Adaptive Priority Mechanisms

Oğuzhan Çelebi* Joel P. Flynn†

Michigan Yale

March 31, 2026

Abstract

We study a model in which an authority wishes to allocate scarce resources to high-scoring agents while caring about the composition of the allocated agents' characteristics. The key challenge facing the authority is uncertainty about the joint distribution of agents' scores and characteristics. We introduce *adaptive priority mechanisms* (APM) that prioritize agents based on their scores and their rank relative to others with the same characteristics. We show that APM uniquely implement the *ex post* optimal allocation. The ubiquitous priority and quota mechanisms are optimal if and *only if* knife-edge conditions on preferences are satisfied—the authority must be risk-neutral or extremely risk-averse over diversity, respectively. With multiple authorities, Deferred Acceptance implements the unique stable matching when all authorities use the optimal APM. We provide a practical roadmap for implementing APM as a market-design solution and illustrate the gains from APM in an application to the Chicago Public Schools system.

*Ross Business School, the University of Michigan, 701 Tappan Street, Ann Arbor, MI 48109 Email: ocelebi@umich.edu

†Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, 06511. Email: joel.flynn@yale.edu

We are grateful to Daron Acemoglu, Mohammad Akbarpour, George-Marios Angeletos, Nick Arnosti, Jonathan Cohen, Viola Corradini, Roberto Corrao, Mert Demirer, Glenn Ellison, Aytek Erdil, Arda Gitmez, Marina Halac, Ömer Karaduman, Elliot Lipnowski, Stephen Morris, Anh Nguyen, Parag Pathak, Charlie Rafkin, Karthik Sastry, Tayfun Sönmez, Alex Teytelboym, Bertan Turhan, Alexander Wolitzky, Bumin Yenmez and seminar participants at the 2022 INFORMS Workshop on Market Design, the 2022 Society for the Advancement of Economic Theory Conference, Michigan Ross, UC Berkeley, The University of Pennsylvania, Microsoft Research, Iowa State University, and the MIT Theory Lunch for helpful comments. We thank Chicago Public Schools for graciously sharing their data and Eryn Heying, Talia Gerstle and Jim Shen for invaluable administrative support. The views expressed here are those of the authors and do not reflect the views of Chicago Public Schools.

1 Introduction

Authorities that allocate resources such as school seats, university courses, refugee visas, and medical supplies often face conflicting objectives. On the one hand, they want to maximize match quality or fairness by allocating resources to the highest-scoring agents according to various criteria such as academic attainment, qualifications (Kornbluth and Kushnir, 2021), likelihood of gaining employment (Ahani, Andersson, Martinello, Teytelboym, and Trapp, 2021; Delacrétaz, Kominers, and Teytelboym, 2023), mortality risk (Pathak, Sönmez, Unver, and Yenmez, 2021), or distance (Dur, Kominers, Pathak, and Sönmez, 2018). On the other hand, they want to achieve diversity across socioeconomic attributes (Aygün and Turhan, 2017; Dur, Pathak, and Sönmez, 2020; Sönmez and Yenmez, 2022a).

To balance these trade-offs, when the use of prices is seen as infeasible or unethical, authorities have broadly used two classes of policies (see Sönmez, 2023, for an extensive review of the policies implemented by authorities in various settings): *quotas*,[1] where a certain portion of the resource is set aside for given groups; and *priorities*, where individuals in given groups receive higher scores.

However, the practical performance of these policies is threatened by uncertainty about the *distribution* of individuals' scores, characteristics, and preferences. For example, Doğan and Erdil (2022) note that "[V]olatile trends of demand for different [degree programs] mean even meeting intake targets can be a difficult juggling act for universities." Exemplifying the substantial financial costs that uncertainty can induce, some UK universities in 2021 offered students £10,000 to study elsewhere after they failed to accurately forecast students' grades (Guardian, 2021). Similarly, Bansak, Lee, Manshadi, Niazadeh, and Paulson (2024) document the critical importance of hard-to-forecast migratory trends in U.S. refugee resettlement, observing that "lack of information about future arrivals presents a major challenge in optimizing matching decisions." Even sophisticated econometric tools provide a limited solution to the problems posed by uncertainty. In the Boston Public Schools system, Pathak and Shi (2021) observe that even state-of-the-art discrete-choice tools lead to "inconsistent performance" in predicting market outcomes due to unpredictable changes in the distribution of applicants' characteristics, finding that "[t]he largest source of error turns out to be in the applicant pool predictions."

In the face of such uncertainty, what mechanism *should* an authority use? Despite the practical importance of this question, we do not know if an authority should use a priority mechanism, a quota mechanism, or something else entirely.

[1] We use *quota* as a general term that includes the widely used reserve policies (see Definition 4). These are a special case of the slot-specific priorities introduced in (Kominers and Sönmez, 2016).

Summary of Main Results. In this paper, we formulate and solve the optimal mechanism design problem of an authority that allocates a resource to agents who are heterogeneous in their scores and belong to different groups. We propose a new class of *adaptive priority mechanisms* (APM) that adjust agents' scores as a function of the number of assigned agents with the same characteristics and that allocate the resource to the set of agents with the highest adjusted scores. With a single authority, we derive an APM that is optimal, implements a unique outcome, and can be specified solely in terms of the *preferences* of the authority (*i.e.,* it is optimal regardless of their beliefs). By contrast, we show that priorities and quotas are optimal if and only if risk aversion over diversity is extremely low or high, respectively. Moreover, optimally set priority and quota policies depend on both the preferences and beliefs of the authority. Thus, the optimal APM improves outcomes, is robust to uncertainty, and requires less information. With many authorities, we show that the combination of the Deferred Acceptance algorithm along with each authority using its optimal APM implements the unique stable allocation. Finally, we provide a practical roadmap for implementing APM as a market design solution. To demonstrate the practicality and potential benefits of adopting APM, we follow this roadmap and provide a proof-of-concept implementation using application and admission data from Chicago Public Schools. We find that the gains from adopting APM have the potential to be large.

Single-Authority Model. We begin our analysis by studying a setting with a single authority that has some amount of a homogeneous resource (*e.g.,* seats at a school, medical resources) that it can allocate to a continuum of agents.[2] Agents differ in their scores (*e.g.,* exam score, clinical need) and discrete attributes (*e.g.,* socioeconomic status, whether they are a frontline health worker). The authority cares about the distributions of scores (through some index such as the average score) and characteristics (such as gender and race) of the agents who are assigned the resource. Thus, the authority's preferences over agents depend on the joint distribution of agents' scores and groups. Motivated by the practical importance of uncertainty, we assume that this distribution is unknown. The authority's problem is to design a *first-best optimal* mechanism: a mechanism that is optimal regardless of their beliefs and implements an *ex post* optimal allocation in all states.

Adaptive Priority Mechanisms. We introduce adaptive priority mechanisms (APM), which proceed in two steps. First, each agent is given an *adaptive priority* that is a function of their own score and the number of agents from the same group to whom the resource is assigned or, equivalently, their rank within their group. Second, APM allocate the resource to agents in order of adaptive priorities, subject to fully allocating the available amount.

[2]In Appendix H, we generalize our analysis and results to a setting with discrete agents.

This class of mechanisms allows the implicit preference for agents from different groups to depend upon the ultimate allocation. When an agent's adaptive priority is increasing in their own score and decreasing in the number of agents with the same attributes that are assigned the resource—a property we call *monotonicity*—the APM implements a unique allocation.

