

PHBS WORKING PAPER SERIES

Visibly Fair Mechanisms

Inácio Bó
University of Macau

Gian Caspari
Leibniz Centre for European Economic
Research

Manshu Khanna
Peking University

September 2025

Working Paper 20250905

Abstract

Priority-based allocation often requires eliminating justified envy, making serial dictatorship (SD) the only non-wasteful direct mechanism with that property. However, SD's outcomes can conflict with the policymaker's objectives. We introduce visible fairness, a framework where fairness is evaluated using coarser information. This is achieved by designing message spaces that strategically conceal information that could render desired allocations unfair. We characterize these mechanisms as generalizations of SD, establish conditions for strategy-proofness, and show how to implement distributional constraints. This creates a new trade-off: achieving distributional goals may require limiting preference elicitation, forgoing efficiency gains even when compatible with the constraints.


Keywords: matching theory, market design, indirect mechanisms

JEL Classification: C78, D47, D78

Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055, China



PHBS
北京大学汇丰商学院



Visibly Fair Mechanisms*

Inácio Bó[†] Gian Caspari[‡] Manshu Khanna[§]

September, 2025

Abstract

Priority-based allocation often requires eliminating justified envy, making serial dictatorship (SD) the only non-wasteful direct mechanism with that property. However, SD's outcomes can conflict with the policymaker's objectives. We introduce visible fairness, a framework where fairness is evaluated using coarser information. This is achieved by designing message spaces that strategically conceal information that could render desired allocations unfair. We characterize these mechanisms as generalizations of SD, establish conditions for strategy-proofness, and show how to implement distributional constraints. This creates a new trade-off: achieving distributional goals may require limiting preference elicitation, forgoing efficiency gains even when compatible with the constraints.

Keywords: Matching Theory, Market Design, Indirect Mechanisms

JEL: C78, D47, D78

*We thank Mustafa Oguz Afacan, Samson Alva, Yuan Gao, Sambuddha Ghosh, Xi Jin, Vikram Manjunath, Debasis Mishra and Yu Zhou for helpful comments. We are grateful for the feedback at several conferences and seminars, including Indian Statistical Institute Delhi, Queen's University of Belfast, University of York, University of Manchester, 14th Conference on Economic Design at University of Essex, 2025 CCBEF-IES Workshop at SWUFE, Greater Bay Area Market Design Workshop at University of Macau, Lisbon Meetings, and SAET Conference.

[†]Department of Economics, Faculty of Social Sciences, University of Macau. Email: inaciobo@um.edu.mo

[‡]Department of Market Design, ZEW — Leibniz Centre for European Economic Research, Mannheim 68161, Germany. Email: gian.caspari@zew.de

[§]Peking University HSBC Business School, Shenzhen 518055, China. Email: manshu@phbs.pku.edu.cn

1 Introduction

Priority-based assignments are pervasive in a wide range of real-world matching contexts, including university admissions, public-sector placements, and cadet-branch allocations in military academies (see, e.g., [Balinski and Sönmez, 1999](#); [Sönmez and Switzer, 2013](#)). In a typical priority-based system, participants are strictly ranked—based on exam scores or a merit list for instance—and are assigned to positions accordingly. The central fairness requirement in these contexts is that no lower-ranked participant should occupy a seat that a higher-ranked participant strictly prefers; otherwise, the latter has a legitimate grievance, known as *justified envy* ([Abdulkadiroğlu and Sönmez, 2003](#)). Such fairness concerns become even more pronounced when the priorities at stake represent strongly protected interests—like property rights or national exam rankings—where even a single instance of justified envy can trigger legal and administrative challenges.¹


Under the standard approach of designing *direct mechanisms*—where each participant reports a complete ranking over positions—*Serial Dictatorship (SD)* is in fact the *only* mechanism that can satisfy non-wastefulness—i.e., not leaving desirable positions unfilled—and no-justified-envy. In SD, the highest-priority participant picks their top choice, the next participant picks from the remaining positions, and so on. This procedure prevents any lower-ranked participant from ending up in a spot that a higher-ranked participant strictly prefers, ensuring no-justified-envy. However, SD can produce allocations that are misaligned with policy goals, such as excessive clustering of top-ranked participants in a small set of elite locations, or undesirable regional or demographic distributions. When considering standard direct mechanisms, there is *no alternative* to SD that both respects strict priorities and operates on full preference lists. Thus, policymakers appear to face a dilemma: given the requirement for “fairness by priority”, how can one construct rules in pursuit of better distributional outcomes?

When we look at real-world priority-based assignments, we see many depart from the fully “direct” approach, as illustrated by the following cases:

1. In the Indian Administrative Services (IAS), officers were formerly assigned

¹One example of these legal challenges is the Federal University of Bahia hiring suit (Brazil, 2025), where a federal judge *blocked* the university from hiring a lower-scoring quota applicant and ordered the single vacancy awarded to the exam’s top scorer ([Agência Estado, 2025](#)). Another was the Italian national residency “fiasco”, in which the Regional Administrative Tribunal of Lazio annulled a ruling that forced higher-ranked doctors to forfeit more-preferred specialties while lower-ranked peers advanced ([Focus.it, 2016](#)). Both rulings treat the harm as a breach of the merit order—i.e., a violation of justified envy—showing that such breaches readily provoke litigation.

Figure 1: Example of a preference ranking in the 2017 IAS Mechanism



to state cadres in a priority-driven process aligned with exam-based merit. This arrangement, which was essentially a serial dictatorship with some modifications, produced undesirable allocations exhibiting *homophily*, i.e., a propensity for officers to serve in or near their home regions. Such geographic clustering was seen as compromising the national integration objective of the service (Thakur, 2023). A 2017 reform imposed a zone-based scheme: all cadres were partitioned into five geographic zones, and each officer now submits a separate ranking of cadres *within* each zone rather than a single list over the entire country (see Figure 1). The revised mechanism guarantees that no lower-ranked officer receives a cadre preferred by a higher-ranked officer *inside the zone where they are ultimately matched*, while the zonal structure itself allows for officers to be more evenly dispersed across India—advancing distributional goals without overriding the preferences participants actually report.

2. Under the U.S. Military Academy matching process of cadets to military branches used in 2006, each cadet (i) ranked the branches and (ii) stated, for every branch, whether they would accept a longer service obligation in exchange for a priority boost (Sönmez and Switzer, 2013; Greenberg *et al.*, 2024). This elicitation did not allow them to report *cross-branch* trade-offs. Omitting this information let the Academy honour the official order-of-merit list while making it possible to steer more cadets into longer commitments that would have been rejected under full preference elicitation. As Sönmez (2024) explains, this design kept such priority violations hidden:

“Several years later, in 2019, I finally learned why the Army initially did not pursue a potential reform of the USMA-2006 mechanism. (...) Any failure of the no-justified-envy axiom rooted in this first issue was also ‘invisible’ to the Army. When a cadet receives his first-choice branch at the increased price but prefers his second choice at the base price, this information was simply unavailable under the strategy space of the USMA-2006 mechanism.”

3. In the Chinese college admissions system, applicants submit a structured rank-order list in which majors are nested under colleges, effectively enforcing a lexicographic hierarchy: once a college is deemed higher-ranked, every major it offers is treated as strictly preferred to any program at a lower-ranked college. Although this structure is known to generate numerous cases of justified envy in practice, none of these can be challenged under the restricted message space (see [Hu *et al.*, 2025](#)).²

As the preceding cases show, limiting what participants may report can hide genuine priority breaches. Following this, we say a mechanism is **visibly fair** when, given the elicited (partial) preferences, no outcome appears to violate priority. For example, if seats are partitioned into zones and participants may only rank seats *within* each zone, allocating by priority inside every zone looks perfectly fair—even though a cross-zone comparison (never elicited) might reveal a lower-ranked participant holding a seat a higher-ranked participant prefers. By restricting the scope of reported preferences, policy makers can pursue goals such as geographic diversity while keeping any latent violations invisible. A closely related idea appears already in [Greenberg *et al.* \(2024\)](#), who introduced the notion of *detectable priority reversals*, a concept that corresponds precisely to visible (un)fairness in the context of the US Army’s branching mechanism.

Inspired by these observations, *we examine the design problem of assigning officers to positions under a strict priority ordering while maintaining visible fairness*. In contrast to standard models that fix a preference-reporting format, in this framework **the policy maker chooses both the message space** (the form of partial preferences agents can report) **and the outcome rule**. Our analysis provides a framework and results on how to configure these two elements together so as to achieve desired policy objectives.

Summary of Results

Our analysis delivers three main sets of results. First, we pin down the precise structure that visible fairness imposes on allocation rules. [Theorem 1](#) shows that any visibly-fair mechanism must operate as an *m-queue allocation*: officers are processed in strict priority order, and each is assigned a state that is undomi-

²As of January 2025, 23 out of 31 provinces in China retain the nested rank-order procedure ([Hu *et al.*, 2025](#)), including Shanghai, Beijing, Tianjin, Hainan, Jiangsu, Fujian, Hubei, Hunan, Guangdong, Heilongjiang, Gansu, Jilin, Anhui, Jiangxi, Guangxi, Shanxi, Henan, Shaanxi, Ningxia, Sichuan, Yunnan, Tibet, and Xinjiang. We provide screenshots from Fujian and Shanghai’s official college-major list sample form in Appendix B.

nated, among the remaining states, within the partial ranking she is allowed to report. When the message space induces a partition of the state space into zones, this characterization yields the more specific results in [Theorems 2 and 3](#). In this setting, the only visibly-fair rules are partitioned priority mechanisms. Moreover, when rankings are permitted across “zones,” the only visibly-fair rules are ranked-partitioned priority mechanisms. While visible fairness implies serial dictatorship (and therefore strategy-proofness) when using direct mechanisms ([Corollary 2](#)), that is not the case for general message spaces. We show in [Theorem 4](#) that strategy-proofness is obtained exactly when the mechanism also satisfies *expressiveness* and (weak) *availability*, two properties that rule out profitable deviations when the mapping from message profiles to outcomes is more general.

Second, we introduce a flexible way to encode distributional goals through *modular upper-bounds*. A quota system is modular when every bound caps groups of officers within an arbitrary subset of states ([Definition 10](#)). Modular bounds induce *zonal message spaces* with a partition of states—formally captured in [Definition 13](#)—where all states subject to the same collection of caps fall in the same “zone.” Building on this structure, the *Modular Priority Mechanism* ([Definition 14](#)) processes officers by priority while dynamically clogging zones whose relevant caps have just filled. [Theorem 6](#) proves that this mechanism simultaneously respects every modular bound, remains visibly fair, and is strategy-proof.

Yet, as [Example 6.1](#) and the impossibility result in [Theorem 7](#) shows, in general, no *static* (one-shot) rule can simultaneously (i) satisfy visible fairness, (ii) respect modular upper-bounds, and (iii) not be Pareto-improved by allocations that respect the caps. The root of the conflict is informational. To guarantee that *every* admissible report keeps the quotas intact, the policy maker must prune the message space in advance, excluding comparisons whose truthful revelation could otherwise force a violation. This secures fairness and feasibility but withholds preference information that would uncover (and implement) mutually beneficial swaps still compatible with the same caps.

Finally, we show that this tension can be resolved using a dynamic mechanism. The *Dynamic Modular Priority Mechanism* ([Definition 16](#)) re-elicits each officer’s preferences after observing earlier assignments, restricting her menu only by the quotas that are now binding. This simple refinement recovers constrained Pareto efficiency, while being strategy-proof ([Theorem 8](#)).

Related Literature

A foundational theory for mechanisms with restricted message spaces comes from [Green and Laffont \(1986\)](#), who study settings where participants are restricted to a limited message space that depends on their true state. Their key insight is that by constraining the information a participant can reveal, the set of implementable outcomes can be expanded beyond what is possible in standard direct mechanisms.

Beyond the already discussed cases of the Indian Administrative Service (IAS) cadre allocation ([Thakur, 2023](#)), the U.S. Military Academy’s cadet-branch matching ([Sönmez and Switzer, 2013](#); [Sönmez, 2013](#); [Greenberg et al., 2024](#)), and Chinese college admissions ([Hu et al., 2025](#)), many real-world mechanisms also limit the extent of preference reporting. In school choice, for example, some systems cap the number of schools an applicant may rank ([Haeringer and Klijn, 2009a](#); [Calsamiglia et al., 2010](#)), while others allow applicants to “bundle” schools into groups without inter-group comparisons ([Huang and Zhang, 2025](#)).

Dynamic procedures often also restrict preferences and by doing so, might reduce complexity ([Pycia and Troyan, 2023](#)). [Bó and Hakimov \(2022\)](#) propose Iterative Deferred Acceptance by letting participants choose from menus, obtaining stable results without demanding complete rankings, and [Bó and Hakimov \(2024\)](#); [Mackenzie and Zhou \(2022\)](#) extend this idea to mechanisms that sequentially offer feasible outcomes—enhancing privacy and performing well in controlled experiments. Even small constraints, like limiting the length of rank-ordered lists, can disrupt classical incentive properties: [Haeringer and Klijn \(2009b\)](#) show how capping the number of ranked schools compromises the usual strategy-proofness of the Deferred Acceptance procedure. Meanwhile, [Caspari and Khanna \(2025\)](#) propose precise conditions for stability and incentives with non-standard preference formats. Collectively, these studies highlight how restricting preference elicitation can open new design possibilities while preserving key desiderata — an insight we leverage in defining and deploying visible fairness. [Decerf et al. \(2024\)](#) is perhaps the most conceptually related work to ours. Their notion of *incontestable assignments* describes an environment where participants cannot fully observe others’ preferences or placements, leaving them unable to identify certain envy issues. This parallels our concept of *visible fairness*, in which certain violations become undetectable. The key difference, however, is that [Decerf et al. \(2024\)](#) derive their informational constraints from the participants’ limited ability to view the full outcome, whereas in our framework, these constraints are intentionally *designed* by the policy maker. Specifically, in our setup the policy maker restricts

what participants can report so as to preclude distributional tensions that might otherwise manifest as visible grievances.³ In addition, while [Decerf *et al.* \(2024\)](#) accommodate a variety of school-specific priorities, our model considers a single, strict priority ranking that orders all participants.

Another paper that considers mechanisms using alternative message spaces is [Cavallo and Dogan \(2024\)](#). The authors analyze Italy’s nationwide teacher-mobility scheme, in which teachers may rank entire municipalities, districts, or provinces—nested geographic units that bundle many schools into a single item on their list. They show that the tie-breaking rule used to resolve these coarse rankings might allocate lower-priority teachers ahead of higher-priority ones, creating detectable instances of justified envy. They also show that these result in legal challenges: Italian courts have repeatedly upheld merit-based claims, and parliamentary testimony records more than 1,000 lawsuits filed each year, on average, over such priority violations.

Partial preferences have also been studied in other contexts within market design. One strand lets participants declare indifference classes directly: [Erdil and Ergin \(2017\)](#); [Manjunath and Westkamp \(2021\)](#); [Andersson *et al.* \(2021\)](#) build mechanisms that treat weak orders — strict ranks punctuated by ties — as the primitives, and then exploit those ties to recover efficiency and strategy-proofness. A second strand considers problems in which not every pair of outcomes can be compared, evaluating which adaptation of standard properties, such as stability, can nonetheless sustain strategy-proofness under suitable conditions ([Caspari and Khanna, 2025](#); [Kuvalekar, 2023](#)). Typically, these frameworks rely on “weak stability,” where participants who are indifferent or indecisive simply cannot block assignments. In contrast, our approach presumes participants *do* have complete preferences but are deliberately constrained from revealing them in full.

