$ be the set of students eligible for a protected category based on minority status. For any $S \subseteq \mathcal{S}$, let $r(S)$ denote the maximum number of protected categories that can be represented by assigning students in S to categories for which they are eligible, with each student assigned to at most one protected category.

We then derive a distributional preference over $2^{\mathcal{S}}$ from $r(\cdot)$: for any $S, S' \subseteq \mathcal{S}$, let $S \succsim S'$ if and only if $|S| = |S'|$ and $r(S) \geq r(S')$. Such a preference reflects a school that cares about how many protected categories are represented in its student body, perhaps because it receives additional funding for each represented category, but does not compare student bodies of unequal size. In this case, $\{s_3, s_4\} \succ \{s_1, s_2\}$, but neither set is comparable to $\{s_1\}$.

Example 4. Consider the four students in Example 1. Using minority status as the first coordinate and low-SES status as the second, define each student's type vector as

$$\tau(s_1) = (1, 1), \qquad \tau(s_2) = (0, 0), \qquad \tau(s_3) = (0, 1), \qquad \tau(s_4) = (1, 0).$$

Suppose the school evaluates cohorts based on the span of the cohort's type vectors. A form of preference for increased diversity as measured by the number of dimensions achieved by the student body.

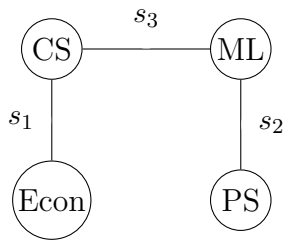
The pair $\{s_1, s_2\}$ has a span of dimension one: s_1 contributes the type vector $(1, 1)$, while s_2 contributes the zero vector. By contrast, $\{s_3, s_4\}$ has a span of dimension two because the vectors $(0, 1)$ and $(1, 0)$ are linearly independent over $\mathbb{R}$. Hence, $\{s_3, s_4\} \succ \{s_1, s_2\}$.

Example 5. Consider a PhD program in computational social science that offers four fields: Computer Science (CS), Machine Learning (ML), Political Science (PS), and Economics (Econ). It wants to admit a cohort that connects as many of these four fields as possible.

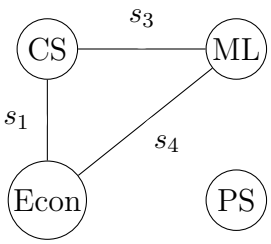
There are four candidates: s_1 with interests in CS and Econ (algorithmic game theory); s_2 with interests in ML and PS (natural language processing for political text); s_3 with interests in CS and ML (algorithmic foundations of machine learning); and s_4 with interests in Econ and ML (econometrics and machine learning).

The cohort $\{s_1, s_2, s_3\}$ connects all four fields, while $\{s_1, s_3, s_4\}$ connects only three of them, as depicted by the graphs in Figure 1. Hence, the PhD program's distributional preference satisfies $\{s_1, s_2, s_3\} \succ \{s_1, s_3, s_4\}$. Importantly, every student in this example is intersectional. Yet the program does not evaluate students only by the number or identity of their fields of interest; it also values the pattern of linkages represented by the cohort.

In our model, $\succsim$ is taken as a primitive, but this does not, in practice, require a school or a clearinghouse to list preferences over all subsets of students. As shown in Example 2, distributional preferences may be derived from student-level distributional scores. More broadly, we show in Section 5 that distributional preferences often admit compact representations through diversity indices, penalty functions, or matroid rank functions. These representations evaluate sets of students through summary statistics of their trait composition, making preferences interpretable and operational.



(i) Cohort $\{s_1, s_2, s_3\}$



(ii) Cohort $\{s_1, s_3, s_4\}$

Figure 1: Each admitted student contributes an undirected edge linking the two fields they bridge. The cohort $\{s_1, s_2, s_3\}$ connects all four fields, whereas $\{s_1, s_3, s_4\}$ connects only three, leaving Political Science isolated.

2.3. Frontier Sets and Choice-Rule Desiderata

A school can only admit a subset of the students that apply to it, and cannot do so in a way that violates its capacity constraint. Thus, the school's distributional goals are constrained by both its applicant pool and its capacity. To formalize these constraints, we introduce two concepts that will be useful in our analysis.

For any set of students $S \subseteq \mathcal{S}$, let

$$\mathrm{NW}(S) \coloneqq \{S' \subseteq S : |S'| = \min\{q, |S|\}\}$$

denote the collection of *non-wasteful* subsets of S: the subsets of S that contain as many students as possible, up to the school's capacity. Let

$$\mathrm{F}(S) \coloneqq \{S' \in \mathrm{NW}(S) : \nexists S'' \in \mathrm{NW}(S) \text{ such that } S'' \succ S'\}$$

denote the *frontier* of S. The frontier consists of all non-wasteful subsets that are undominated, or maximal, with respect to the distributional preference. For any set $S \subseteq \mathcal{S}$, $\mathrm{NW}(S)$ is a finite, non-empty collection of subsets of students. Since every finite nonempty preorder has at least one maximal element, $\mathrm{F}(S)$ is nonempty.[1]

We now turn to our main question: **How should a school allocate its limited number of seats in a way that balances its distributional goals and priority considerations?** We address this question by imposing a set of normative choice desiderata, formalized as axioms, on how a school chooses which students to admit.

Formally, the school's *choice rule* is a function $\mathrm{Ch} : 2^{\mathcal{S}} \to 2^{\mathcal{S}}$ that maps each set of applicants $S \subseteq \mathcal{S}$ into a set of admitted students $\mathrm{Ch}(S)$ such that

[1] Note that even when $S = \varnothing$, $\mathrm{NW}(S) = \mathrm{F}(S) = \{\varnothing\}$ are nonempty sets.

(*a*) $\mathrm{Ch}(S) \subseteq S$, and

(*b*) $|\mathrm{Ch}(S)| \leq q$.

Our first axiom imposes a weak efficiency constraint: no seat may remain empty while an eligible applicant is available.

Axiom 1. A choice rule Ch is **non-wasteful** if, for every set $S \subseteq \mathcal{S}$,

$$|\mathrm{Ch}(S)| = \min\{q, |S|\}.$$

Non-wastefulness precludes choice rules that may purposefully leave seats vacant, such as quota-based rules that admit students of a given trait only up to a specified cap. Nevertheless, the axiom is natural in our primary applications, such as public school choice and civil service hiring, which operate under statutory or funding constraints that require seats to be filled whenever acceptable applicants are available.[2]

For each set S, non-wastefulness requires $\mathrm{Ch}(S)$ to belong to $\mathrm{NW}(S)$. However, when the applicant size $|S|$ exceeds the school's capacity q, $\mathrm{NW}(S)$ is not a singleton. Therefore, non-wastefulness alone does not generally determine which students should be admitted. In such cases, our second axiom requires the school to select an undominated, or maximal, set with respect to its distributional objective while holding fixed the number of admitted students.

Axiom 2. A choice rule Ch is **distributionally maximal** if, for every set $S \subseteq \mathcal{S}$ and $S' \subseteq S$,

$$|S'| = |\mathrm{Ch}(S)| \implies S' \not\succ \mathrm{Ch}(S).$$

The first two axioms are jointly equivalent to requiring the chosen set to be in the frontier.

Lemma 1. *A choice rule* Ch *is non-wasteful and distributionally maximal if and only if* $\mathrm{Ch}(S) \in \mathrm{F}(S)$ *for every set* $S \subseteq \mathcal{S}$.

The lemma establishes that the frontier is the relevant benchmark for evaluating non-wasteful and distributionally maximal choice rules. However, the school's distributional preference $\succsim$ may leave many subsets indifferent or incomparable, so $\mathrm{F}(S)$ may contain multiple sets. Our third axiom requires the priority ranking to guide the school's choice among such sets. It imposes a fairness requirement: the school cannot reject a higher-priority student in favor of a lower-priority admitted student whenever admitting the higher-priority student instead would weakly improve the distributional outcome.

[2]Unacceptable applicants can be removed from $\mathcal{S}$ without affecting the analysis.

Axiom 3. A choice rule Ch is **free of justified envy** if, for every $S \subseteq \mathcal{S}$, $s \in \text{Ch}(S)$, and $s' \in S \setminus \text{Ch}(S)$,

$$\Big(\text{Ch}(S) \setminus \{s\}\Big) \cup \{s'\} \succsim \text{Ch}(S) \implies s \; \pi \; s'.$$

The three axioms above are the central desiderata for the single-school choice problem: the school must fill its seats guided by its distributional objective while also respecting priorities whenever doing so is compatible with the distributional objective. Our single-school characterization in Section 3 focuses primarily on these three requirements.

Our fourth, and final, axiom plays a limited role in single-school choice problems, but it is central to our market-design analysis in Section 4. The axiom imposes a consistency condition: processing the applicant pool all at once or in batches should not change the admitted set.

Axiom 4. A choice rule Ch is **path independent** if, for every $S, S' \subseteq \mathcal{S}$,

$$\text{Ch}(S \cup S') = \text{Ch}\left(\text{Ch}(S) \cup S'\right).$$

Path independence is a basic rationality axiom, and thus normatively appealing even in single-school choice problems (Echenique and Yenmez, 2015); it also underlies existence and structural results in matching markets (Blair, 1988; Alkan and Gale, 2003; Fleiner, 2003), making it a central axiom in two-sided market design problems.

2.4. Distributional Choice Rule

We now introduce the greedy choice rule of interest, which we refer to as the *distributional choice rule* and denote by Ch^{π}.

Distributional Choice Rule:

Input: A set of students $S \subseteq \mathcal{S}$. Let $k \coloneqq |S|$. Label students in S such that $s_1 \; \pi \; s_2 \; \pi \; \cdots \; \pi \; s_k$.

Initialization: Set $S_0 \coloneqq \varnothing$.

Step $i \in \{1, \ldots, k\}$**:** If $S_{i-1} \cup \{s_i\} \subseteq S'$ for some $S' \in \text{F}(S)$, set $S_i \coloneqq S_{i-1} \cup \{s_i\}$. Otherwise, set $S_i \coloneqq S_{i-1}$.

Output: $\text{Ch}^{\pi}(S) \coloneqq S_k$.

Let us illustrate how the distributional choice rule combines priority and distributional preferences using Example 2 and Example 3. Suppose the school has capacity $q = 2$ and the priority ranking is $s_1 \; \pi \; s_2 \; \pi \; s_3 \; \pi \; s_4$.

Consider first the score-based preference in Example 2. Among non-wasteful sets, the maximal total score is 3, with $\mathrm{F}(\mathcal{S}) = \Big\{\{s_1, s_4\}\Big\}$. The algorithm first considers s_1. Since s_1 belongs to the frontier set $\{s_1, s_4\}$, the rule admits s_1. It next sequentially considers s_2 and s_3, and rejects both because neither $\{s_1, s_2\}$ nor $\{s_1, s_3\}$ belongs to the frontier. Finally, the rule considers s_4 and admits her since $\{s_1, s_4\}$ is in the frontier. Hence, $\mathrm{Ch}^{\pi}(\mathcal{S}) = \{s_1, s_4\}$.

Next, consider the distributional preference in Example 3. In this case, the school values the number of protected categories that can be represented when assigning each admitted student to at most one category. The frontier consists of all non-wasteful sets that represent both protected categories. Therefore,

$$\mathrm{F}(\mathcal{S}) = \Big\{\{s_1, s_3\}, \{s_1, s_4\}, \{s_3, s_4\}\Big\}.$$

The algorithm first considers and admits s_1, since s_1 is contained in some frontier set. It then considers and rejects s_2, since no frontier set contains s_2. Next, it considers and admits s_3, since the set $\{s_1, s_3\}$ is in the frontier. Once these two students have been admitted, no additional student can be added while remaining contained in a two-student frontier set. Thus, $\mathrm{Ch}^{\pi}(\mathcal{S}) = \{s_1, s_3\}$.

The two examples have the same students, capacity, and priority ranking. Yet, each distributional preference systematically induces a different frontier and therefore a different choice. More broadly, at each step of the algorithm, Ch^{π} greedily admits the highest-priority student whose admission remains compatible with eventually reaching some set on the frontier. In fact, by construction, the output of the algorithm lies in $\mathrm{F}(S)$ for each applicant set S. As a result, by Lemma 1, Ch^{π} is non-wasteful and distributionally maximal.

The rule is also free of justified envy. To see why, consider two students s and s' with $s \in \mathrm{Ch}^{\pi}(S)$ and $s' \in S \setminus \mathrm{Ch}^{\pi}(S)$. Suppose, toward a contradiction, that $s' \; \pi \; s$ and

$$\Big(\mathrm{Ch}^{\pi}(S) \setminus \{s\}\Big) \cup \{s'\} \succsim \mathrm{Ch}^{\pi}(S).$$

Since $\mathrm{Ch}^{\pi}(S) \in \mathrm{F}(S)$, the second condition implies that $\Big(\mathrm{Ch}^{\pi}(S) \setminus \{s\}\Big) \cup \{s'\} \in \mathrm{F}(S)$. Indeed, $\Big(\mathrm{Ch}^{\pi}(S) \setminus \{s\}\Big) \cup \{s'\}$ is non-wasteful and weakly preferred to a set in the frontier, so it cannot be strictly dominated by any non-wasteful subset of S. Moreover, because $s' \; \pi \; s$, student s' is considered before s in the algorithm. Let i be the step at which s' is considered. Since s' is rejected, $S_i = S_{i-1}$. Moreover, because s is considered later, $S_{i-1} \subseteq \mathrm{Ch}^{\pi}(S) \setminus \{s\}$. Hence, $S_{i-1} \cup \{s'\} \subseteq \Big(\mathrm{Ch}^{\pi}(S) \setminus \{s\}\Big) \cup \{s'\} \in \mathrm{F}(S)$. However, this implies that the algorithm would have admitted s' at step i, leading to a contradiction. Therefore Ch^{π} is free of justified envy.

Thus, Ch^{π} is non-wasteful, distributionally maximal, and free of justified envy. However,

showing that Ch^{π} is the unique choice rule satisfying these three axioms or that it satisfies path independence requires imposing additional structure on the distributional preference. We identify that structure in the following section.

3. Choice with Distributional Preferences

The normative axioms involve a given distributional preference, but some discipline on these preferences is needed. We identify two structural properties of distributional preferences under which our axioms uniquely characterize the distributional choice rule for a fixed applicant pool. We then introduce a third property that, together with the first two, extends this characterization across applicant pools and ensures path independence.

The structural properties introduced below restrict $\succsim$ only over sets of cardinality q. This is without loss for our analysis. If $|S| \leq q$, non-wastefulness requires $\text{Ch}(S) = S$, so neither the distributional preference nor the priority ranking affects the choice. If $|S| > q$, every non-wasteful subset of S has cardinality q, so all relevant distributional comparisons are among sets of cardinality q. Consequently, when specifying applications in Section 5, it is sufficient to define the distributional objective over sets of size q.

The first property—the *upper-bound property*—ensures that incomparability among sets of cardinality q does not create a decision deadlock. If two such sets are incomparable, then their union must contain a non-wasteful subset that is strictly preferred to at least one of them.

Definition 1. The distributional preference $\succsim$ satisfies the **upper-bound property** if, whenever two sets $S, S' \subseteq \mathcal{S}$ are incomparable under $\succsim$ and satisfy $|S| = |S'| = q$, there exists $S'' \in \text{NW}(S \cup S')$ such that $S'' \succ S$ or $S'' \succ S'$.

By definition, the frontier consists of non-wasteful sets that are maximal: no other non-wasteful set strictly dominates them. Maximality alone, however, does not imply that a given frontier set is comparable to every other non-wasteful set. As formalized by the following proposition, the upper-bound property rules out such incomparability and guarantees that the frontier coincides with the collection of non-wasteful sets that are not merely maximal but *optimal*, i.e., most-preferred with respect to the distributional preference among all non-wasteful sets.

Proposition 1. *The distributional preference satisfies the upper-bound property if and only if, for each set $S \subseteq \mathcal{S}$, the following hold:*

(*i*) *For every $S', S'' \in \text{F}(S)$, we have $S' \sim S''$.*

(*ii*) *For every* $S' \in \mathrm{F}(S)$ *and every* $S'' \in \mathrm{NW}(S) \setminus \mathrm{F}(S)$*, we have* $S' \succ S''$.

The second property—the *maximizer property*—imposes a symmetric exchange condition on optimal non-wasteful sets.

Definition 2. The distributional preference $\succsim$ satisfies the **maximizer property** if, for every distinct $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, whenever $S \succsim S''$ and $S' \succsim S''$ hold for every $S'' \in \mathrm{NW}(S \cup S')$, there exist $s \in S \setminus S'$ and $s' \in S' \setminus S$ such that

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

The maximizer property says that optimal non-wasteful sets are not isolated from one another. If two distinct sets S and S' of cardinality q are optimal among the non-wasteful subsets of their union, then there exists a student in $S \setminus S'$ and a student in $S' \setminus S$ who can be exchanged so that each exchanged set is distributionally indifferent to the corresponding original set. Such an exchange moves each set one step toward the other by increasing their overlap. Under the upper-bound property and Proposition 1, the maximizer property implies that if $S, S' \in \mathrm{F}(S \cup S')$, then the exchanged sets remain in $\mathrm{F}(S \cup S')$. In this sense, the maximizer property is a discrete connectedness condition on the frontier, analogous to the connectedness of the set of optima in continuous convex optimization problems.

Proposition 2. *Suppose that the distributional preference satisfies the upper-bound property. Then the maximizer property is satisfied if and only if, for every set* $S \subseteq \mathcal{S}$*,* $\mathrm{F}(S)$ *forms the bases of a matroid on the ground set* S*.*

Together with the upper-bound property, the maximizer property ensures that the frontier forms the bases of a matroid, which in turn supports the greedy characterization of the distributional choice rule. To see that this structure does not hold without the maximizer property, consider the following example.

Example 6. Let $\mathcal{S} = \{s_1, s_2, s_3, s_4\}$ and $q = 2$. Consider a distributional preference where $\{s_1, s_4\} \sim \{s_2, s_3\} \succ T$ for every other $T \in \mathrm{NW}(\mathcal{S})$. This distributional preference satisfies the upper-bound property but the maximizer property fails: $\{s_1, s_4\}$ and $\{s_2, s_3\}$ are both undominated within $\mathrm{NW}(\mathcal{S})$, yet since $\{s_1, s_4\} \cap \{s_2, s_3\} = \varnothing$, any swap of a student between them yields either $\{s_1, s_2\}$, $\{s_1, s_3\}$, $\{s_2, s_4\}$, or $\{s_3, s_4\}$, all of which are strictly dominated. Consequently, $\mathrm{F}(\mathcal{S}) = \{\{s_1, s_4\}, \{s_2, s_3\}\}$ does not form the bases of a matroid.[3]

The maximizer property is therefore necessary and sufficient for the matroidal structure of the frontier and, as the following theorem shows, for the distributional choice rule to be

[3]Specifically, the base exchange axiom fails; see Appendix A.

the unique choice rule that is non-wasteful, distributionally maximal, and free of justified envy.