We show that a particular monotone APM is first-best optimal. Under this optimal APM, an agent's priority is equal to their own score plus their marginal contribution to diversity utility. Intuitively, this mechanism equates the benefits and costs of allocating to the marginal agent, regardless of the ultimate joint distribution of agents' scores and groups. Moreover, this APM can be described and communicated to stakeholders *ex ante* as a function of the authority's preferences, without any reference to its beliefs or hypothetical states of the world. This has substantial practical benefits over *ad hoc* and *ex post* adjustments of priority and quota policies, which would be illegal in many countries, including the UK.

(Sub)Optimality of Priorities and Quotas. We next establish that the widely used priority and quota mechanisms are generally dominated by APM. More precisely, we find that priorities and quotas are first-best optimal if and only if (i) the authority is risk-neutral over diversity, in which case priorities are optimal, or (ii) the authority is extremely risk-averse over diversity, in which case quotas are optimal. The role of risk-aversion is best seen as arising from a trade-off between the relative strengths of priority and quota mechanisms. On the one hand, by mandating a minimal level of representation from underrepresented groups, quotas *guarantee* a level of diversity. On the other hand, as relatively more underrepresented agents receive the resource in the states in which they have relatively higher scores, priorities *positively select* these agents. Adaptive priority mechanisms optimally exploit the guarantee effects of quotas and the positive selection effects of priorities, and are always optimal. In a price-theoretic sense, this is analogous to ideas that govern the trade-off between price and quantity regulation (Weitzman, 1974) and the fact that a demand function would outperform both alternatives (for a literature review of demand functions, see Rostek and Yoon, 2023).

Multiple Authorities and Stable Allocations. While the single-authority model is relevant for studying settings with a single resource, in many markets there are multiple authorities who control heterogeneous resources over which agents have heterogeneous preferences. We generalize our analysis to this setting and show that there is a unique stable allocation. We characterize this unique stable allocation and show that a mechanism is consistent with stability if and only if it coincides with the single-authority-optimal APM. Furthermore, if all authorities use the optimal APM to determine the set of admitted agents, then the widely used Deferred Acceptance (DA) algorithm implements the unique stable matching.

APM as a Practical Market Design Solution. We conclude the paper by providing a practical, three-step roadmap for implementing APM as a market design solution. First, understand the preferences of stakeholders over allocations using standard stated- or revealed-preference methodologies. A key benefit of the optimal APM is that it only requires elicitation of preferences and not beliefs, which the optimal design of priority and quota mechanisms would require. Second, estimate the optimal APM and communicate this to stakeholders. We show that this can be achieved via a simple tabular format, in which rows correspond to groups, columns correspond to an agent's score rank within their group, and the entries correspond to score boosts. Third, implement the optimal APM by running the Deferred Acceptance algorithm. This implements a stable allocation and makes it dominant for agents to report their preferences truthfully.

Finally, we perform a proof-of-concept exercise to benchmark the benefits from APM. Concretely, we apply our practical roadmap using application and admission data from 2013-2017 on the selective exam schools of Chicago Public Schools (CPS), a setting also empirically studied by Angrist, Pathak, and Zárate (2019) and Ellison and Pathak (2021). CPS uses a reserve system to increase the admissions of underrepresented groups. Moreover, we document that there is substantial variation in the joint distribution of students' scores and characteristics over time. This justifies our focus on the importance of uncertainty and implies that APM *must* generate gains for the authority. Estimating preference parameters to best rationalize the pursued reserve policy, we find that the gains from using the optimal APM are equivalent to eliminating 37.5% of the loss to CPS' payoffs from failing to admit a diverse class of students. This gain is 2.3 times larger than the estimated gain from an actual 2012 policy reform that increased the size of all reserves. This exercise shows both that APM could be practically implemented and that the gains from so doing may be considerable.

Related Literature. In the market design literature on matching with affirmative action concerns following Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu (2005), many papers have studied the axiomatic properties of various affirmative action policies (see *e.g.,* Kojima, 2012; Hafalir, Yenmez, and Yildirim, 2013; Ehlers, Hafalir, Yenmez, and Yildirim, 2014; Echenique and Yenmez, 2015; Doğan, 2016; Kominers and Sönmez, 2016; Goto, Kojima, Kurata, Tamura, and Yokoo, 2017; Kamada and Kojima, 2017, 2018; Erdil and Kumano, 2019; Imamura, 2020). These include specific studies of alternative quota-like mechanisms that are used in many countries, including India (Aygün and Turhan, 2020; Sönmez and Yenmez, 2022a,b), Germany (Westkamp, 2013) and Brazil (Aygün and Bó, 2021).

In this paper, we instead pursue the methodological approach of mechanism design and welfare economics by analyzing optimal mechanisms from the perspective of an authority with some given preferences over allocations. Chan and Eyster (2003) share this perspective in their analysis of the costs and benefits of banning affirmative action.[3] In this vein, Çelebi and Flynn (2022) analyzes the narrower problem of how to optimally coarsen agents' scores into priorities. This analysis nevertheless restricted authorities to using a priority mechanism that does not consider agents' characteristics and implementing only allocations that are stable with respect to these priorities. Thus, our focus on comparing priorities, quotas, and optimal mechanisms distinguishes our analysis from the previous literature, which study the properties of each policy in isolation.[4]

Another key distinguishing feature of our modeling approach is the explicit consideration that market designers face uncertainty about the joint distribution of agents' scores, characteristics, and preferences. While this has not been studied in the existing literature, a plethora of evidence demonstrates the sizable practical challenges that such uncertainty poses (see Pathak and Shi, 2021; Doğan and Erdil, 2022; Bansak, Lee, Manshadi, Niazadeh, and Paulson, 2024, and the preceding discussion).

Overall, by taking a mechanism design approach and explicitly modeling uncertainty, we arrive at a new and practical market design solution for two-sided matching markets that dominates the existing alternatives, including priority and quota mechanisms.

Finally, by showing that Deferred Acceptance with adaptive priority mechanisms implements the unique stable allocation in the presence of endogenous preferences, we extend results from Azevedo and Leshno (2016) and Abdulkadiroğlu, Che, and Yasuda (2015) that show that Deferred Acceptance implements the unique stable allocation with fixed preferences. As a byproduct, our paper contributes to a literature that studies the uniqueness of stable allocations in large markets, complementing the result of Che, Kim, and Kojima (2019) that there is a unique stable allocation when the distribution of agents has full support and authorities' choice functions are generated by submodular quotas. Our uniqueness result concerns another important class of economies in which authorities have separable preferences over scores and diversity of a form that is commonly assumed (see *e.g.,* Chan and Eyster, 2003; Ellison and Pathak, 2021; Dessein, Frankel, and Kartik, 2023).

Outline. Section 2 exemplifies our main results. Section 3 studies optimal mechanisms with a single authority. Section 4 extends the model to include many authorities. Section 5 presents a general roadmap for implementing APM in practice and, as a proof-of-concept, applies this to Chicago Public Schools. Section 6 concludes. The proofs are in Appendix B.

[3]Other analyses of this issue include Temnyalov (2023), who characterizes when efficient assignment requires differential treatment across groups; and Shi (2022), who studies the optimal design of priority rules within standard allocation mechanisms under a planner objective based on agents' utilities.

[4]In a similar spirit, Combe, Dur, Tercieux, Terrier, and Ünver (2022) propose new mechanisms to overcome problems associated with unequal distribution of experienced teachers in schools and quantify the improvements compared to benchmarks.