The design of matching markets with distributional constraints through quota systems has emerged as a critical area of research in market design, balancing equity objectives with efficiency and stability considerations ([Echenique and Yenmez, 2015](#); [Abdulkadiroglu and Grigoryan, 2023](#)).⁴ [Kamada and Kojima \(2015\)](#)

³In many real-life applications, such as the IAS hiring and public sector hiring contests in Italy and Brazil, transparency requirements imply the public disclosure of information such as exam papers, scoring sheets, and even interview recordings. These are not considered private and must be accessible to ensure administrative and social oversight ([Autorità Nazionale Anticorruzione, 2025](#); [Controladoria-Geral da União, 2023](#); [Abizada and Bó, 2021](#)). Under these informational circumstances, incontestability might become equivalent to standard elimination of justified envy, and therefore under single priority imply serial dictatorship.

⁴Practical implementations in education markets reveal both the potential and complexity of quota systems. [Combe *et al.* \(2022\)](#) quantified these trade-offs through France’s teacher assignment reforms, where experience-based distribution constraints reduced novice teacher con-

introduced the idea of matching with distributional constraints, showing that conventional stable matching algorithms can break down under strict regional quotas. To address these deficiencies, they introduced new mechanisms that ensure such constraints are respected while preserving or improving upon stability, efficiency, and incentive alignment. Subsequent research refines and generalizes these insights: for instance, [Kamada and Kojima \(2018\)](#) identify structural conditions enabling strategy-proof and stable mechanisms under distributional constraints, and their more recent work ([Kamada and Kojima, 2024](#)) extends stability ideas to increasingly nuanced affirmative action policies.⁵

At the core of this field lies the tension between rigid distributional quotas and the flexible preferences of participants. [Fragiadakis and Troyan \(2017\)](#) demonstrated this through military cadet matching, where static reservation systems created inefficiencies by locking seats for specific groups prematurely. Their dynamic quota mechanism represented a paradigm shift, adjusting reservation targets based on revealed preferences while maintaining strategy-proofness. This approach inspired subsequent innovations like the Adaptive Deferred Acceptance (ADA) mechanism by [Goto *et al.* \(2017\)](#), which introduced hereditary constraints—rules where satisfying a constraint automatically satisfies all its subsets. The ADA mechanism’s success in Japanese medical residency matching showed that carefully designed constraints need not sacrifice core market principles like nonwastefulness and strategy-proofness.

Structure of the paper. In Section 2, we introduce the model and definitions, covering partial preferences, feasible allocations, and our message-space framework. Section 3 then characterizes visibly fair mechanisms, identifying them as queue-allocation variants and examining applications such as zonal message spaces. Section 4 turns to incentives, specifying exactly when these mechanisms are strategy-proof via the conditions of expressiveness and availability. Sections 5 and 6 present our results on distributional objectives, including modular upper-bounds and the Modular Priority Mechanism, alongside an analysis of efficiency that leads to our Dynamic Modular Priority Mechanism. Section 7 concludes. All

centrations in disadvantaged schools by 18% without significant efficiency losses. [Combe *et al.* \(2025\)](#) extended this through a reassignment algorithm that prioritized understaffed schools, demonstrating how temporal flexibility in constraints (allowing multi-year adjustment periods) could mitigate short-term displacement costs.

⁵Additional contributions include [Aziz *et al.* \(2019\)](#), who introduce the principle of “cutoff stability” for diversity-constrained matching, and [Kojima *et al.* \(2020\)](#), who identify conditions ensuring that distributional constraints do not undermine substitutability in job-matching markets.

proofs are relegated to Appendix A.

2 Model and Definitions

A problem consists of:

1. a finite set of **officers** $I = \{i_1, i_2, \dots, i_n\}$,
2. a finite set of **states** $S = \{s_1, s_2, \dots, s_m\}$,
3. a **capacity** for each state $(q_s)_{s \in S}$, such that $\sum_{s \in S} q_s \geq n$,
4. a strict **preference** (*asymmetric, complete, and transitive*) for each officer $(\succ_i)_{i \in I}$ over states S ,⁶ and
5. a **priority** ranking π of officers I , where officer i is ranked higher than officer j if $\pi(i) < \pi(j)$.

For a given problem, the goal is to produce an allocation of officers to states. Formally, an **allocation** $a = (a_i)_{i \in I}$ is a list specifying a state $a_i \in S$ for each officer $i \in I$. An allocation is **feasible** if, for each $s \in S$, we have $|\{i \in I : a_i = s\}| \leq q_s$. We denote the set of all feasible allocations by $\mathcal{A}$. Furthermore, without loss of generality, we assume that officers with lower subscripts have a higher priority, i.e., $\pi(i_1) < \pi(i_2) < \dots < \pi(i_n)$. While the allocation decision is based on officers' reported preferences and assigned priorities, in our setup, officers do not necessarily communicate their full preferences directly but instead provide partial preference information from a menu of partial preferences available.

More specifically, let M_i denote the **message space** for officer $i \in I$. Each **message** $m_i \in M_i$ is an *irreflexive* and *acyclic* binary relation $\succ_{m_i}$ over the set of states S .⁷ Throughout, " $\succ_{m_i}$ " will be read as "preferred under the message m_i ." We say that states s and s' are **comparable** under m_i if $s = s'$ or if m_i ranks one strictly above the other, i.e., $s \succ_{m_i} s'$ or $s' \succ_{m_i} s$. We denote message space profiles and message profiles by $M = (M_i)_{i \in I}$ and $m = (m_i)_{i \in I}$.

⁶We denote by $\succsim_i$ the associated weak preference—that is, $s \succsim_i s' \iff s \succ_i s' \text{ or } s = s'$.

⁷Binary relation $\succ_{m_i}$ on S is *irreflexive* if for all $s \in S$, $\neg(s \succ_{m_i} s)$. Binary relation $\succ_{m_i}$ on S is *acyclic* if for all $s, s' \in S$, for all $K \in \mathbb{N}$, and for all $s^0, \dots, s^K \in S$,

$$[s = s^0 \text{ and } s^{k-1} \succ_{m_i} s^k \text{ for all } k \in \{1, \dots, K\} \text{ and } s^K = s'] \implies \neg(s' \succ_{m_i} s).$$

We denote by $\succsim_{m_i}$ the associated weak binary relation.

Definition 1. A message space M_i satisfies **richness** if, whenever there exists a message $m_i \in M_i$ and two states $s, s' \in S$ such that $s \succ_{m_i} s'$ and there is no state $y \in S$ where $s \succ_{m_i} y \succ_{m_i} s'$, then there also exists a message $m'_i \in M_i$ such that:

- for every pair $(s_1, s_2) \in S \times S$ with $\{s_1, s_2\} \neq \{s, s'\}$, $s_1 \succ_{m_i} s_2$ if and only if $s_1 \succ_{m'_i} s_2$, and
- $s' \succ_{m'_i} s$.

In words, whenever the designer allows an officer to send a message that ranks s above s' , there also exists a message $m'_i \in M_i$ that preserves every other pairwise comparison of m_i , but reverses the comparison of s and s' . The condition involving the third state y prevents this reversal from violating transitivity.

Henceforth, we assume each message space M_i satisfies richness, unless noted otherwise. This property guarantees that the message space never forces the officer to reveal a comparison in one direction without permitting the symmetric comparison in the opposite direction.

Having defined both allocations and messages, we now formally define a mechanism. A **M-mechanism** is a function from message profiles to allocations, $\psi : M \rightarrow A$. We will use the shorthand **mechanism** when the space of message profiles is clear from the context.

3 Visibly Fair Mechanisms

3.1 Visible Fairness

The key notions that we introduce in this paper is that of visible fairness.

Definition 2. An allocation a is **visibly unfair under** m if for some $i \in I$ either

- i) there is a $j \in I$ such that $a_i \neq a_j$, $\pi(i) < \pi(j)$, and $a_j \succ_{m_i} a_i$, or
- ii) there is a $s \in S$ such that $a_i \neq s$, $|\{i \in I : a_i = s\}| < q_s$, and $s \succ_{m_i} a_i$.

A mechanism ψ is **visibly fair** if there does not exist $m \in M$ such that $\psi(m)$ is visibly unfair under m .

At first sight, visible (un)fairness appears to be a combination of standard non-wastefulness and elimination of justified envy. The distinction, however, lies in the fact that m_i , in general, is incomplete, and therefore some existing wastefulness

or justified envy is “invisible” due to the limits to the expression of the associated preferences imposed by the message space.

Define the set $G(X, m_i)$ of all m_i -**maximal elements** of $X \subseteq S$ by

$$G(X, m_i) = \{s \in X : \neg(s' \succ_{m_i} s) \text{ for all } s' \in X\}.$$

Since $\succ_{m_i}$ is acyclic and S is finite, $G(X, m_i) \neq \emptyset$.⁸ We next introduce a new family of mechanisms.

Definition 3. A mechanism ψ is a **m-queue allocation mechanism** if $\psi(m)$ is the outcome produced by the following procedure:

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): $a_k = s^k \in G(S^k, m_k)$. If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

The m -queue allocation mechanism is the natural partial-preference analogue of serial dictatorship. Starting from the highest-ranked officer, they are matched to undominated states, given previously assigned ones. The difference here being, of course, that there might be multiple undominated states, and therefore some selection criterion between them must also be defined.

Theorem 1. *A mechanism ψ is visibly fair if and only if it is a m -queue allocation mechanism.*

When considering complete message spaces, Theorem 1 gives us the following corollary:

Corollary 1. *If for every $i \in I$, every $m_i \in M_i$ and every $s, s' \in S$, states s and s' are comparable under m_i , then Serial Dictatorship is the unique visibly fair M -mechanism.*

Theorem 1 provides a full characterization of visibly fair mechanisms: any such mechanism must be an m -queue allocation mechanism. While this definition is compact, it captures an extensive class of mechanisms. Designing an m -queue mechanism involves two layers of choice—first, specifying the message spaces officers can use, and second, given a message space, determining which of the undominated (m_i -maximal) states each officer is matched to. This latter

⁸For a proof, see [Bossert and Suzumura \(2010, Theorem 2.6\)](#).

choice can depend on exogenous policy parameters or on the overall message profile. In contrast, as indicated in the corollary above, when officers are allowed to submit complete preferences over all states, visible fairness alone pins down a unique mechanism—serial dictatorship. The characterization in Theorem 1 therefore illustrates how relaxing the message space, dramatically expands the set of mechanisms that satisfy visible fairness, enabling greater flexibility in accommodating policy goals beyond those implementable via serial dictatorship.

3.2 Two Special Message Spaces

We now turn to two natural families of message spaces, which are used in practice and serve as the foundation for implementing distributional objectives in Section 5.

3.2.1 Zonal Message Space

We now introduce a family of message spaces that induce a partition of the set of states into “zones.” Formally, suppose a message space M_i is such that the states S can be partitioned into disjoint subsets $Z = \{z_1, \dots, z_\ell\}$ with $\bigcup_{j=1}^\ell z_j = S$ and $z_j \cap z_{j'} = \emptyset$ for $j \neq j'$. We call M_i a **zonal message space** if within each subset z_j , any two states are comparable under any $m_i \in M_i$, while any two states in different subsets are never comparable under any $m_i \in M_i$. That is, zonal message spaces have the following properties:

Within-zone completeness: For every strict total order R on S , there exists a message $m_i \in M_i$ such that, for every zone $z_j \in Z$ and all states $s, s' \in z_j$ with $s \neq s'$,

$$s R s' \iff s \succ_{m_i} s'.$$

In other words, within each zone, the message space is rich enough to permit any ranking of states within that zone.

Across-zone incomparability: For every pair of zones $(z_j, z_{j'})$ with $j \neq j'$, and for every message $m_i \in M_i$, no two states in z_j and $z_{j'}$ are ever ordered under $\succ_{m_i}$. That is, the message space never allows an officer to rank a state in z_j relative to a state in $z_{j'}$.

From an officer’s perspective, submitting a message in a zonal message space amounts to choosing a complete (strict) ordering over states within each zone but leaving no comparison defined across zones.

Example 3.1. Suppose the states S are $\{s_1, s_2, s_3, s_4\}$. One could design a message space that effectively partitions these states into two zones:

$$z_1 = \{s_1, s_2\}, \quad z_2 = \{s_3, s_4\}.$$

Officers would be required to provide complete rankings among $\{s_1, s_2\}$ and among $\{s_3, s_4\}$, but they would never be allowed to compare s_1 (or s_2) with s_3 (or s_4). Consequently, any message in this space is of the form:

- Within z_1 , rank s_1 above s_2 (or vice versa).
- Within z_2 , rank s_3 above s_4 (or vice versa).
- Across z_1 and z_2 , no ordering is possible.

Although the zone partition can be entirely endogenously derived from how comparisons are restricted in the message space, a mechanism designer may also choose to impose such a structure normatively, for example to ensure no cross-zone comparisons are made (as in the Indian Civil Services, where certain sets of states are grouped into “zones”). From a theoretical standpoint, both perspectives are equivalent: a zonal message space is simply one in which states are fully comparable within each zone and never comparable across zones.

Notice that given any Z , i , zonal message space M_i , and a preference over states $\succ_i$, there is exactly one message $m_i \in M_i$ such that for every pair of states s, s' , $s \succ_i s' \iff s \succ_{m_i} s'$.

Suppose each officer has a zonal message space. Consider a **zone selection function** $\mathcal{C}_i : 2^S \times M \rightarrow Z$, which is a function such that for any $X \subseteq S$, $X \cap \mathcal{C}_i(X, m) \neq \emptyset$.⁹

Definition 4. A mechanism ψ is a **partitioned priority mechanism** if there exists a zone selection function profile $(\mathcal{C}_i)_{i \in I}$ such that for any message profile m , $\psi(m)$ is the outcome produced by the following procedure:

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): $a_k \in G(S^k, m_k) \cap \mathcal{C}_k(S^k, m) = \{s^k\}$.¹⁰ If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

⁹That is, as long as there are states with spare capacity available, $\mathcal{C}_i$ must choose a zone with at least one of them.

¹⁰It is easy to see that within each zone, there is a unique m_i -maximal element among any set of remaining states with spare capacity, since all states in a zone are comparable under m_i .

Zone selection functions constitute the essential component of the definition that results in the large variety of these mechanisms. They indicate, for each profile of messages, which zone will be used to determine an officer's outcome. This zone can depend on some exogenous parameter, on the allocations of higher-ranked officers and/or the preferences stated by other officers, as well as her own message. Once a zone is determined, however, the state that will be matched to the officer depends only on her preferences between the remaining states in that zone.¹¹

Theorem 2. *For a zonal message space M , ψ is visibly fair if and only if it is a partitioned priority mechanism.*

3.2.2 Zonal Message Space with Ranking over the Zones

We now enrich the idea of a zonal message space by allowing officers to impose a strict ordering *across* the zones. As before, let $Z = \{z_1, \dots, z_\ell\}$ be a partition of S into disjoint zones. Let, moreover, $\max(X, m_i) \equiv G(X, m_i)$ and $\min(X, m_i) \equiv \{s \in X : \nexists s' \in X, s \succ_{m_i} s'\}$. A **zonal message space with ranking over the zones** is a message space M_i in which:

Within each zone, states are comparable just as in the standard zonal case. That is, for every strict total order R on S , there exists a message $m_i \in M_i$ such that, for every zone $z_j \in Z$ and all states $s, s' \in z_j$ with $s \neq s'$,

$$s R s' \iff s \succ_{m_i} s'.$$

Across zones, each message $m_i \in M_i$ implies a complete ranking $\triangleright_{m_i}$ over zones Z s.t.

$$z_i \triangleright_{m_i} z_j \iff \max(z_i, m_i) \succ_{m_i} \min(z_j, m_i).$$

No other preferences among states can be expressed in these message spaces.