Theorem 1. *Suppose that the distributional preference satisfies the upper-bound property. Then the following statements are equivalent:*

(*i*) *The distributional preference satisfies the maximizer property.*

(*ii*) *For every priority ranking* π*,* Ch^{π} *is the unique choice rule that is non-wasteful, distributionally maximal, and free of justified envy.*

Hence, while the distributional choice rule satisfies our first three axioms without any additional assumptions on the distributional preference, it is the unique such choice rule under the upper-bound and maximizer properties. In other words, any rule that differs from Ch^{π} must violate at least one of the three axioms. Under the same two properties, our next result strengthens the uniqueness property to a dominance result: At any applicant set where a rule differs from Ch^{π}, it must either leave seats unfilled, select a strictly worse distributional outcome, or choose a set of students that are inferior according to the priority ranking. To formalize this last notion, we extend the priority ranking over individual students to a partial order over sets of students.

Definition 3. Let $S, S' \subseteq \mathcal{S}$. We say that S **priority dominates** S' if $|S| \geq |S'|$ and, for every $k \leq |S'|$, the k-th highest-priority student in S is either identical to, or has higher priority than, the k-th highest-priority student in S'.

Priority domination extends the school's priority ranking from individual students to sets of students. A set S priority dominates S' if it contains at least as many students as S' and if, rank by rank, it contains weakly higher-priority students.

Theorem 2. *Suppose that the distributional preference satisfies the upper-bound property and the maximizer property. Let* Ch *be a choice rule such that* $\mathrm{Ch}(S) \neq \mathrm{Ch}^{\pi}(S)$ *for some* $S \subseteq \mathcal{S}$*. Then at least one of the following holds:*

(*i*) ***Wastefulness:*** $|\mathrm{Ch}^{\pi}(S)| > |\mathrm{Ch}(S)|$.

(*ii*) ***Distributional Inferiority:*** $\mathrm{Ch}^{\pi}(S) \succ \mathrm{Ch}(S)$.

(*iii*) ***Priority Inferiority:*** $\mathrm{Ch}^{\pi}(S)$ *priority dominates* $\mathrm{Ch}(S)$.

The results so far concern choice from a fixed applicant pool. The upper-bound and maximizer properties ensure that, within each applicant pool, the frontier has the exchange structure needed for the first three axioms to uniquely select the greedy distributional choice rule. Path independence—our last choice axiom—requires more: processing an applicant pool

all at once or in batches must lead to the same admitted set. Equivalently, the choice rule must be consistent across varying applicant pools. This consistency requirement is not only desirable for a single-school choice problem, it is also foundational for deferred-acceptance implementation in centralized matching markets.

For a choice rule to be path independent, the frontiers associated with different applicant pools must also be "coherent," something that the upper-bound and maximizer properties do not ensure by themselves. The next property imposes exactly this cross-pool coherence.

Definition 4. The distributional preference $\succsim$ satisfies the **coherence property** if, for every distinct $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$ and every $\hat{s} \in S \setminus S'$ satisfying $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$, whenever

(i) $S \succ S''$ for every $S'' \in \mathrm{NW}(S \cup S')$ with $\hat{s} \notin S''$,
(ii) $S \succsim S''$ for every $S'' \in \mathrm{NW}(S \cup S')$ with $\hat{s} \in S''$, and
(iii) $S' \succsim S''$ for every $S'' \in \mathrm{NW}\left((S \cup S') \setminus \{\hat{s}\}\right)$,

then there exist $s \in (S \setminus S') \setminus \{\hat{s}\}$ and $s' \in S' \setminus S$ such that

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

The economic content of coherence can be understood as a comparative-statics property of the frontier. Let $T \coloneqq S \cup S'$ and suppose $\hat{s}$ satisfies the antecedent of the coherence property. This means that $\hat{s}$ is an **essential** student in T: given an applicant set T, the school cannot attain a distributionally maximal set without admitting student $\hat{s}$. Equivalently, every set in $\mathrm{F}(T)$ contains $\hat{s}$. Once $\hat{s}$ is removed, every set in the frontier of T becomes infeasible. Thus, the frontier of the reduced pool $T \setminus \{\hat{s}\}$ shifts and must be disjoint from the frontier of T. Coherence requires this shift in the frontier to be "continuous in the exchange of students."

In particular, suppose S is optimal among the non-wasteful sets of T, and S' is optimal among the non-wasteful sets of $T \setminus \{\hat{s}\}$. Then coherence requires some student in $(S \setminus S') \setminus \{\hat{s}\}$ to be exchangeable with some student in $S' \setminus S$ while leaving both resulting sets distributionally indifferent to their corresponding original sets. Thus, while removing an essential student shifts the frontier, it cannot make the new frontier isolated (in the exchange sense) from the original one.

Theorem 3. *Suppose that the distributional preference satisfies the upper-bound property. Then the following statements are equivalent:*

(i) *The distributional preference satisfies the maximizer and coherence properties.*
(ii) *For every priority ranking* π, Ch^{π} *satisfies path independence.*

Thus, given the upper-bound property, the maximizer and coherence properties are both necessary and sufficient for path independence of the distributional choice rule. The maximizer property gives each frontier the matroid structure needed for greedy choice within a fixed applicant pool. Coherence imposes the additional frontier consistency condition needed across different applicant pools. If either property fails, one can construct a priority ranking for which the greedy distributional rule selects a student from the larger pool but rejects that same student after another applicant is removed, violating path independence. For instance, the failure of path independence in Example 1 can be attributed to a failure of the maximizer property.[4]

3.1. Revealed Priorities

We now consider a setting in which the priority ranking is not taken as primitive. Instead, priority comparisons are inferred from the school's choice behavior. Suppose a selected student s is chosen over a rejected student s' even though replacing s with s' would weakly improve the distributional outcome. Such a choice reveals s to have higher priority than s'. The next axiom requires these revealed priority comparisons to be acyclic.

Axiom 5. A choice rule Ch satisfies the **revealed priority axiom** if there are no sequences of applicant sets $\{S_k\}_{k=1}^K$ and students $\{s_k\}_{k=1}^K$ such that, for every $k \in \{1, \dots, K\}$,

(i) $s_k \in \mathrm{Ch}(S_k)$,

(ii) $s_{k+1} \in S_k \setminus \mathrm{Ch}(S_k)$, and

(iii) $\Big(\mathrm{Ch}(S_k) \setminus \{s_k\}\Big) \cup \{s_{k+1}\} \succsim \mathrm{Ch}(S_k)$,

where $s_{K+1} = s_1$.

The condition $s_{K+1} = s_1$ closes the sequence into a cycle. Thus, the axiom rules out choice behavior that reveals s_1 above s_2, s_2 above s_3, and so on, while also revealing s_K above s_1. In this sense, rather than starting from an exogenously given priority ranking, we infer one from the school's choice behavior. The axiom then requires these inferred priority comparisons to be acyclic, which is precisely what is needed for them to be extended to a strict linear priority ranking.

Theorem 4. *Suppose that the distributional preference satisfies the upper-bound property and the maximizer property. Then a choice rule* Ch *is non-wasteful, distributionally maximal, and satisfies the revealed priority axiom if and only if* $\mathrm{Ch} = \mathrm{Ch}^\pi$ *for some priority ranking* π.

[4] In Example 1, both $\{s_1, s_2\}$ and $\{s_3, s_4\}$ are in the frontier. However, for any $s \in \{s_1, s_2\}$ and $s' \in \{s_3, s_4\}$, at least one of the two exchanged sets must be either $\{s_2, s_3\}$ or $\{s_2, s_4\}$. Both sets represent only one protected category and are therefore strictly dominated by the frontier sets $\{s_1, s_2\}$ and $\{s_3, s_4\}$. As a result, the maximizer property does not hold.

This result complements Theorem 1 by showing that the distributional choice rule does not depend on taking priorities as exogenously fixed. When a priority ranking is given, non-wastefulness, distributional maximality, and freedom from justified envy uniquely identify Ch^{π}. When priorities are instead inferred from behavior, non-wastefulness, distributional maximality, and acyclicity of revealed priorities again force the choice rule to be a distributional choice rule for some priority ranking.

Together, Theorems 1–4 characterize the distributional choice rule at the single-school level. Extending this characterization to centralized markets introduces a new challenge because multiple schools' choice rules, each based on its own priority ranking and distributional preference, must interact through a single matching mechanism. Section 4 shows how these structural properties support implementation in centralized matching markets through the deferred-acceptance mechanism.

4. Market Design with Distributional Preferences

This section extends our single-school analysis to centralized matching markets. The market consists of a finite set of schools $\mathcal{C}$ and a finite set of students $\mathcal{S}$. Each school $c \in \mathcal{C}$ has a capacity $q_c \in \mathbb{Z}_{++}$, a priority ranking π_c over $\mathcal{S}$, and a distributional preference $\succsim_c$ over subsets of students. Each student $s \in \mathcal{S}$ has a preference relation P_s, which is a strict linear order over $\mathcal{C} \cup \{\varnothing\}$, where $\varnothing$ denotes the student's outside option. The corresponding weak order is denoted by R_s. We denote the set of all student preference relations by $\mathcal{P}$ and a preference profile by $P \coloneqq (P_s)_{s\in\mathcal{S}} \in \mathcal{P}^{\mathcal{S}}$. Thus, a matching market is summarized by a tuple $(\mathcal{S}, \mathcal{C}, P, (\succsim_c)_{c\in\mathcal{C}}, (\pi_c)_{c\in\mathcal{C}}, (q_c)_{c\in\mathcal{C}})$.

Throughout, school capacities, priority rankings, and distributional preferences are common knowledge. On the other hand, student preferences are private information, which motivates the incentive requirements we impose below.

A matching is a function $\mu : \mathcal{S} \to \mathcal{C} \cup \{\varnothing\}$, where $\mu(s) \in \mathcal{C} \cup \{\varnothing\}$ represents the school (or outside option) that is matched to $s \in \mathcal{S}$, and $\mu^{-1}(c) \subseteq \mathcal{S}$ represents the subset of students that are matched to school $c \in \mathcal{C}$. We impose that a matching cannot exceed any school's capacity, i.e., $|\mu^{-1}(c)| \leq q_c$ for every $c \in \mathcal{C}$. Let $\mathcal{M}$ denote the set of all matchings.

We say a school $c \in \mathcal{C}$ is acceptable for student s if $c \; P_s \; \varnothing$.

Axiom 6. A matching μ is **individually rational** if every student is matched either to an acceptable school or to the outside option; that is, $\mu(s) \; R_s \; \varnothing$ for every $s \in \mathcal{S}$.

We now adapt the single-school axioms to the matching environment.

Axiom 7. A matching μ is **non-wasteful** if, for every $c \in \mathcal{C}$ and $s \in \mathcal{S}$,

$$c \; P_s \; \mu(s) \implies \left|\mu^{-1}(c)\right| = q_c.$$

Non-wastefulness requires that no school leave a seat empty while a student who wants that seat remains assigned elsewhere or unmatched.

Axiom 8. A matching μ is **distributionally maximal** if, for every $c \in \mathcal{C}$ and every $S \subseteq \{s \in \mathcal{S} : c \; R_s \; \mu(s)\}$, we have

$$|S| = \left|\mu^{-1}(c)\right| \implies S \not\succ_c \mu^{-1}(c).$$

The set $\{s \in \mathcal{S} : c \; R_s \; \mu(s)\}$ consists of the students who weakly prefer school c to their assignment under μ. In short, this is the set of students who are willing to attend school c under the matching μ. Distributional maximality requires each school to be matched to a set of students that is distributionally undominated among the students who are willing to attend it, holding fixed the number of students assigned to the school.

Axiom 9. A matching μ is **free of justified envy** if, for every $c \in \mathcal{C}$ and distinct students $s, s' \in \mathcal{S}$ with $\mu(s) = c$ and $c \; P_{s'} \; \mu(s')$,

$$\left(\mu^{-1}(c) \setminus \{s\}\right) \cup \{s'\} \succsim_c \mu^{-1}(c) \implies s \; \pi_c \; s'.$$

This definition integrates distributional objectives with priority rankings. In standard school choice, a student has justified envy if she prefers a school to her assignment and has higher priority than some admitted student. In our framework, the priority claim is conditional on the distributional objective. If replacing an admitted student s with a student s' who wants school c would weakly improve the school's distributional outcome, then s must have higher priority than s'. Thus, a lower-priority student can be matched to a school over a higher-priority student only when replacing the former with the latter would fail to weakly improve that school's distributional outcome.

A matching mechanism is a function $\phi : \mathcal{P}^{\mathcal{S}} \to \mathcal{M}$ that maps each profile of student preferences P to a matching $\phi(P)$. We write $\phi_s(P)$ for the assignment of student s under the matching $\phi(P)$.

We say that a mechanism ϕ is non-wasteful if, for every preference profile $P \in \mathcal{P}^{\mathcal{S}}$, the matching $\phi(P)$ is non-wasteful. We define individually rational, distributionally maximal, and free-of-justified-envy mechanisms analogously.

Axiom 10. A mechanism ϕ is **strategy-proof** if no student can benefit from misreporting her preferences: for every $P \in \mathcal{P}^{\mathcal{S}}$, every $s \in \mathcal{S}$, and every $P'_s \in \mathcal{P}$,

$$\phi_s(P) \; R_s \; \phi_s(P'_s, P_{-s}).$$

We now define the deferred-acceptance mechanism induced by a profile of school choice rules. Let $(\mathrm{Ch}_c)_{c\in\mathcal{C}}$ be a profile of choice rules, one for each school.

Deferred-Acceptance Algorithm (DA) based on $(\mathrm{Ch}_c)_{c\in\mathcal{C}}$
Input: A profile of student preferences $P = (P_s)_{s\in\mathcal{S}}$.

Step 1: Each student $s \in \mathcal{S}$ proposes to their most preferred acceptable school, if any. Let S_1^c be the set of proposals received by school c. School c tentatively holds the set $\mathrm{Ch}_c(S_1^c)$ and permanently rejects $S_1^c \setminus \mathrm{Ch}_c(S_1^c)$. If no students are rejected, the algorithm terminates.

Step $k \geq 2$**:** Each student $s \in \mathcal{S}$ rejected at Step $k-1$ proposes to their most preferred acceptable school from which they have not yet been rejected, if any. Let S_k^c be the union of new proposals to c and the students tentatively held by c from the previous step. School c tentatively holds $\mathrm{Ch}_c(S_k^c)$ and permanently rejects $S_k^c \setminus \mathrm{Ch}_c(S_k^c)$. If no students are rejected, the algorithm terminates.

Output: The final tentative assignments constitute the matching outcome.

The **DA matching mechanism** based on $(\mathrm{Ch}_c)_{c\in\mathcal{C}}$ is the mechanism whose outcome is determined by the algorithm above for any given preference profile. In the market generated by distributional preferences, the relevant choice-rule profile is $(\mathrm{Ch}_c^{\pi_c})_{c\in\mathcal{C}}$, where each school uses its distributional choice rule.

Theorem 5. *Suppose that, for every school $c \in \mathcal{C}$, the distributional preference $\succsim_c$ satisfies the upper-bound property, the maximizer property, and the coherence property. A matching mechanism is non-wasteful, distributionally maximal, individually rational, strategy-proof, and free of justified envy if and only if it is the deferred-acceptance mechanism based on the distributional choice rules $(\mathrm{Ch}_c^{\pi_c})_{c\in\mathcal{C}}$.*

Theorem 5 is the mechanism-level counterpart of the choice-theoretic results in Section 3. It shows that the same structural properties—upper-bound, maximizer, and coherence—that characterize single-school choice also characterize implementation in centralized markets. If each school uses its distributional choice rule, deferred acceptance produces matchings that are individually rational, non-wasteful, distributionally maximal, and free of justified envy,

while preserving strategy-proofness for students. Conversely, any mechanism satisfying these desiderata must coincide with the deferred-acceptance mechanism based on the schools' distributional choice rules. Thus, the upper-bound, maximizer, and coherence properties provide a unified foundation for centralized implementation of distributional objectives.

5. Applications

In this section, we apply our framework to several classes of distributional preferences. The purpose is twofold. First, we show how familiar policies, including reserves and floors-and-ceilings policies, can be written as choice rules derived from distributional preferences that satisfy the upper-bound, maximizer, and coherence properties. Hence, these familiar policies are instances of our distributional choice rule for particular distributional preferences. Second, we show how the same framework accommodates distributional objectives that are not naturally expressed as reserves, especially objectives involving overlapping or intersectional identities.

We focus on sets $S \subseteq \mathcal{S}$ with $|S| = q$ without loss of generality. The upper-bound, maximizer, and coherence properties are defined for q-sized sets, and thus, it is sufficient to specify how a school's distributional preference ranks sets of size q. In particular, two distributional preferences that agree on all comparisons among size-q sets induce the same frontier sets and the same distributional choice rule, even if they disagree on sets of other cardinalities.

It is useful to formalize the attributes that determine a school's distributional preferences. To that end, let $\mathcal{T} \coloneqq \{t_1, \ldots, t_K\}$ denote a finite set of distributionally relevant identities or protected traits. Each student $s \in \mathcal{S}$ has a type $\tau(s) \in \{0,1\}^K$, which is a vector with $\tau_k(s) = 1$ if student s has trait $t_k \in \mathcal{T}$ and $\tau_k(s) = 0$ otherwise. In Example 1, $K = 2$ with $t_1 =$ minority status and $t_2 =$ low SES, so $\tau(s_1) = (1,1)$, $\tau(s_2) = (0,0)$, $\tau(s_3) = (0,1)$, and $\tau(s_4) = (1,0)$. The quantity $\|\tau(s)\|_1 = \sum_k \tau_k(s)$ captures how many traits a student s has, and hence represents the student's intersectionality.

5.1. Score-Based Preferences

We begin with preferences derived from student-level scores, which is a useful application of our framework to both separable and nonseparable distributional preferences.

Let $f : \mathcal{S} \to \mathbb{R}$ assign a distributional score to each student. Examples of distributional scores include:

- **Trait-based scoring:** $f(s) = \sigma^t(\tau(s))$, where $\sigma^t : \{0,1\}^K \to \mathbb{R}$ maps a student's type vector into a score. For example, given weights $\lambda = (\lambda_k)_{k=1}^K \in [0,1]^K$ satisfying

$\sum_k \lambda_k = 1$ and a type vector $\tau \in \{0,1\}^K$, the trait-based score could be $\sigma^t(\tau) = \lambda \cdot \tau$. Hence, λ_k measures the marginal contribution of trait t_k to a student's overall score and captures how much the school values that trait.

- **Merit-based scoring:** $f(s) = \sigma^m(s)$, where $\sigma^m(s) \in \mathbb{R}$ is an exam score of student s.
- **Multidimensional scoring:** $f(s) = g(\sigma^t(\tau(s)), \sigma^m(s))$, where the aggregator function g allows for arbitrary interactions between trait- and merit-based scores.

For each set $S \subseteq \mathcal{S}$ with $|S| = q$, label the students in $S = \{s_1, \ldots, s_q\}$ so that $f(s_1) \geq \cdots \geq f(s_q)$. With some abuse of notation, define $f(S) \coloneqq (f(s_1), \ldots, f(s_q))$ as the profile of distributional scores sorted in descending order. We consider two different distributional preferences that compare any two sets of size q based on their respective profiles of scores.