2 Comparing Mechanisms: An Example

The Setting. A single school has capacity $q < 1$. Students are of unit total measure, have scores in $[0, 1]$, and are either minority or majority students. The authority has linear-quadratic preferences $\xi : \mathbb{R}^2 \to \mathbb{R}$ over students' total scores $\bar{s}$ and the measure of admitted minority students x:

$$\xi(\bar{s}, x) = \bar{s} + \gamma \left(x - \frac{\beta}{2} x^2 \right) \tag{1}$$

where $\gamma \geq 0$ indexes their general preference for admitting minority students and $\beta \geq 0$ indexes the degree of risk aversion regarding the measure of admitted minority students.

The minority students are of measure κ and have scores that are uniform over $[0, 1]$. The majority students are of measure $1 - \kappa$ and all have common underlying score $\omega \in [\underline{\omega}, \overline{\omega}] \subseteq [0, 1]$ with distribution Λ. The score of the majority students, ω, parameterizes how well they score relative to the minority students. Finally, we assume that the affirmative action preference is neither too small nor too large with the following: $\min\{\kappa, q\} > \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega} - \underline{\omega})$, $\kappa(1 - \underline{\omega}) < \frac{1+\gamma-\overline{\omega}}{\frac{1}{\kappa}+\gamma\beta}$. These conditions ensure that optimal affirmative action policies affect minority admissions in all states while not allocating the entire resource to minority students.

The authority can implement an APM, a priority mechanism, or a quota mechanism. An APM increases the scores of minority students by $A(y)$ when y other minority students are admitted, does not change the scores of majority students, and allocates seats to the students with highest transformed scores.[5] An (additive) priority mechanism $\alpha \in \mathbb{R}_+$ increases uniformly the scores of minority students by α. The authority then admits the highest-scoring measure q students. A quota policy $Q \in [0, \min\{\kappa, q\}]$ sets aside measure Q of the capacity for the minority students. The highest-scoring minority students of measure Q are first allocated to quota slots, and all other agents are then admitted to the residual $q - Q$ places according to the underlying score.[6]

We illustrate how these three policies prioritize minority students in Figure 1. By definition, priority mechanisms award a constant score boost of α. Quota mechanisms give enough points to always ensure admission until measure Q is reached and then give no advantage. APM allow any pattern of prioritization as a function of minority admissions (we plot only the optimal APM, which turns out to be linear in this context).

[5] Formally, this mechanism allocates seats to $x(\omega)$ minority students and $q - x(\omega)$ majority students, where $s(x(\omega)) + A(x(\omega)) = \omega$, and $s(x(\omega))$ denotes the score of the marginal minority student when the highest-scoring $x(\omega)$ minority students are admitted.

[6] This corresponds to a precedence order that processes quota slots first. In this example, one can show that processing quota slots last is equivalent to setting a priority.

Figure 1: How Priorities, Quotas, and APM Prioritize Minority Students

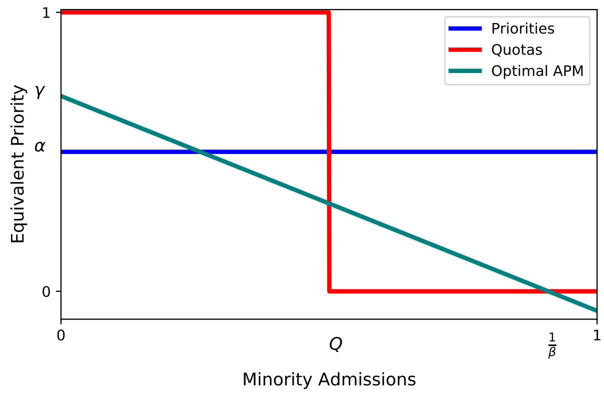

Note: Illustration of the equivalent priority given to a minority student as a function of the measure of admitted minority students under: the optimal APM (see Proposition 1), a priority mechanism α, and a quota mechanism Q.

Comparing Mechanisms. Let the authority's expected utility be V^* under any optimal (expected utility maximizing) mechanism, V_A under an optimal adaptive priority mechanism, V_P under an optimal priority mechanism, and V_Q under an optimal quota mechanism. The following proposition characterizes the relationships between these mechanisms:

Proposition 1. *The following statements are true:*

1. *The APM $A(y) = \gamma(1 - \beta y)$ is optimal, $V^* = V_A$*
2. *The comparative advantage of priorities over quotas is given by:*
$$\Delta \equiv V_P - V_Q = \frac{\kappa}{2}\left(1 - \kappa\gamma\beta\right)\mathrm{Var}[\omega] \tag{2}$$
3. *The comparative advantage of APM over priorities and quotas is given by:*
$$\Delta^* \equiv \min\{V^* - V_P, V^* - V_Q\} = \begin{cases} \frac{1}{2}\left(\kappa\gamma\beta\right)^2 \frac{\kappa \mathrm{Var}[\omega]}{1+\kappa\gamma\beta}, & \kappa\gamma\beta \leq 1, \\ \frac{1}{2}\frac{\kappa \mathrm{Var}[\omega]}{1+\kappa\gamma\beta}, & \kappa\gamma\beta > 1. \end{cases} \tag{3}$$

We now develop intuition for the comparative advantage of priorities over quotas. First, observe that a quota of Q admits measure Q minority students in all states of the world. However, a priority policy induces variability in the measure of admitted minority students. We call the gain to quota policies in eliminating this variation the *guarantee effect* and find mathematically that it is equal to $\frac{\kappa}{2}\left(1 + \kappa\gamma\beta\right)\mathrm{Var}[\omega]$ in payoff terms.

Second, the optimal priority policy admits more minority students when minority students score relatively well and fewer when minority students score relatively poorly. To demonstrate this, we show that minority admissions in state ω under the optimal priority policy are $x(\alpha, \omega) = \bar{x}(\alpha) + \varepsilon(\omega)$ where $\bar{x}(\alpha) = \kappa(1 + \alpha - \mathbb{E}[\omega])$ and $\varepsilon(\omega) = \kappa\left(\mathbb{E}[\omega] - \omega\right)$. Thus, in the states where minority students score relatively better ($\omega < \mathbb{E}[\omega]$), we have that

Figure 2: Comparative Statics for the Positive Selection and Guarantee Effects

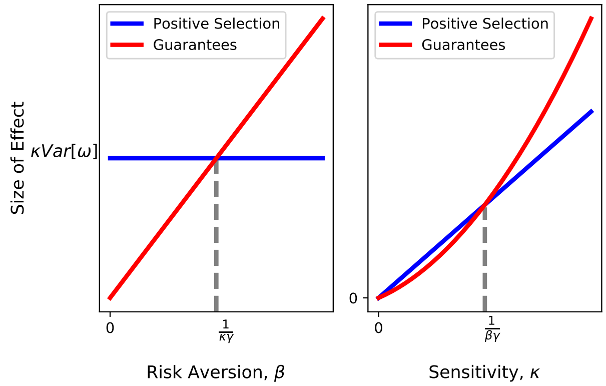

Note: Illustration of the comparative statics for the trade-offs between priority and quota mechanisms. Positive Selection plots the positive selection effect, $\kappa\text{Var}[\omega]$, and Guarantee plots the guarantee effect, $\frac{\kappa}{2}(1+\kappa\gamma\beta)\,\text{Var}[\omega]$. As per Equation 2 in Proposition 1, priorities dominate quotas if and only if $1 \geq \kappa\gamma\beta$, where the point of indifference is denoted by the dashed grey line.