In words, these rankings augment zonal message spaces in a minimal sense: they allow officers to express some information about how they rank zones in a weakest sense: by ranking zone z above z' , they are saying that the best state in z is preferred over the worst in z' . Notice that if this was not the case, we would have z being ranked above z' while every state in z' is preferred to every state in

¹¹Notice, moreover, that the definition of the mechanism requires that $\mathcal{C}_k$ chooses a zone with states with spare capacity, which by assumption always exists.

z .¹²

Notice, moreover, that Zonal Message Space with Ranking over the Zones do not satisfy the richness condition we introduced in definition 1: you can express $\max(z_k, m_i) \succ_{m_i} \min(z_j, m_i)$, but not $\min(z_j, m_i) \succ_{m_i} \max(z_k, m_i)$.

Example 3.2. Suppose the states S are $\{s_1, s_2, s_3, s_4\}$. One could design a message space that effectively partitions these states into two zones:

$$z_1 = \{s_1, s_2\}, \quad z_2 = \{s_3, s_4\}.$$

Officers' messages contain complete rankings among $\{s_1, s_2\}$ and among $\{s_3, s_4\}$, and a ranking over zones $\triangleright_{m_i}$. Consequently, any message in this space is of the form:

- Within z_1 , rank s_1 above s_2 (or vice versa).
- Within z_2 , rank s_3 above s_4 (or vice versa).
- In addition to these:
 - If $s_1 \succ_{m_i} s_2$ and $s_3 \succ_{m_i} s_4$, either s_1 is ranked above s_4 ($z_1 \triangleright_{m_i} z_2$), or s_3 above s_2 ($z_2 \triangleright_{m_i} z_1$).
 - If $s_1 \succ_{m_i} s_2$ and $s_4 \succ_{m_i} s_3$, either s_1 is ranked above s_3 ($z_1 \triangleright_{m_i} z_2$), or s_4 above s_2 ($z_2 \triangleright_{m_i} z_1$).
 - If $s_2 \succ_{m_i} s_1$ and $s_3 \succ_{m_i} s_4$, either s_2 is ranked above s_4 ($z_1 \triangleright_{m_i} z_2$), or s_3 above s_1 ($z_2 \triangleright_{m_i} z_1$).
 - If $s_2 \succ_{m_i} s_1$ and $s_4 \succ_{m_i} s_3$, either s_2 is ranked above s_3 ($z_1 \triangleright_{m_i} z_2$), or s_4 above s_1 ($z_2 \triangleright_{m_i} z_1$).
- No other preferences across states are possible.

Suppose each officer i has a zonal message space with ranking over the zones, as described above. Define a **ranked zone selection function** $\mathcal{C}_i : 2^S \times M \rightarrow Z$ to be a mapping that, for each subset of states $X \subseteq S$ and each message profile m , selects a zone $\mathcal{C}_i(X, m) \in Z$ such that:

¹²To see this, suppose that $Z = \{z_1, z_2\}$, $z_1 = \{s_1, s_2\}$ and $z_2 = \{s_3, s_4\}$. Let m_i be such that $s_1 \succ_{m_i} s_2$ and $s_3 \succ_{m_i} s_4$. If a ranking $z_1 \triangleright_{m_i} z_2$ is associated with $s_1 \succ_{m_i} s_4$, we have this additional comparison between states and nothing else. But if it was associated with $s_4 \succ_{m_i} s_1$, this would imply that $s_3 \succ_{m_i} s_4$, $s_4 \succ_{m_i} s_1$, $s_1 \succ_{m_i} s_2$, thus making s_3 the m_i -maximal element in $\{s_1, s_2, s_3, s_4\}$, which would be at odds with any reasonable interpretation of what $z_1 \triangleright_{m_i} z_2$ implies for preferences among these states.

1. $X \cap \mathcal{C}_i(X, m) \neq \emptyset$, and
2. either

$$X \cap \mathcal{C}_i(X, m) \neq \{\min(\mathcal{C}_i(X, m), m_i)\},$$

or there is no zone $z \in Z$ with $z \triangleright_{m_i} \mathcal{C}_i(X, m)$ and $\max(z, m_i) \in X$.

A ranked zone selection function restricts which zone will be assigned to an officer on the basis of the limits visible fairness imply given the ranking over zones: if a zone z contains only the state deemed as the least preferred in z , no zone z' for which $z' \triangleright_{m_i} z$ can have spare capacity in its most-preferred state.

Definition 5. A mechanism ψ is a **ranked partitioned priority mechanism** if there exists a ranked zone selection function profile $(\mathcal{C}_i)_{i \in I}$ such that for any message profile m , $\psi(m)$ is the outcome produced by the following procedure:

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): $a_k \in G(S^k, m_k) \cap \mathcal{C}_k(S^k, m) = \{s^k\}$.¹³ If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

Theorem 3. For zonal message space with ranking over zones, ψ is visibly fair if and only if it is a ranked partitioned priority mechanism.

The presence of ranking over zones implies some restrictions on the zones that the zone selection function can determine, as shown in the example below.

Example 3.3. There are three states $S = \{s_1, s_2, s_3\}$ (capacity 1 each) and two zones

$$z_1 = \{s_1\}, \quad z_2 = \{s_2, s_3\}.$$

¹³Ranked zone selection function makes sure that within the selected zone, m_k -maximal element exists. To see why, let $z = \mathcal{C}_k(S^k, m)$ be the selected zone. We have two cases from the ranked zone selection function definition:

Case 1: $S^k \cap z \neq \emptyset$ and $S^k \cap z \neq \{\min(z, m_k)\}$

Then $S^k \cap z$ contains non-minimal elements of zone z . Since all states in a zone are comparable under m_k , the m_k -maximal state within $S^k \cap z$ is unique. By the across-zone ranking property, only $\min(z, m_k)$ can be dominated by states in different zones, so the m_k -maximal state within $S^k \cap z$ is also m_k -maximal in all of S^k . Therefore, $G(S^k, m_k) \cap \mathcal{C}_k(S^k, m)$ is non-empty and a singleton.

Case 2: $S^k \cap z \neq \emptyset$ and no zone $z' \triangleright_{m_i} z$ has $\max(z', m_k) \in S^k$

The condition ensures that all maximal elements of higher-ranked zones are unavailable in S^k . Therefore, m_k -maximal state within $S^k \cap z$ faces no domination from higher-ranked zones and is m_k -maximal in S^k . Therefore, again $G(S^k, m_k) \cap \mathcal{C}_k(S^k, m)$ is non-empty and a singleton.

There are two officers, i_1 and i_2 . We will consider visibly fair mechanisms in which the message space is zonal with rankings over these zones for both officers.

Officers i_1 and i_2 both submit the same message:

$$z_1 \triangleright_{m_i} z_2, \quad s_2 \succ_{m_i} s_3.$$

Because z_2 contains *two* available states, the ranked-zone selection function could, in this scenario, place i_1 in *either* zone z_1 or z_2 . Suppose that it places on z_2 . Given m_1 , i_1 is matched to s_2 .

Now z_1 still has s_1 free, and z_2 only s_3 . Since z_1 remains vacant and is the *highest-ranked* for i_2 , the mechanism must choose z_1 and assign s_1 . Selecting z_2 would contradict $z_1 \triangleright_{m_2} z_2$ while z_1 still offers an available seat, and is therefore not allowed.

Ordinary partitioned priority mechanisms (without the zone ranking) would leave the planner free to swap i_2 between z_1 and z_2 , illustrating how adding cross-zone orderings tightens the designer's hands, in comparison.

4 Incentives

To analyze incentives, we first define what it means for an officer to report truthfully. The idea is straightforward: given an officer's preference over states, their report is truthful if the underlying preference information in the submitted message aligns with their actual preferences. For a given officer $i \in I$, let $\mathcal{Q}_i$ be the **set of all preferences** over S . Officer i 's message m_i is a **truthful message** for a preference $\succsim_i \in \mathcal{Q}_i$, if for all $s, s' \in S$,

$$s \succ_{m_i} s' \implies s \succ_i s'.$$

Note that our richness assumption on the message space M_i ensures the existence of a truthful message for every officer. That is, for every preference $\succsim_i \in \mathcal{Q}_i$, there exists at least one message $m_i \in M_i$ that is truthful.

When it comes to incentives, a key desideratum is that an officer who submits a truthful message should never receive a less preferred allocation than if they were to report any other message. Formally, a mechanism ψ is **strategy-proof** if for all $i \in I$, $\succsim_i \in \mathcal{Q}_i$, for any $m_i \in M_i$ that is a truthful message for $\succsim_i$, and any other message $\hat{m}_i \in M_i$ we have

$$\psi(m)_i \succsim_i \psi(\hat{m}_i, m_{-i})_i.$$

Strategy-proofness implies, therefore, that if an officer has multiple truthful messages, then they cannot lead to different outcomes. We next define two conditions that are sufficient for a visibly fair mechanism to be strategy-proof. Weakening one of the two conditions is also necessary for a visibly fair mechanism to be strategy-proof.

The first condition, expressiveness, requires that whenever an officer changes her message and thereby obtains another assignment, this new assignment must be comparable to the officer's originally assigned state under the original message. Note that, the condition allows for the new assignment to be identical to the original assignment and that a state is always comparable to itself.

Definition 6. Let ψ be a mechanism and let m be a message profile. The interest into state $s \in S$ is **expressed** by officer i under m if s and $\psi(m)_i$ are comparable under m_i .

A mechanism ψ satisfies **expressiveness** if for all $i \in I$, any $m_i \in M_i$ and any other message $\hat{m}_i \in M_i$ interest into state $\psi(\hat{m}_i, m_{-i})_i$ was *expressed* by officer i under m .

The second condition, availability, requires that whenever an officer changes her message and thereby obtains another assignment, this new assignment must always correspond to a state that was already available to her under the original message. Note that, the condition allows for the new assignment to be identical to the original assignment in which case the state is trivially available under the original message.

Definition 7. Let ψ be a mechanism and let m be a message profile. A state $s \in S$ is **available** to officer i under m if $|\{j \in I : \psi(m)_j = s\} \cap \{j \in I : \pi(j) < \pi(i)\}| < q_s$.

A mechanism ψ satisfies **availability** if for all $i \in I$, any $m_i \in M_i$ and any other message $\hat{m}_i \in M_i$, we have that state $\psi(\hat{m}_i, m_{-i})_i$ is *available* to officer i under m .

Interestingly, availability is too strong of a condition for ensuring strategy-proofness. Indeed, while availability ensures that an officer cannot manipulate the availability of states by submitting a different message, weak availability only requires that an officer cannot manipulate the availability of weakly preferred states, evaluated at the original message, relative to her assignment under the original message. This condition, together with expressiveness is necessary for a visibly fair mechanism to be strategy-proof.

Definition 8. A mechanism ψ satisfies **weak availability** if for all $i \in I$, $\succsim_i \in \mathcal{Q}_i$, for any $m_i \in M_i$ that is a truthful message for $\succsim_i$, and any other message $\hat{m}_i \in M_i$ s.t. $\psi(\hat{m}_i, m_{-i})_i \succsim_i \psi(m)_i$, we have that state $\psi(\hat{m}_i, m_{-i})_i$ was *available* to officer i under m .

Now we are ready to formally state our main result on incentives.

Theorem 4. *A visibly fair mechanism is strategy-proof if and only if it satisfies expressiveness and weak availability.*

The following corollary is immediate:

Corollary 2. *A visibly fair mechanism is strategy-proof if it satisfies expressiveness and availability.*

To give some intuition behind the results we give two examples. The first one is a strategy-proof mechanism violating availability — but satisfying weak availability and expressiveness.

Example 4.1 (Strategy-proof mechanism violating availability). Consider a problem with two officers $I = \{i_1, i_2\}$. Without loss of generality, we let officer i_1 have higher priority than i_2 , i.e., $\pi(i_1) < \pi(i_2)$. There are two states $S = \{s_1, s_2\}$, each with capacity $q_s = 1$. Consider a mechanism ψ where i_1 can only submit a single message m_{i_1} without any preference information and i_2 can either submit message $m_{i_2} : s_1 \succ_{m_{i_2}} s_2$ or $m'_{i_2} : s_2 \succ_{m'_{i_2}} s_1$. Finally, let $\psi(m_{i_1}, m_{i_2}) = (s_2, s_1)$ and $\psi(m_{i_1}, m'_{i_2}) = (s_1, s_2)$.

Clearly, the mechanism is strategy-proof as i_1 cannot influence the outcome by submitting a different message, and i_2 always gets her top choice when submitting preferences truthfully. Moreover, the mechanism is visibly fair as i_1 gives no preference information and i_2 gets her top choice.

It is easy to see that availability is violated as, e.g., consider message m_i and message m'_i . Note that, m'_i leads to a different outcome $\psi(m_{i_1}, m'_{i_2})_{i_2} = \{s_2\} \neq \psi(m)_{i_2} = \{s_1\}$ which is not available under m .

Note that weak availability is not violated as under m_i which is a truthful message for $s_1 \succ_{i_2} s_2$ we have $\psi(m)_{i_2} \succ_{i_2} \psi(m_{i_1}, m'_{i_2})_{i_2}$, and analogous under m'_{i_1} which is a truthful message for $s_2 \succ_{i_2} s_1$ we have $\psi(m_{i_1}, m'_{i_2})_{i_2} \succ_{i_2} \psi(m)_{i_2}$.

Overall this example illustrates how availability is too strong a requirement for strategy-proofness.

The second example shows a mechanism satisfying expressiveness but violating weak availability and thus strategy-proofness.

Example 4.2 (Expressive but not weakly available mechanism). Consider a problem with two officers $I = \{i_1, i_2\}$, where, without loss of generality, officer i_1 has higher priority than i_2 , i.e., $\pi(i_1) < \pi(i_2)$. There are two states $S = \{s_1, s_2\}$, each with capacity $q_s = 1$. Consider a mechanism ψ where i_1 can only submit a single message m_{i_1} without any preference information and i_2 can either submit message $m_{i_2} : s_1 \succ_{m_{i_2}} s_2$ or $m'_{i_2} : s_2 \succ_{m'_{i_2}} s_1$. Finally, let $\psi(m_{i_1}, m_{i_2}) = (s_1, s_2)$ and $\psi(m_{i_1}, m'_{i_2}) = (s_2, s_1)$.

Clearly, the mechanism is not strategy-proof as i_2 always gets his second choice when submitting a truthful message and can get his first choice by simply reporting the opposite message. The mechanism is visibly fair as i_1 does not give any preference information, and i_2 only reports the preferred state, which is given to the higher priority officer i_1 .

The mechanism also satisfies expressiveness as under any message i_2 gives full preference information, while i_1 has a single message automatically satisfying expressiveness. At the same time the mechanism violates weak availability as e.g. under m_{i_2} which is a truthful message for $s_1 \succ_{i_2} s_2$ and another message m'_{i_2} , where $\psi(m'_{i_1}, m_{i_2})_{i_2} \succ_{i_2} \psi(m_{i_1}, m_{i_2})_{i_2}$ we have that $\psi(m'_{i_1}, m_{i_2})_{i_2}$ is not available under m .

Finally, an alternative condition that ensures a mechanism is strategy-proof, without requiring visible fairness, is defined next. Coherence requires that whenever an officer changes her message and thereby obtains a different assignment, this new assignment is not in the set of undominated states — consisting of both the new and original assignment — evaluated at the original message. Note that this condition only applies when the two assignments are distinct states.

Definition 9. A mechanism satisfies **coherence** if for any message $m_i \in M_i$ and any other message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \neq \psi(m_i, m_{-i})_i$ we have $\psi(\hat{m}_i, m_{-i})_i \notin G(\{\psi(m_i, m_{-i})_i\} \cup \{\psi(\hat{m}_i, m_{-i})_i\}, m_i)$.

Indeed, the following result follows almost immediately:

Theorem 5. *A mechanism is strategy-proof if and only if it satisfies coherence.*

4.1 Two special message spaces revisited

In this section we show that, neither the class of partition priority mechanism nor the ranked partition priority mechanisms are always strategy-proof.