The first distributional preference, which we call the *index distributional preference*, compares sets by applying an index to their score profiles. Formally, let $\Psi : \mathbb{R}^q \to \mathbb{R}$ be a *diversity index*. We assume that Ψ is strictly increasing in each coordinate, so that replacing one student with a higher-scoring student strictly improves the index.

Given two sets $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, we then define the index distributional preference by

$$S \succsim S' \quad \iff \quad \Psi(f(S)) \geq \Psi(f(S')).$$

Note that the index distributional preference is complete. Hence, the upper-bound property holds vacuously. We show that this class of distributional preference satisfies all three structural properties.

Proposition 3. *Suppose the index Ψ is strictly increasing. Then the index distributional preference satisfies the upper-bound property, the maximizer property, and the coherence property.*

This class of preferences accommodates the *additive distributional preference* in Example 2 by specifying

$$\Psi(f(S)) = \sum_{s \in S} f(s).$$

Since Ψ in this case is additively separable, the value of admitting a student is independent of which other students have already been admitted.

More generally, Ψ may be nonseparable, so that the marginal value of a student's score can depend on the scores of the other admitted students. For example, suppose $f(s) \geq 0$ for all $s \in \mathcal{S}$. Given a set $S \subseteq \mathcal{S}$, let $F(\cdot|S) : \mathbb{R}_+ \to [0,1]$ represent its *score survival function* defined as

$$F(m|S) \coloneqq \frac{|\{s \in S : f(s) \geq m\}|}{q}$$

for each $m \in \mathbb{R}_+$. Finally, let the index be given by

$$\Psi(f(S)) = \int_0^\infty g(m, F(m|S))dm,$$

where $g : \mathbb{R}_+ \times [0,1] \to \mathbb{R}$ is measurable, strictly increasing in its second argument, satisfies $g(m,0) = 0$ for every $m \in \mathbb{R}_+$, and is such that the integral is finite for every $S \subseteq \mathcal{S}$. Different specifications of the function g correspond to different distributional preferences. The dependence of g on m determines which score thresholds the school values most: placing greater weight on lower thresholds rewards broad representation, while placing greater weight on higher thresholds rewards cohorts containing students with especially high scores. The shape of g in its second argument determines how the school values the fraction of the cohort whose scores exceed each threshold. For example, concavity generates diminishing returns to additional students above a given threshold, while convexity places increasing value on having a larger share of the cohort above that threshold. Since the index depends on the entire survival function of the cohort, it can be nonseparable across students even though it is derived from individual scores.

Such nonseparability is often difficult to accommodate in matching markets because it can conflict with substitutability, and hence with the stability and implementation properties required of centralized mechanisms (Echenique and Yenmez, 2015). Here, however, strict monotonicity of the index implies that the frontier consists of the q-student sets with the highest feasible score profile. Thus, this class of preferences preserves the exchange and coherence structure needed for implementation on the frontier while allowing the shape of the index to affect how non-frontier sets are ranked.

The second distributional preference, which we call the *pointwise distributional preference*, compares sets based on coordinate-wise domination of the profile of scores. Formally, given $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, we define the pointwise distributional preference by

$$S \succsim S' \quad \iff \quad f(S) \geq f(S'),$$

where for $x, y \in \mathbb{R}^q$, $x \geq y$ if $x_i \geq y_i$ for all $i \in \{1, \ldots, q\}$.

Unlike the index distributional preference, the pointwise distributional preference is incomplete. Thus, the upper-bound property is substantive here.

Proposition 4. *The pointwise distributional preference satisfies the upper-bound property, the maximizer property, and the coherence property.*

While the index and pointwise distributional preferences differ as binary relations, they induce identical frontiers. If $|S| \leq q$, then $\mathrm{F}(S) = \{S\}$. If $|S| > q$, then $\mathrm{F}(S)$ consists of all q-

student subsets of S whose profile of scores is equal to the vector of the q highest distributional scores available in S. Thus, the two preferences lead to the same distributional choice rule even though one is complete and the other is not. This illustrates that what matters for our theory is how the preferences shape the frontier. How preferences compare off-frontier sets is not important.

5.2. Matroidal Distributional Preferences

We next consider distributional preferences derived from matroid rank functions. This class is useful because many familiar policies, such as standard reserve policies as well as overlapping reserves (Sönmez and Yenmez, 2020, 2022), have a natural matroidal structure. In these cases, the rank of a set measures how much of the school's distributional objective can be satisfied by the students in the set.

Let $M = (\mathcal{S}, \mathcal{I})$ be a matroid on the student set $\mathcal{S}$, and let $\mathrm{r}_M : 2^{\mathcal{S}} \to \mathbb{Z}_+$ denote its rank function: for any $S \subseteq \mathcal{S}$, $\mathrm{r}_M(S)$ is the cardinality of the largest independent subset of S.[5]

Given $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, we define the *matroidal distributional preference* induced by M by

$$S \succsim S' \quad \Longleftrightarrow \quad \mathrm{r}_M(S) \geq \mathrm{r}_M(S').$$

In words, the independent sets of the matroid specify which groups of students can be "counted" toward meeting the school's distributional objective. The rank $\mathrm{r}_M(S)$ is the largest number of students in S that can be counted in this sense. Thus, one set is preferred to another when it contains a larger subset of students that can count toward that objective.

Proposition 5. *The matroidal distributional preference satisfies the upper-bound property, the maximizer property, and the coherence property.*

The proposition shows that matroidal distributional preferences are nested in our framework. This is useful because, even when the matroid is not the primitive object governing the school's distributional preference, it can still be viewed as a compact representation of how the primitive distributional preference compares sets.

We now present several classes of matroidal distributional preferences.

5.2.1 Nonintersectional Matroids

Let us first consider settings in which each student has at most one of the K protected traits. This arises, for example, when

[5] See Appendix A for a formal definition of the matroidal rank function.

(*a*) the protected traits are mutually exclusive, such as $\mathcal{T} = \{\text{in-state}, \text{out-of-state}\}$,

(*b*) each combination of identities is treated as a distinct protected trait, such as $\mathcal{T} = \{\text{minority-only}, \text{low-SES-only}, \text{minority-and-low-SES}\}$, or

(*c*) each student is assigned to one of her possibly multiple protected traits.

We model such settings via partition matroids, which are particularly useful for formalizing standard reserve policies where each student qualifies for the reserved seats associated with at most one protected trait.

For each trait $t_k \in \mathcal{T}$, let

$$\mathcal{S}_k \coloneqq \{s \in \mathcal{S} : \tau_k(s) = 1\}$$

be the set of students with trait t_k. Suppose that these sets are pairwise disjoint, so each student has at most one protected trait. We allow for some students to have no protected traits, so it is possible that $\bigcup_{k=1}^{K} \mathcal{S}_k \subseteq \mathcal{S}$.

Let $R_k \in \mathbb{Z}_+$ denote the number of reserved seats for students with trait t_k, with $\sum_k R_k \leq q$. The associated partition matroid has independent sets

$$\mathcal{I} = \left\{ I \subseteq \bigcup_{k=1}^{K} \mathcal{S}_k : |I \cap \mathcal{S}_k| \leq R_k \text{ for every } k \in \{1, \ldots, K\} \right\}.$$

Thus, students with no protected trait do not contribute to the matroid rank. The rank function is

$$\mathrm{r}_M(S) = \sum_{k=1}^{K} \min\left\{ |S \cap \mathcal{S}_k|, R_k \right\}.$$

For example, suppose $\mathcal{T} = \{\text{in-state}, \text{out-of-state}\}$, with reserves $R_1 = 80$ and $R_2 = 20$. Consider a set S with 90 in-state students, 15 out-of-state students, and 10 international students. Since international status is not one of the protected traits in $\mathcal{T}$, international students do not contribute to the matroid rank. Thus, $\mathrm{r}_M(S) = \min\{80, 90\} + \min\{20, 15\} = 95$.

Thus, $\mathrm{r}_M(S)$ counts how many reserved seats can be filled by students in S, and the resulting frontier consists of the non-wasteful sets that maximize matroid rank. The distributional choice rule induced by partition matroids corresponds to a standard reserve policy (Hafalir et al., 2013; Ehlers et al., 2014; Echenique and Yenmez, 2015), which prioritizes sets of students that fill the maximum number of reserved seats.

In Example 1, we considered two scenarios: one in which the SES reserve is processed first, and one in which the minority reserve is processed first. These two scenarios can be represented by partition-matroid distributional preferences after assigning intersectional students to a single protected trait ex-ante. If intersectional students are assigned to the low-SES trait,

the induced distributional choice rule corresponds to the scenario in which the SES reserve is processed first. If instead intersectional students are assigned to the minority trait, the induced distributional choice rule corresponds to the scenario in which the minority reserve is processed first.

The example highlights one drawback of partition matroids; they require each student to belong to at most one type. Intersectional students with multiple protected traits must be assigned to a single identity, which fails to capture the full complexity of their background. This motivates the intersectional matroid structures we examine next.

5.2.2 Matroids for Intersectional Identities

We next consider different classes of matroidal distributional preferences that accommodate intersectional identities.

Transversal matroids and overlapping reserves. When a student has multiple protected traits, she may be eligible for the reserved seats associated with more than one trait. Partition matroids handle this multiple eligibility by assigning the student to one of her identities ex-ante. Here, we instead consider transversal matroids, which formalize overlapping reserves: a student may be eligible ex-ante for the reserved seats of multiple traits, but, once admitted, she counts toward the reserved seats of at most one trait. Such overlapping, or horizontal, reserves are studied by Sönmez and Yenmez (2020, 2022).

For each protected trait $t_k \in \mathcal{T}$, recall that $\mathcal{S}_k$ is the set of students who have that trait, and R_k is the number of reserved seats associated with the trait, with $\sum_k R_k \leq q$. Unlike partition matroids, we no longer require the sets $\mathcal{S}_k$ to be pairwise disjoint. For each trait t_k, index the reserved seats by $\ell = 1, \ldots, R_k$. Define

$$\mathcal{R} \coloneqq \bigcup_{k=1}^{K} \{(\ell, k) : \ell \in \{1, \ldots, R_k\}\}.$$

A set $I \subseteq \mathcal{S}$ is an independent set of the transversal matroid if there exists an injective map $\eta : I \to \mathcal{R}$ such that, whenever $\eta(s) = (\ell, k)$, we have $\tau_k(s) = 1$. In words, each student in an independent set can be assigned to a distinct reserve seat for a trait she has. The injectivity of η ensures that no two students are assigned to the same seat, while the construction of $\mathcal{R}$ ensures that at most R_k students can be assigned to the reserved seats for trait t_k.

The rank function of this transversal matroid measures the maximum number of reserved seats that can be filled by students in a given set. In particular, whenever an intersectional student is eligible for several reserve seats, the matroid optimizes which seat she fills. The

induced matroidal distributional preference therefore compares sets by the maximum number of reserve seats they can fill. This is how the distributional preference in Example 3 is derived.

Notice that when the sets $\mathcal{S}_k$ are disjoint, we recover the independent sets of partition matroids. Hence, partition matroids are a special case of transversal matroids. More broadly, rather than assigning intersectional students to one of their potentially numerous traits ex-ante, the transversal matroid allows intersectional students to remain eligible for multiple reserved seats while ensuring that each admitted student counts toward at most one reserve seat.

Vector matroids. Vector matroids capture a distributional objective under which the school values cohorts containing linearly non-redundant combinations of protected traits.

Recall that each student s is associated with a type $\tau(s) \in \{0,1\}^K$, which we view as a vector in $\mathbb{R}^K$. A set $I \subseteq \mathcal{S}$ is independent in the vector matroid if the vectors $\{\tau(s) : s \in I\}$ are linearly independent over $\mathbb{R}$. For a student $s \notin S$, adding s to S increases its rank precisely when

$$\tau(s) \notin \operatorname{span}_{\mathbb{R}}\{\tau(s') : s' \in S\}.$$

The associated rank function is

$$\mathrm{r}_M(S) = \dim\left(\operatorname{span}_{\mathbb{R}}\{\tau(s) : s \in S\}\right).$$

Thus, the induced matroidal distributional preference compares cohorts according to the maximum number of linearly independent trait profiles they contain. Since $\mathrm{r}_M(S) \leq K$, the school values at most K mutually non-redundant profiles, regardless of cohort size.

In Example 4, the pair $\{s_1, s_2\}$ has rank one because $\tau(s_2)$ is the zero vector and therefore adds no dimension to the span of $\tau(s_1)$. In contrast, the pairs $\{s_1, s_3\}$, $\{s_1, s_4\}$, and $\{s_3, s_4\}$ each have rank two because the two type vectors are linearly independent and span $\mathbb{R}^2$. Thus, unlike the score-based preference in Example 2, which assigns an explicit premium to intersectionality, the vector-matroid preference assigns no intrinsic premium to an intersectional student: she increases rank only when her trait profile adds a dimension not already generated by the other students in the cohort.

Graphic matroids. In some settings, the school may not value each trait in isolation, but rather care about the linkages across traits. For example, a student with expertise in machine learning and economics may be valued differently from one whose expertise links machine learning to biology, even if both are interdisciplinary students. This objective reflects the idea that bridging gaps between otherwise distinct traits can create value for the group as a

whole (Burt, 1992, 2004). Graphic matroids capture this structure.

To model this, suppose each student has exactly two protected traits, so that $\|\tau(s)\|_1 = 2$ for every $s \in \mathcal{S}$. We can then interpret $\mathcal{T}$ as a set of vertices of a multigraph with each student naturally corresponding to an edge between the two traits she has. Formally, if $\tau_k(s) = \tau_{k'}(s) = 1$ with $k \neq k'$, define the edge contributed by the student as $e_s = \{t_k, t_{k'}\}$. For a set $S \subseteq \mathcal{S}$, let

$$E(S) \coloneqq \{e_s : s \in S\}$$

denote the corresponding multiset of edges. The multigraph is then $G = (\mathcal{T}, E(\mathcal{S}))$.

A set $I \subseteq \mathcal{S}$ is an independent set of the graphic matroid if the multigraph $(\mathcal{T}, E(I))$ contains no cycle. Equivalently, the students in I add new links among traits without creating a redundancy. The rank function of this matroid is (Oxley, 2006, Section 1.1)

$$\mathrm{r}_M(S) = K - c(E(S)),$$

where $c(E(S))$ is the number of connected components of the graph $(\mathcal{T}, E(S))$, counting isolated vertices as their own components.

The induced matroidal distributional preference therefore compares sets by the maximum number of non-redundant trait linkages represented by students in the set. Thus, the school may prefer a cohort that links otherwise separate traits over one whose students all connect traits within the same already-connected cluster.

Example 5 and Figure 1 illustrate such a graphic-matroid distributional preference. In that example, the cohort $\{s_1, s_2, s_3\}$ contributes edges that connect all four fields into a single component, giving $\mathrm{r}_M(\{s_1, s_2, s_3\}) = 4 - 1 = 3$. By contrast, the cohort $\{s_1, s_3, s_4\}$ contributes edges CS–Economics, CS–ML, and Economics–ML, which form a cycle, while leaving Political Science as an isolated component, giving $\mathrm{r}_M(\{s_1, s_3, s_4\}) = 4 - 2 = 2$. The school therefore prefers $\{s_1, s_2, s_3\}$ to $\{s_1, s_3, s_4\}$. In particular, adding s_2 rather than s_4 to $\{s_1, s_3\}$ is preferable because s_2 introduces a bridge to Political Science, while s_4 adds another connection within the already connected CS–Economics–ML cluster.

Gammoid matroids. In some settings, the school may care not only about which traits are linked, but also about the paths through which those linkages arise. For example, a medical school may treat two students specializing in primary care as distributionally different if one is a first-generation student who attended a public undergraduate university, while the other comes from a rural hometown and worked as a research assistant before applying to medical school. Even though both students are linked to the same specialty, their backgrounds and prior experiences trace different paths to that specialty, and those paths

may shape how they contribute within the admitted cohort. Gammoid matroids capture this structure.

To model this, define a directed graph $G = (\mathcal{T}, E)$ on the trait space $\mathcal{T}$. Let $\mathcal{T}^o \subseteq \mathcal{T}$ be the set of source nodes, representing student backgrounds, and let $\mathcal{T}^d \subseteq \mathcal{T}$ be the set of terminal nodes, representing career outcomes. We allow for the possibility that some nodes are neither source nor terminal nodes, so $\mathcal{T}^o \cup \mathcal{T}^d \subseteq \mathcal{T}$. Each student $s \in \mathcal{S}$ is associated with a source node in $\mathcal{T}^o$, and the directed graph determines which terminal nodes in $\mathcal{T}^d$ can be reached from that source. If multiple students share the same source node or need to be linked to the same terminal node, one can use student-specific copies of the relevant nodes and enlarge the vertex set to accommodate such copies.

A set $I \subseteq \mathcal{S}$ is an independent set of the gammoid matroid if there exist $|I|$ vertex-disjoint paths from the source nodes associated with students in I to terminal nodes. For a set $S \subseteq \mathcal{S}$, the rank function $\mathrm{r}_M(S)$ equals the maximum number of students in S that can be linked from their source nodes to terminal nodes by vertex-disjoint paths. Thus, the induced matroidal distributional preference compares sets according to the maximum number of directed source-to-terminal linkages that can be represented without redundancy.

5.3. Floors and Ceilings

We next consider floors-and-ceilings policies. These policies specify lower and upper bounds on the number of admitted students from each protected trait. They are commonly written directly as choice rules. We show how they can instead be rationalized by distributional preferences over sets of students.

For this subsection, suppose that each student has exactly one protected trait.[6] For each trait $t_k \in \mathcal{T}$ and each set $S \subseteq \mathcal{S}$, let

$$n_k(S) \coloneqq \sum_{s \in S} \tau_k(s)$$

denote the number of students in S with that trait, and let $n(S) \coloneqq (n_1(S), \ldots, n_K(S))$ denote the set's trait-composition. A floors-and-ceilings policy specifies, for each trait t_k, a lower bound $L_k \in \mathbb{Z}_+$ and an upper bound $U_k \in \mathbb{Z}_+$, with $L_k \leq U_k$ (Ehlers et al., 2014). We say that a set $S \subseteq \mathcal{S}$ is *feasible* if $L_k \leq n_k(S) \leq U_k$ for each $k \in \{1, \ldots, K\}$. Otherwise, we say S is infeasible.

A natural first attempt is to define a dichotomous distributional preference, which ranks

[6] One can allow for students to have *at most* one protected trait by treating students with no protected trait as belonging to a residual unprotected category t_0.

sets solely by whether they are feasible. For sets $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, define

$$S \succsim S' \quad \Longleftrightarrow \quad S \text{ is feasible or } S' \text{ is infeasible.}$$

This captures a school that treats the bounds as hard constraints: all feasible sets are mutually indifferent and strictly dominate all infeasible sets, while all infeasible sets are mutually indifferent to one another, regardless of how severely they violate the constraints.

Proposition 6. *The dichotomous distributional preference satisfies the upper-bound property and the maximizer property.*

The upper-bound property is satisfied because the dichotomous distributional preference is complete. The maximizer property follows because all feasible sets are mutually indifferent and admit symmetric exchanges, as do all infeasible sets.