$\varepsilon(\omega) > 0$ and $x(\alpha, \omega) > \bar{x}(\alpha)$. We call this effect the *positive selection* effect and find that this benefits a priority policy by $-\text{Cov}[\omega, \varepsilon(\omega)] = \kappa\text{Var}[\omega]$ in payoff terms.

The preference between priority and quota mechanisms is determined by which of the guarantee and positive selection effects dominates. We illustrate how risk aversion and the measure of minority students affect these opposing effects in Figure 2. If the authority is close enough to risk-neutral (*i.e.*, $\frac{1}{\kappa\gamma} > \beta$), then priorities are strictly preferred as positive selection dominates guarantees. If the authority is sufficiently risk-averse (*i.e.*, $\frac{1}{\kappa\gamma} < \beta$), then quotas are strictly preferred as the guarantee effects dominate positive selection. The threshold for risk aversion scales inversely with the measure of minority students κ. This is because κ determines the *sensitivity* of minority admissions to a small change in the priority of minority students. As a result, higher κ favors quota policies by increasing the magnitude of the guarantee effect relative to the positive selection effect. Finally, the extent of uncertainty $\text{Var}[\omega]$ may intensify an underlying preference but never determines which regime is preferred.

An APM is optimal and overcomes the limitations posed by both priorities and quotas. In this case, the optimal APM is linear in the measure of admitted minority students, with slope given by the authority's risk aversion over minority admissions, awarding each minority student a subsidy equivalent to their marginal contribution to the diversity preferences of the authority. This allows the adaptive priorities to optimally balance the positive selection and guarantee effects, and implement the first-best allocation in every state. In Figure 3, we illustrate how the losses from priority mechanisms and quota mechanisms vary with risk aversion and sensitivity. As risk aversion moves, the loss from priority and quota policies relative to the optimum is greatest when the authority is indifferent between the two regimes.

Figure 3: Comparative Statics for the Losses from Priorities and Quotas

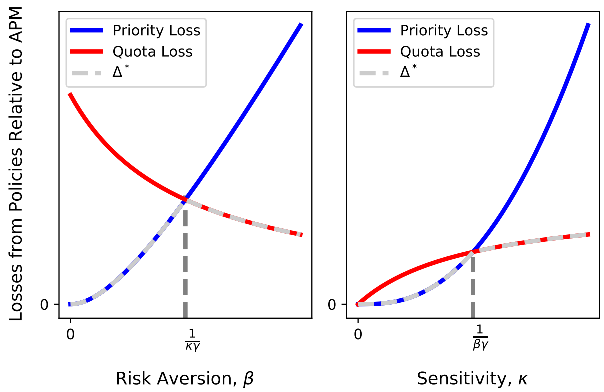

Note: Illustration of the comparative statics for the losses from optimal priority and quota policies relative to the optimal APM (as presented in Equation 3 in Proposition 1). The lower envelope of the losses, Δ^*, corresponds to the comparative advantage of the optimal APM over priorities and quotas. The point of indifference between priorities and quotas is denoted by the dashed grey line.

The loss from restricting to priority or quota policies is zero when the authority is risk-neutral or there is no uncertainty regarding relative scores, and decreases as the authority becomes extremely risk-averse. As sensitivity increases, the scope for affirmative action increases and so the gains from APM also increase. Thus, we should expect there to be large gains from switching to APM precisely when there is substantial uncertainty, authorities have intermediate levels of risk aversion, or the scope for implementing affirmative action is significant. A key advantage of the optimal APM is that it depends only on the school's preferences and not its beliefs. This contrasts with the belief-dependence of the optimal priority and quota policies.[7]

A Price-Theoretic Intuition. This comparison of mechanisms echoes that of *prices vs. quantitities* by Weitzman (1974). In fact, there is a formal mapping between the two.[8] Intuitively, the positive selection effect is equivalent to the effect that price regulation gives rise to the greatest production in states where the firm's marginal cost is lowest. Moreover, the guarantee effect is equivalent to the ability of quantity regulation to stabilize the level of production. An APM corresponds in the Weitzman (1974) setting to a regulator setting neither a price nor a quantity, but completely specifying the optimal demand curve. Thus, the comparison of mechanisms for allocating goods without prices boils down to similar trade-offs between well-understood price-based mechanisms for goods allocation.

[7]The optimal quota policy is given by $Q^* = (1 + \gamma - \mathbb{E}[\omega])/(\frac{1}{\kappa} + \gamma\beta)$, while the optimal priority policy sets the expected measure of minorities to Q^*. In this simple setting, the policies depend on Λ through $\mathbb{E}[\omega]$.

[8]This mapping holds under the following correspondence between Weitzman's (1974) parameters and ours: the inverse cost curvature $C''^{-1} \mapsto \kappa$, the benefit curvature $B'' \mapsto -\gamma\beta$, and the variance of marginal costs $\mathrm{Var}[\omega] \mapsto \mathrm{Var}[\alpha(\theta)]$. Under this mapping, our comparative advantage expressions coincide exactly.

3 Optimal Mechanisms with a Single Authority

We begin our general analysis by studying the resource allocation problem of a single authority. In this context, we define APM and derive an optimal APM that attains the first-best. We moreover provide necessary and sufficient conditions for the optimality of the ubiquitous priority and quota mechanisms and find that they are extremely restrictive, implying that there are likely gains from switching to APM.

3.1 Model

An authority allocates a single resource of measure $q \in (0,1)$. Agents differ in their type $\theta \in \Theta = [0,1] \times \mathcal{M}$ comprising their scores $s \in [0,1]$ and the group to which they belong, $m \in \mathcal{M}$, where their score denotes their suitability for the resource and $\mathcal{M}$ is a finite set comprising potential attributes such as race, gender, or socioeconomic status. We denote the score and group of any type θ by $s(\theta)$ and $m(\theta)$, respectively.

As we have motivated, in many market design contexts, accurately forecasting the distribution of students' types is hard (see Pathak and Shi, 2021; Doğan and Erdil, 2022; Bansak, Lee, Manshadi, Niazadeh, and Paulson, 2024, and the discussion in the introduction). To model this important feature of these settings, we introduce uncertainty about the distribution of types. The authority's uncertainty is parameterized by $\omega \in \Omega$, where Ω is the set of all distributions over Θ that admit a density. The authority believes that ω has distribution $\Lambda \in \Delta(\Omega)$. We denote the measure of types by F_ω with density f_ω in state $\omega \in \Omega$.[9]

An allocation $\mu : \Theta \to \{0,1\}$ specifies for any type $\theta \in \Theta$ whether they are assigned to the resource.[10] Two allocations μ and μ' are *essentially the same* if they coincide up to a measure zero set. The set of possible allocations is $\mathcal{U}$. An allocation is feasible if it allocates no more than measure q of the resource. A mechanism is a function $\phi : \Omega \to \mathcal{U}$ that returns a feasible allocation for any possible measure of types.

As we have also motivated, authorities often have preferences over scores and diversity. To model this, we define the aggregate score index of any allocation as:

$$\bar{s}_h(\mu, \omega) = \int_\Theta \mu(s,m) h(s) \mathrm{d}F_\omega(s,m) \tag{4}$$

for some continuous, strictly increasing function $h : [0,1] \to \mathbb{R}_+$, which determines the extent to which the authority values agents with higher scores. To capture diversity, we compute

[9] Formally, we mean that $f_\omega(s,m) = \frac{\partial}{\partial s} F_\omega(s,m)$ exists for all $s \in [0,1]$ and $m \in \mathcal{M}$.