In [Example 4.3](#) we show that mechanisms using zonal message spaces can fail weak availability and expressiveness, and therefore not be strategy-proof. Moreover, in [Example 4.4](#) we show that zonal message spaces can be sufficient conditions for visibly fair mechanisms to fail expressiveness, and therefore strategy-proofness.

Example 4.3 (Partitioned-priority can violate strategy-proofness). Consider a problem with three officers $I = \{i_1, i_2, i_3\}$. There are three states $S = \{s_1, s_2, s_3\}$, each with capacity $q_s = 1$. The zonal message space for all officers consists of two zones: $z_1 = \{s_1, s_2\}$ and $z_2 = \{s_3\}$. Suppose each officer's true preference is $s_1 \succ_i s_2 \succ_i s_3$. For each officer i , let m_i be the message that ranks $s_1 \succ_{m_i} s_2$, and let m'_i reverse it: $s_2 \succ_{m'_i} s_1$.

Consider a zone selection function $\mathcal{C}$ that includes the following:

$$\begin{aligned} \mathcal{C}_{i_1}(S, m) &= z_1, & \mathcal{C}_{i_1}(S, (m_{-i_3}, m'_{i_3})) &= z_1, \\ \mathcal{C}_{i_2}(\{s_2, s_3\}, m) &= z_1, & \mathcal{C}_{i_2}(\{s_2, s_3\}, (m_{-i_3}, m'_{i_3})) &= z_2, \\ \mathcal{C}_{i_3}(\{s_3\}, m) &= z_2, & \mathcal{C}_{i_3}(\{s_2\}, (m_{-i_3}, m'_{i_3})) &= z_1. \end{aligned}$$

Notice that these conditions are consistent with zone selection function that induces a partitioned priority mechanism. Consider the following successful manipulation for i_3 : Under truthful report m , the allocation is (s_1, s_2, s_3) ; but when i_3 flips her internal ranking ($m_{i_3} \rightarrow m'_{i_3}$), the mechanism assigns (s_1, s_3, s_2) , giving i_3 a preferred assignment as $s_2 \succ_{i_3} s_3$.

Weak availability is violated as e.g. s_2 is available under m'_{i_3} but not under m_{i_3} which is a truthful message for $\succ_{i_3}$ under which s_2 is preferred to s_3 .

Similarly, expressiveness is violated as e.g. officer i_3 does not express any preference information regarding s_2 and s_3 , even though s_2 is allocated to i_3 under message m'_{i_3} and s_3 is allocated to i_3 under message m_{i_3} .

Finally, we give an example that illustrates that a ranked partition priority mechanism might not be strategy-proof due to its violation of expressiveness.

Example 4.4 (Ranked-partitioned priority can violate Expressiveness).

Consider a setting with three officers $I = \{i_1, i_2, i_3\}$. The set of states is $S = \{s_1, s_2, s_3\}$, each with a capacity of $q_s = 1$. The message space is *zonal* for all officers and requires officers to rank two zones: $z_1 = \{s_1, s_2\}$ and $z_2 = \{s_3\}$. Each officer's message must provide a full ranking over the states within zone z_1 , i.e., between s_1 and s_2 , and additionally indicate whether $z_1 \triangleright_{m_i} z_2$ or $z_2 \triangleright_{m_i} z_1$.

True preferences for i_1 : $s_1 \succ_{i_1} s_2 \succ_{i_1} s_3$. This yields a *unique* truthful message $m_{i_1}^\triangleright : (s_1 \succ s_2 \triangleright s_3)$,¹⁴ ensuring i_1 obtains s_1 in any visibly fair mechanism.

Two preferences for i_2 :

1. First, suppose i_2 has $s_3 \succ_{i_2} s_2 \succ_{i_2} s_1$. The unique truthful message $m_{i_2} : (s_3 \triangleright s_2 \succ s_1)$ forces i_2 to end up with s_3 under any visibly fair, strategy-proof mechanism. To see this, suppose that i_2 is assigned s_2 instead. Then, by submitting $m_{i_2}' : s_3 \triangleright s_1 \succ s_2$, any visibly fair mechanism must assign s_3 to i_2 , leading to a successful manipulation.
2. Next, consider $\succ_{i_2} : s_2 \succ s_3 \succ s_1$. Here, *two* truthful messages are possible. One message $m_{i_2} : (s_3 \triangleright s_2 \succ s_1)$ again, as we have just argued, assigns s_3 to i_2 . Another message $m_{i_2}'' : (s_2 \succ s_1 \triangleright s_3)$ must assign s_2 to i_2 .

Here, expressiveness is violated as e.g. officer i_2 does not express any preference information regarding s_2 and s_3 given message m_{i_2} , even though s_2 is allocated to i_2 under message m_{i_2}'' and s_3 is allocated to i_2 under message m_{i_2} . On the other hand, weak availability does not pose a problem in the above example.

Given [Theorem 4](#), combined with [Example 4.3](#) and [Example 4.4](#) the following corollary is immediate:

Corollary 3. *Consider the zonal message space with and without rankings, then both the partitioned priority mechanism and ranked partition priority mechanism might not be strategy-proof.*

5 Achieving Distributional Objectives

When using direct mechanisms, our earlier discussion shows that only serial dictatorship (SD) achieves visible fairness. By eliciting less information about preferences—through carefully designed restricted message spaces—we expand the set of allocations that can be deemed visibly fair for a given problem. This relaxation provides the policy maker with additional flexibility, allowing for the implementation of a broader array of allocation rules that still adhere to this notion of fairness.

¹⁴A message $m_{i_1}^\triangleright : s_1 \succ s_2 \triangleright s_3$ is our abbreviation for the preference message with $s_1 \succ_{m_{i_1}} s_2$ and $z_1 \triangleright_{m_{i_1}} z_2$.

There are, in principle, many distinct distributional objectives that can be accommodated within this broader framework. In this section, we introduce one family of such objectives, which we denote *Modular Upper-Bounds*, and give complete instructions on how to design mechanisms that are visibly fair, strategy-proof, and respect these bounds. Modular upper-bounds model distributional objectives by imposing limits on the number of officers of certain types assigned to specific subsets of states.

5.1 Modular Upper-Bounds

We extend our original model by saying that each officer i has a type t from a finite set of types T , where t_i denotes the type of officer i . The distributional goals of the designer are modeled through type-specific modular upper-bounds, where for a set of types, a collection of upper-bounds specifies limits on the allocation of officers with those types to subsets of states.

Definition 10. A **modular upper-bound system** is a finite collection

$$H = \{ (\Xi_h, S_h, k_h) \},$$

where for each element $h \in H$.¹⁵

- $\emptyset \neq \Xi_h \subseteq T$ is the set of types covered by the quota,
- $S_h \subseteq S$ is a subset of states, and
- $k_h \in \mathbb{N}$ is the ceiling.

For every type $t \in T$ we write

$$H^t := \{ (\Xi_h, S_h, k_h) \in H : t \in \Xi_h \}.$$

For every state $s \in S$ and type $t \in T$, we write its **upper-bound signature** as

$$H^{s,t} = \{ (\Xi_h, S_h, k_h) \in H : s \in S_h, t \in \Xi_h \}.$$

Definition 11. An allocation $a \in \mathcal{A}$ **respects the modular upper-bounds** H if, for every $(\Xi_h, S_h, k_h) \in H$:

$$|\{ i \in I : a_i \in S_h, t_i \in \Xi_h \}| \leq k_h.$$

¹⁵In our notation, H contains a set of upper-bounds, each of which represented by the letter h . In this context, Ξ_h for example, is the first component of h .

We say that the modular upper-bound (Ξ_h, S_h, k_h) is **binding** at allocation $a \in \mathcal{A}$ if the constraint is satisfied with equality.

Since in our model officers cannot be left unmatched, we need to guarantee that these upper-bounds are compatible with that restriction while using visibly fair mechanisms. Formally, hereafter we will restrict our attention to modular upper-bound systems that satisfy the following property of *sequential solvency*.

Definition 12. For any allocation $a \in \mathcal{A}$, officer $i \in I$, and modular upper bound $h \in H$, define $n_s := |\{j \neq i : a_j = s\}|$ and $n_h := |\{j \neq i : a_j \in S_h, t_j \in \Xi_h\}|$. The modular upper-bounds system H satisfies **sequential solvency** if and only if for all $i \in I$ and for all $a \in \mathcal{A}$ that respect H ,

$$\exists s \in S \text{ s.t. } \begin{cases} n_s < q_s, \text{ and} \\ h \in H^{s, t_i} \implies n_h < k_h \quad \forall h \in H. \end{cases}$$

The intuition for the definition above is simple. Regardless of which capacities or collection of upper-bounds bind, there will always be a compatible state for every remaining officer. That is, while matching officers one at a time, modular upper-bounds can restrict *where* officers are matched, but *not whether* they are matched. Since it relies on the particular number of agents of each type in I , it allows for interesting and practical constraints, as we will show in examples that will follow.¹⁶

Example 5.1. Consider a problem with five states $S = \{s_1, s_2, s_3, s_4, s_5\}$, each having capacity $q_s = 1$. There are two officer types: t_1 and t_2 .

- For type t_1 , the upper-bound system is

$$H^{t_1} = \left\{ (\{t_1\}, \{s_1, s_2, s_3\}, 2) \right\},$$

meaning that at most 2 type- t_1 officers may be assigned to states in $\{s_1, s_2, s_3\}$.

- For type t_2 , the upper-bound system is

$$H^{t_2} = \left\{ (\{t_2\}, \{s_3, s_4, s_5\}, 1) \right\},$$

meaning that at most 1 type- t_2 officer may be assigned to states in $\{s_3, s_4, s_5\}$.

¹⁶Notice that it is crucial that the definition depends on the profile of types of officers. Otherwise, the upper-bounds would have to be satisfied when all agents have the same type, making only bounds that never bind compatible with not leaving officers unmatched.

The literature has proposed several ways to formalize “quota-type” constraints.¹⁷ An example of a very *permissive* notion is the hereditary family of Goto *et al.* (2017): write a matching as a vector that counts, for every state, how many officers of each type are assigned there; a subset of vectors is *hereditary* when it is closed under coordinate-wise decrements. Any system of pure ceilings clearly has this property, so every modular upper-bound instance fits inside the hereditary domain. A *tighter* specification is the hierarchical (laminar) system analysed by (Kamada and Kojima, 2015, 2018): here the subsets that carry quotas must form a tree—any two are either disjoint or one contains the other. Laminar caps are useful when the policy maker wants, say, regional ceilings that line up neatly with district ceilings, but they rule out overlapping constraints such as “no more than ten officers in the Northeastern states *and* no more than eight in the coastal states.” All laminar systems are modular, yet the converse is false.

5.2 Modular-induced Message Spaces

To design mechanisms that respect modular upper-bounds, we define the *Modular-induced Message Spaces*. These are zonal message spaces where states are partitioned into zones based on the upper-bounds, grouping together states involved in the same upper-bounds.

Definition 13. For any type $t \in T$, define an equivalence relation $\sim_t$ on S such that, for all $s, s' \in S$,

$$s \sim_t s' \quad \text{if and only if} \quad H^{s,t} = H^{s',t}.$$

The **Modular-induced Message Space** associated with the modular upper-bounds H^t is the zonal message space M_t with zones $Z = \{z_1^t, z_2^t, \dots, z_k^t\}$, where each zone z_j^t is an equivalence class under $\sim_t$, that is,

$$z_j^t = \{s \in S : s \sim_t s_j\},$$

for some representative state $s_j \in S$.

This construction ensures that:

- i) Zones are disjoint and partition the set of states: $\bigcup_j z_j^t = S$ and $z_j^t \cap z_{j'}^t = \emptyset$ for $j \neq j'$.

¹⁷An incomplete list includes Echenique and Yenmez (2015), Kamada and Kojima (2015), Goto *et al.* (2017), Kamada and Kojima (2018), Aziz *et al.* (2019), Kojima *et al.* (2020), and Kamada and Kojima (2024).

- ii) All states within the same zone are involved in exactly the same set of upper-bounds for the type t . Therefore, if some upper-bound is binding for some state in a zone, then it binds for all states in that zone.

5.3 Modular Prioritized Allocation Mechanism

Fix a finite set of agents I and a strict priority order π on I . Let H be a given modular upper-bound system for a finite set of types T . Each type $t \in T$ induces a *modular-induced message space* M_t , as in Definition 13, where for each type $t \in T$, S is partitioned into zones $Z^t = \{z_1^t, z_2^t, \dots\}$ according to equivalence classes of states under the same upper-bound signature.

Thus, for an officer i of type t_i , the mechanism *offers* the zonal space M_{t_i} , with associated zones $z_1^{t_i}, z_2^{t_i}, \dots$, requiring her to submit a message $m_i \in M_{t_i}$. That message ranks all states within each zone but cannot compare states across different zones.

As part of the mechanism design, each officer i is also assigned an *exogenous* ranking

$$z_1^{t_i} \blacktriangleright_i z_2^{t_i} \blacktriangleright_i \dots \blacktriangleright_i \dots$$

over the same zones, independent of the message m_i . These exogenous rankings—which can encode policy priorities such as emphasizing certain zones first or last—do not depend on agents' reports.

Definition 14. Given a strict priority ranking π , a modular upper-bound system H , an exogenous zone ranking $\blacktriangleright_i$ for each officer i , and a profile of messages $m = (m_i)_{i \in I}$, each m_i in the modular-induced message space M_{t_i} , the **Modular Priority Mechanism** ψ proceeds as follows:

Initialization:

- For each $s \in S$, set remaining capacity $q_s^{rem} = q_s$.
- For each type t and zone z_ℓ^t , set a flag $B_\ell^t = \text{False}$.
- Set $a_i = \emptyset$ for all $i \in I$.

Quotas update procedure:

We next describe the procedure that updates the zone flags below:

- For each $(\Xi_h, S_h, k_h) \in H$:

- Let $N_h = |\{i \in I : a_i \in S_h \text{ and } t_i \in \Xi_h\}|$.
- If $N_h = k_h$, then for every $t \in \Xi_h$ and ℓ such that $S_h \cap z_\ell^t \neq \emptyset$, set $B_\ell^t = \text{True}$.

Sequential Assignment: Process the officers in the order $(i_1, \dots, i_n)$. For each $k = 1, \dots, n$:

1. Let t_{i_k} be the type of officer i_k . Her message m_{i_k} partitions S into zones $z_1, \dots, z_K$.
2. Starting from the top-ranked zone z_1 under $\blacktriangleright_{i_k}$, find the first zone z_ℓ such that:
 - $B_{z_\ell}^{t_{i_k}} = \text{False}$. (No modular upper-bound for type t_{i_k} is yet fully binding in z_ℓ .)
 - There is a state $s \in z_\ell$ for which $q_s^{rem} > 0$.
3. If no such zone is found, set $a_{i_k} = \emptyset$, and move to i_{k+1} . Otherwise:
 - Let s^* be the *most-preferred* state of i_k *within* z_ℓ (according to m_{i_k}) that still has $q_{s^*}^{rem} > 0$.
 - Set $a_{i_k} = s^*$ and reduce capacity $q_{s^*}^{rem} \leftarrow q_{s^*}^{rem} - 1$.
 - Apply the quotas update procedure.

Outcome: After processing all agents $i_1, \dots, i_n$, the mechanism outputs the allocation $a = (a_i)_{i \in I}$.

The Modular Priority Mechanism is, therefore, a partitioned priority mechanism in which each officer type is associated with a zonal message space of states that share the same upper-bound constraints. An officer's final assignment is determined by two key factors: (1) a counter that tracks remaining capacity for the relevant states, and (2) a flag indicating whether any upper-bound restrictions in that zone have become binding. Because all states in a given zone are governed by the same set of constraints, a single triggered bound applies uniformly across the entire zone. Below, we present the main result for this mechanism and two examples that illustrate how it operates in practice.¹⁸

¹⁸In [Appendix B](#) we show that the modular upper-bounds system presented in [Example 5.2](#) and [Example 5.3](#) satisfy sequential solvency.