However, this strict bifurcation has an important limitation: the dichotomous distributional preference can fail the coherence property. Consequently, the induced choice rule can violate path independence, as the following example demonstrates.

Example 7. Suppose $q = 3$ and there are two protected traits, t_1 and t_2. The bounds are given by $L_1 = 2$, $L_2 = 0$, $U_1 = U_2 = 3$. Let $\mathcal{S} = \{s_1, s_2, s_3, s_4, s_5\}$, where s_4 and s_5 have trait t_1, while s_1, s_2, and s_3 have trait t_2. Suppose the priority ranking is $s_1 \ \pi \ s_2 \ \pi \ s_3 \ \pi \ s_4 \ \pi \ s_5$.

The frontier of the entire applicant pool is

$$\mathrm{F}(\mathcal{S}) = \{\{s_1, s_4, s_5\}, \{s_2, s_4, s_5\}, \{s_3, s_4, s_5\}\},$$

so the distributional choice rule yields $\mathrm{Ch}^{\pi}(\mathcal{S}) = \{s_1, s_4, s_5\}$.

Now consider the applicant pool after s_5 is removed. In this case, the lower bound for trait t_1 cannot be satisfied, so all subsets of the applicant pool are infeasible. The frontier of the reduced applicant pool is

$$\mathrm{F}(\mathcal{S} \setminus \{s_5\}) = \{\{s_1, s_2, s_3\}, \{s_1, s_2, s_4\}, \{s_1, s_3, s_4\}, \{s_2, s_3, s_4\}\},$$

so the distributional choice rule yields $\mathrm{Ch}^{\pi}(\mathcal{S} \setminus \{s_5\}) = \{s_1, s_2, s_3\}$. Because s_4 is admitted from the larger pool but rejected after s_5 is removed, the choice rule violates substitutability, and hence path independence.

The problem is that the dichotomous preference treats all constraint violations as equivalent. A set that violates a bound by one student is ranked the same as a set that violates a bound by 10 students. To recover path independence, the distributional preference must compare violations more finely.

We now define a graded floors-and-ceilings distributional preference. For each trait t_k and each set $S \subseteq \mathcal{S}$, define the error function

$$\varepsilon_k(n_k(S)) = \begin{cases} -\eta(L_k - n_k(S)) & \text{if } n_k(S) < L_k, \\ 0 & \text{if } L_k \leq n_k(S) \leq U_k, \\ -(n_k(S) - U_k) & \text{if } n_k(S) > U_k, \end{cases} \tag{1}$$

where $\eta \geq 0$ is a penalty parameter.[7] The error for trait t_k is zero when the number of students with that trait lies within the desired lower and upper bounds. Falling below the floor is penalized at rate η, while exceeding the ceiling is penalized at rate one.

For sets $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, define the floors-and-ceilings distributional preference by

$$S \succsim S' \quad \Longleftrightarrow \quad \sum_{k=1}^{K} \varepsilon_k(n_k(S)) \geq \sum_{k=1}^{K} \varepsilon_k(n_k(S')).$$

As with the dichotomous distributional preference, the floors-and-ceilings distributional preference ranks sets that satisfy all floor and ceiling bounds as distributionally maximal. Where the two preferences differ, however, is in how they handle constraint violations. Unlike the dichotomous distributional preference, here, the school treats the bounds as soft constraints, and trades off floor and ceiling violations based on the parameter η.

Proposition 7. *For any $\eta \geq 0$, the floors-and-ceilings distributional preference satisfies the upper-bound property, the maximizer property, and the coherence property.*

Because the floors-and-ceilings distributional preference satisfies all three structural properties, the results above imply that the induced distributional choice rule satisfies all four choice axioms, in particular path independence. When $\eta > q$, the induced distributional choice rule closely parallels the standard floors-and-ceilings choice rule (Ehlers et al., 2014): under this parameterization, a single floor violation is more costly than any possible collection of ceiling violations within a q-student set, so the school first minimizes floor violations, and then, subject to that, minimizes ceiling violations.

Returning to Example 7, consider the floors-and-ceilings distributional preference and suppose $\eta > 0$. Notice that the ceiling bounds are never violated, so an error arises only when the floor bounds are violated. In the original applicant pool $\mathcal{S}$, the frontier is unchanged:

$$\mathrm{F}(\mathcal{S}) = \{\{s_1, s_4, s_5\}, \{s_2, s_4, s_5\}, \{s_3, s_4, s_5\}\},$$

since these are exactly the sets with total error of zero. Thus, the distributional choice rule

[7] It is possible to further generalize the error function by choosing a parameter $\eta_k \geq 0$ for each $k = 1, \ldots, K$.

again selects $\mathrm{Ch}^{\pi}(\mathcal{S}) = \{s_1, s_4, s_5\}$. The difference appears after s_5 is removed. In the reduced pool $\mathcal{S} \setminus \{s_5\}$, no set can satisfy the lower bound $L_1 = 2$, but any set containing s_4 has one unit of floor violation, while any set omitting s_4 has two units of floor violation. Hence the frontier of the reduced pool is

$$\mathrm{F}(\mathcal{S} \setminus \{s_5\}) = \{\{s_1, s_2, s_4\}, \{s_1, s_3, s_4\}, \{s_2, s_3, s_4\}\},$$

and the distributional choice rule selects $\mathrm{Ch}^{\pi}(\mathcal{S} \setminus \{s_5\}) = \{s_1, s_2, s_4\}$. Thus, s_4 remains admitted after s_5 is removed. The path-independence failure in Example 7 disappears because the floors-and-ceilings preference distinguishes between different infeasible sets by preferring sets that come closer to meeting the violated bound.

5.4. Discrete Concavity

We conclude the application section by connecting the structural properties of distributional preferences to discrete convex analysis. The previous examples were defined through score indices, particular policies, or matroidal representations. Here, we instead begin with a value function over a cohort's trait-composition and impose a discrete exchange condition on that value function. This gives a general way to verify the maximizer and coherence properties for distributional preferences that depend on the composition of the admitted cohort.

Suppose each student belongs to exactly one category in $\mathcal{T}$. For each set $S \subseteq \mathcal{S}$, recall $n(S) \coloneqq (n_1(S), \ldots, n_K(S))$ denotes the set's trait-composition. Let $\Phi : \mathbb{Z}_+^K \to \mathbb{R}$ be a distributional value function over trait-compositions. For sets $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, define the induced distributional preference by

$$S \succsim S' \quad \Longleftrightarrow \quad \Phi(n(S)) \geq \Phi(n(S')).$$

Since this preference is represented by a real-valued function on q-student sets, it is complete. Hence, the upper-bound property holds vacuously. The substantive restrictions are the maximizer and coherence properties, which require the value function Φ to have an appropriate exchange structure. To that end, we employ an ordinal version of M-concavity that is adapted to sets of fixed cardinality. Let e_k denote the k-th standard basis vector in $\mathbb{Z}^K$.

Definition 5. The value function $\Phi : \mathbb{Z}_+^K \to \mathbb{R}$ satisfies **q-ordinal concavity** if, for any trait-compositions $x, y \in \mathbb{Z}_+^K$ with $\sum_{k=1}^K x_k = \sum_{k=1}^K y_k = q$, and for every $k \in \{1, \ldots, K\}$ such that $x_k > y_k$, there exists $k' \in \{1, \ldots, K\}$ with $x_{k'} < y_{k'}$ such that at least one of the following holds:

(*i*) $\Phi(x - e_k + e_{k'}) > \Phi(x)$,

(*ii*) $\Phi(y + e_k - e_{k'}) > \Phi(y)$, or

(*iii*) $\Phi(x - e_k + e_{k'}) = \Phi(x)$ and $\Phi(y + e_k - e_{k'}) = \Phi(y)$.

In economic terms, this condition is a discrete concavity requirement: the school cannot strictly prefer both of two cohorts to the intermediate compositions obtained by moving them one step toward each other.[8] When one cohort has more students with trait t_k and fewer students with trait $t_{k'}$ than another, exchanging these traits moves each cohort's composition one step toward the other. The q-ordinal concavity condition requires that such an exchange either strictly improves the value of one cohort or preserves the value of both.

Proposition 8. *Suppose that Φ satisfies q-ordinal concavity. Then the distributional preference induced by Φ satisfies the upper-bound property, the maximizer property, and the coherence property.*

A simple and important class is separable discrete-concave value functions

$$\Phi(x) = \sum_{k=1}^{K} \phi_k(x_k),$$

where, for each $k \in \{1, \ldots, K\}$, the function $\phi_k : \mathbb{Z}_+ \to \mathbb{R}$ has weakly decreasing marginal increments, i.e., $\phi_k(x_k + 1) - \phi_k(x_k)$ is weakly decreasing in x_k.[9] Since the marginal value of admitting an additional student of a given trait decreases as that trait becomes more represented in the cohort, such a separable value function Φ satisfies q-ordinal concavity.

Notice that the floors-and-ceilings preference in the previous section is a special case of a preference induced by a q-ordinally concave value function where each ϕ_k is given by (1). In that case, the marginal increments are given by

$$\phi_k(x_k + 1) - \phi_k(x_k) = \begin{cases} \eta & \text{if } x_k < L_k \\ 0 & \text{if } L_k \leq x_k < U_k \\ -1 & \text{if } x_k \geq U_k, \end{cases}$$

which is indeed weakly decreasing in x_k. Thus, Proposition 8 immediately implies Proposition 7.

More broadly, q-ordinal concavity imposes the same exchange logic on trait-compositions that the maximizer and coherence properties impose on the exchange of students across

[8] This definition is adapted from the ordinal concavity studied by Hafalir et al. (2026) and restricted to fixed-cardinality sets.

[9] In fact, when each ϕ_k has weakly decreasing marginal increments, the separable value function Φ satisfies the fixed-cardinality exchange inequality associated with M-concavity. This cardinal exchange condition implies q-ordinal concavity.

cohorts. Hence, unlike the score-based preferences in Section 5.1, which rely on monotonicity properties of the index function, this class relies on local exchange properties of the value function. It therefore provides an alternative route for verifying that a distributional preference satisfies the three structural properties needed for implementation.

6. Conclusion

This paper develops a preference-based approach to distributional objectives in school choice and related problems. We deviate from the literature by taking as primitive the school's distributional preference over student cohorts, rather than encoding a school's distributional objectives directly through reserves, quotas, or institutional choice rules. Our approach accommodates overlapping and intersectional identities and makes it possible to ask when a distributional objective can be reconciled with priority rankings, consistent choice across applicant pools, and centralized matching.

We identify an upper-bound property and two exchange properties. The upper-bound property ensures that, within each applicant pool, all distributionally maximal non-wasteful cohorts are indifferent and strictly dominate the remaining non-wasteful cohorts. Conditional on this property, the maximizer property holds if and only if, for every priority ranking, the greedy distributional choice rule is the unique choice rule satisfying non-wastefulness, distributional maximality, and freedom from justified envy. Conditional on the same property, the maximizer and coherence properties hold if and only if, for every priority ranking, that rule satisfies path independence. When every school's distributional preference satisfies all three structural properties, deferred acceptance based on these rules is the unique mechanism satisfying the corresponding three axioms together with individual rationality and strategy-proofness.

These results also delineate where the characterizations fail. Conditional on the upper-bound property, the failure of the maximizer property implies that there exists a priority ranking for which another choice rule satisfies the three choice axioms, while the failure of either exchange property implies that there exists a priority ranking for which the distributional choice rule fails path independence. Such failures undermine the choice-theoretic structure supporting the deferred-acceptance characterization.

The applications show that both familiar and less conventional distributional objectives fit within the framework. Standard and overlapping reserves arise from matroidal distributional preferences, while a graded formulation of floors and ceilings also satisfies the three structural properties. The framework additionally accommodates score-based preferences that assign different values to intersectional profiles, preferences that value linearly non-redundant com-

binations of traits, and preferences that value linkages across traits. It therefore provides preference foundations for canonical policies while supplying structured formulations for distributional objectives beyond conventional reserve systems.

Our analysis suggests a preference-first approach to policy design. Rather than selecting a reserve policy and subsequently investigating its properties, a designer can begin with the institution's distributional preference and ask whether it satisfies the structural properties supporting the desired choice and matching guarantees. The framework thereby separates the normative question of which cohorts are distributionally desirable from the implementation question of which objectives can be reconciled with individual priorities and centralized matching.

Appendix

Appendix A. Useful Results from Matroid Theory

We first state some equivalent characterizations for the bases of a matroid, which we use in our proofs. Fix a ground set $\mathcal{E}$. A collection of sets $\mathcal{B} \subseteq 2^{\mathcal{E}}$ forms the bases of a matroid if and only if the following hold:

(B1) *Non-emptiness*: $\mathcal{B} \neq \varnothing$.

(B2) *Exchange*: For every $B, B' \in \mathcal{B}$ and $e \in B \setminus B'$, there exists $e' \in B' \setminus B$ such that $(B \setminus \{e\}) \cup \{e'\} \in \mathcal{B}$.

One may replace (B2) with the following *strong simultaneous exchange* property (Oxley, 2006, Lemma 2.2.2): for any $B, B' \in \mathcal{B}$ and $e \in B \setminus B'$, there exists $e' \in B' \setminus B$ such that

$$(B \setminus \{e\}) \cup \{e'\} \in \mathcal{B} \quad \text{and} \quad (B' \setminus \{e'\}) \cup \{e\} \in \mathcal{B}.$$

Similarly, one may replace the exchange axiom (B2) with the following *weak simultaneous exchange* property (Murota, 2003, Theorem 4.3): for any distinct $B, B' \in \mathcal{B}$, there exists $e \in B \setminus B'$ and $e' \in B' \setminus B$ such that

$$(B \setminus \{e\}) \cup \{e'\} \in \mathcal{B} \quad \text{and} \quad (B' \setminus \{e'\}) \cup \{e\} \in \mathcal{B}.$$

Given a matroid $M = (\mathcal{E}, \mathcal{I})$, where $\mathcal{I}$ is its collection of independent sets, and a priority ranking π on $\mathcal{E}$, the **standard matroid greedy algorithm** initializes $G = \varnothing$ and scans the elements of $\mathcal{E}$ in π-order. When an element e is considered, the algorithm adds e to G if and only if $G \cup \{e\} \in \mathcal{I}$. The algorithm terminates with a maximal independent set, hence a basis. We call this set the **greedy basis** of M under π and denote it by G_M^π. We use the following classic result regarding greedy algorithms on matroids.

Lemma 2 (Gale, 1968)**.** *Let $\mathcal{E}$ be a ground set, let $M = (\mathcal{E}, \mathcal{I})$ be a matroid on $\mathcal{E}$, and let π be a priority ranking over $\mathcal{E}$. Then G_M^π priority dominates every $I \in \mathcal{I}$.*

Given a matroid $M = (\mathcal{E}, \mathcal{I})$, its rank function $\mathrm{r}_M : 2^{\mathcal{E}} \to \mathbb{Z}_+$ is defined by

$$\mathrm{r}_M(E) \coloneqq \max_{I \subseteq E,\ I \in \mathcal{I}} |I|$$

for each $E \subseteq \mathcal{E}$. We refer to $\mathrm{r}_M(\mathcal{E})$ as the **matroid rank** of M.

We will use a quotient relation between two matroids on a common ground set. In our application, the relevant quotients have ranks that differ by one. Given two matroids M^c

and M^d on a finite common ground set $\mathcal{E}$, we say that M^c is a **quotient** of M^d if for every $E' \subseteq E \subseteq \mathcal{E}$,

$$\mathrm{r}_{M^c}(E) - \mathrm{r}_{M^c}(E') \leq \mathrm{r}_{M^d}(E) - \mathrm{r}_{M^d}(E'). \tag{2}$$

Furthermore, if $\mathrm{r}_{M^d}(\mathcal{E}) = \mathrm{r}_{M^c}(\mathcal{E}) + 1$, we say that M^c is a **rank-one quotient** of M^d.

There is an equivalent characterization of rank-one quotients via an elementary construction using deletion and contraction operations on matroids. Formally, given a matroid $M = (\mathcal{E}, \mathcal{I})$ and an element $e \in \mathcal{E}$, let

$$\mathcal{I}^{\setminus e} \coloneqq \{I \in \mathcal{I} : I \subseteq \mathcal{E} \setminus \{e\}\},$$

and for e satisfying $\{e\} \in \mathcal{I}$, let

$$\mathcal{I}^{/e} \coloneqq \{I \subseteq \mathcal{E} \setminus \{e\} : I \cup \{e\} \in \mathcal{I}\}.$$

The first collection consists of the independent sets of M that do not contain e. The second consists of the sets obtained by taking an independent set of M that contains e and removing e. Both $\mathcal{I}^{\setminus e}$ and $\mathcal{I}^{/e}$ define a collection of independent sets on the reduced ground set $\mathcal{E} \setminus \{e\}$. Moreover, by the hereditary property, $\mathcal{I}^{/e} \subseteq \mathcal{I}^{\setminus e}$.

The matroid $M^{\setminus e} \coloneqq (\mathcal{E} \setminus \{e\}, \mathcal{I}^{\setminus e})$ is the **deletion of** e **from** M, while the matroid $M^{/e} \coloneqq (\mathcal{E} \setminus \{e\}, \mathcal{I}^{/e})$ is the **contraction of** M **by** e.

Given two matroids M^c and M^d on a finite common ground set $\mathcal{E}$, M^c is a rank-one quotient of M^d if and only if there exists an element $e \notin \mathcal{E}$ and a matroid $M = (\mathcal{E} \cup \{e\}, \mathcal{I})$ with $\{e\} \in \mathcal{I}$ such that $M^c = M^{/e}$ is the contraction of M by e (hence "c" for contraction) and $M^d = M^{\setminus e}$ is the deletion of e from M (hence "d" for deletion) (Oxley, 2006, Section 7.3).

The following two auxiliary facts are standard consequences of the theory of matroid quotients and flag matroids. We include direct proofs because the particular formulations used here are tailored to our application.

Lemma 3. *Let M^c and M^d be matroids on the same finite ground set $\mathcal{E}$. Then M^c is a quotient of M^d if and only if $G^{\pi}_{M^c} \subseteq G^{\pi}_{M^d}$ for every priority ranking π on $\mathcal{E}$.*

Proof. **($\Rightarrow$).** Suppose M^c is a quotient of M^d. Then (2) holds. Fix a priority ranking π on $\mathcal{E}$, and write the elements of $\mathcal{E}$ as

$$e_1 \; \pi \; e_2 \; \pi \; \cdots \; \pi \; e_{|\mathcal{E}|}.$$

For each $k \in \{1, \ldots, |\mathcal{E}|\}$, let $E_k \coloneqq \{e_1, \ldots, e_k\}$ and $E_0 \coloneqq \varnothing$. For an arbitrary matroid M on $\mathcal{E}$, the standard matroid greedy algorithm admits e_k if and only if $\mathrm{r}_M(E_k) - \mathrm{r}_M(E_{k-1}) = 1$.