[10] Formally, μ is a measurable function with respect to the Borel $\sigma-$algebra of the product topology in Θ.

the measure of agents of each group allocated the resource $x(\mu, \omega) = \{x_m(\mu, \omega)\}_{m \in \mathcal{M}}$ as:

$$x_m(\mu, \omega) = \int_{[0,1]} \mu(s, m) f_\omega(s, m) \mathrm{d}s \tag{5}$$

To separate the roles of scores and diversity, in our main analysis, we impose that their utility function over these dimensions is given by $\xi : \mathbb{R}^{|\mathcal{M}|+1} \to \mathbb{R}$, where:

$$\xi\left(\bar{s}_h, x\right) \equiv g\left(\bar{s}_h + \sum_{m \in \mathcal{M}} u_m(x_m)\right) \tag{6}$$

for some continuous, strictly increasing function $g : \mathbb{R} \to \mathbb{R}$ and continuously differentiable and concave functions $u_m : \mathbb{R} \to \mathbb{R}$ for all $m \in \mathcal{M}$. We also assume that the authority always prefers to allocate the entire resource.[11]

The preference of the authority is a monotone transformation of a quasi-linear utility index comprised of scores and a diversity preference. Intuitively, u_m determines the preference for assigned agents of group m, with its concavity following from a preference for diversity.[12] The function g contributes to risk preferences as well as the complementarity/substitutability of scores and diversity (if g is convex (concave) at a point, then scores and diversity are complements (substitutes) at that point). In Appendix D, we relax the separability between scores and diversity, the additive separability of u_m across groups, and concavity, and show that the optimality of APM continues to hold; Section 3.5 summarizes these results.

We define the value of a mechanism ϕ under distribution Λ as the authority's expected utility of the allocations induced by that mechanism:

$$\Xi(\phi, \Lambda) = \int_\Omega \xi(\bar{s}_h(\phi(\omega), \omega), x(\phi(\omega), \omega)) \mathrm{d}\Lambda(\omega) \tag{7}$$

We say that a mechanism is first-best optimal if it maximizes the authority's expected utility for all possible *distributions of* measures of agents' characteristics.

Definition 1 (First-Best Optimality). *A mechanism ϕ^* is first-best optimal if:*

$$\Xi(\phi^*, \Lambda) = \sup_\phi \Xi(\phi, \Lambda) \tag{8}$$

for all $\Lambda \in \Delta(\Omega)$.

This is a demanding property for a mechanism to possess as it requires a mechanism to implement an *ex post* optimal allocation in all states of the world.

[11] A necessary and sufficient condition for this is: $h(0) + u'_m(q) \geq 0$ for all $m \in \mathcal{M}$.

[12] Our specification allows the designer to have different preferences for allocating the resource to agents from different groups. For example, this allows for a designer with affirmative action motives who prefers to assign the resource to some particular group m: $u'_m(x_m) > u'_{m'}(x_m)$ for all x or a designer who prefers a balanced composition of allocated agents: $u'_m(x_m) = u'_{m'}(x_m)$ for all $m \in \mathcal{M}$ and x.

3.2 Adaptive Priority Mechanisms

Toward deriving a first-best optimal mechanism, we introduce APM. To this end, we first introduce an *adaptive priority policy* $A = \{A_m\}_{m\in\mathcal{M}}$, where $A_m : \mathbb{R} \times [0,1] \to \mathbb{R}$. The adaptive priority policy assigns priority $A_m(y_m, s)$ to an agent with score s in group m when measure y_m of agents of the same group is allocated the object. Given an adaptive priority policy, an APM implements allocations in the following way:

Definition 2 (Adaptive Priority Mechanism)**.** *An adaptive priority mechanism, induced by an adaptive priority A, implements an allocation μ in state ω if the following are satisfied:*

1. *Allocations are in order of priorities: $\mu(\theta) = 1$ if and only if for all θ' with $\mu(\theta') = 0$,*

$$A_{m(\theta)}(x_{m(\theta)}(\mu,\omega), s(\theta)) > A_{m(\theta')}(x_{m(\theta')}(\mu,\omega), s(\theta')) \tag{9}$$

2. *The resource is fully allocated:*

$$\sum_{m\in\mathcal{M}} x_m(\mu,\omega) = q \tag{10}$$

With some abuse of terminology, we will often refer to an APM as the adaptive priority A that induces it. By way of illustration, we provide a simple example of the flexibility of APM to act like a hybrid of priority and quota policies.

Example 1. Let $\mathcal{M} = \{m, n\}$ and the capacity be $q = 0.5$. We consider the adaptive priority policy $A = \{A_m, A_n\}$ given by:

$$A_m(x,s) = s, \qquad A_n(x,s) = \begin{cases} s+1 & \text{if } x \leq 0.1 \\ s+0.1 & \text{if } x \in (0.1, 0.25) \\ s & \text{if } x \geq 0.25 \end{cases} \tag{11}$$

This leaves the score of group m agents unchanged and gives agents of group n a score boost of: 1 if less than measure 0.1 group n agents is assigned, 0.1 if between measure 0.1 and 0.25 group n agents is assigned, and no score boost at all if measure greater than 0.25 group n agents is assigned.

To understand the properties of this adaptive priority policy, observe that the highest-scoring measure 0.1 group n agents are guaranteed the resource, even in states where they score poorly. Therefore, A_n practically embeds a quota of 0.1. For admissions levels between 0.1 and 0.25, the APM acts like a priority policy and boosts the scores of group n agents by 0.1, increasing the admissions of group n when group n agents score moderately well. For admissions levels beyond 0.25, group n agents are given no further advantage. Thus, when diversity is attained, this APM "phases out" and no longer advantages any group. $\triangle$

At this point, we have not established that a given APM implements any allocation at all, or that it implements a unique allocation (*i.e.,* it may not even be a mechanism). However, there is a natural subclass of APM that do implement a unique allocation: those that are monotone. An APM A is *monotone* when (i) $A_m(\cdot, s)$ is a decreasing function for all $m \in \mathcal{M}, s \in [0,1]$ and (ii) $A_m(y_m, \cdot)$ is a strictly increasing function for all $m \in \mathcal{M}, y_m \in \mathbb{R}$.[13]

Proposition 2. *Any Monotone APM A implements an essentially unique allocation.*

Moreover, the unique outcome of a monotone APM can be implemented by a simple algorithm:[14]

Algorithm 1 (Algorithm for Implementation of APM)**.** *The APM algorithm proceeds in the following four steps:*

1. *For each θ, define*
$$\overline{x}(\theta) = \int_{s(\theta)}^{1} f_\omega(s, m(\theta))\mathrm{d}s \tag{12}$$
as the measure of agents who have higher scores than θ and belong to the same group.
2. *Construct a ranking of the agents as*
$$R(\theta) = A_{m(\theta)}(\overline{x}(\theta), s(\theta)) \tag{13}$$
3. *Define the cutoff ranking for the agents as $\overline{R}_\omega$ by*
$$\int_\Theta \mathbb{I}\{R(\theta) \geq \overline{R}_\omega\}\mathrm{d}F_\omega(\theta) = q \tag{14}$$
4. *Allocate the resource to all θ with $R(\theta) \geq \overline{R}_\omega$.*

Intuitively, this algorithm works by ranking all agents by their score within each group m and assigning agents in order of their transformed scores evaluated at the measure of *higher-score* agents of the same group. Informally, the algorithm moves down the ranking of agents until the resource is exhausted. This algorithm allows APM to be equivalently communicated as making an agent's priority depend on either (i) the measure of admitted agents from each group, or (ii) each agent's rank within their group.