Theorem 6. *The Modular Priority Mechanism is visibly fair, strategy-proof, and respects modular upper-bounds.*

Example 5.2 (Distributing officers across two regions). Consider four states $S = \{s_1, s_2, s_3, s_4\}$, which are partitioned administratively into two regions: $R_1 = \{s_1, s_2\}$ and $R_2 = \{s_3, s_4\}$, and each $s \in S$ having capacity $q_s = 2$. There are 8 officers $I = \{i_1, \dots, i_8\}$, and an officer type corresponds to the region they are originally from: $t_i \in \{1, 2\}$.

The policy objective is that at most 50% of the jobs in a region are taken by local officers. For each $r \in \{1, 2\}$, the modular upper-bound is, therefore:

$$H^r = \{(\{r\}, R_r, 2)\},$$

so that for an officer of type r the cap applies only to states in R_r . For any given region r , denote by $\hat{r}$ the other region in T .

For an officer i of type r , the upper-bound signature is

$$H^{s,r} = \begin{cases} \{(\{r\}, R_r, 2)\}, & \text{if } s \in r, \\ \emptyset, & \text{if } s \in \hat{r}. \end{cases}$$

Thus, the induced equivalence relation partitions S into two zones:

$$z_1^r = \{s_1, s_2\}, \quad z_2^r = \{s_3, s_4\}.$$

Officers therefore use the same zonal message space, where they submit complete rankings over states within each zone (without comparing states across z_1^r and z_2^r).

Exogenous zone ranking: each officer of type r has the fixed ordering

$$\begin{aligned} z_1^r &\blacktriangleright_i z_2^r & \text{if } t_i = 1 \\ z_2^r &\blacktriangleright_i z_1^r & \text{if } t_i = 2 \end{aligned}$$

so that the officer's own region's zone is ranked above other states.

A modular-priority mechanism enforcing these bounds proceeds as in Definition 14, using the zonal message spaces M_1 and M_2 and exogenous rankings over zones $(\blacktriangleright_i)_{i \in I}$.

Example 5.3 (Distributing doctors across regions and urban/rural divides). Consider nine hospitals $S = \{s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, s_9\}$ partitioned into regions

$$R_1 = \{s_1, s_2, s_3\}, \quad R_2 = \{s_4, s_5, s_6\}, \quad R_3 = \{s_7, s_8, s_9\},$$

with each hospital having capacity $q_s = 4$. In each region, the first hospital is *rural* (s_1, s_4, s_7) and the others are *urban* ($s_2, s_3, s_5, s_6, s_8, s_9$). There are 27 doctors, 9 of each type $t \in \{1, 2, 3\}$, indicating their home region.

Two types of modular upper-bounds apply:

$$H^t = \{(\{t\}, R_t, 6)\} \quad \text{for each } t = 1, 2, 3, \quad \text{and} \quad H^U = \{(\{1, 2, 3\}, U, 19)\}$$

where $U = \{s_2, s_3, s_5, s_6, s_8, s_9\}$ is the set of all urban hospitals, and $(U, 19)$ is a universal cap across all doctors.

A doctor of type 1, for instance, faces the following *signatures*:

$$H^{s,1} = \begin{cases} \{(\{1\}, R_1, 6)\}, & \text{if } s \in R_1 \text{ is rural (here, } s_1), \\ \{(\{1\}, R_1, 6), (\{1, 2, 3\}, U, 19)\}, & \text{if } s \in R_1 \text{ is urban (here, } s_2, s_3), \\ \{(\{1, 2, 3\}, U, 19)\}, & \text{if } s \text{ is urban but not in } R_1 \text{ (e.g. } s_5, s_6, s_8, s_9), \\ \emptyset, & \text{if } s \text{ is rural outside } R_1 \text{ (e.g. } s_4, s_7). \end{cases}$$

Hence the modular-induced partition for a type-1 doctor is:

$$z_1^1 = \{s_1\}, \quad z_2^1 = \{s_2, s_3\}, \quad z_3^1 = \{s_4, s_7\}, \quad z_4^1 = \{s_5, s_6, s_8, s_9\}.$$

Within each zone, the doctor ranks hospitals *fully*, yet makes no cross-zone comparisons. Message spaces following the same principle are constructed for the remaining types.

Exogenous zone ranking: each doctor of type 1 has the fixed ordering

$$z_2^1 \blacktriangleright_i z_1^1 \blacktriangleright_i z_4^1 \blacktriangleright_i z_3^1,$$

so that hospitals in her own region's zone are ranked above other states, and urban hospitals are ranked above rural. Exogenous zone rankings for other types of doctors are defined analogously.

A modular-priority mechanism enforces these caps by assigning each type- t doctor to her top feasible and non-binding hospital within the partition induced by $(R_t, 6)$ and $(U, 19)$.

A modular-priority mechanism enforcing these caps proceeds as in Definition 14, using the zonal message spaces defined above and the exogenous rankings over zones $(\blacktriangleright_i)_{i \in I}$.

As noted previously, the specification of the exogenous rankings over zones $(\blacktriangleright_i)_{i \in I}$ does not affect the theoretical properties established in our results. However, if we know that, for instance, most doctors typically prefer urban hospitals over rural ones, defining the rankings as in the examples above tends to improve efficiency of the final allocations, without harming the other objectives. Naturally, if the actual preferences of participants substantially diverge from such assumptions, employing this approach could lead to less desirable outcomes.

It is also worth noting that the mechanisms introduced above—and the overarching notion of visible fairness—are ill-suited to implementing *within*-state affirmative action quotas. Such policies, which often appear in the matching literature on diversity (e.g. majority quotas or type-specific ceilings in each state), could be encoded as modular upper-bounds with one state in each upper-bound. However, a scenario in which, for example, each state can admit a maximum number of agents of a certain type, the resulting zonal message space would have one state per zone. This, of course, would eliminate a role for preferences in the allocation. If we attempted to compensate for this by using zonal message spaces with ranking over zones, we would easily conclude that no visibly fair mechanism with these characteristics would respect upper bounds.

6 Efficiency

As the following example demonstrates, the Modular Prioritized Allocation Mechanism does not guarantee efficiency.

Example 6.1. Consider two officers $I = \{i_1, i_2\}$ with priority $\pi(i_1) < \pi(i_2)$ and of the same type t , and two states $S = \{s_1, s_2\}$ with capacities $q_{s_1} = 2$ and $q_{s_2} = 1$. Let there be a modular upper-bound of one officer of type t allowed in s_1 . Thus, the modular-induced message space partitions states into two zones: $z_1 = \{s_1\}$ and $z_2 = \{s_2\}$. Let the exogenous ranking of zones be $z_1 \blacktriangleright_i z_2$ for both officers.

Suppose that officers' true preferences are as follows:

$$\begin{aligned} \succ_{i_1}: \quad & s_2 \succ s_1, \\ \succ_{i_2}: \quad & s_1 \succ s_2. \end{aligned}$$

Under the Modular Prioritized Allocation Mechanism, the resulting allocation is $a = (a_{i_1}, a_{i_2}) = (s_1, s_2)$.

Example 6.1 shows that it is possible to improve the efficiency of a mechanism that remains visibly fair and satisfies all modular upper-bounds. Specifically, because officer i_1 's assignment to either s_1 or s_2 does not violate any bounds, it is unnecessary to separate those two states into different zones in i_1 's message space. Allowing the officer to place both states in the same zone incorporates her actual preferences more fully, thus leading to a more efficient assignment.

We now introduce a “second-best” notion of efficiency tailored to settings with modular upper-bounds.

Definition 15. An allocation a is **constrained Pareto efficient** if:

- i) a respects the upper-bounds, and
- ii) there is no other allocation a' that respects the modular upper-bounds such that for every officer i we have $a'_i \succsim_i a_i$.

A **mechanism** is constrained Pareto efficient if, for any problem, when agents submit any *truthful message*, the outcome is constrained Pareto efficient.

In other words, an allocation is constrained Pareto efficient if there is no alternative allocation that respects both capacity and modular upper-bound constraints and strictly Pareto improves upon a (with respect to the agents' true preferences). A mechanism is constrained Pareto efficient if, whenever agents submit truthful messages, the resulting allocation is constrained Pareto efficient *under their full (true) preferences*. The following result demonstrates that, in general, no static mechanism can simultaneously achieve constrained Pareto efficiency and visible fairness.

Theorem 7. *Not all modular upper-bound constraints admit a (static) mechanism that is simultaneously visibly fair and constrained Pareto efficient.*

The negative result in Theorem 7 is not, of course, universal. Standard SD satisfies those properties for modular upper-bound constraints that never bind. On the other hand, constant mechanisms can trivially respect constraints that result in a single feasible allocation which, by its uniqueness, would be constrained Pareto efficient.

Theorem 7, however, underscores a key tension in designing visibly fair mechanisms that must also adhere to non-trivial constraints such as modular upper-bounds. At its core, this tension arises from the message-space design: enforcing both fairness and quota requirements demands that the space be curated to preempt scenarios where certain preference reports would necessarily violate upper-bounds. Consequently, to avoid such violations, the mechanism must collect more restricted preference information than might otherwise be desirable. This reduction in elicited information, while upholding fairness and preserving the bounds, can lead to efficiency losses because the mechanism may lack the information needed to detect and implement mutually beneficial reallocations that remain compliant with all constraints.

Nonetheless, if one can condition each officer's reported preferences on the assignments of higher-ranked officers, it becomes possible to obtain second-best outcomes. We exploit this insight to propose a dynamic mechanism, described next.

Definition 16. The **Dynamic Modular Priority Mechanism** ψ operates as follows:

Initialization:

- For each officer $i \in I$, set $a_i^0 = \emptyset$.
- For each type $t \in T$, initialize the set of binding modular upper-bounds $B_0^t = \emptyset$.

Quotas update procedure:

We next describe the procedure that updates the set of binding upper-bounds:

- Let $(B_k^t)_{t \in T}$ be the current set of binding modular upper-bounds.
- For each $(\Xi_h, S_h, k_h) \in H$:
 - Let $N_h = |\{i \in I : a_i \in S_h \text{ and } t_i \in \Xi_h\}|$.
 - If $N_h = k_h$, then for every $t \in \Xi_h$, add h to B_k^t .

Sequential Assignment: Process the officers in the order $(i_1, \dots, i_n)$. For each $k = 1, \dots, n$:

1. Let M_{i_k} be a zonal message space that partitions S into two zones:

$$z_1 = \left\{ s \in S : \left(\forall (\Xi_h, S_h, k_h) \in B_k^{t_{i_k}}, s \notin S_h \right) \quad \text{and} \quad |\{j < k : a_j = s\}| < q_s \right\}$$

$$z_2 = S \setminus z_1$$

2. Elicit from officer i_k a message $m_k \in M_{i_k}$. Let s^* be officer i_k 's most-preferred state in zone z_1 , according to m_{i_k} .
3. Set $a_{i_k} = s^*$.
4. Follow the quotas update procedure and proceed to the next officer i_{k+1} .

Outcome: After processing all officers $i_1, \dots, i_n$, the assignment $a = (a_i)_{i \in I}$ is the mechanism's final allocation.

The Dynamic Modular Priority Mechanism, therefore, sequentially assigns officers, dynamically tailoring the menu of available states for each officer based on previous assignments. At each officer's turn, the mechanism constructs a zone containing all states with remaining capacity that do not belong to any subset of states where a modular upper-bound for the officer's type has already become binding. The officer is then matched to her most-preferred state among these feasible alternatives. By construction, this ensures visible fairness since each officer always obtains their top choice from that zone. Furthermore, by dynamically adjusting the zones and enabling officers to fully express their preferences within these constraints, the mechanism attains constrained efficiency: any allocation that improves an officer's assignment without harming others would necessarily violate at least one modular upper-bound. The following proposition summarizes these properties, including incentives.¹⁹

Theorem 8. *The Dynamic Modular Priority Mechanism is visibly fair, constrained Pareto efficient, strategy-proof, and respects modular upper-bounds.*

¹⁹While the definitions of visible fairness and respecting modular upper bounds carry over directly without introducing additional notation, strategy-proofness warrants further clarification in our dynamic setting. Because the Dynamic Modular Priority Mechanism allows officers' strategies to depend on both history and the structure of the message spaces (i.e., the composition of the zones), our original, static definition of strategy-proofness does not explicitly cover this complexity. Nevertheless, the informal argument used in the proof of Theorem 8 presented in Appendix A remains sound: no matter how other players behave over time, an officer cannot secure a better outcome by misreporting, so truth-telling continues to be a weakly dominant strategy.

It is also worth noting that the efficiency gains achieved with the Dynamic Modular Priority Mechanism come at a time cost. Whereas the standard procedure requires only a single round of simultaneous preference elicitation, the dynamic approach proceeds sequentially, allocating officers one by one. While this added complexity may be negligible in small markets, it can become impractical in larger settings, where the number of officers is substantial.

7 Conclusion

We explored how designing mechanisms that restrict the preferences participants can report helps reconcile multiple policy objectives—particularly distributional constraints—with fairness principles grounded in strict priority orders. We introduced the concept of visible fairness, in which a mechanism never produces an allocation that appears to violate a participant’s priority based on the partially observed preferences. Theorem 1 demonstrated that every visibly fair mechanism operates as a variant of serial dictatorship adapted to partial preferences. We further showed that this framework can accommodate diverse distributional objectives by employing modular upper-bounds, which limit how many individuals of certain types can be placed in specified subsets of positions. Central to this approach are modular-induced message spaces, which prevent participants from specifying cross-group preference comparisons that would otherwise undermine these quotas.

We then characterized what makes these mechanisms incentive-compatible, by parsing out two critical conditions—expressiveness and (weak) availability—that together ensure no participant can profit by misreporting her partial preferences. While restricting the scope of preference reporting can, in principle, worsen efficiency, we introduced a dynamic modular-priority framework that sequentially elicits partial preferences and updates feasibility constraints in real time, thus recovering the best possible Pareto outcomes subject to capacity and distributional limits. Overall, this paper highlights how deliberately constraining participants’ message spaces provides policy makers with new levers for achieving desired policy goals while preserving fairness and incentive properties in priority-based allocation settings.

References

- ABDULKADIROĞLU, A. and GRIGORYAN, A. (2023). Market design for distributional objectives in allocation problems: An axiomatic approach. *Working paper*.
- ABDULKADIROĞLU, A. and SÖNMEZ, T. (2003). School choice: A mechanism design approach. *American Economic Review*, **93** (3), 729–747.
- ABIZADA, A. and BÓ, I. (2021). Hiring from a pool of workers. *Journal of Economic Behavior & Organization*, **186**, 576–591.
- AGÊNCIA ESTADO (2025). Justiça manda universidade contratar 1^a colocada no lugar de negra aprovada no sistema de cota. Federal Court decision, 10^a Vara Cível/BA, 17 Dec 2024, reported 8 Apr 2025.
- ANDERSSON, T., CSEH, Á., EHLERS, L. and ERLANSON, A. (2021). Organizing Time Exchanges: Lessons from Matching Markets. *American Economic Journal: Microeconomics*, **13** (1), 338–373.
- AUTORITÀ NAZIONALE ANTICORRUZIONE (2025). Obblighi di pubblicazione concernenti i bandi di concorso (art. 19, d.lgs. 33/2013). FAQ aggiornate al 5 maggio 2025.
- AZIZ, H., GASPERS, S., SUN, Z. and WALSH, T. (2019). From matching with diversity constraints to matching with regional quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, p. 377–385.
- BALINSKI, M. and SÖNMEZ, T. (1999). A tale of two mechanisms: student placement. *Journal of Economic theory*, **84** (1), 73–94.
- BÓ, I. and HAKIMOV, R. (2022). The iterative deferred acceptance mechanism. *Games and Economic Behavior*, **135**, 411–433.
- and HAKIMOV, R. (2024). Pick-an-Object Mechanisms. *Management Science*, **70** (7), 4693–4721.
- BOSSERT, W. and SUZUMURA, K. (2010). *Consistency, choice, and rationality*. Harvard University Press.