Suppose $e_k \in G^{\pi}_{M^c}$. Then, applying (2) to $E_{k-1} \subseteq E_k$ and using the fact that a rank function can increase by at most one when a single element is added to a set, we have

$$1 = \mathrm{r}_{M^c}(E_k) - \mathrm{r}_{M^c}(E_{k-1}) \leq \mathrm{r}_{M^d}(E_k) - \mathrm{r}_{M^d}(E_{k-1}) \leq 1.$$

Thus, $e_k \in G^{\pi}_{M^d}$. Since the priority ranking was fixed arbitrarily, we conclude that $G^{\pi}_{M^c} \subseteq G^{\pi}_{M^d}$ for all π on $\mathcal{E}$.

($\Leftarrow$). Suppose $G^{\pi}_{M^c} \subseteq G^{\pi}_{M^d}$ for every priority ranking π on $\mathcal{E}$. We verify (2). Fix $E' \subseteq E \subseteq \mathcal{E}$. Choose a priority ranking π on $\mathcal{E}$ that ranks the elements of E' first, then the elements of $E \setminus E'$, and then the elements of $\mathcal{E} \setminus E$, with arbitrary order within each block.

When restricted to the first block E', greedy on M^d under π produces a maximal independent subset of E' in M^d of cardinality $\mathrm{r}_{M^d}(E')$. After the elements of E' and $E \setminus E'$ have all been processed, the selected elements inside E form a maximal independent subset of E in M^d of cardinality $\mathrm{r}_{M^d}(E)$. Hence

$$|G^{\pi}_{M^d} \cap E'| = \mathrm{r}_{M^d}(E') \quad \text{and} \quad |G^{\pi}_{M^d} \cap E| = \mathrm{r}_{M^d}(E).$$

Therefore $\left|G^{\pi}_{M^d} \cap (E \setminus E')\right| = \mathrm{r}_{M^d}(E) - \mathrm{r}_{M^d}(E')$. The same argument for M^c gives

$$|G^{\pi}_{M^c} \cap (E \setminus E')| = \mathrm{r}_{M^c}(E) - \mathrm{r}_{M^c}(E').$$

Since $G^{\pi}_{M^c} \subseteq G^{\pi}_{M^d}$, we obtain

$$\mathrm{r}_{M^c}(E) - \mathrm{r}_{M^c}(E') = |G^{\pi}_{M^c} \cap (E \setminus E')| \leq |G^{\pi}_{M^d} \cap (E \setminus E')| = \mathrm{r}_{M^d}(E) - \mathrm{r}_{M^d}(E').$$

This verifies (2). Hence, M^c is a quotient of M^d. □

Lemma 4. *Let M^c and M^d be matroids on a finite set $\mathcal{E}$, and suppose that M^c is a rank-one quotient of M^d. Then, for every basis B^c of M^c and every basis B^d of M^d with $B^c \setminus B^d \neq \varnothing$, there exist $e^c \in B^c \setminus B^d$ and $e^d \in B^d \setminus B^c$ such that $(B^c \setminus \{e^c\}) \cup \{e^d\}$ is a basis of M^c and $(B^d \setminus \{e^d\}) \cup \{e^c\}$ is a basis of M^d.*

Proof of Lemma 4. By the deletion-contraction characterization of rank-one quotients, there exists an element $e \notin \mathcal{E}$ and a matroid $M = (\mathcal{E} \cup \{e\}, \mathcal{I})$ with $\{e\} \in \mathcal{I}$ such that $M^c = M^{/e}$ and $M^d = M^{\setminus e}$. Then B^c is a basis of M^c if and only if $B^c \cup \{e\}$ is a basis of M, and B^d is a basis of M^d if and only if B^d is a basis of M and $e \notin B^d$.

Let B^c be a basis of M^c and let B^d be a basis of M^d satisfying $B^c \setminus B^d \neq \varnothing$. Take any $e^c \in B^c \setminus B^d$. Then, $B^c \cup \{e\}$ and B^d are each a basis of M, and $e^c \in (B^c \cup \{e\}) \setminus B^d$.

By the strong basis exchange property (e.g., Oxley, 2006, Lemma 2.2.2), there exists $e^d \in B^d \setminus (B^c \cup \{e\}) = B^d \setminus B^c$ such that both

$$\Big((B^c \cup \{e\}) \setminus \{e^c\}\Big) \cup \{e^d\} \quad \text{and} \quad (B^d \setminus \{e^d\}) \cup \{e^c\}$$

are each a basis of M as well.

The first basis contains e, so contracting e from it gives $(B^c \setminus \{e^c\}) \cup \{e^d\}$, which is a basis of M^c. The second basis does not contain e, so it is a basis of M^d. This proves the lemma. □

Appendix B. Proofs of Lemmas and Propositions

Proof of Lemma 1.

(⇐). Suppose that $\mathrm{Ch}(S) \in \mathrm{F}(S)$ for each set $S \subseteq \mathcal{S}$. By the definition of the frontier, $\mathrm{F}(S) \subseteq \mathrm{NW}(S)$, so $\mathrm{Ch}(S) \in \mathrm{NW}(S)$. Hence $|\mathrm{Ch}(S)| = \min\{q, |S|\}$, and Ch is non-wasteful. Moreover, any set $S' \subseteq S$ with $|S'| = |\mathrm{Ch}(S)|$ is also non-wasteful, so $S' \in \mathrm{NW}(S)$. By definition of the frontier and the fact that $\mathrm{Ch}(S) \in \mathrm{F}(S)$, we conclude that $S' \not\succ \mathrm{Ch}(S)$. Thus, Ch is also distributionally maximal.

(⇒). Suppose that Ch is non-wasteful and distributionally maximal. Consider any set $S \subseteq \mathcal{S}$. Non-wastefulness implies $|\mathrm{Ch}(S)| = \min\{q, |S|\}$, so $\mathrm{Ch}(S) \in \mathrm{NW}(S)$. Since Ch is distributionally maximal, there is no $S' \subseteq S$ with $|S'| = |\mathrm{Ch}(S)|$ such that $S' \succ \mathrm{Ch}(S)$. Equivalently, there is no $S' \in \mathrm{NW}(S)$ such that $S' \succ \mathrm{Ch}(S)$. Hence $\mathrm{Ch}(S) \in \mathrm{F}(S)$. □

Proof of Proposition 1.

(⇒). Suppose the distributional preference satisfies the upper-bound property. Consider any set $S \subseteq \mathcal{S}$. If $|S| \leq q$, then $\mathrm{NW}(S) = \mathrm{F}(S) = \{S\}$, so Part (i) of the proposition holds trivially. Moreover, $\mathrm{NW}(S) \setminus \mathrm{F}(S) = \varnothing$, so Part (ii) of the proposition holds vacuously. Hence, for the remainder of the proof, suppose $|S| > q$.

Part (i): Let $S', S'' \in \mathrm{F}(S)$. Toward a contradiction, suppose that S' and S'' are incomparable. By the upper-bound property, there exists $T \in \mathrm{NW}(S' \cup S'')$ such that $T \succ S'$ or $T \succ S''$. Since $S' \cup S'' \subseteq S$ and $|S' \cup S''| \geq q$, we have $T \in \mathrm{NW}(S)$. This contradicts either $S' \in \mathrm{F}(S)$ or $S'' \in \mathrm{F}(S)$. Thus S' and S'' must be comparable. Additionally, because neither strictly dominates the other, $S' \sim S''$.

Part (ii): Let $S' \in \mathrm{F}(S)$ and let $S'' \in \mathrm{NW}(S) \setminus \mathrm{F}(S)$. Since $S'' \notin \mathrm{F}(S)$, there exists $T_1 \in \mathrm{NW}(S)$ such that $T_1 \succ S''$. If $T_1 \in \mathrm{F}(S)$, then by Part (i) of the proposition, $S' \sim T_1$, and therefore $S' \succ S''$ by transitivity.

If $T_1 \notin \mathrm{F}(S)$, then we can repeat the above argument to obtain $T_2 \in \mathrm{NW}(S)$ with $T_2 \succ T_1$. Iteratively, whenever $T_k \notin \mathrm{F}(S)$, there exists $T_{k+1} \in \mathrm{NW}(S)$ such that $T_{k+1} \succ T_k$. This yields a sequence $T_1, T_2, \ldots$ in $\mathrm{NW}(S)$ such that

$$T_k \succ T_{k-1} \succ \ldots \succ T_1 \succ S''.$$

Since $\mathrm{NW}(S)$ is a finite set, and the strict part of a preorder is transitive and irreflexive, the sequence construction cannot continue indefinitely. Hence, there exists some K with $T_K \in \mathrm{F}(S)$; otherwise the sequence construction could be continued. By construction, $T_K \succ S''$, and by Part (i) of the proposition, $S' \sim T_K$. Therefore, by transitivity, $S' \succ S''$.

($\Leftarrow$). Suppose Parts (i) and (ii) of the proposition hold. Consider any two sets $S, S' \subseteq \mathcal{S}$ such that S and S' are incomparable and that $|S| = |S'| = q$. Hence, the antecedent for the upper-bound property holds.

We first show that $S \notin \mathrm{F}(S \cup S')$ and $S' \notin \mathrm{F}(S \cup S')$. Indeed, if both were in $\mathrm{F}(S \cup S')$, then Part (i) would imply that $S \sim S'$, contradicting incomparability. Similarly, if only one of them is in $\mathrm{F}(S \cup S')$, say $S \in \mathrm{F}(S \cup S')$ and $S' \notin \mathrm{F}(S \cup S')$, then Part (ii) would imply that $S \succ S'$, again contradicting incomparability. Hence, $S, S' \in \mathrm{NW}(S \cup S') \setminus \mathrm{F}(S \cup S')$.

We already know that $\mathrm{F}(S \cup S')$ is non-empty. Moreover, for every $T \in \mathrm{F}(S \cup S')$, Part (ii) implies that $T \succ S$ and $T \succ S'$, which establishes the existence of a set $T \in \mathrm{NW}(S \cup S')$ that dominates either S or S' as desired. □

Proof of Proposition 2.

($\Rightarrow$). Suppose the maximizer property holds. Consider any set $S \subseteq \mathcal{S}$. We shall show that $\mathrm{F}(S)$ satisfies the bases axioms: non-emptiness (B1) and exchange (B2) (see Appendix A).

By construction, $\mathrm{F}(S)$ is non-empty, so (B1) is satisfied. If $\mathrm{F}(S)$ is a singleton, then the exchange condition is satisfied vacuously. Thus, suppose $|\mathrm{F}(S)| > 1$. Then $|S| > q$, so every set in $\mathrm{F}(S)$ has cardinality q.

Take two distinct sets $S', S'' \in \mathrm{F}(S)$. Since the distributional preference satisfies the upper-bound property, Proposition 1 implies $S' \sim S''$. Moreover, for every $T \in \mathrm{NW}(S)$, Proposition 1 implies $S' \succsim T$ and $S'' \succsim T$. Since $\mathrm{NW}(S' \cup S'') \subseteq \mathrm{NW}(S)$, we have $S' \succsim T$ and $S'' \succsim T$ for all $T \in \mathrm{NW}(S' \cup S'')$. By the maximizer property, there exist $s' \in S' \setminus S''$ and $s'' \in S'' \setminus S'$ such that

$$(S' \setminus \{s'\}) \cup \{s''\} \sim S' \quad \text{and} \quad (S'' \setminus \{s''\}) \cup \{s'\} \sim S''.$$

Both exchanged sets belong to $\mathrm{NW}(S)$. Since each is indifferent to a frontier set, Proposition 1 implies that both exchanged sets also belong to $\mathrm{F}(S)$. Hence, $\mathrm{F}(S)$ satisfies the weak

simultaneous exchange property, which is equivalent to (B2) by Murota (2003, Theorem 4.3). Therefore, $\mathrm{F}(S)$ forms the bases of a matroid.

($\Leftarrow$). Consider two distinct sets $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$ such that $S \succsim S''$ and $S' \succsim S''$ for all $S'' \in \mathrm{NW}(S \cup S')$. By definition, $S, S' \in \mathrm{F}(S \cup S')$. Since $\mathrm{F}(S \cup S')$ forms the bases of a matroid for ground set $S \cup S'$, it satisfies the weak simultaneous exchange property: there exist $s \in S \setminus S'$ and $s' \in S' \setminus S$ such that

$$(S \setminus \{s\}) \cup \{s'\} \in \mathrm{F}(S \cup S') \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S \cup S').$$

Since the upper-bound property is satisfied, by Proposition 1, all frontier sets are indifferent to each other. Specifically, we have

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

Hence, $\succsim$ also satisfies the maximizer property. □

Proof of Proposition 3. We will use the following observation repeatedly: for any $S \subseteq \mathcal{S}$, $s \in S$, and $s' \notin S$ with $f(s') > f(s)$, we have $f\Big((S \setminus \{s\}) \cup \{s'\}\Big) > f(S)$, where for $x, y \in \mathbb{R}^q$, $x > y$ if $x_i \geq y_i$ for all $i \in \{1, \ldots, q\}$ and a strict inequality in at least one dimension.

Upper-bound Property: For any $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, either $\Psi(f(S)) \geq \Psi(f(S'))$ so that $S \succsim S'$, or $\Psi(f(S)) \leq \Psi(f(S'))$ so that $S' \succsim S$. Since no two sets of size q are incomparable, the upper-bound property holds vacuously.

Maximizer Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and suppose that $S \succsim S''$ and $S' \succsim S''$ for every $S'' \in \mathrm{NW}(S \cup S')$.

Take any $s \in S \setminus S'$ and any $s' \in S' \setminus S$. We claim that $f(s) = f(s')$. Suppose first that $f(s) > f(s')$. Then $f\Big((S' \setminus \{s'\}) \cup \{s\}\Big) > f(S')$, and by strict monotonicity of Ψ, we get $\Psi\Big(f((S' \setminus \{s'\}) \cup \{s\})\Big) > \Psi(f(S'))$. This implies $(S' \setminus \{s'\}) \cup \{s\} \succ S'$, contradicting the assumption that S' weakly dominates every element of $\mathrm{NW}(S \cup S')$. Similarly, if $f(s') > f(s)$, an analogous argument would yield $(S \setminus \{s\}) \cup \{s'\} \succ S$, contradicting the assumption that S weakly dominates every element of $\mathrm{NW}(S \cup S')$.

Therefore, $f(s) = f(s')$. Exchanging equal-score students leaves the profile of scores unchanged for both sets, so

$$\Psi\Big(f((S' \setminus \{s'\}) \cup \{s\})\Big) = \Psi(f(S')) \quad \text{and} \quad \Psi\Big(f((S \setminus \{s\}) \cup \{s'\})\Big) = \Psi(f(S)).$$

Consequently,

$$(S' \setminus \{s'\}) \cup \{s\} \sim S' \quad \text{and} \quad (S \setminus \{s\}) \cup \{s'\} \sim S.$$

This proves the maximizer property.

Coherence Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and let $\hat{s} \in S \setminus S'$ with $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$. Suppose that the three antecedents in Definition 4 hold.

Choose any $s \in (S \setminus S') \setminus \{\hat{s}\}$ and any $s' \in S' \setminus S$. Such an s' exists because $|S| = |S'|$ and $S \setminus S'$ is nonempty, which implies that $S' \setminus S$ is also nonempty. Since $s \neq \hat{s}$, the set $(S \setminus \{s\}) \cup \{s'\}$ contains $\hat{s}$ and belongs to $\mathrm{NW}(S \cup S')$. By condition (ii), $S \succsim (S \setminus \{s\}) \cup \{s'\}$. If $f(s') > f(s)$, then $\Psi\Big(f((S \setminus \{s\}) \cup \{s'\})\Big) > \Psi(f(S))$, contradicting the preceding weak dominance. Hence $f(s) \geq f(s')$.

Similarly, $(S' \setminus \{s'\}) \cup \{s\}$ belongs to $\mathrm{NW}((S \cup S') \setminus \{\hat{s}\})$. By condition (iii), $S' \succsim (S' \setminus \{s'\}) \cup \{s\}$. If $f(s) > f(s')$, then $\Psi\Big(f((S' \setminus \{s'\}) \cup \{s\})\Big) > \Psi(f(S'))$, contradicting the preceding weak dominance. Hence $f(s') \geq f(s)$.

Therefore, $f(s) = f(s')$. Exchanging equal-score students leaves the profile of scores unchanged for both sets, so

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

This proves the coherence property. □

Proof of Proposition 4.

Upper-bound Property: Let $S, S' \subseteq \mathcal{S}$ satisfy $|S| = |S'| = q$, and suppose that S and S' are incomparable. Let $T \subseteq S \cup S'$ be a set of q students with the highest f-values in $S \cup S'$. Then $T \in \mathrm{NW}(S \cup S')$. Moreover, by construction, T pointwise weakly dominates every set in $\mathrm{NW}(S \cup S')$, and in particular $T \succsim S$ and $T \succsim S'$. If neither comparison were strict, then we would have $T \sim S$ and $T \sim S'$, which would imply $S \sim S'$, contradicting incomparability. Therefore, T strictly dominates either S or S'. Hence, the upper-bound property holds.

Maximizer and Coherence Properties: The arguments are analogous to those for Proposition 3, with pointwise dominance of the ordered vector of f-values in place of strict monotonicity of Ψ. □

Proof of Proposition 5. Fix a matroid $M = (\mathcal{S}, \mathcal{I})$. We use the following standard exchange property of rank functions: for any two sets $X, Y \subseteq \mathcal{S}$ with $|X| = |Y|$ and any $x \in X \setminus Y$, there exists $y \in Y \setminus X$ such that

$$\mathrm{r}_M(X) + \mathrm{r}_M(Y) \leq \mathrm{r}_M\Big((X \setminus \{x\}) \cup \{y\}\Big) + \mathrm{r}_M\Big((Y \setminus \{y\}) \cup \{x\}\Big). \tag{3}$$

This is the fixed-cardinality exchange property for matroid rank functions, which is equivalent to r_M being M-concave when restricted to a domain of equi-cardinal sets (Shioura, 2012).

Upper-bound Property: For any $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, either $\mathrm{r}_M(S) \geq \mathrm{r}_M(S')$ so that $S \succsim S'$, or $\mathrm{r}_M(S') \geq \mathrm{r}_M(S)$ so that $S' \succsim S$. Since no two sets of size q are incomparable, the upper-bound property holds vacuously.

Maximizer Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and suppose that $S \succsim S''$ and $S' \succsim S''$ for every $S'' \in \mathrm{NW}(S \cup S')$. Equivalently, $\mathrm{r}_M(S) \geq \mathrm{r}_M(S'')$ and $\mathrm{r}_M(S') \geq \mathrm{r}_M(S'')$ for every $S'' \in \mathrm{NW}(S \cup S')$.

Choose any $s \in S \setminus S'$. By (3), there exists $s' \in S' \setminus S$ such that

$$\mathrm{r}_M(S) + \mathrm{r}_M(S') \leq \mathrm{r}_M\left((S \setminus \{s\}) \cup \{s'\}\right) + \mathrm{r}_M\left((S' \setminus \{s'\}) \cup \{s\}\right).$$

However, both S and S' maximize the rank function over $\mathrm{NW}(S \cup S')$, and both exchanged sets belong to $\mathrm{NW}(S \cup S')$. Consequently, the above inequality must bind, and

$$\mathrm{r}_M\left((S \setminus \{s\}) \cup \{s'\}\right) = \mathrm{r}_M(S) \quad \text{and} \quad \mathrm{r}_M\left((S' \setminus \{s'\}) \cup \{s\}\right) = \mathrm{r}_M(S').$$

Thus

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

This proves the maximizer property.