[13]Observe that monotone adaptive priority mechanisms are fair in the sense that they preserve the ranking of agents within any group and assign higher priority to an agent whenever there are fewer agents from their group who are allocated the resource.

[14]Formally, when we consider an Adaptive Priority *Mechanism*, we are studying any selection from the set of allocations that the APM implements. As monotone APMs implement an essentially unique allocation, this is without loss of optimality. When we refer to the "unique" allocation, we refer to the cutoff allocation defined in the proof of Proposition 2 and which is implemented by Algorithm 1.

3.3 Adaptive Priority Mechanisms Achieve the First-Best

Having shown that monotone APM implement a unique allocation and provided an algorithm to compute this allocation, we now show that a certain, monotone APM is first-best optimal:

Theorem 1. *The following APM is monotone and first-best-optimal:*

$$A^*_m(y_m, s) \equiv h^{-1}(h(s) + u'_m(y_m)) \tag{15}$$

Moreover, if a mechanism is first-best-optimal, then it implements essentially the same allocations as A^*.

To gain intuition for the form of this mechanism, suppose that the authority has linear utility over scores $h(s) \equiv s$. In this case, $A^*_m(y_m, s) = s + u'_m(y_m)$, so an agent in group m is awarded a boost of $u'_m(y_m)$ when there are y_m higher-scoring agents of the same group, their direct marginal contribution to the diversity preferences of the authority. This is optimal because this boost precisely trades off the marginal benefit of additional representation with the marginal costs of reduced scores. Moreover, failing to award this precise level of boost would result in a suboptimal allocation. Thus, any optimal mechanism must be essentially identical to the optimal APM we have characterized. To generalize this beyond linear utility of scores, consider the following observation: we can map agents' scores from s to $h(s)$, and consider the optimal boost in this space. As h is monotone, this preserves the ordinal structure of the optimal allocation, and the authority has linear preferences over $h(s)$. Thus, in this transformed space, the optimal boost remains additive and given by $u'_m(y_m)$. To find the optimal transformed score in the original space, we simply invert the transformation h and apply it to the optimal score in the transformed space, yielding the formula for the optimal mechanism in Theorem 1.

The logic underlying this result is simple and echoes standard consumer preference theory: just as it is best to equate the marginal rates of substitution between different goods for a consumer, an authority should set allocations to equate the marginal rates of substitution between merit and diversity across all groups. Theorem 1 shows that this can be achieved through a simple generalization of priority mechanisms by allowing the priority boost to depend on the representation of each group.

In addition to its optimality, the APM characterized in Theorem 1 has several practical benefits over alternative mechanisms (which we discuss at length in Section 5) and is important as an intermediate step in our theoretical results that characterize stable allocations with multiple authorities in Section 4.

3.4 (Sub)Optimality of Priorities and Quotas

We have shown that APM are optimal. However, the primary classes of mechanisms that have been used in practice are priority and quota mechanisms. Therefore, it is important to understand whether (and when) these mechanisms are also optimal. We now establish that APM generally provide a strict improvement over priority and quota mechanisms and characterize when priority and quota mechanisms attain optimality.

We first define priority and quota mechanisms. A *priority policy* $P : \Theta \to [0, 1]$ awards an agent of type $(s, m) \in \Theta$ a priority $P(s, m)$, that depends on their score and group.

Definition 3 (Priority Mechanisms)**.** *A priority mechanism, induced by a priority policy P, allocates the resource in order of priorities until measure q has been allocated, with ties broken uniformly and at random.*

We define a *quota policy* as (Q, D), where $Q = \{Q_m\}_{m\in\mathcal{M}}$ and $D : \mathcal{M} \cup \{R\} \to \{1, 2, \ldots, |\mathcal{M}| + 1\}$ is a bijection. The vector Q reserves measure of the capacity Q_m for agents in group m, with residual capacity $Q_R = q - \sum_{m\in\mathcal{M}} Q_m$ open to agents of all types. The bijection D (the precedence order) gives the order in which the groups are processed.

Definition 4 (Quota Mechanisms)**.** *A quota mechanism, induced by a quota policy (Q, D), proceeds by allocating the measure $Q_{D^{-1}(k)}$ agents from group $D^{-1}(k)$ (if there are sufficient agents from this group) to the resource in ascending order of k, and in descending order of score within each k. If there are insufficiently many agents of any group to fill the quota, the residual capacity is allocated to a final round in which all agents are eligible.*

We now characterize when priority and quota mechanisms are (sub)optimal. To do this, we first provide some definitions. Authority preferences are *non-trivial* if for all $m, n \in \mathcal{M}$, we have that:

$$h(1) + u'_n(0) > h(0) + u'_m(q) \tag{16}$$

Intuitively, the authority's preferences are non-trivial when their concerns for representation of certain groups do not always outweigh the consideration of scores.[15] The authority is *risk-neutral* over diversity if for all $m \in \mathcal{M}$, $u'_m : [0, q] \to \mathbb{R}$ is constant, *i.e.,* there are constant marginal returns to admitting more agents from all groups. If there are decreasing marginal returns, then the authority's preferences feature risk aversion. We define extremely risk-averse preferences as follows. Let $\tilde{u}$ and $\tilde{h}$ be functions describing diversity and score

[15] Note that failure of non-triviality means there exists m and n such that $h(1) + u'_n(0) \leq +h(0) + u'_m(q)$, *i.e.,* a group n agent with the maximum score is less preferred than a group m agent with the minimum score even when all of the entire capacity is allocated to group m agents.

preferences, and let $\{x_m^{\text{tar}}\}_{m\in\mathcal{M}}$ be a vector of target allocation levels. Moreover, assume that these satisfy: (i) $\tilde{u}'_m(x_m) = 0$ for all $x_m > x_m^{\text{tar}}$ (ii) $\tilde{u}'_m(x_m) \geq \tilde{h}(1) - \tilde{h}(0)$ for $x_m \leq x_m^{\text{tar}}$ and (iii) $\sum_{m\in\mathcal{M}} x_m^{\text{tar}} \leq q$. Intuitively, an authority whose preferences are represented by $\tilde{u}$ and $\tilde{h}$ is very risk-averse as the condition that $\tilde{u}'_m(x_m) \geq \tilde{h}(1) - \tilde{h}(0)$ implies that the loss from being below the target level for a group x_m^{tar} dominates any benefit from increased scores, effectively making the authority infinitely risk-averse to missing the target. We say that the authority is *extremely risk-averse* if the authority's preferences over the optimal allocations can be represented by $(\tilde{u}, \tilde{h})$.[16]

The following result characterizes the first-best (sub)optimality of priorities and quotas:

Theorem 2. *Suppose that the authority has non-trivial preferences. The following statements are true:*

1. *There exists a first-best optimal priority mechanism if and only if the authority is risk-neutral. Moreover, this mechanism is given by* $P(s, m) = h^{-1}(h(s) + u'_m)$.
2. *There exists a first-best optimal quota mechanism if and only if the authority is extremely risk-averse. Moreover, this mechanism is given by* $Q_m = x_m^{tar}$ *and* $D(R) = |\mathcal{M}| + 1$.