- CALSAMIGLIA, C., HAERINGER, G. and KLIJN, F. (2010). Constrained School Choice: An Experimental Study. *American Economic Review*, **100** (4), 1860–1874.
- CASPARI, G. and KHANNA, M. (2025). Nonstandard choice in matching markets. *International Economic Review*, **66** (2), 757–786.
- CAVALLO, M. and DOGAN, B. (2024). When geography shapes preferences: Redesigning teacher assignment in Italy. *Working paper*.
- COMBE, J., DUR, U., TERCIEUX, O., TERRIER, C. and UNVER, M. U. (2025). Market design for distributional objectives in (re) assignment: An application to improve the distribution of teachers in schools. *Working Paper*.
- , TERCIEUX, O. and TERRIER, C. (2022). The design of teacher assignment: Theory and evidence. *The Review of Economic Studies*, **89** (6), 3154–3222.
- CONTROLADORIA-GERAL DA UNIÃO (2023). *Parecer nº 1103/2023/CGRAI/DIRAI/SNAI/CGU: Lei de Acesso à Informação – Recurso de 3ª instância no processo 23546.053930/2023-07*. Parecer técnico 1103/2023, Controladoria-Geral da União, Brasília, assinado eletronicamente em 8 set. 2023.
- DECERF, B., HAERINGER, G. and DER LINDEN, M. V. (2024). *Incontestable Assignments*. Working paper, Working Paper.
- ECHENIQUE, F. and YENMEZ, M. B. (2015). How to Control Controlled School Choice. *American Economic Review*, **105** (8), 2679–2694.
- ERDIL, A. and ERGIN, H. (2017). Two-sided matching with indifferences. *Journal of Economic Theory*, **171**, 268–292.
- FOCUS.IT (2016). Medicina: concorso specializzazioni 2014, tar accoglie ricorso. TAR Lazio (Rome) ruling on the 2014 national residency match, 20 Jun 2016.
- FRAGIADAKIS, D. and TROYAN, P. (2017). Improving matching under hard distributional constraints. *Theoretical Economics*, **12** (2), 863–908.
- GOTO, M., KOJIMA, F., KURATA, R., TAMURA, A. and YOKOO, M. (2017). Designing Matching Mechanisms under General Distributional Constraints. *American Economic Journal: Microeconomics*, **9** (2), 226–262.

- GREEN, J. R. and LAFFONT, J.-J. (1986). Partially Verifiable Information and Mechanism Design. *The Review of Economic Studies*, **53** (3), 447–456.
- GREENBERG, K., PATHAK, P. A. and SÖNMEZ, T. (2024). Redesigning the us army’s branching process: A case study in minimalist market design. *American Economic Review*, **114** (4), 1070–1106.
- HAERINGER, G. and KLIJN, F. (2009a). Constrained school choice. *Journal of Economic Theory*, **144** (5), 1921–1947.
- and — (2009b). Constrained school choice. *Journal of Economic Theory*, **144** (5), 1921–1947.
- HU, X., YAO, L. and ZHANG, J. (2025). Reforming china’s two-stage college admission: An experimental study. *Working Paper*.
- HUANG, L. and ZHANG, J. (2025). Bundled school choice. *Working Paper*.
- KAMADA, Y. and KOJIMA, F. (2015). Efficient Matching under Distributional Constraints: Theory and Applications. *American Economic Review*, **105** (1), 67–99.
- and — (2018). Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition: Matching with constraints. *Theoretical Economics*, **13** (2), 761–793.
- and — (2024). Fair Matching under Constraints: Theory and Applications. *The Review of Economic Studies*, **91** (2), 1162–1199.
- KOJIMA, F., SUN, N. and YU, N. N. (2020). Job Matching under Constraints. *American Economic Review*, **110** (9), 2935–2947.
- KUVALEKAR, A. (2023). Matching with Incomplete Preferences. *Working paper*.
- MACKENZIE, A. and ZHOU, Y. (2022). Menu mechanisms. *Journal of Economic Theory*, **204**, 105511.
- MANJUNATH, V. and WESTKAMP, A. (2021). Strategy-proof exchange under trichotomous preferences. *Journal of Economic Theory*, **193**, 105197.
- PYCIA, M. and TROYAN, P. (2023). A Theory of Simplicity in Games and Mechanism Design. *Econometrica*, **91** (4), 1495–1526.

- SÖNMEZ, T. (2013). Bidding for army career specialties: Improving the rotc branching mechanism. *Journal of Political Economy*, **121** (1), 186–219.
- SÖNMEZ, T. (2024). Minimalist Market Design: A Framework for Economists with Policy Aspirations. *Working Paper*.
- SÖNMEZ, T. and SWITZER, T. B. (2013). Matching with (branch-of-choice) contracts at the united states military academy. *Econometrica*, **81** (2), 451–488.
- THAKUR, A. (2023). Matching in the Civil Service: A Market Design Approach to Public Administration and Development. *Working Paper*.

A Appendix: Proofs

A.1 Proof of Theorem 1

A mechanism ψ is visibly fair if and only if it is a m -queue allocation mechanism.

Proof. Part 1: A m -queue allocation mechanism ψ is visibly fair.

Case 1. Consider any $m \in M$ and suppose by contradiction that for individual i there exists $\psi(m)_j \succ_i \psi(m)_i$, where $\pi(i) < \pi(j)$.

By construction of the mechanism, suppose individual i is assigned at step k ; note that both $\psi(m)_j \in S^k$ and $\psi(m)_i \in S^k$ as $\pi(i) < \pi(j)$. It follows, by $\psi(m)_j \succ_i \psi(m)_i$ that $\psi(m)_i \notin G(S^k, m_k)$ — a contradiction.

Case 2. Consider any $m \in M$ and suppose by contradiction that for individual i there exists $s \succ_i \psi(m)_i$, where $|\{i \in I : \psi(m)_i = s\}| < q_s$.

By construction of the mechanism, suppose individual i is assigned at step k ; note that both $s \in S^k$ and $\psi(m)_i \in S^k$ as $|\{i \in I : \psi(m)_i = s\}| < q_s$. It follows, by $s \succ_i \psi(m)_i$ that $\psi(m)_i \notin G(S^k, m_k)$ — a contradiction.

Part 2: A visibly fair mechanism is a m -queue allocation mechanism

Take an arbitrary visibly fair mechanism ψ and fix a message profile $m \in M$. Denote the resulting allocation by $a = \psi(m)$. We now *construct* for this profile the m -queue allocation procedure that replicates a .

Define set S^k for each officer $k \in I$ as follows:

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): Set $s^k = a_k$. If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

Notice that for officer k , the set S^k consists of only those states that still have remaining capacities to accommodate k after the allocation of higher-ranking officers is taken into account. That is, each state $s \in S \setminus S^k$ is filled by higher ranking offers, $|\{i \in I : i < k \text{ and } a_i = s\}| = q_s$. Now we must have that $a^k \in G(S^k, m_k)$, otherwise there exists $s \in S^k$ such that $s \succ_{m_k} a_k$, which makes the outcome visibly unfair as a higher ranking offer does not occupy this state.

An m -queue mechanism working down the priority list, using sets $\{S^k\}_{k=1}^n$ defined above and the procedure given in [Definition 3](#), would give every officer the same assignment as ψ . Since the construction can be repeated for every profile m , ψ is an m -queue allocation mechanism. $\square$

A.2 Proof of [Theorem 2](#)

For a zonal message space M , ψ is visibly fair if and only if it is a partitioned priority mechanism.

Proof. Part 1: A partitioned priority mechanism is visibly fair.

The proof follows the same logic as in Theorem 1, Part 1, and is omitted for brevity.

Part 2: A visibly fair mechanism is a partitioned priority mechanism.

Fix an arbitrary message profile $m \in M$ and let $a = \psi(m)$ be the allocation produced by any visibly fair mechanism ψ . We construct zone selection functions $\{C_k\}_{k=1}^n$ and sets $\{S^k\}_{k=1}^n$ that would reproduce the same allocation step by step.

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): If $a_k \in z$, set $C_k(S^k, m) = z$. If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

After performing the construction for every k we have defined all zone selection functions $\{C_k\}_{k=1}^n$ and sets $\{S^k\}_{k=1}^n$.

Notice that for officer k , the set S^k consists of only those states that still have remaining capacities to accommodate k after the allocation of higher-ranking officers is taken into account. That is, each state $s \in S \setminus S^k$ is filled by higher ranking offers, $|\{i \in I : i < k \text{ and } a_i = s\}| = q_s$. Now we must have that $a^k \in G(S^k, m_k) \cap C_k(S^k, m)$, otherwise there exists $s \in S^k$ such that $s \succ_{m_k} a_k$, which makes the outcome visibly unfair as a higher ranking offer does not occupy this state.

A partitioned priority mechanism working down the priority list, using zone selection functions $\{C_k\}_{k=1}^n$, sets $\{S^k\}_{k=1}^n$ defined above and the procedure given in [Definition 4](#), would assign every officer the same post as ψ .

Since the construction can be repeated for every profile m , ψ is a partitioned priority mechanism. $\square$

A.3 Proof of [Theorem 3](#)

For zonal message space with ranking over zones, ψ is visibly fair if and only if it is a ranked partitioned priority mechanism.

Proof. Part 1: A ranked partitioned priority mechanism is visibly fair.

The proof follows the same logic as in [Theorem 1](#), Part 1, and is omitted for brevity.

Part 2: A visibly fair mechanism is a ranked partitioned priority mechanism.

[Theorem 2](#)'s constructed zone selection function and argument apply here as well. Fix an arbitrary message profile $m \in M$ and let $a = \psi(m)$ be the allocation produced by any visibly fair mechanism ψ . We construct zone selection functions $\{C_k\}_{k=1}^n$ and sets $\{S^k\}_{k=1}^n$ that would reproduce the same allocation step by step.

Step 0: Set $S^1 = S$.

Step k ($1 \leq k \leq n$): If $a_k \in z$, set $\mathcal{C}_k(S^k, m) = z$. If the number of officers assigned to s^k reaches q_{s^k} , that is $|\{i \in I : i \leq k \text{ and } a_i = s^k\}| = q_{s^k}$, then $S^{k+1} \equiv S^k \setminus \{s^k\}$. Otherwise, $S^{k+1} = S^k$.

After performing the construction for every k we have defined all zone selection functions $\{C_k\}_{k=1}^n$ and sets $\{S^k\}_{k=1}^n$.

Notice that for officer k , the set S^k consists of only those states that still have remaining capacities to accommodate k after the allocation of higher-ranking officers is taken into account. That is, each state $s \in S \setminus S^k$ is filled by higher ranking offers, $|\{i \in I : i < k \text{ and } a_i = s\}| = q_s$. Now we must have that $a^k \in G(S^k, m_k) \cap \mathcal{C}_k(S^k, m)$, otherwise there exists $s \in S^k$ such that $s \succ_{m_k} a_k$, which makes the outcome visibly unfair as a higher ranking offer does not occupy this state.

Also, see that the constructed zone selection function is a ranked zone selection function. Suppose not, then for officer $k \in I$, $S^k \cap \mathcal{C}_k(S^k, m) = \{\min(\mathcal{C}_k(S^k, m), m_k)\}$ and there is a zone $z \in Z$ such that $z \triangleright_k \mathcal{C}_k(S^k, m)$ and $\max(m_k, z) \in S^k$. Since

$\max(m_k, z) \succ_{m_k} \min(\mathcal{C}_k(S^k, m), m_k)$, $a^k = \min(\mathcal{C}_k(S^k, m_k), m_k)$ is visibly unfair for officer k under m .

A ranked partitioned priority mechanism working down the priority list, using ranked zone selection functions $\{\mathcal{C}_k\}_{k=1}^n$, sets $\{S^k\}_{k=1}^n$ defined above and the procedure given in [Definition 5](#), would assign every officer the same post as ψ .

Since the construction can be repeated for every profile m , ψ is a ranked partitioned priority mechanism. □

A.4 Proof of [Theorem 4](#)

A visibly fair mechanism is strategy-proof if and only if it satisfies expressiveness and weak availability.

Proof. **If part a.** A visibly fair mechanism satisfying weak availability is strategy-proof if it satisfies expressiveness. We prove the contrapositive statement: If a visibly fair mechanism satisfying weak availability is not strategy-proof it does not satisfy expressiveness.

1. By assumption (violation of strategy-proofness), there exists a truthful message m_i and a message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \succ_i \psi(m)_i$.
2. By weak availability $\psi(\hat{m}_i, m_{-i})_i$ is available under m .
3. The expressed information $\psi(\hat{m}_i, m_{-i})_i \succ_{m_i} \psi(m)_i$, together with $\psi(\hat{m}_i, m_{-i})_i$ being available under m , would lead to a violation of visible fairness. Thus, we have found a violation of expressiveness.

If part b. A visibly fair mechanism satisfying expressiveness is strategy-proof if it satisfies weak availability. We prove the contrapositive statement: If a visibly fair mechanism satisfying expressiveness is not strategy-proof it does not satisfy weak availability.

1. By assumption (violation of strategy-proofness), there exists a truthful message m_i and a message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \succ_i \psi(m)_i$.
2. By expressiveness the information $\psi(\hat{m}_i, m_{-i})_i \succ_{m_i} \psi(m)_i$ is available under m .
3. The availability of $\psi(\hat{m}_i, m_{-i})_i$ under m , together with the expressed information $\psi(\hat{m}_i, m_{-i})_i \succ_{m_i} \psi(m)_i$, would lead to a violation of visible fairness. Thus, we have found a violation of weak availability.

Only if part a. A visibly fair mechanism satisfies expressiveness if it is strategy-proof. We prove the contrapositive statement: If a visibly fair mechanism does not satisfy expressiveness then it is not strategy-proof.

1. Note that for the same state expressiveness is always satisfied. As the mechanism does not satisfy expressiveness, there exist officer i and messages m_i and $\hat{m}_i$ such that $\psi(m)_i \neq \psi(\hat{m}_i, m_{-i})_i$, but neither $\psi(\hat{m}_i, m_{-i})_i \succ_{m_i} \psi(m)_i$ nor $\psi(m)_i \succ_{m_i} \psi(\hat{m}_i, m_{-i})_i$.
2. Consider any preference $\succsim_i^* \in Q_i$ s.t. $\psi(\hat{m}_i, m_{-i}) \succ_i^* \psi(m)$ and $s \succ_i^* s'$ whenever $s \succ_{m_i} s'$. By construction, m_i is a truthful message for preference $\succsim_i^* \in Q_i$. Moreover, there exists another message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \succ_i \psi(m)_i$, i.e., a violation of strategy-proofness.

Only if part b.

A visibly fair mechanism satisfies weak availability if it is strategy-proof. We prove the contrapositive statement: If a visibly fair mechanism does not satisfy weak availability then it is not strategy-proof.

1. As the mechanism does not satisfy weak availability, there exists $m_i \in M_i$ that is truthful for a preference $\succsim_i \in Q_i$ and a message $m'_i \in M_i \setminus \{m_i\}$ s.t. $\psi(\hat{m}_i, m_{-i})_i \succsim_i \psi(m)_i$; but state $\psi(\hat{m}_i, m_{-i})_i$ is not available to officer i under m .
2. It follows that $\psi(\hat{m}_i, m_{-i})_i \neq \psi(m)_i$, and therefore we have a successful manipulation, as $\psi(m'_i, m_{-i})_i \succ_i \psi(m)_i$.

□

Proof of Theorem 5

A mechanism is strategy-proof if and only if it satisfies coherence.