Coherence Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and let $\hat{s} \in S \setminus S'$ with $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$. Suppose that the three antecedents in Definition 4 hold.

Choose any $s \in (S \setminus S') \setminus \{\hat{s}\}$. By (3) applied to S and S', there exists $s' \in S' \setminus S$ such that

$$\mathrm{r}_M(S) + \mathrm{r}_M(S') \leq \mathrm{r}_M\left((S \setminus \{s\}) \cup \{s'\}\right) + \mathrm{r}_M\left((S' \setminus \{s'\}) \cup \{s\}\right).$$

Since $s \neq \hat{s}$, the set $(S \setminus \{s\}) \cup \{s'\}$ contains $\hat{s}$ and belongs to $\mathrm{NW}(S \cup S')$. By condition (ii), $S \succsim (S \setminus \{s\}) \cup \{s'\}$, or equivalently, $\mathrm{r}_M(S) \geq \mathrm{r}_M((S \setminus \{s\}) \cup \{s'\})$. Similarly, $(S' \setminus \{s'\}) \cup \{s\}$ belongs to $\mathrm{NW}((S \cup S') \setminus \{\hat{s}\})$. By condition (iii), $S' \succsim (S' \setminus \{s'\}) \cup \{s\}$, or equivalently, $\mathrm{r}_M(S') \geq \mathrm{r}_M((S' \setminus \{s'\}) \cup \{s\})$.

Together with the rank-exchange inequality, all the above inequalities must bind, i.e.,

$$\mathrm{r}_M\left((S \setminus \{s\}) \cup \{s'\}\right) = \mathrm{r}_M(S) \quad \text{and} \quad \mathrm{r}_M\left((S' \setminus \{s'\}) \cup \{s\}\right) = \mathrm{r}_M(S').$$

Thus

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

This proves the coherence property. □

Proof of Proposition 6.

Upper-bound Property: For any $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, there are four possibilities. Both are feasible, in which case $S \sim S'$; both infeasible, in which case $S \sim S'$; S is feasible while S' is infeasible, in which case $S \succ S'$; or S' is feasible while S is infeasible, in which case $S' \succ S$. Since no two sets of size q are incomparable, the upper-bound property holds vacuously.

Maximizer Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and suppose that $S \succsim S''$ and $S' \succsim S''$ for every $S'' \in \mathrm{NW}(S \cup S')$.

If S is infeasible, then no set in $\mathrm{NW}(S \cup S')$ is feasible; otherwise a feasible set would strictly dominate S. Hence, every set in $\mathrm{NW}(S \cup S')$ is infeasible, and any exchange between S and S' leaves both exchanged sets indifferent to the originals. Thus, the maximizer property holds.

Now suppose that S is feasible. Since $S' \succsim S$, the set S' is also feasible. Thus, $n_k(S) \in [L_k, U_k]$ and $n_k(S') \in [L_k, U_k]$ for all $k = 1, \ldots, K$.

Suppose $n_k(S) = n_k(S')$ for all k. Then there exist $s \in S \setminus S'$ and $s' \in S' \setminus S$ with $\tau(s) = \tau(s')$. Thus, $n_k(S) = n_k((S \setminus \{s\}) \cup \{s'\})$ and $n_k(S') = n_k((S' \setminus \{s'\}) \cup \{s\})$ for all k, so both exchanged sets remain feasible and hence indifferent to the originals.

Next, suppose $n_{k'}(S) \neq n_{k'}(S')$ for some k'. Without loss of generality, let $n_{k'}(S) > n_{k'}(S')$. Hence, there exists a student $s \in S \setminus S'$ with $\tau_{k'}(s) = 1$. Moreover, there exists $k'' \neq k'$ such that $n_{k''}(S) < n_{k''}(S')$; otherwise, because each student has only one protected trait, $|S| = \sum_k n_k(S) > \sum_k n_k(S') = |S'|$, which contradicts the fact that $|S| = |S'| = q$. Thus, there exists a student $s' \in S' \setminus S$ with $\tau_{k''}(s') = 1$.

Let us consider the exchanged set $(S \setminus \{s\}) \cup \{s'\}$. Note that

$$n_{k'}((S \setminus \{s\}) \cup \{s'\}) = n_{k'}(S) - 1 \geq n_{k'}(S') \geq L_{k'}$$

where the first inequality follows because $n_{k'}(S) > n_{k'}(S')$ by assumption, and the last inequality follows because S' is feasible. Similarly,

$$n_{k''}((S \setminus \{s\}) \cup \{s'\}) = n_{k''}(S) + 1 \leq n_{k''}(S') \leq U_{k''},$$

where the first inequality follows because $n_{k''}(S) < n_{k''}(S')$, and the last inequality follows because S' is feasible. For all other $k \notin \{k', k''\}$, $n_k((S \setminus \{s\}) \cup \{s'\}) = n_k(S)$. Thus, $(S \setminus \{s\}) \cup \{s'\}$ is also feasible. Consequently, $S \sim (S \setminus \{s\}) \cup \{s'\}$.

An analogous argument shows that $(S' \setminus \{s'\}) \cup \{s\}$ is also feasible and hence, $S' \sim (S' \setminus \{s'\}) \cup \{s\}$. Thus, the maximizer property holds. □

Proof of Proposition 7. The floor-and-ceiling distributional preference is an example of a

distributional preference that is induced by a q-ordinally concave value function. Hence, Proposition 7 follows from Proposition 8. □

Proof of Proposition 8.

Upper-bound Property: For any $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, either $\Phi(n(S)) \geq \Phi(n(S'))$, in which case $S \succsim S'$ or $\Phi(n(S')) \geq \Phi(n(S))$, in which case $S' \succsim S$. Since no two sets of size q are incomparable, the upper-bound property holds vacuously.

Maximizer Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and suppose that $S \succsim S''$ and $S' \succsim S''$ for every $S'' \in \text{NW}(S \cup S')$.

Choose any $s \in (S \setminus S')$ and suppose the student has trait t_k. If there exist $s' \in S' \setminus S$ such that $\tau_k(s') = 1$, then exchanging these students leaves trait-compositions unchanged. Hence,

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

Next, suppose no such student exists in $S' \setminus S$, which implies that $n_k(S) > n_k(S')$. By q-ordinal concavity, there exists $k' \neq k$ with $n_{k'}(S) < n_{k'}(S')$ such that at least one of the three conditions in Definition 5 holds. Since $n_{k'}(S) < n_{k'}(S')$, there exists $s' \in S' \setminus S$ with trait $t_{k'}$.

However, also notice that both S and S' are maximal in $\text{NW}(S \cup S')$, so we must also have $\Phi(n(S)) \geq \Phi(n((S \setminus \{s\}) \cup \{s'\}))$ and $\Phi(n(S')) \geq \Phi(n((S' \setminus \{s'\}) \cup \{s\}))$. In other words, Condition (i) and Condition (ii) of Definition 5 cannot hold. Thus, Condition (iii) of Definition 5 must hold, which is equivalent to

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

This proves the maximizer property.

Coherence Property: Let $S, S' \subseteq \mathcal{S}$ be distinct sets with $|S| = |S'| = q$, and let $\hat{s} \in S \setminus S'$ with $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$. Suppose that the three antecedents in Definition 4 hold.

Choose any $s \in (S \setminus S') \setminus \{\hat{s}\}$ and suppose the student has trait t_k. If there exist $s' \in S' \setminus S$ such that $\tau_k(s') = 1$, then exchanging these students leaves trait-compositions unchanged. Hence,

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

Next, suppose no such student exists in $S' \setminus S$, which implies that $n_k(S) > n_k(S')$. By q-ordinal concavity, there exists $k' \neq k$ with $n_{k'}(S) < n_{k'}(S')$ such that at least one of the three conditions in Definition 5 holds. Since $n_{k'}(S) < n_{k'}(S')$, there exists $s' \in S' \setminus S$ with trait $t_{k'}$.

Since $s \neq \hat{s}$, the set $(S \setminus \{s\}) \cup \{s'\}$ contains $\hat{s}$ and belongs to $\mathrm{NW}(S \cup S')$. By condition (ii), $S \succsim (S \setminus \{s\}) \cup \{s'\}$, or equivalently, $\Phi(n(S)) \geq \Phi(n((S \setminus \{s\}) \cup \{s'\}))$. Hence, Condition (i) of Definition 5 is not satisfied.

Similarly, $(S' \setminus \{s'\}) \cup \{s\}$ belongs to $\mathrm{NW}((S \cup S') \setminus \{\hat{s}\})$. By condition (iii), $S' \succsim (S' \setminus \{s'\}) \cup \{s\}$, or equivalently, $\Phi(n(S')) \geq \Phi(n((S' \setminus \{s'\}) \cup \{s\}))$. Hence, Condition (ii) of Definition 5 is also not satisfied.

Thus, Condition (iii) of Definition 5 must hold, which is equivalent to

$$(S \setminus \{s\}) \cup \{s'\} \sim S \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \sim S'.$$

This proves the coherence property. □

Appendix C. Proofs of Theorems

Proof of Theorem 1. We first show that (i) implies (ii). Let the distributional preference satisfy the maximizer property. Fix a priority ranking π and consider the distributional choice rule Ch^π. In Section 2.4, we have already discussed that Ch^π is non-wasteful, distributionally maximal, and free of justified envy. Thus, we need only show that Ch^π is the unique choice rule that satisfies these three axioms.

To that end, let Ch be any choice rule that is non-wasteful, distributionally maximal, and free of justified envy. By Lemma 1, $\mathrm{Ch}(S) \in \mathrm{F}(S)$ for every $S \subseteq \mathcal{S}$. Suppose, toward a contradiction, that there exists $S \subseteq \mathcal{S}$ such that $\mathrm{Ch}(S) \neq \mathrm{Ch}^\pi(S)$. Such a set S is necessarily non-empty.

Since the distributional preference satisfies the upper-bound and maximizer properties, Proposition 2 implies that $\mathrm{F}(S)$ forms the bases of a matroid. By the weak simultaneous exchange property (see Appendix A), there exist $s \in \mathrm{Ch}(S) \setminus \mathrm{Ch}^\pi(S)$ and $s' \in \mathrm{Ch}^\pi(S) \setminus \mathrm{Ch}(S)$ such that

$$(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \in \mathrm{F}(S) \quad \text{and} \quad (\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S).$$

By Proposition 1, $(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \sim \mathrm{Ch}(S)$ and $(\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\} \sim \mathrm{Ch}^\pi(S)$.

Since Ch is free of justified envy, $(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \sim \mathrm{Ch}(S)$ implies that $s \;\pi\; s'$. At the same time, since Ch^π is also free of justified envy, $(\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\} \sim \mathrm{Ch}^\pi(S)$ implies $s' \;\pi\; s$. However, as the priority ranking π is a strict linear order on $\mathcal{S}$, we cannot have both $s \;\pi\; s'$ and $s' \;\pi\; s$. Therefore, no such set S exists, and $\mathrm{Ch} = \mathrm{Ch}^\pi$.

We now show that (ii) implies (i). For each priority ranking π, let Ch^π be the unique choice rule that is non-wasteful, distributionally maximal, and free of justified envy. Toward

a contradiction, suppose that the maximizer property does not hold. Since the upper-bound property is satisfied, Proposition 2 implies that there exists a set $T \subseteq \mathcal{S}$ such that $\mathrm{F}(T)$ does not form the bases of a matroid with ground set T.

Since $\mathrm{F}(T)$ is non-empty, it must violate the matroid base-exchange axiom (B2), i.e., there exist distinct sets $S, S' \in \mathrm{F}(T)$ and some $i \in S \setminus S'$ such that for every $s' \in S' \setminus S$, we have

$$(S \setminus \{i\}) \cup \{s'\} \notin \mathrm{F}(T).$$

Among all tuples (T, S, S', i) that violate (B2), choose the tuple for which $|S \setminus S'|$ is minimal.

Since $S \in \mathrm{F}(T)$ and $(S \setminus \{i\}) \cup \{s'\} \in \mathrm{NW}(T) \setminus \mathrm{F}(T)$ for every $s' \in S' \setminus S$, Proposition 1 implies

$$S \succ (S \setminus \{i\}) \cup \{s'\} \qquad \text{for every } s' \in S' \setminus S. \tag{4}$$

Moreover, $|S \setminus S'| \geq 2$. Indeed, if $S \setminus S' = \{i\}$, then exchanging i with the unique student $s' \in S' \setminus S$ yields $(S \setminus \{i\}) \cup \{s'\} = S' \sim S$, which contradicts (4).

Claim 1. *For every $j \in (S \setminus S') \setminus \{i\}$ and every $k \in S' \setminus S$,*

$$S' \succ (S' \setminus \{k\}) \cup \{j\}. \tag{5}$$

Proof. Toward a contradiction, suppose there exists some such j and k such that $S' \not\succ (S' \setminus \{k\}) \cup \{j\}$. Let $\widetilde{S}' := (S' \setminus \{k\}) \cup \{j\}$. Since $S' \in \mathrm{F}(T)$ and $\widetilde{S}' \in \mathrm{NW}(T)$, by Proposition 1, the fact that $S' \not\succ \widetilde{S}'$ implies that $\widetilde{S}' \in \mathrm{F}(T)$.

Moreover, $i \in S \setminus \widetilde{S}'$, and for every $\tilde{s} \in \widetilde{S}' \setminus S = (S' \setminus S) \setminus \{k\}$, (4) gives

$$S \succ (S \setminus \{i\}) \cup \{\tilde{s}\},$$

which implies that $(S \setminus \{i\}) \cup \{\tilde{s}\} \notin \mathrm{F}(T)$. Thus, the tuple $(T, S, \widetilde{S}', i)$ also violates the base-exchange condition. However, $\left|S \setminus \widetilde{S}'\right| = |S \setminus S'| - 1$, contradicting the fact that the tuple (T, S, S', i) was chosen to minimize $|S \setminus S'|$. Therefore (5) holds. □

We now construct a priority ranking π as follows: Rank students in $S \cap S'$ first, then students in $(S \setminus S') \setminus \{i\}$, then students in $S' \setminus S$, then i, and finally rank all students outside $S \cup S'$ last. Within each block, the students are ranked in arbitrary order. Thus, for every $\tilde{s} \in S \cap S'$, $s \in (S \setminus S') \setminus \{i\}$, $s' \in S' \setminus S$, and $t \in T \setminus (S \cup S')$, we have

$$\tilde{s} \;\pi\; s \;\pi\; s' \;\pi\; i \;\pi\; t.$$

We first verify that $\mathrm{Ch}^{\pi}(T) = S$. The greedy algorithm admits every student in $S \cap S'$,

since this set is contained in the frontier set S. It then admits every student in $(S \setminus S') \setminus \{i\}$, since the tentative set remains contained in S. At that point, the tentative set is $S \setminus \{i\}$. For any $s' \in S' \setminus S$, adding s' to the tentative set would produce $(S \setminus \{i\}) \cup \{s'\}$, which is not in $\mathrm{F}(T)$ by (4). Hence, every student in $S' \setminus S$ is rejected. Finally, i is considered and admitted because $(S \setminus \{i\}) \cup \{i\} = S \in \mathrm{F}(T)$. Once i is admitted, the tentative set is S, which has cardinality q; hence no later student can be admitted. Thus $\mathrm{Ch}^\pi(T) = S$.

Define a new choice rule Ch by

$$\mathrm{Ch}(T') = \begin{cases} S' & \text{if } T' = T, \\ \mathrm{Ch}^\pi(T') & \text{otherwise.} \end{cases}$$

Because $\mathrm{Ch}(T) = S'$ whereas $\mathrm{Ch}^\pi(T) = S$, $\mathrm{Ch} \neq \mathrm{Ch}^\pi$.

Notice that for every input $T' \subseteq \mathcal{S}$, $\mathrm{Ch}(T') \in \mathrm{F}(T')$. Indeed, for input $T' \neq T$, $\mathrm{Ch}(T') = \mathrm{Ch}^\pi(T')$, and $\mathrm{Ch}^\pi(T') \in \mathrm{F}(T')$ by construction. For input $T' = T$, $\mathrm{Ch}(T) = S' \in \mathrm{F}(T)$. Therefore, by Lemma 1, Ch is non-wasteful and distributionally maximal.

It remains to verify that Ch is free of justified envy. For every input $T' \neq T$, this follows because Ch coincides with Ch^π, which is free of justified envy. Hence, it suffices to check the input $T' = T$.

Let $k \in \mathrm{Ch}(T) = S'$ and $j \in T \setminus S'$ be such that $(S' \setminus \{k\}) \cup \{j\} \succsim S'$ so that the antecedent of the free-of-justified-envy axiom holds. We show that $k \; \pi \; j$. There are four cases.

Case 1: $j = i$. Since i is ranked below every student in S', we have $k \; \pi \; j$.

Case 2: $j \in (S \setminus S') \setminus \{i\}$ *and* $k \in S \cap S'$. By construction of π, every student in $S \cap S'$ is ranked above every student in $(S \setminus S') \setminus \{i\}$. Hence $k \; \pi \; j$.

Case 3: $j \in (S \setminus S') \setminus \{i\}$ *and* $k \in S' \setminus S$. By (5), $S' \succ (S' \setminus \{k\}) \cup \{j\}$, which contradicts the maintained hypothesis $(S' \setminus \{k\}) \cup \{j\} \succsim S'$. Thus this case cannot occur.

Case 4: $j \in T \setminus (S \cup S')$. By construction of π, every student in S' is ranked above every student in $T \setminus (S \cup S')$. Hence $k \; \pi \; j$.

In every case in which the antecedent of the free-of-justified-envy axiom is satisfied, the selected student k has higher priority than the rejected student j. Therefore Ch is free of justified envy.

We have therefore constructed a choice rule $\mathrm{Ch} \neq \mathrm{Ch}^\pi$ that is non-wasteful, distributionally maximal, and free of justified envy. This contradicts the initial assertion that Ch^π is the unique such choice rule for every priority ranking. To avoid such a contradiction, the maximizer property must hold. □

Proof of Theorem 2. Let Ch, π, and S be as in the statement of Theorem 2. Suppose that condition (i) does not hold, i.e., $|\mathrm{Ch}^\pi(S)| \leq |\mathrm{Ch}(S)|$. We show that either condition (ii) or condition (iii) holds.

Because $\mathrm{Ch}(S) \subseteq S$, we have $|\mathrm{Ch}(S)| \leq |S|$. Additionally, all choice rules must respect the capacity, so $|\mathrm{Ch}(S)| \leq q$. Hence, $|\mathrm{Ch}(S)| \leq \min\{q, |S|\}$. Since Ch^π is non-wasteful, we have $\min\{q, |S|\} = |\mathrm{Ch}^\pi(S)| \leq |\mathrm{Ch}(S)|$, implying that $|\mathrm{Ch}(S)| = \min\{q, |S|\}$.

If $|S| \leq q$, then $|\mathrm{Ch}^\pi(S)| = |\mathrm{Ch}(S)| = |S|$, and hence $\mathrm{Ch}^\pi(S) = \mathrm{Ch}(S) = S$, contradicting the premise that $\mathrm{Ch}(S) \neq \mathrm{Ch}^\pi(S)$. Therefore, $|S| > q$. We thus have $|\mathrm{Ch}^\pi(S)| = |\mathrm{Ch}(S)| = q$, implying that $\mathrm{Ch}(S) \in \mathrm{NW}(S)$.