That risk-neutrality and high risk aversion are sufficient for the optimality of priority and quota mechanisms is intuitive. With risk-neutral preferences, priorities can perfectly balance the score and diversity goals without regard for the state of the world. This is because, under risk-neutrality, there is a constant "exchange rate" between the two: how the authority compares any two agents does not depend on the final allocation and thus can be specified *ex ante* by a priority policy. If the authority is extremely risk-averse as to the prospect of failing to assign x_m^{tar} agents from group m, then a quota allows them to always achieve this target level of allocation in all states of the world while minimally sacrificing scores. It is less obvious that risk-neutrality and high risk aversion are necessary. We prove this result by constructing adversarial measures of agents for each preference profile that render any priority or quota mechanism suboptimal unless the authority is risk-neutral or extremely risk-averse, respectively. Importantly, this result also shows that the only optimal quota mechanisms are those that process open slots last.

This result highlights the fragility of priority mechanisms to uncertainty absent the strong assumption of risk-neutrality over diversity. Intuitively, this is because they feature no guarantees as to how many agents of different groups will be assigned. Indeed, the unfortunate interaction between priority mechanisms and unforeseen market realizations has led to public

[16] More formally, this means that there exists $(\tilde{u}, \tilde{h})$ such that the optimal allocation under (u, h) is also optimal under $(\tilde{u}, \tilde{h})$ for all $\omega \in \Omega$.

backlash against priority mechanisms. For example, in the Vietnamese university admissions system, which combines exam scores with priority boosts for students from disadvantaged groups, a year of unexpectedly easy exams led to "top-scoring students missing out on the opportunity to attend their university of choice" and generated backlash against the system (Tuoi Tre News, 2017).

Moreover, our result highlights that quota mechanisms similarly fail to achieve the first-best away from high levels of risk aversion as they do not take advantage of the potential for positive selection. Our quantitative analysis in Section 5 in the context of quota mechanisms in Chicago Public Schools suggests that the variation in the distribution of characteristics across years generates meaningful welfare gains from switching to APM.

To formalize the connection between uncertainty and the importance of the adaptability of APM, we consider a setting with *no uncertainty*, where Λ is a Dirac measure. In this context, we say that a mechanism is optimal without uncertainty if it is a utility maximizer.

Proposition 3. *If there is no uncertainty, then there exist optimal priority and quota mechanisms.*

This result shows that if an authority is certain about the market, then appropriately constructed priority and quota mechanisms would be optimal. This formalizes the idea that the suboptimality of priority and quota mechanisms stems from their inability to adapt to the state.

As we have described, uncertainty about the distribution of agents' characteristics is an important feature of even well-established markets. Aggregate shocks to students' examination performance (Doğan and Erdil, 2022), refugees' migratory trends (Bansak, Lee, Manshadi, Niazadeh, and Paulson, 2024), and a host of other factors (Pathak and Shi, 2021) generate substantial fluctuations in distributions from year-to-year across a wide range of contexts. Thus, absent the strong conditions on authority preferences that we have characterized in Theorem 2, APM strictly dominate priority and quota mechanisms.

3.5 Extensions: Beyond Separable Preferences, Noisy and Endogenous Scores, and Discrete Economies

Beyond Separable Preferences. Our main analysis made assumptions on the separability of (i) diversity preferences across groups and (ii) score and diversity preferences. These assumptions are standard in the literature on affirmative action concerns (see e.g., Athey, Avery, and Zemsky, 2000; Chan and Eyster, 2003; Ellison and Pathak, 2021) and can moreover be justified by the axiomatization developed by Çelebi (2023): if preferences satisfy appropriate adaptations of responsiveness (in the sense of Roth, 1985), substitutes (in the sense of Roth, 1984), and acyclicity (which is a strengthening of transitivity, proposed by Tversky, 1964), then the assumptions underlying Equation 6 are satisfied.

These points notwithstanding, we explore the robustness of our results to relaxing the separability and concavity embedded in Equation 6 in Appendix D. First, we show that a modified APM remains optimal when preferences are non-separable over diversity. That is, when $\sum_{m\in\mathcal{M}} u_m(x_m)$ is replaced with $u(x)$, the optimal APM now replaces $u'_m(x_m)$ with the partial derivative $u^{(m)}(x)$ (Proposition 5). Among other things, this allows our model to capture preferences in situations with overlapping group membership (Aygün and Bó, 2021; Sönmez and Yenmez, 2022a), *e.g,* when people have different genders and belong to different socioeconomic groups.[17] Second, we show that allowing for a non-separable preference between scores and diversity leads to the optimality of a modified APM, which now must be made to condition on the aggregate score $\bar{s}_h$ in addition to the vector of admissions levels across groups x (Proposition 6). Intuitively, the optimal APM in this setting replaces the marginal diversity utility with the marginal rate of substitution between diversity and scores. These extensions show that the core economic insight on the first-best optimality of APM extends to richer preference specifications. Third, we relax the concavity of u_m and show that, although the optimal allocation is no longer unique, any optimal allocation is implemented by the optimal APM (Proposition 7).

Noisy and Endogenous Scores. In Appendix E, we extend our baseline model to account for the fact that scores may either be a noisy measure underlying ability or scores may be endogenously generated by agents by undertaking costly effort. The noisiness and endogeneity of scores may raise concerns about the applicability of our results if authorities care about ability and not scores themselves. However, we show that APM continue to generate optimal outcomes in both cases. With noisy scores, intuitively, the APM must replace the observed score with the conditional expectation of ability given the score (Proposition 9). Thus, with the caveat that computing expected abilities given scores may be more complex than simply using scores, this shows that the economics of APM extend to noisy measures of ability. With endogenous scores, we show that if more able agents find it easier to achieve higher scores (a standard single-crossing condition), then we can leverage a theorem of Mailath (1987) to show that there exists a separating equilibrium in which the authority uses the optimal APM as in Theorem 1 and all agents separate (Proposition 8). Thus, with endogenous scores, the optimal APM continues to be first-best optimal.

[17] As a simple example, if people can be men m or women w and rich r or poor p (so the groups are $\mathcal{M} = \{wp, wr, mp, mr\}$) and the authority cares about increasing the representation of women and poor people, then the utility function could be given by $u(x) = \hat{u}(x_{wp} + x_{wr}, x_{mp} + x_{wp})$ and our analysis would apply so long as u is concave.

Discrete Economies. In Appendix H, we translate our analysis and results to the discrete context. Concretely, we establish the optimality of APM, characterize the (sub)-optimality of priorities and quotas, and—looking ahead to the multi-authority context—show that agent-proposing Deferred Acceptance, when combined with the optimal APM, implements the agent-optimal stable allocation.

4 Stable Mechanisms with Multiple Authorities

The single-authority model is relevant for many resource allocation contexts, such as the medical resource allocation problem of a hospital and the allocation of (homogeneous) government jobs to candidates. However, in other settings such as school or university admissions, multiple authorities must decide upon their admissions policies and rules. In this section, we generalize our single authority model to a setting with multiple authorities. We define *stability* in this setting and show that there is a unique stable allocation. Moreover, we show that a mechanism is consistent with stability if and only if it coincides with the *single-authority-optimal* APM from Theorem 1. Finally, we show that when authorities use the optimal APM, the widely used Deferred Acceptance algorithm implements the unique stable matching, which makes the use of APM practical in a multiple-authority setting.