Proof. If. A mechanism is strategy-proof if it satisfies coherence. We prove the contrapositive statement: If a mechanism is not strategy-proof it does not satisfy coherence.

By assumption, there exists a truthful message m_i and a message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \succ_i \psi(m)_i$. As $\psi(\hat{m}_i, m_{-i})_i$ is strictly preferred we have $\psi(\hat{m}_i, m_{-i})_i \neq \psi(m)_i$. As m_i is truthful we cannot have $\psi(m) \succ_{m_i} \psi(\hat{m}_i, m_{-i})$,

and therefore $\psi(\hat{m}_i, m_{-i}) \in G(\{\psi(m)\} \cup \{\psi(\hat{m}_i, m_{-i})\}, m_i)$, i.e., a violation of coherence.

Only if. A mechanism satisfies coherence if it is strategy-proof. We prove the contrapositive statement: If a mechanism does not satisfy coherence then it is not strategy-proof.

By assumption, there exists a message m_i and a message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i}) \in G(\{\psi(m)\} \cup \{\psi(\hat{m}_i, m_{-i})\}, m_i)$. Consider any preference $\succsim_i^* \in Q_i$ s.t. $\psi(\hat{m}_i, m_{-i}) \succ_i^* \psi(m)$ and $s \succ_i^* s'$ whenever $s \succ_{m_i} s'$. By construction, m_i is a truthful message for preference $\succsim_i^* \in Q_i$. Moreover, there exists another message $\hat{m}_i \in M_i$ such that $\psi(\hat{m}_i, m_{-i})_i \succ_i \psi(m)_i$, i.e., a violation of strategy-proofness. $\square$

A.5 Proof of Theorem 6

The Modular Priority Mechanism (MPM) is visibly fair, strategy-proof and respects modular upper-bounds.

Proof. First, it should be clear that the Modular Priority Mechanism is a partitioned priority mechanism: each agent, following the priority order, is associated with a zonal message space, and is matched to the most-preferred state from a zone that still has states with spare capacity. Since the modular upper-bound system H satisfies sequential solvency, at every step of the MPM an officer has a state available for which no upper-bound is binding and with spare capacity.

Next, we show that, for any given set of type-specific modular upper-bound systems and exogenous rankings $(\blacktriangleright_i)_{i \in I}$, the Modular Priority Mechanism, represented by the function ψ , is strategy-proof.

Notice first that, by construction, an officer i cannot, by submitting a different message, change the assignment of any officer $j < i$. Since ψ only produces feasible outcomes, this implies that ψ satisfies availability.

Since the zone z from which i 's assignment will be drawn from depends on the assignment of officers with higher priority and the ranking $\blacktriangleright_i$, both unaffected by m_i , the zone to which i will be assigned during the execution of the mechanism cannot be changed by changes in her message. Let $s = \psi(m)$. Since the zone from which i 's outcome will be drawn will always be in z regardless of the message i sends, $\psi(m'_i, m_{-i}) \in z$ for any $m'_i \in M_i$. Since both s and $\psi(m'_i, m_{-i})$ belong to zone z , it must be the case that s and $\psi(m'_i, m_{-i})_i$ are comparable under m_i . Therefore, ψ satisfies expressiveness, and by Corollary 2, strategy-proofness.

Finally, we need to show that the Modular Priority Mechanism respects modular upper-bounds. That is, the final allocation $\psi(m) = a$ satisfies, for every upper-bound $(\Xi_h, S_h, k_h) \in H$:

$$|\{i \in I : a_i \in S_h, t_i \in \Xi_h\}| \leq k_h.$$

We prove this by induction on the steps $k = 1, 2, \dots, n$ of the mechanism. In order to facilitate notation and comprehension, we will denote by a^k the “tentative” allocation a by the end of step k of the mechanism.

Base Case ($k = 0$):

At initialization, no officers have been assigned. That is, $a_i^0 = \emptyset$ for all $i \in I$. Therefore, for any upper-bound $(\Xi_h, S_h, k_h) \in H$, we have:

$$|\{i \in I : a_i^0 \in S_h, t_i \in \Xi_h\}| = 0 \leq k_h.$$

Thus, the allocation trivially respects all modular upper-bounds at step $k = 0$.

Inductive Step:

Assume that at step $k - 1$, the current assignment a^{k-1} respects all modular upper-bounds; that is, for every upper-bound $(\Xi_h, S_h, k_h) \in H$:

$$N_h^{k-1} = |\{i \in I : a_i^{k-1} \in S_h, t_i \in \Xi_h\}| \leq k_h.$$

We need to show that after assigning officer i_k at step k , the updated assignment a^k also respects all modular upper-bounds.

Consider officer i_k of type t_{i_k} . According to the mechanism, i_k is assigned to a state s^* within a zone $z_{\ell^*}^{t_{i_k}}$ where the upper-bounds for type t_{i_k} involving states in that zone are not yet binding before i_k ’s assignment—this is tracked by the flag $B_{\ell^*}^{t_{i_k}} = \text{False}$.

That is, for every upper-bound $(\Xi_h, S_h, k_h) \in H$ for which $s^* \in S_h$:

$$|\{i \in I : a_i^{k-1} \in S_h, t_i \in \Xi_h\}| < k_h.$$

Since all of these constraints are strictly below their upper-bounds after assigning i_k to s^* they will remain below their bounds or hit them. None of them will go beyond their limits. Therefore, the resulting assignment a^k respects all modular upper-bounds.

Conclusion:

In both cases, the assignment a^k at step k respects all modular upper-bounds. By induction, the final allocation a^n produced by the mechanism satisfies:

$$|\{i \in I : a_i^n \in S_h, t_i \in \Xi_h\}| \leq k_h,$$

for every every upper-bound $h \in H$.

□

A.6 Proof of Theorem 7

Not all modular upper-bound constraints admit a (static) mechanism that is simultaneously visibly fair and constrained Pareto efficient.

Proof. Consider a problem with three officers $I = \{i_1, i_2, i_3\}$ (with priority $\pi(i_1) < \pi(i_2) < \pi(i_3)$) and two states $S = \{s_1, s_2\}$ with capacity $q_{s_1} = q_{s_2} = 2$. Assume that all officers are of the same type t and that there is a single modular upper-bound $(\{t\}, \{s_1\}, 1)$; that is, at most one officer of type t may be assigned to s_1 .

Thus, any allocation that *respects the upper-bounds* must assign at most one officer to s_1 . We now show that no visibly fair mechanism in this setup is constrained Pareto efficiency.

Since there are two states, each officer's message space admits only two possibilities regarding the comparison between s_1 and s_2 : (i) The officer can express a strict preference between s_1 and s_2 , or (ii) The officer cannot express any comparison between s_1 and s_2 .

Next, note that no mechanism that elicits preferences of more than one officer between s_1 and s_2 and is visibly fair respects modular upper-bounds. To see that, assume that two officers express that $s_1 \succ_i s_2$. Visible fairness requires that s_1 does not have any empty slot. But this would violate modular upper-bounds.

Any visibly fair mechanism that respects modular upper-bounds, therefore, will not elicit preferences for at least two officers.

Let's denote the overall message-space configuration by a triple (X_1, X_2, X_3) , where $X_k = Y$ if officer i_k can express a preference between s_1 and s_2 , and $X_k = N$ otherwise.

We now examine the four cases in which we do not elicit preferences for at least two officers and show that in each case, *either* the mechanism must return an assignment that violates the modular upper-bound (i.e., assigns two or more

officers to s_1), *or* the outcome is not constrained Pareto efficient (because it is Pareto dominated by another allocation that still respects the upper-bound).

Case 1: $(X_1, X_2, X_3) = (Y, N, N)$. Suppose i_1 reports $s_2 \succ s_1$. By visible fairness she is assigned s_2 . By visible fairness and modular upper-bounds, the mechanism must assign i_2 and i_3 to s_1 and s_2 . Suppose it assigns i_2 to s_1 and i_3 to s_2 . If their true preferences are such that $s_2 \succ_{i_2} s_1$ and $s_1 \succ_{i_3} s_2$, an assignment that swaps their matches Pareto dominates the outcome, while still respecting modular upper-bounds. The same reasoning can be applied if the mechanism swaps the assignments of i_2 and i_3 .

Case 2: $(X_1, X_2, X_3) = (N, N, Y)$. Suppose that i_3 reports $s_2 \succ_{i_3} s_1$. There are two cases to consider: (i) one of i_1, i_2 is matched to s_1 , or (ii) both i_1 and i_2 are matched to s_2 . In the first case, let without loss of generality i_1 be matched to s_2 and i_2 be matched to s_1 and. Then, if i_1 prefers s_1 and i_2 prefers s_2 , swapping their assignments would Pareto dominate the outcome produced by the mechanism while still respecting modular upper-bounds.

In the second case, both i_1 and i_2 are matched to s_2 . Suppose that one of them prefer s_1 to s_2 . Then swapping that agent's assignment with i_3 would Pareto dominate the outcome produced by the mechanism while still respecting modular upper-bounds.

Case 3: $(X_1, X_2, X_3) = (N, Y, N)$. Suppose that i_2 reports $s_2 \succ_{i_2} s_1$. Then, visible fairness requires that she is assigned to s_2 . By modular upper-bounds, i_1 and i_3 are matched to s_1 and s_2 . Suppose that both prefer the other's assignment. Then, swapping their matches would Pareto dominate the outcome produced by the mechanism while still respecting modular upper-bounds.

Case 4: $(X_1, X_2, X_3) = (N, N, N)$.

The mechanism must choose an assignment of two officers to s_2 and one to s_1 solely by some fixed rule. Let i be the officer matched to s_1 and i', i'' be matched to s_2 (their precise identities do not matter in this case). Suppose that i prefers s_2 to s_1 , and i' prefers s_1 to s_2 . Then, swapping their matches would Pareto dominate the outcome produced by the mechanism while still respecting modular upper-bounds.

In every case, the following occurs: either the outcome produced by the mechanism (which is constrained to be *upper-bound respecting*) violates the modular upper-bound (by assigning more than one officer to s_1), or there exists another

allocation that both respects the upper-bounds and yields a strict Pareto improvement with respect to the officers' true preferences. Therefore, for these modular upper-bounds, no mechanism exists that is simultaneously visibly fair and is constrained Pareto efficient. □

A.7 Proof of Theorem 8

The Dynamic Modular Priority Mechanism is visibly fair, constrained Pareto efficient, strategy-proof, and respects modular upper-bounds.

Proof. Part 1: Visible fairness.

In the Dynamic Modular Priority Mechanism, officers are processed in strict priority order $i_1, i_2, \dots, i_n$. For each officer i_k , the mechanism partitions the state space S into two zones: z_1 and z_2 . Officer i_k is matched to her most-preferred state in z_1 . Therefore, there is no state that i_k ranks above her assignment. This is true for every officer, and therefore the resulting assignment is visibly fair.

Part 2: Respects modular upper-bounds.

We show by induction on the assignment order that for any modular upper-bound (Ξ_h, S_h, k_h) , the number of officers with type in Ξ_t assigned to states in S_h never exceeds k_h .

Base Case. Before any assignment (step 0), no officer is assigned and the count is $0 \leq k_h$.

Inductive Step. Suppose that after assigning officers $i_1, \dots, i_{k-1}$, the upper-bound (Ξ_h, S_h, k_h) is not violated. When officer i_k (of type t_{i_k}) is considered, if $t_{i_k} \notin \Xi_h$ the claim is unaffected. If $t_{i_k} \in \Xi_h$ and i_k is assigned a state s with $s \in S_h$, then by construction s belongs to zone z_1 for i_k and the current count of officers with type $t \in \Xi_h$ in S_h is strictly lower than k_h . If after assignment the count reaches k_h , the mechanism updates B^t by adding h ; hence, any subsequent officer of type t will have S_h excluded from her zone z_1 . Thus, by induction, no upper-bound is ever violated.

Part 3: Strategy-proof.

We show that no officer can obtain a strictly better outcome by misreporting her preferences. In the mechanism, the available set of states for officer i_k is determined solely by:

1. The assignments of all higher-priority officers $i_1, \dots, i_{k-1}$, and
2. The set $B_k^{t_{i_k}}$ of binding modular upper-bounds for type t_{i_k} .

Neither of these depends on the message m_{i_k} provided by officer i_k . Thus, her feasible set—the zone z_1 —remains fixed regardless of her report. Within z_1 , if m_{i_k} truthfully reflects her preference ordering, she is assigned her top available choice. Any misreport would merely permute her ranking over the same fixed set z_1 and cannot result in an outcome strictly preferred (by her true preference) to her truthful assignment. Therefore, the mechanism is strategy-proof.

Part 4: Constrained Pareto efficient.

Assume for contradiction that there exists an allocation a' which Pareto dominates a and also respects all modular upper-bounds. Let i^* be the highest-priority officer for whom $a'_{i^*} \succ_{i^*} a_{i^*}$; that is, for every officer $j < i^*$, we have $a'_j = a_j$.

Consider officer i^* . Under the Dynamic Modular Priority Mechanism, when i^* was processed, she was assigned a_{i^*} as her most-preferred available state in zone z_1 . If $a'_{i^*} \succ_{i^*} a_{i^*}$, then the state a'_{i^*} must have been available when i^* was considered. There are two possibilities:

1. *Capacity constraint:* The state a'_{i^*} might have been unavailable because its capacity was already exhausted by officers $j < i^*$. However, since $a'_j = a_j$ for all $j < i^*$, the capacity allocated in a and a' is identical for states assigned to higher-priority officers. Thus, if a'_{i^*} is available in a' , it must have been available for i^* in a as well.
2. *Binding Upper-Bound:* The other reason for a'_{i^*} not to be in i^* 's zone z_1 under a is if a'_{i^*} belonged to a set S_h for which the corresponding upper-bound (Ξ_h, S_h, k_h) was already binding at the time of i^* 's assignment. However, if a'_{i^*} were in such a set and a' respects the upper-bound, then a'_{i^*} could not be assigned to i^* without causing the total number of type t_{i^*} officers in S_h to exceed k_h .

In either case, if a'_{i^*} is strictly better for i^* than a_{i^*} , then the state a'_{i^*} must have been available for i^* under the mechanism; hence, the mechanism would have assigned a'_{i^*} to i^* rather than a_{i^*} . This contradiction shows that no allocation a' can Pareto dominate a while still respecting all modular upper-bounds. Therefore, any allocation a' that Pareto dominates a must violate some modular upper-bound. $\square$

A.8 More on efficiency

This section explores the relationship between visible efficiency and fairness in allocation mechanisms. An allocation is visibly efficient under message profile m

if no alternative allocation visibly Pareto dominates it (meaning no reallocation would make all affected agents better off according to their reported preferences). The key results show that visibly fair allocations are always visibly efficient, but the converse doesn't hold. When comparing visible efficiency to true Pareto efficiency (based on actual preferences), every Pareto efficient allocation is visibly efficient under truthful messages, but visible efficiency doesn't guarantee Pareto efficiency—even visibly fair allocations can be Pareto inefficient. Furthermore, when message $\hat{m}$ contains more preference information than m , allocations that are visibly fair or efficient under $\hat{m}$ remain so under m , but not vice versa, demonstrating that visible fairness and efficiency depend critically on the message spaces.

An allocation a is **visibly efficient under** m if it is not visible that it is Pareto dominated by another allocation, that is, there is no allocation $a' \in \mathcal{A}$ such that for all $i \in I$ with $a'_i \neq a_i$, we have $a'_i \succ_{m_i} a_i$. A mechanism ψ is **visibly efficient** if $\psi(m)$ is visibility efficient for all $m \in M$.

Theorem 9. *For any message $m \in M$, the following statements are true.*

1. *Every visibly fair allocation is visibly efficient.*
2. *A visibly efficient allocation may not be visibly fair.*

Proof. Statement 1: Fix a message profile $m \in M$. Suppose $a \in \mathcal{A}$ is a visibly fair allocation that is not visibly efficient under m . Then there exists another allocation $a' \in \mathcal{A}$ such that for all $i \in I$ with $a'_i \neq a_i$, we have $a'_i \succ_{m_i} a_i$.