If $\mathrm{Ch}(S) \notin \mathrm{F}(S)$, then Proposition 1 implies $\mathrm{Ch}^\pi(S) \succ \mathrm{Ch}(S)$, because $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S)$ and $\mathrm{Ch}(S) \in \mathrm{NW}(S) \setminus \mathrm{F}(S)$. Hence condition (ii) holds.

Now suppose $\mathrm{Ch}(S) \in \mathrm{F}(S)$. By Proposition 2, $\mathrm{F}(S)$ forms the bases of a matroid. In this matroid, a set is independent if and only if it is contained in some frontier set, which is exactly the algorithm used by Ch^π at each step. Hence $\mathrm{Ch}^\pi(S)$ is the greedy basis of this matroid under π. Since $\mathrm{Ch}(S) \in \mathrm{F}(S)$ is also a basis, and hence an independent set, Lemma 2 implies that $\mathrm{Ch}^\pi(S)$ priority dominates $\mathrm{Ch}(S)$. Hence condition (iii) holds.

Thus, whenever condition (i) fails, either condition (ii) or condition (iii) holds. Therefore at least one of the three conditions in Theorem 2 must hold. □

Proof of Theorem 3. We first show that (i) implies (ii). Suppose the distributional preference satisfies the maximizer and coherence properties. We show that, for every priority ranking π, the distributional choice rule Ch^π satisfies path independence. We use the following characterization of path independence due to Aizerman and Malishevski (1981): a choice rule Ch is path independent if and only if, for every $S \subseteq \mathcal{S}$ and every $s \in S$,

- (a) **Irrelevance of rejected students (IRS):** If $s \notin \mathrm{Ch}(S)$, then $\mathrm{Ch}(S \setminus \{s\}) = \mathrm{Ch}(S)$.
- (b) **Substitutability:** If $s \in \mathrm{Ch}(S)$, then for any $\hat{s} \in S \setminus \{s\}$, $s \in \mathrm{Ch}(S \setminus \{\hat{s}\})$.

We shall show that the distributional choice rule Ch^π satisfies IRS and substitutability, and hence, is path independent.

IRS. Consider an applicant set $S \subseteq \mathcal{S}$, and suppose $s \in S \setminus \mathrm{Ch}^\pi(S)$. Since Ch^π is non-wasteful, this implies $|S| > q$. Hence both S and $S \setminus \{s\}$ have non-wasteful sets of cardinality q.

By Lemma 1, $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S)$. Since $s \notin \mathrm{Ch}^\pi(S)$, the set $\mathrm{Ch}^\pi(S)$ belongs to $\mathrm{NW}(S \setminus \{s\})$. Moreover, because $\mathrm{Ch}^\pi(S)$ is maximal in $\mathrm{NW}(S)$, it remains maximal in the smaller set $\mathrm{NW}(S \setminus \{s\})$, implying $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S \setminus \{s\})$.

Next, consider the applicant set $S \setminus \{s\}$. By Lemma 1, $\mathrm{Ch}^\pi(S \setminus \{s\}) \in \mathrm{F}(S \setminus \{s\})$. Since both $\mathrm{Ch}^\pi(S \setminus \{s\})$ and $\mathrm{Ch}^\pi(S)$ belong to $\mathrm{F}(S \setminus \{s\})$, Proposition 1 implies $\mathrm{Ch}^\pi(S \setminus \{s\}) \sim \mathrm{Ch}^\pi(S)$.

Consequently, $\mathrm{Ch}^\pi(S \setminus \{s\}) \in \mathrm{F}(S)$; otherwise, $\mathrm{Ch}^\pi(S)$ would strictly dominate $\mathrm{Ch}^\pi(S \setminus \{s\})$ by Proposition 1, contradicting the indifference.

By Proposition 2, $\mathrm{F}(S \setminus \{s\})$ forms the bases of a matroid on $S \setminus \{s\}$. Under this matroid, $\mathrm{Ch}^\pi(S \setminus \{s\})$ is the greedy basis, while $\mathrm{Ch}^\pi(S)$ is also a basis. Applying Lemma 2, $\mathrm{Ch}^\pi(S \setminus \{s\})$ priority dominates $\mathrm{Ch}^\pi(S)$.

Similarly, $\mathrm{F}(S)$ forms the bases of a matroid on S. Under this matroid, $\mathrm{Ch}^\pi(S)$ is the greedy basis, while $\mathrm{Ch}^\pi(S \setminus \{s\})$ is also a basis. Applying Lemma 2 again, $\mathrm{Ch}^\pi(S)$ priority dominates $\mathrm{Ch}^\pi(S \setminus \{s\})$.

Since the two sets priority dominate one another and π is a strict linear order, they must be equal. Hence Ch^π satisfies IRS.

Substitutability. Let $S \subseteq \mathcal{S}$, let $s \in \mathrm{Ch}^\pi(S)$, and let $\hat{s} \in S \setminus \{s\}$. If $\hat{s} \notin \mathrm{Ch}^\pi(S)$, then IRS gives $\mathrm{Ch}^\pi(S \setminus \{\hat{s}\}) = \mathrm{Ch}^\pi(S)$, and the claim follows. Hence, suppose that $\hat{s} \in \mathrm{Ch}^\pi(S)$. If $|S \setminus \{\hat{s}\}| \leq q$, then non-wastefulness gives $\mathrm{Ch}^\pi(S \setminus \{\hat{s}\}) = S \setminus \{\hat{s}\}$, so the claim again follows trivially. Thus, also assume that $|S \setminus \{\hat{s}\}| > q$.

Set $A := \mathrm{Ch}^\pi(S)$ and $B := \mathrm{Ch}^\pi(S \setminus \{\hat{s}\})$. Suppose, toward a contradiction, that $s \notin B$. We shall consider two cases and show each case leads to a contradiction. Since $|S \setminus \{\hat{s}\}| > q$ and Ch^π is non-wasteful, we have $|A| = |B| = q$. Additionally, $s, \hat{s} \in A \setminus B$ and $s \in (A \setminus B) \setminus \{\hat{s}\}$, so $(A \setminus B) \setminus \{\hat{s}\} \neq \varnothing$.

Case 1: $B \notin \mathrm{F}(S)$. Since $A \in \mathrm{F}(S)$ and $B \notin \mathrm{F}(S)$, Proposition 1 gives $A \succ B$. Recall that $|S \setminus \{\hat{s}\}| > q$ and $(A \setminus B) \setminus \{\hat{s}\} \neq \varnothing$. Thus, the domain restrictions for the coherence property are satisfied. We verify the three antecedents of the coherence property for the pair A, B and the student $\hat{s}$.

Take any $T \in \mathrm{NW}(A \cup B)$ with $\hat{s} \notin T$. Since $|A| = |B| = q$ and $(A \cup B) \setminus \{\hat{s}\} \subseteq S \setminus \{\hat{s}\}$, we have $T \in \mathrm{NW}(S \setminus \{\hat{s}\})$. Since $B \in \mathrm{F}(S \setminus \{\hat{s}\})$, Proposition 1 gives $B \succsim T$. Combined with $A \succ B$ and transitivity, we conclude that $A \succ T$ for all $T \in \mathrm{NW}(A \cup B)$ with $\hat{s} \notin T$. Hence, the first antecedent holds.

Since A is maximal in $\mathrm{NW}(S)$ and $A \cup B \subseteq S$, A remains maximal in the smaller set $\mathrm{NW}(A \cup B)$, i.e., $A \in \mathrm{F}(A \cup B)$. By Proposition 1, $A \succsim T$ for all $T \in \mathrm{NW}(A \cup B)$, so the second antecedent holds.

Similarly, since B is maximal in $\mathrm{NW}(S \setminus \{\hat{s}\})$ and $(A \cup B) \setminus \{\hat{s}\} \subseteq S \setminus \{\hat{s}\}$, B remains maximal in the smaller set $\mathrm{NW}((A \cup B) \setminus \{\hat{s}\})$, i.e., $B \in \mathrm{F}((A \cup B) \setminus \{\hat{s}\})$. By Proposition 1, $B \succsim T$ for all $T \in \mathrm{NW}((A \cup B) \setminus \{\hat{s}\})$, so the third antecedent holds.

We can thus apply the coherence property: there exists $t \in (A \setminus B) \setminus \{\hat{s}\}$ and $t' \in B \setminus A$ such that

$$(A \setminus \{t\}) \cup \{t'\} \sim A \quad \text{and} \quad (B \setminus \{t'\}) \cup \{t\} \sim B.$$

Since $A \in \mathrm{F}(S)$ and $(A \setminus \{t\}) \cup \{t'\} \in \mathrm{NW}(S)$, the first indifference and Proposition 1 imply $(A \setminus \{t\}) \cup \{t'\} \in \mathrm{F}(S)$. Additionally, notice that $(B \setminus \{t'\}) \cup \{t\} \in \mathrm{NW}(S \setminus \{\hat{s}\})$, because $B \setminus \{t'\} \subseteq B \subseteq S \setminus \{\hat{s}\}$ and $t \neq \hat{s}$. Since $B \in \mathrm{F}(S \setminus \{\hat{s}\})$, the second indifference and Proposition 1 imply $(B \setminus \{t'\}) \cup \{t\} \in \mathrm{F}(S \setminus \{\hat{s}\})$.

By Proposition 2, $\mathrm{F}(S \setminus \{\hat{s}\})$ forms the bases of a matroid on $S \setminus \{\hat{s}\}$. Under this matroid, B is the greedy basis, while $(B \setminus \{t'\}) \cup \{t\}$ is also a basis. Applying Lemma 2, B priority dominates $(B \setminus \{t'\}) \cup \{t\}$. Because these two sets differ only by replacing t' with t, this priority dominance implies $t' \; \pi \; t$.

Similarly, by Proposition 2, $\mathrm{F}(S)$ forms the bases of a matroid on S. Under this matroid, A is the greedy basis, while $(A \setminus \{t\}) \cup \{t'\}$ is also a basis. Applying Lemma 2, A priority dominates $(A \setminus \{t\}) \cup \{t'\}$. Because these two sets differ only by replacing t with t', this priority dominance implies $t \; \pi \; t'$. However, this contradicts the asymmetry of the strict linear order π.

Case 2: $B \in \mathrm{F}(S)$. By Proposition 2, $\mathrm{F}(S)$ forms the bases of a matroid on S. Since A and B are both in $\mathrm{F}(S)$ and since $s \in A \setminus B$, the strong basis exchange property gives some $s' \in B \setminus A$ such that

$$(A \setminus \{s\}) \cup \{s'\} \in \mathrm{F}(S) \quad \text{and} \quad (B \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S).$$

Since $B \setminus \{s'\} \subseteq B \subseteq S \setminus \{\hat{s}\}$ and $s \in (A \setminus B) \setminus \{\hat{s}\}$, we have that $(B \setminus \{s'\}) \cup \{s\} \subseteq S \setminus \{\hat{s}\}$. Additionally, since $|B| = |(B \setminus \{s'\}) \cup \{s\}| = q$, we conclude that $(B \setminus \{s'\}) \cup \{s\} \in \mathrm{NW}(S \setminus \{\hat{s}\})$. In fact, since $(B \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S)$, it remains maximal in $\mathrm{NW}(S \setminus \{\hat{s}\})$. Thus, $(B \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S \setminus \{\hat{s}\})$.

By Proposition 2, $\mathrm{F}(S \setminus \{\hat{s}\})$ forms the bases of a matroid on $S \setminus \{\hat{s}\}$. Under this matroid, B is the greedy basis, while $(B \setminus \{s'\}) \cup \{s\}$ is also a basis. Applying Lemma 2, B priority dominates $(B \setminus \{s'\}) \cup \{s\}$. Because these two sets differ only by replacing s' with s, this priority dominance implies $s' \; \pi \; s$.

On the other hand, under the matroid whose bases are $\mathrm{F}(S)$, A is the greedy basis, while $(A \setminus \{s\}) \cup \{s'\}$ is also a basis. Applying Lemma 2, A priority dominates $(A \setminus \{s\}) \cup \{s'\}$. Because these two sets differ only by replacing s with s', this priority dominance implies $s \; \pi \; s'$. However, this contradicts the asymmetry of the strict linear order π. Again we obtain a contradiction.

The two cases exhaust all possibilities, so the supposition that $s \notin \mathrm{Ch}^{\pi}(S \setminus \{\hat{s}\})$ is false. As a result, Ch^{π} satisfies substitutability. Since Ch^{π} also satisfies IRS, the characterization of Aizerman and Malishevski (1981) implies that Ch^{π} satisfies path independence.

We now prove that (ii) implies (i). Suppose that, for every priority ranking π, the distri-

butional choice rule Ch^π satisfies path independence. We first show that the distributional preference satisfies the maximizer property. Then, taking the maximizer property as a given, we show that the distributional preference also satisfies the coherence property.

Path independence $\implies$ *maximizer.* Toward a contradiction, suppose that the maximizer property does not hold. Since the upper-bound property is satisfied, Proposition 2 implies that there exists a set $T \subseteq \mathcal{S}$ such that $\text{F}(T)$ does not form the bases of a matroid with ground set T.

Since $\text{F}(T)$ is non-empty, it must violate the matroid base-exchange axiom (B2), i.e., there exist distinct sets $S, S' \in \text{F}(T)$ and some $s \in S \setminus S'$ such that for every $s' \in S' \setminus S$, we have

$$(S \setminus \{s\}) \cup \{s'\} \notin \text{F}(T).$$

Choose a priority ranking π that ranks all students in $S \setminus \{s\}$ first, then all students in $S' \setminus S$, then s, and finally all students in $T \setminus (S \cup S')$, with an arbitrary ranking across students within each block. Under this priority ranking, the distributional choice rule applied to T selects $\text{Ch}^\pi(T) = S$. Indeed, every student in $S \setminus \{s\}$ is admitted because the set selected so far remains contained in the frontier set S. Each student $s' \in S' \setminus S$ is rejected because $(S \setminus \{s\}) \cup \{s'\}$ has cardinality q and is not in $\text{F}(T)$. Finally, s is admitted because $S \in \text{F}(T)$.

Since Ch^π is path independent, it satisfies IRS and substitutability. Because no student in $T \setminus (S \cup S')$ belongs to $\text{Ch}^\pi(T) = S$, repeated application of IRS gives $\text{Ch}^\pi(T) = \text{Ch}^\pi(S \cup S') = S$. Now remove s from $S \cup S'$. By substitutability, every student in $S \setminus \{s\}$ must remain chosen, so $S \setminus \{s\} \subseteq \text{Ch}^\pi((S \cup S') \setminus \{s\})$.

Because $S' \subseteq (S \cup S') \setminus \{s\}$, the reduced applicant pool has at least q students, so $|\text{Ch}^\pi((S \cup S') \setminus \{s\})| = q$. Hence, there exists some $s^* \in S' \setminus S$ such that

$$\text{Ch}^\pi((S \cup S') \setminus \{s\}) = (S \setminus \{s\}) \cup \{s^*\}.$$

By construction, the distributional choice rule always picks a set in the frontier. Thus, we must have $(S \setminus \{s\}) \cup \{s^*\} \in \text{F}((S \cup S') \setminus \{s\})$.

Additionally, notice that $S' \in \text{F}((S \cup S') \setminus \{s\})$; because S' is maximal in T, it remains so in the reduced pool $(S \cup S') \setminus \{s\}$. By Proposition 1, $S' \sim (S \setminus \{s\}) \cup \{s^*\}$. Using the fact that $S' \in \text{F}(T)$ and $(S \setminus \{s\}) \cup \{s^*\} \in \text{NW}(T)$, we once again apply Proposition 1 to get that $(S \setminus \{s\}) \cup \{s^*\} \in \text{F}(T)$. However, this contradicts the assumption that, for every $s' \in S' \setminus S$,

$$(S \setminus \{s\}) \cup \{s'\} \notin \text{F}(T).$$

Hence, to avoid the contradiction, the maximizer property must hold.

Path independence + maximizer $\Longrightarrow$ coherence. Take $S, S' \subseteq \mathcal{S}$ with $|S| = |S'| = q$, and take $\hat{s} \in S \setminus S'$ satisfying $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$. Suppose S, S' and $\hat{s}$ satisfy the three antecedents in Definition 4:

$$S \succ T \text{ for every } T \in \mathrm{NW}(S \cup S') \text{ with } \hat{s} \notin T, \tag{6}$$

$$S \succsim T \text{ for every } T \in \mathrm{NW}(S \cup S') \text{ with } \hat{s} \in T, \tag{7}$$

$$S' \succsim T \text{ for every } T \in \mathrm{NW}\Big((S \cup S') \setminus \{\hat{s}\}\Big). \tag{8}$$

Let $U := S \cup S'$ and $W := U \setminus \{\hat{s}\}$. Clearly, $q \leq |W| \leq |U|$. By (6) and (7), $S \in \mathrm{F}(U)$: for every $T \in \mathrm{NW}(U)$, either $\hat{s} \notin T$ and $S \succ T$ by (6), or $\hat{s} \in T$ and $S \succsim T$ by (7). Thus, no element of $\mathrm{NW}(U)$ strictly dominates S. Moreover, every set in $\mathrm{F}(U)$ contains $\hat{s}$: if $T \in \mathrm{NW}(U)$ and $\hat{s} \notin T$, then $S \succ T$ by (6), so $T \notin \mathrm{F}(U)$. Finally, (8) implies $S' \in \mathrm{F}(W)$, because $S' \in \mathrm{NW}(W)$ and S' weakly dominates every element of $\mathrm{NW}(W)$.

By Proposition 2, $\mathrm{F}(U)$ forms the bases of a matroid on U. Denote this matroid by M, and note that it has rank q. Since every basis of M contains $\hat{s}$, contracting M by $\hat{s}$ yields a matroid on W whose bases are

$$\mathcal{B}^c := \{T \setminus \{\hat{s}\} : T \in \mathrm{F}(U)\}.$$

Denote this matroid by M^c, and note that M^c has rank $q - 1$.

Separately, by Proposition 2, $\mathrm{F}(W)$ forms the bases of a matroid on W. Denote this matroid by M^d, and note that M^d has rank q.

Fix an arbitrary priority ranking ρ on W. Extend ρ to a priority ranking π on U by ranking $\hat{s}$ above every element of W and ordering the elements of W according to ρ. Let $G^{\rho}_{M^c}$ and $G^{\rho}_{M^d}$ denote the greedy basis under ρ of matroids M^c and M^d, respectively.

Because every set in $\mathrm{F}(U)$ contains $\hat{s}$, and because $\hat{s}$ has the highest priority under π, Ch^{π} applied to U first admits $\hat{s}$. Then, the choice rule proceeds along W in ρ-order, admitting at each step the highest-ρ element compatible with extending to a set in $\mathrm{F}(U)$. Thus, once $\hat{s}$ has been admitted, the choice rule is precisely the greedy algorithm on M^c under ρ, so

$$\mathrm{Ch}^{\pi}(U) = \{\hat{s}\} \cup G^{\rho}_{M^c}. \tag{9}$$

Notice that $\mathrm{Ch}^{\pi} = \mathrm{Ch}^{\rho}$ on W since ρ and π coincide on W. When applying Ch^{π} to W, the choice rule proceeds along W in ρ-order, admitting at each step the highest-ρ element compatible with extending to a set in $\mathrm{F}(W)$. Hence,

$$\mathrm{Ch}^{\pi}(W) = G^{\rho}_{M^d}. \tag{10}$$

By assumption, Ch^π satisfies path independence. Thus,

$$\mathrm{Ch}^\pi(U) = \mathrm{Ch}^\pi(W \cup \{\hat{s}\}) = \mathrm{Ch}^\pi(\mathrm{Ch}^\pi(W) \cup \{\hat{s}\})$$

By (9) and (10), we get

$$\{\hat{s}\} \cup G^\rho_{M^c} = \mathrm{Ch}^\pi(G^\rho_{M^d} \cup \{\hat{s}\}) \subseteq \{\hat{s}\} \cup G^\rho_{M^d},$$

where the set inclusion follows by definition of a choice rule. Therefore, $G^\rho_{M^c} \subseteq G^\rho_{M^d}$. Since ρ was arbitrary, we have that $G^\rho_{M^c} \subseteq G^\rho_{M^d}$ for every priority ranking ρ on W. By Lemma 3, we can conclude that M^c is a quotient of M^d. Moreover, recall that M^c has rank $q-1$ and M^d has rank q. Thus, M^c is a rank-one quotient of M^d.

Now observe that $S \setminus \{\hat{s}\} \in \mathcal{B}^c$, $S' \in \mathrm{F}(W)$, and $(S \setminus S') \setminus \{\hat{s}\} \neq \varnothing$ by assumption. Applying Lemma 4, there exist $s \in (S \setminus \{\hat{s}\}) \setminus S'$ and $s' \in S' \setminus S$ such that

$$((S \setminus \{\hat{s}\}) \setminus \{s\}) \cup \{s'\} \in \mathcal{B}^c \quad \text{and} \quad (S' \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(W).$$

By the definition of $\mathcal{B}^c$, adding $\hat{s}$ to the first set yields $(S \setminus \{s\}) \cup \{s'\} \in \mathrm{F}(U)$. Since $S \in \mathrm{F}(U)$, Proposition 1 implies $(S \setminus \{s\}) \cup \{s'\} \sim S$. Since $S' \in \mathrm{F}(W)$ and $(S' \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(W)$, Proposition 1 also implies $(S' \setminus \{s'\}) \cup \{s\} \sim S'$. Thus the conclusion of the coherence property holds. Therefore, the distributional preference satisfies the coherence property. □

Proof of Theorem 4.

(⇐). Consider Ch^π for some priority ranking π. By construction, $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S)$ for every $S \subseteq \mathcal{S}$. Therefore, by Lemma 1, Ch^π is non-wasteful and distributionally maximal.

We show that Ch^π satisfies the revealed priority axiom. Let $S \subseteq \mathcal{S}$, $s \in \mathrm{Ch}^\pi(S)$, and $s' \in S \setminus \mathrm{Ch}^\pi(S)$, and suppose that

$$\Big(\mathrm{Ch}^\pi(S) \setminus \{s\}\Big) \cup \{s'\} \succsim \mathrm{Ch}^\pi(S).$$

Since $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S)$ and the swapped set belongs to $\mathrm{NW}(S)$, Proposition 1 implies that $\Big(\mathrm{Ch}^\pi(S) \setminus \{s\}\Big) \cup \{s'\} \in \mathrm{F}(S)$. By Proposition 2, $\mathrm{F}(S)$ forms the bases of a matroid. Since $\mathrm{Ch}^\pi(S)$ is the greedy basis of this matroid under π, Lemma 2 implies that $\mathrm{Ch}^\pi(S)$ priority dominates $\Big(\mathrm{Ch}^\pi(S) \setminus \{s\}\Big) \cup \{s'\}$. These two sets differ only by replacing s with s', so priority domination implies $s \; \pi \; s'$.

Now suppose, toward a contradiction, that there exist sequences $\{S_k\}_{k=1}^K$ and $\{s_k\}_{k=1}^K$ satisfying the three conditions of the revealed priority axiom. The preceding argument implies

$s_k \pi s_{k+1}$ for every $k \in \{1, \dots, K\}$. Hence,

$$s_1 \pi s_2 \pi \cdots \pi s_K \pi s_1.$$

By transitivity, this implies $s_1 \pi s_1$, contradicting the asymmetry of π. Thus, Ch^π satisfies the revealed priority axiom.

($\Rightarrow$). Suppose that Ch is non-wasteful, distributionally maximal, and satisfies the revealed priority axiom. By Lemma 1, $\mathrm{Ch}(S) \in \mathrm{F}(S)$ for every $S \subseteq \mathcal{S}$.

Define a binary relation ρ on $\mathcal{S}$ as follows: $s \rho s'$ if there exists $S \subseteq \mathcal{S}$ such that $s \in \mathrm{Ch}(S)$, $s' \in S \setminus \mathrm{Ch}(S)$, and $(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \succsim \mathrm{Ch}(S)$. The revealed priority axiom implies that ρ is acyclic. Therefore, by a generalization of Szpilrajn's lemma (Chambers and Echenique, 2016, Theorem 1.5), ρ can be completed to a priority ranking over $\mathcal{S}$. Let π be such a completion, and let Ch^π be the associated distributional choice rule. By construction, $\mathrm{Ch}^\pi(S) \in \mathrm{F}(S)$ for every $S \subseteq \mathcal{S}$.

Toward a contradiction, suppose that there exists $S \subseteq \mathcal{S}$ such that $\mathrm{Ch}(S) \neq \mathrm{Ch}^\pi(S)$. Since the distributional preference satisfies the upper-bound and maximizer properties, Proposition 2 implies that $\mathrm{F}(S)$ forms the bases of a matroid. By the strong basis exchange property, there exist $s \in \mathrm{Ch}(S) \setminus \mathrm{Ch}^\pi(S)$ and $s' \in \mathrm{Ch}^\pi(S) \setminus \mathrm{Ch}(S)$ such that

$$(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \in \mathrm{F}(S) \quad \text{and} \quad (\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S).$$

By Proposition 1, $(\mathrm{Ch}(S) \setminus \{s\}) \cup \{s'\} \sim \mathrm{Ch}(S)$. Therefore, by the definition of ρ, we have $s \rho s'$. Since π extends ρ, it follows that $s \pi s'$.

On the other hand, since $(\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\} \in \mathrm{F}(S)$, Lemma 2 implies that $\mathrm{Ch}^\pi(S)$ priority dominates $(\mathrm{Ch}^\pi(S) \setminus \{s'\}) \cup \{s\}$ under π. These two sets differ only by replacing s' with s, so priority domination implies $s' \pi s$. However, this contradicts our previous conclusion that $s \pi s'$. Hence no such S exists, and therefore $\mathrm{Ch} = \mathrm{Ch}^\pi$ for some priority ranking π. □

Proof of Theorem 5. We use the deferred-acceptance characterization of Doğan et al. (2025). Their result applies to profiles of choice rules that are characterized by *solitary* axioms and that also satisfy path independence and *size monotonicity*. We define these terms shortly. They show that the corresponding mechanism-level extensions of those solitary axioms, together with individual rationality and strategy-proofness, characterize the DA mechanism based on the underlying choice rules.

To that end, fix a school $c \in \mathcal{C}$. A choice rule Ch_c satisfies **size monotonicity** if $|\mathrm{Ch}_c(S')| \leq |\mathrm{Ch}_c(S)|$ for every $S, S' \subseteq \mathcal{S}$ with $S' \subseteq S$. A **solitary** axiom is a correspon-

dence $\psi_c : 2^{\mathcal{S}} \rightrightarrows 2^{\mathcal{S}}$ such that for each applicant set S, the set $\psi_c(S) \subseteq 2^S$ is the set of all subsets of S that satisfy a particular choice axiom at school c. For our purposes, the relevant solitary axioms are non-wastefulness, distributional maximality, and freedom from justified envy. A choice rule Ch_c satisfies a solitary axiom ψ_c if $\text{Ch}_c(S) \in \psi_c(S)$ for every $S \subseteq \mathcal{S}$.

Given a matching μ, define the **demand** for c at μ by

$$D_c(\mu) \coloneqq \{s \in \mathcal{S} : c \; R_s \; \mu(s)\}.$$

Since preferences are reflexive in their weak form, $\mu^{-1}(c) \subseteq D_c(\mu)$. We say that μ satisfies the **extension** of a solitary axiom ψ_c if $\mu^{-1}(c) \in \psi_c(D_c(\mu))$.

We now show that the three choice rule axioms in Theorem 1 are solitary and that their extensions coincide with the corresponding matching axioms.

Non-wasteful: For every $S \subseteq \mathcal{S}$, define

$$\psi_c^{nw}(S) \coloneqq \{S' \subseteq S : |S'| = \min\{q_c, |S|\}\}.$$

Then a choice rule Ch_c is non-wasteful if and only if $\text{Ch}_c(S) \in \psi_c^{nw}(S)$ for every $S \subseteq \mathcal{S}$ (Axiom 1), so this is a solitary axiom.

Recall that a matching μ must respect a school's capacity, i.e., $|\mu^{-1}(c)| \leq q_c$. A matching μ satisfies the extension if and only if $|\mu^{-1}(c)| = \min\{q_c, |D_c(\mu)|\}$. This condition trivially holds if $|\mu^{-1}(c)| = q_c$. Because $\mu^{-1}(c) \subseteq D_c(\mu)$, if $|\mu^{-1}(c)| < q_c$, the extension is satisfied if and only if $D_c(\mu) = \mu^{-1}(c)$. Since student preferences are strict, a student $s \notin \mu^{-1}(c)$ belongs to $D_c(\mu)$ if and only if $c \; P_s \; \mu(s)$. Therefore, the extension is equivalent to the matching non-wastefulness condition (Axiom 7): if some student strictly prefers c to their assignment, then c must have no empty seats.

Distributionally maximal: For every $S \subseteq \mathcal{S}$, define

$$\psi_c^{pd}(S) \coloneqq \{S'' \subseteq S : \nexists S' \subseteq S \text{ such that } |S'| = |S''| \text{ and } S' \succ_c S''\}.$$

Then Ch_c is distributionally maximal if and only if $\text{Ch}_c(S) \in \psi_c^{pd}(S)$ for every $S \subseteq \mathcal{S}$ (Axiom 2), so this is a solitary axiom.

A matching μ satisfies the extension if and only if $\mu^{-1}(c) \in \psi_c^{pd}(D_c(\mu))$. Equivalently, there exists no subset $S' \subseteq D_c(\mu)$ with $|S'| = |\mu^{-1}(c)|$ such that $S' \succ_c \mu^{-1}(c)$. This is exactly the matching axiom requiring μ to be distributionally maximal (Axiom 8).

Free from justified envy: For every $S \subseteq \mathcal{S}$, define

$$\psi_c^{ne}(S) := \left\{ S' \subseteq S : \begin{array}{l} \text{for every } s' \in S' \text{ and } s'' \in S \setminus S', \\ (S' \setminus \{s'\}) \cup \{s''\} \succsim_c S' \text{ implies } s' \; \pi_c \; s'' \end{array} \right\}.$$

Then Ch_c is free of π_c-justified envy if and only if $\mathrm{Ch}_c(S) \in \psi_c^{ne}(S)$ for every $S \subseteq \mathcal{S}$ (Axiom 3), so this is a solitary axiom.

A matching μ satisfies the extension if and only if $\mu^{-1}(c) \in \psi_c^{ne}(D_c(\mu))$. Equivalently, for every $s \in \mu^{-1}(c)$ and $s' \in D_c(\mu) \setminus \mu^{-1}(c)$,

$$(\mu^{-1}(c) \setminus \{s\}) \cup \{s'\} \succsim_c \mu^{-1}(c) \implies s \; \pi_c \; s'.$$

Because preferences are strict, $s' \in D_c(\mu) \setminus \mu^{-1}(c)$ is equivalent to $c \; P_{s'} \; \mu(s')$. Hence this extension is exactly the matching condition requiring freedom from π_c-justified envy (Axiom 9).

We have thus far shown that for each school $c \in \mathcal{C}$, non-wastefulness, distributional maximality, and freedom from justified envy are all solitary choice axioms. Since $\succsim_c$ satisfies upper-bound and maximizer properties, by Theorem 1, these solitary axioms uniquely characterize $\mathrm{Ch}_c^{\pi_c}$. As $\mathrm{Ch}_c^{\pi_c}$ is non-wasteful, it selects $\min\{q_c, |S|\}$ students from each set S. Hence, the distributional choice rule also satisfies size monotonicity. Since $\succsim_c$ also satisfies the coherence property, by Theorem 3, $\mathrm{Ch}_c^{\pi_c}$ satisfies path independence.

Moreover, the three solitary axioms extend exactly to the corresponding non-wastefulness, distributional maximality, and freedom from justified envy matching axioms. Thus, applying Doğan et al. (2025) yields that a matching mechanism satisfies the three extended axioms, individual rationality, and strategy-proofness if and only if it is the DA mechanism based on the choice rule profile $(\mathrm{Ch}_c^{\pi_c})_{c \in \mathcal{C}}$. □

References

Abdulkadiroğlu, A. and A. Grigoryan (2025): "Market Design for Distributional Objectives in Allocation Problems: An Axiomatic Approach," SSRN working paper.

Aizerman, M. A. and A. V. Malishevski (1981): "General theory of best variants choice: Some aspects," *IEEE Transactions on Automatic Control*, 26, 1030–1040.

Alcalde, J. and S. Barberà (1994): "Top dominance and the possibility of strategy-proof stable solutions to matching problems," *Economic Theory*, 4, 417–435.

Alkan, A. and D. Gale (2003): "Stable schedule matching under revealed preference," *Journal of Economic Theory*, 112, 289 – 306.

Alves, M. J. and C. V. Willie (1987): "Controlled choice assignments: A new and more effective approach to school desegregation," *The Urban Review*, 19, 67–88.

Aygün, O. and I. Bó (2021): "College admission with multidimensional privileges: The Brazilian affirmative action case," *American Economic Journal: Microeconomics*, 13, 1–28.

Aygün, O. and B. Turhan (2017): "Large-Scale Affirmative Action in School Choice: Admissions to IITs in India," *American Economic Review*, 107, 210–13.

——— (2020): "Dynamic reserves in matching markets," *Journal of Economic Theory*, 188, 105069.

Blair, C. (1988): "The Lattice Structure of the Set of Stable Matchings with Multiple Partners," *Mathematics of Operations Research*, 13, 619–628.

Bonet, C., N. Arnosti, and J. Sethuraman (2024): "Explainable Affirmative Action," in *Proceedings of the 25th ACM Conference on Economics and Computation*, New York, NY: Association for Computing Machinery, EC '24, 310.

Burt, R. S. (1992): *Structural Holes: The Social Structure of Competition*, Harvard University Press.

——— (2004): "Structural Holes and Good Ideas," *American Journal of Sociology*, 110, 349–399.

Çelebi, O. (2024): "Diversity Preferences and Affirmative Action," Working paper.

Chambers, C. P. and F. Echenique (2016): *Revealed preference theory*, vol. 56, Cambridge University Press.

Correa, J., R. Epstein, J. Escobar, I. Rios, B. Bahamondes, C. Bonet, N. Epstein, N. Aramayo, M. Castillo, A. Cristi, and B. Epstein (2019): "School Choice in Chile," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, New York, NY: Association for Computing Machinery, EC '19, 325–343.

Doğan, B. and A. Erdil (2025): "Widening access in university admissions," *Journal of Economic Theory*, 224, 105973.

Doğan, B., K. Imamura, and M. B. Yenmez (2025): "Market design with deferred acceptance: A recipe for characterizations," *Journal of Economic Theory*, 228, 106057.

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

Dur, U., P. A. Pathak, and T. Sönmez (2020): "Explicit vs. Statistical Preferential Treatment in Affirmative Action: Theory and Evidence from Chicago's Exam Schools," *Journal of Economic Theory*, 187, 104996.

Echenique, F., T. Mekonnen, and M. B. Yenmez (2024): "Diversity in Choice as Majorization," Working paper.

Echenique, F. and M. B. Yenmez (2015): "How to Control Controlled School Choice," *American Economic Review*, 105, 2679–2694.

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

Ehlers, L. and B. Klaus (2016): "Object allocation via deferred-acceptance: Strategy-proofness and comparative statics," *Games and Economic Behavior*, 97, 128–146.

Fleiner, T. (2003): "A Fixed-Point Approach to Stable Matchings and Some Applications," *Mathematics of Operations Research*, 28, 103–126.

Fragiadakis, D. and P. Troyan (2017): "Improving matching under hard distributional constraints," *Theoretical Economics*, 12, 863–908.

Gale, D. (1968): "Optimal assignments in an ordered set: An application of matroid theory," *Journal of Combinatorial Theory*, 4, 176–180.

Hafalir, I. E., F. Kojima, M. B. Yenmez, and K. Yokote (2026): "Market Design with Distributional Objectives," *Mathematics of Operations Research*, articles in Advance.

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

Imamura, K. (2025): "Meritocracy versus Diversity," *Journal of Economic Theory*, 228, 106047.

Kamada, Y. and F. Kojima (2015): "Efficient Matching under Distributional Constraints: Theory and Applications," *American Economic Review*, 105, 67–99.

——— (2017): "Stability concepts in matching with distributional constraints," *Journal of Economic Theory*, 168, 107–142.

——— (2018): "Stability and strategy-proofness for matching with constraints: a necessary and sufficient condition," *Theoretical Economics*, 13, 761–794.

Kojima, F. and M. Manea (2010): "Axioms for Deferred Acceptance," *Econometrica*, 78, 633–653.

Kojima, F., A. Tamura, and M. Yokoo (2018): "Designing matching mechanisms under constraints: An approach from discrete convex analysis," *Journal of Economic Theory*, 176, 803 – 833.

Kumano, T. and M. Kurino (2022): "Quota Adjustment Process," Working paper.

Manzini, P. and M. Mariotti (2007): "Sequentially Rationalizable Choice," *American Economic Review*, 97, 1824–1839.

Murota, K. (2003): *Discrete Convex Analysis*, Society for Industrial and Applied Mathematics.

Oxley, J. G. (2006): *Matroid theory*, Oxford University Press, USA.

Shioura, A. (2012): "Matroid rank functions and discrete concavity," *Japan Journal of Industrial and Applied Mathematics*, 29, 535–546.

Sönmez, T. and M. B. Yenmez (2020): "Affirmative Action with Overlapping Reserves," Boston College Working Paper 990.

——— (2022): "Affirmative Action in India via Vertical, Horizontal, and Overlapping Reservations," *Econometrica*, 90, 1143–1176.

Westkamp, A. (2013): "An analysis of the German university admissions system," *Economic Theory*, 53, 561–589.

Yokote, K., I. E. Hafalir, F. Kojima, and M. B. Yenmez (2025): "Rationalizing Path-Independent Choice Rules," in *Proceedings of the 26th ACM Conference on Economics and Computation*, Association for Computing Machinery, 1019.