Let i^* be the highest ranking officer among the ones that are allocated to a different state under a and a' and let s^* be her assigned state under a' . That is, for

$$\bar{I} := \{i \in I : a_i \neq a'_i\}, \quad i^* := \arg \min_{i \in \bar{I}} \pi(i), \quad \text{and} \quad a'_{i^*} = s^*.$$

The allocation a must be visibly unfair under m . This is because either $a_i = s^*$ for some $i \in \bar{I}$, or $|\{i \in I : a_i = s^*\}| < q_{s^*}$. Yet $s^* \succ_{m_{i^*}} a_{i^*}$.

Statement 2: Let $I = \{i_1, i_2\}$, $S = \{s_1, s_2\}$, and $q_{s_1} = q_{s_2} = 1$. Additionally, assume zonal message space for both officers, $Z = \{z_1\}$ with $z_1 = \{s_1, s_2\}$. For preferences $s_1 \succ_{i_1} s_2$ and $s_1 \succ_{i_2} s_2$, the allocation $(a_1, a_2) = (s_2, s_1)$ is visibly efficient but not visibly fair under the truthful message.

□

In contrast, an allocation a is **Pareto efficient** if it is not Pareto dominated by another allocation, that is, there is no allocation $a' \in \mathcal{A}$ such that for all $i \in I$ with $a'_i \neq a_i$, we have $a'_i \succ_i a_i$. A mechanism ψ is **Pareto efficient** if $\psi(m)$ is Pareto efficient for all $m \in M$.

Theorem 10. *For any truthful message $m \in M$, the following statements are true.*

1. *Every Pareto efficient allocation is visibly efficient.*
2. *A visibly efficient allocation may not be Pareto efficient.*
3. *A visibly fair allocation may not be Pareto efficient.*

Proof. Statement 1: Consider an allocation $a \in \mathcal{A}$ that is not visibly efficient for some truthful message $m \in M$. This implies there is another allocation $a' \in \mathcal{A}$ such that for all $i \in I$ with $a'_i \neq a_i$, we have $a'_i \succ_{m_i} a_i$, and therefore $a'_i \succ_i a_i$ (m is a truthful message). Thus, a is not Pareto efficient.

Statement 2 and 3: Let $I = \{i_1, i_2\}$, $S = \{s_1, s_2, s_3\}$, and $q_{s_1} = q_{s_2} = q_{s_3} = 1$. Additionally, assume a zonal message space, $Z = \{z_1, z_2\}$ with $z_1 = \{s_1, s_2\}$ and $z_2 = \{s_3\}$. For preferences $s_3 \succ_{i_1} s_1 \succ_{i_1} s_2$ and $s_1 \succ_{i_2} s_3 \succ_{i_2} s_2$, the allocation $(a_1, a_2) = (s_1, s_3)$ is visibly fair and visibly efficient under the truthful message, as s_1 and s_3 are not comparable. However, it is not Pareto efficient. $\square$

Message $\hat{m}$ **contains more preference information than** m , if for all $i \in I$:

$$s \succ_{m_i} s' \implies s \succ_{\hat{m}_i} s'.$$

Theorem 11. *Suppose message $\hat{m}$ contains more preference information than m . Then the following statements are true.*

1. *Every visibly fair allocation under $\hat{m}$ is also visibly fair under m .*
2. *A visibly fair allocation under m may not be visibly fair under $\hat{m}$.*
3. *Every visibly efficient allocation under $\hat{m}$ is also visibly efficient under m .*
4. *A visibly efficient allocation under m may not be visibly efficient under $\hat{m}$.*

Proof. Statement 1: There are two cases: (i) Allocation a is not visibly fair under m because there exist some $i \in I$ such that there is a $j \in I$ such that $a_i \neq a_j$, $\pi(i) < \pi(j)$, and $a_j \succ_{m_i} a_i$. Since $a_j \succ_{\hat{m}_i} a_i$, a cannot be visibly fair under $\hat{m}$.

(ii) Allocation a is not visibly fair under m because there exist some $i \in I$ such that there is a $s \in S$ such that $a_i \neq s$, $|\{i \in I : a_i = s\}| \leq q_s$, and $s \succ_{m_i} a_i$. Since $s \succ_{\hat{m}_i} a_i$, a cannot be visibly fair under $\hat{m}$.

Statement 3: We use the contrapositive. If allocation a is not visibly efficient under m , then there exists another allocation a' such that for all $i \in I$ with $a'_i \neq a_i$,

we have $a'_i \succ_{m_i} a_i$. Since $a'_i \succ_{\hat{m}_i} a_i$ for all such $i \in I$, a cannot be visibly efficient under $\hat{m}$.

Statement 2 and 4: Let $I = \{i_1, i_2\}$, $S = \{s_1, s_2, s_3\}$, and $q_{s_1} = q_{s_2} = q_{s_3} = 1$. Additionally, assume a zonal message space for every officer, $Z = \{z_1, z_2\}$ with $z_1 = \{s_1, s_2\}$ and $z_2 = \{s_3\}$. For preferences $s_3 \succ_{i_1} s_1 \succ_{i_1} s_2$ and $s_1 \succ_{i_2} s_3 \succ_{i_2} s_2$, the allocation $(a_1, a_2) = (s_1, s_3)$ is visibly fair and visibly efficient under the truthful message m as s_1 and s_3 are not comparable.

Consider the zonal message space $Z = \{S\}$. The truthful message $\hat{m} = (\succ_i)_{i \in I}$ contains more preference information than m . However, allocation $(a_1, a_2) = (s_1, s_3)$ is not visibly fair nor visibly efficient under $\hat{m}$. $\square$

B Appendix: Proofs that examples satisfy sequential solvency

Proof that Example 5.2 satisfies sequential solvency.

Assume, for contradiction, that there exists an officer i who, given some feasible placement a_{-i} of the other 7 officers, has *no* admissible state. Without loss of generality let $t_i = R_1$.

- If any seat in R_2 were vacant it would be admissible for i (because the bound $(R_1, 2)$ does not apply there), contradicting the assumption. Hence both states in R_2 are fully occupied (4 officers).
- The local cap in R_2 is $(R_2, 2)$; therefore at most two of the four occupants can be of type R_2 . So at least $4 - 2 = 2$ of them are of type R_1 .
- Among the remaining 7 officers, at most one R_1 -type can still be in R_1 (since two are already in R_2). Thus $(R_1, 2)$ is *not* binding in R_1 .
- Only 7 officers are placed, but R_1 has capacity $q_{s_1} + q_{s_2} = 4$; hence at least one seat in R_1 is empty.

But then, $(R_1, 2)$ is *not* binding in R_1 and there is spare capacity in R_1 , contradicting i not having an admissible state.

Proof that Example 5.3 satisfies sequential solvency.

Suppose, for contradiction, that some type-1 doctor i and feasible placement a_{-i} of the other 26 doctors leave i without an admissible hospital. This can only happen if:

- (i) all hospitals in R_2 and R_3 are full; or
- (ii) the only vacancies there are urban and the urban cap (19) is binding.

Case (i): $Q(R_2) + Q(R_3) = 24$, so at most 24 of the others are outside R_1 . Then at most 2 are in R_1 , leaving ≥ 10 free seats there; the local cap $(R_1, 6)$ is not binding, so i could be placed in R_1 .

Case (ii): With 19 urban doctors, the remaining $26 - 19 = 7$ occupy rural seats, leaving $12 - 7 = 5$ rural vacancies overall. If all rural seats in R_2 and R_3 were full, they would require 8 rural occupants, contradicting the total of 7. Thus some rural hospital in R_2 or R_3 has a vacancy admissible to i .

In both cases we contradict the assumption; hence sequential solvency holds.

C Appendix: Indirect Message Spaces

C.1 Preference Elicitation in All India Services

The 2017 Cadre Allocation Policy for India’s All India Services, including the Indian Administrative Service (IAS), Indian Police Service (IPS), and Indian Forest Service (IFoS), introduces a zonal system that divides all states and union territories into five geographical zones, requiring candidates to first indicate their zone preferences in descending order, followed by cadre preferences within each preferred zone. For illustrative purposes, we include a screenshot of the submitted preferences from 2020 IFoS examination.

C.2 Rank-order lists in Chinese College Admissions

23 out of 31 provinces in China implement the structured rank-order list system, in which majors are effectively nested under colleges, as noted by [Hu *et al.* \(2025\)](#). These provinces include: Shanghai, Beijing, Tianjin, Hainan, Jiangsu, Fujian, Hubei, Hunan, Guangdong, Heilongjiang, Gansu, Jilin, Anhui, Jiangxi, Guangxi, Shanxi, Henan, Shaanxi, Ningxia, Sichuan, Yunnan, Tibet, and Xinjiang. For illustrative purposes, we include a screenshot of the official college-major preference form from Shanghai.

C.3 Reserve Officer Training Corps (ROTC) Mechanism

[Sönmez \(2013\)](#)’s model of cadet-branch matching problem consists of

1. a finite set of cadets $I = \{i_1, i_2, \dots, i_n\}$,
2. a finite set of branches $B = \{b_1, b_2, \dots, b_m\}$,
3. a vector of branch capacities $q = (q_b)_{b \in B}$,
4. a set of “terms” $T = \{t_1, \dots, t_k\}$,
5. a list of cadet preferences $P = (P_i)_{i \in I}$ over $(B \times T) \cup \{\emptyset\}$, and
6. a list of base priority rankings $\pi = (\pi_b)_{b \in B}$.

The ROTC mechanism is not direct. Instead, each cadet submits a ranking of branches $\succ'_i$, and he can sign a branch-of-choice contract for any of his top three choices under $\succ'_i$

Rank No.	Roll_No	Name of Candidate	Date of Birth	Gender	Education Qualificatoin	Home_State (Yes/No)	Home State	PREFERENCE																									Addl. Preference for PH						
								Zone I					Zone II					Zone III					Zone IV					Zone V											
								AG	HR	MP	PB	RJ	UK	Pref. for Zones	BH	JH	OD	UP	Pref. for Zones	CHH	GJ	MP	MS	Pref. for Zones	AM	MN	NG	SK	TR	WB	Pref. for Zones	AP		KT	KR	TN	TG		
1	1303818	SCORAL BEN K R	28.09.1993	M	M.Sc.	YES	KERALA	3	4	5	2	6	3	1	4	3	2	1	4	2	3	4	1	2	5	1	5	6	3	4	2	1	4	2	1	5	3		
2	1210900	GOBBILLA VIDYADHARI	07.12.1994	F	B.TECH	YES	ANDHRA PRADESH	3	5	6	1	4	3	2	4	3	2	1	4	2	3	1	4	2	5	3	6	5	1	4	2	1	1	5	3	4	2		
3	848230	PALLUVAL VISHNU VARDHAN REDDY	01.07.1994	M	B.TECH	YES	ANDHRA PRADESH	3	1	5	3	6	4	2	5	4	2	1	3	2	3	4	2	1	4	2	5	6	1	4	3	1	1	5	3	4	2		
4	1005056	K A V PRASAD REDDY	28.08.1993	M	M.B.A.	YES	TELANGANA	3	3	5	4	6	2	1	4	4	2	1	2	2	4	3	2	1	5	2	5	6	1	4	3	1	2	5	4	3	1		
5	1134380	KALPESH KUMAR SHARMA	13.04.1994	M	B.TECH	YES	RAJASTHAN	1	6	4	3	5	1	2	4	4	2	1	3	2	3	4	1	2	5	1	5	6	2	4	3	3	2	1	5	4	2		
6	66319422	JONG PAVANKUMAR APPASAHEB	11.06.1988	M	M.Sc.	YES	MAHARASHTRA	4	4	5	3	6	1	2	3	4	2	1	3	1	3	4	2	1	5	2	5	6	1	4	3	2	2	1	5	4	3		
7	319662	GURHARSH SINGH	13.07.1989	M	M.Sc.	YES	HIMACHAL PRADESH	1	3	5	1	4	6	2	4	3	2	1	4	2	3	4	1	2	5	2	3	6	1	4	5	3	2	1	4	3	5		
8	6905519	GAURAV JAIN	04.02.1992	M	M.TECH	YES	MADHYA PRADESH	4	4	5	3	6	1	2	3	4	2	1	3	1	2	3	1	4	5	2	5	6	3	4	1	2	3	1	4	5	2		
9	4199127	SHREYAS SRIVASTAVA	17.06.1989	M	B.TECH	YES	MADHYA PRADESH	2	3	5	1	6	4	2	4	4	3	1	2	1	3	4	1	2	5	2	5	6	1	4	3	3	4	1	5	3	2		
10	5607027	DHIVYA N	02.11.1997	F	B.TECH	YES	TAMIL NADU	2	4	6	1	5	3	2	4	3	2	1	4	3	4	3	2	1	5	2	5	6	3	4	1	1	4	2	3	1	5		
11	803894	CHAVAN SUHAS MADHUKAR	28.10.1990	M	M.TECH	YES	MAHARASHTRA	4	3	5	4	6	1	2	3	4	3	2	1	4	1	3	4	2	1	5	2	5	6	1	4	3	2	2	1	5	4	3	
12	1118577	RAAMKRISHNA SARAN	25.06.1996	M	B.TECH	YES	RAJASTHAN	1	6	4	3	5	6	1	2	3	3	2	1	4	1	2	3	4	1	2	5	2	5	6	1	4	3	2	1	5	4	2	
13	400024	SHUBHAM BAJAJ	20.02.1997	M	M.TECH	YES	MADHYA PRADESH	4	4	3	5	6	1	2	3	3	2	1	4	1	2	3	1	4	2	5	3	5	6	1	4	2	2	3	1	5	4	2	
14	16625781	VIKALPA N VISHWAKARMA	02.02.1992	F	M.TECH	YES	MAHARASHTRA	4	4	6	3	5	5	1	2	1	3	3	2	1	4	1	3	4	2	1	5	3	5	6	1	2	4	2	5	1	3	4	2
15	6119017	MRIDULA SINGH	01.03.1997	F	M.TECH	YES	UTTAR PRADESH	3	6	4</																													

2023 年上海市普通高等学校招生考生志愿表
表 4 - 本科普通批次（样表）

高考报名号				姓名	性别	所在报名区	选考科目	应试语种			
批 次	院校专业组代码	院校专业组名称	专业代码	专业名称	专业代码	专业名称	专业代码	专业名称	专业代码	专业名称	是否专业调剂
本科普通批次	1										
	2										
	3										
	4										
	5										
	6										
	7										
	8										
	9										
	10										
	11										
	12										
	13										
	14										
	15										
	16										
	17										
	18										
	19										
	20										
	21										
	22										
	23										
	24										

注：1. 考生在本科普通批次最多可填报 24 个院校专业组志愿，每个院校专业组内最多可填报 4 个专业志愿。
2. 本科普通批次实行平行志愿，投档时，按考生高考成绩总分从高到低排序，逐分、逐个地按其填报的院校专业组志愿顺序依次检索，一旦投档则不再继续检索，实行一次投档。
3. 考生在每个院校专业组志愿中均须填写“是否专业调剂”代码，其中“1”表示全愿意，“2”表示全不愿意，“3”表示除高收费专业外其他愿意，“4”表示除医科专业外其他愿意，考生在该栏必须填写“1”到“4”之间的数字，其中“3”“4”可以并列填写。专业调剂只能在所投档的院校专业组内开展。
4. 考生须在志愿填报系统中按规定输入志愿表信息。志愿表须一式三份打印，并由考生本人签字（其他人签字无效），作为投档录取依据。

考生签名： 日期： 年 月 日

Source: <https://www.shmeea.edu.cn/page/08000/20230407/17353.html>