4.1 The Multi-Authority Model

There are authorities $c \in \mathcal{C} = c_0 \cup \bar{\mathcal{C}} = \{c_0, c_1, \ldots, c_{|\mathcal{C}|-1}\}$ with capacities q_c, where c_0 is a dummy authority that corresponds to an agent going unmatched. The agents differ in their authority-specific scores, the group to which they belong, and their preferences over the authorities, $\succ$. We index agents by their type $\theta = (s, m, \succ) \in [0,1]^{|\mathcal{C}|} \times \mathcal{M} \times \mathcal{R} = \Theta$, where $\mathcal{R}$ is set of all complete, transitive, and strict preference relations over $\mathcal{C}$. For each type θ, $s_c(\theta)$ denotes the score of θ at authority c and $m(\theta)$ denotes the group of θ. From now, to economize on notation, we suppress indexing states by $\omega \in \Omega$ and let the measure of types be F, with density f.[18] We assume that f has full support over Θ (*i.e.,* $f > 0$) and that $F(\Theta)$ is less than the capacity of c_0 and greater than the capacity of $\bar{\mathcal{C}}$.

Each authority has preferences over the agents they are assigned of the form introduced in the previous section:

$$\xi_c(\bar{s}_{h_c}, x_c) = g_c\left(\bar{s}_{h_c} + \sum_{m \in \mathcal{M}} u_{m,c}(x_{m,c})\right) \tag{17}$$

[18]Formally, this density is given by $f(s, m, \succ) = \frac{\partial}{\partial s} F(s, m, \succ)$.

where the extent to which they care about risk g_c, scores h_c, and diversity $\{u_{m,c}\}_{m\in\mathcal{M}}$ are potentially specific to each authority.

A matching is a function $\mu : \mathcal{C} \cup \Theta \to 2^{\Theta} \cup \mathcal{C}$ where $\mu(\theta) \in \mathcal{C}$ is the authority that any type θ is assigned and $\mu(c) \subseteq \Theta$ is the set of agents that is assigned to authority c.[19] Given a matching μ, $\bar{s}_{h_c,c}(\mu)$ and $x_c(\mu) = \{x_{m,c}(\mu)\}_{m\in\mathcal{M}}$ denote the score indices and measures of agents from different groups matched to c at μ. We say that c *prefers* μ to μ', which we denote by $\mu \succ_c \mu'$, if $\xi_c(\bar{s}_{h_c,c}(\mu), x_c(\mu)) > \xi_c(\bar{s}_{h_c,c}(\mu'), x_c(\mu'))$. Toward representing a matching as a lower-dimensional object, we moreover define a cutoff matching as one in which agents are assigned to the authority that they most prefer among the set of authorities in which their score clears a group-specific threshold:

Definition 5. *A matching μ is a cutoff matching if there exist cutoffs $S = \{S_{m,c}\}_{m\in\mathcal{M},c\in\mathcal{C}}$ such that $\mu(\theta) = c$ if (i) $s_c(\theta) \geq S_{m(\theta),c}$ and (ii) for all c' with $c' \succ_\theta c$, $s_{c'}(\theta) < S_{m(\theta),c'}$.*

Given S, the *demand* of an agent θ is their favorite authority among those for which they clear the cutoff:

$$D^{\theta}(S) = \{c : s_c(\theta) \geq S_{m(\theta),c} \text{ and } c \succeq_\theta c' \text{ for all } c' \text{ with } s_{c'}(\theta) \geq S_{m(\theta),c'})\} \tag{18}$$

The *aggregate demand* for authority c is the set of agents who demand it $D_c(S) = \{\theta : D^{\theta}(S) = c\}$, while $\tilde{D}_c(S_{-c}) = D_c((0,\ldots,0), S_{-c})$ returns the set of all agents who would demand c if offered admission when other authorities' cutoffs are S_{-c}.

4.2 Characterization of Stable Allocations

We first characterize the set of stable allocations. Our context presents two challenges in this regard. First, the priorities that are typically used to define stability are not primitives of our model. Therefore, to define stability, we will use the preferences of the authorities induced by Equation 17. Second, unlike discrete models, a single agent does not affect the preferences of an authority. Therefore, we need to consider a positive mass of agents to define blocking.

For each matching μ, authority $c \neq c_0$, and two sets of agents $\tilde{\Theta}$ and $\hat{\Theta}$, we let $\hat{\mu}_{(\hat{\Theta},\tilde{\Theta},c,\mu)}$ denote the matching that maps $\hat{\Theta}$ to c and $\tilde{\Theta}$ to c_0 and otherwise coincides with μ.[20] A set of

[19] The mathematical definition of a matching for the continuum economy we study follows Azevedo and Leshno (2016) and requires that μ satisfies the following four properties: (i) for all $\theta \in \Theta$, $\mu(\theta) \in \mathcal{C}$; (ii) for all $c \in \mathcal{C}$, $\mu(c) \subseteq \Theta$ is measurable and $F(\mu(c)) \leq q_c$; (iii) $c = \mu(\theta)$ iff $\theta \in \mu(c)$; (iv) for any $c \in \mathcal{C}$, the set $\{\theta \in \Theta : c \succ_\theta \mu(\theta)\}$ is open.

[20] Formally,

$$\hat{\mu}_{(\hat{\Theta},\tilde{\Theta},c,\mu)}(\theta) = \begin{cases} c_0 & \text{if } \theta \in \tilde{\Theta} \\ c & \text{if } \theta \in \hat{\Theta} \\ \mu(\theta) & \text{otherwise} \end{cases} \tag{19}$$

agents $\hat{\Theta}$ *blocks* matching μ at authority c by $\tilde{\Theta}$ if (i) for all $\theta \in \hat{\Theta}$, $c \succ_\theta \mu(\theta)$, (ii) $\tilde{\Theta} \subseteq \mu(c)$, (iii) $F(\tilde{\Theta}) = F(\hat{\Theta})$, and (iv) $\hat{\mu}_{(\hat{\Theta},\tilde{\Theta},c,\mu)} \succ_c \mu$. A matching μ is *not blocked* if there does not exist such a $(\hat{\Theta}, \tilde{\Theta}, c)$. A matching μ satisfies *within-group fairness* if for all $\theta, \theta' \in \Theta$ such that $m(\theta') = m(\theta)$ and $s_{\mu(\theta)}(\theta') > s_{\mu(\theta)}(\theta)$, it holds that $\mu(\theta') \succeq_{\theta'} \mu(\theta)$.[21] A matching μ is *stable* if it satisfies within-group fairness, is not blocked, and all non-dummy authorities fill their capacity. The following result establishes that there exists a unique stable matching and that this is a cutoff matching.

Theorem 3. *There is a unique stable matching. This matching is a cutoff matching.*

This result complements existing results on the uniqueness of stable allocations in continuum economies under full support (Azevedo and Leshno (2016) for responsive preferences and Che, Kim, and Kojima (2019) for submodular quotas). Theorem 3 extends the scope of known economies in which there is a unique stable allocation by considering economies in which agents belong to different socioeconomic groups and the preferences of authorities depend non-linearly on their admissions of various groups. Importantly, our setting results in the authorities having preferences over various sets of agents that are endogenous to the composition of the admitted agents.

We now formally describe these complications and how we resolve them by using our characterization of the single authority optimal APM from Theorem 1. First, imagine that there is only one group of agents $|\mathcal{M}| = 1$, so that authorities' preferences are determined by the scores of the agents, as in Azevedo and Leshno (2016). Given a set of cutoffs S_{-c}, a cutoff t_c clears the market for c if $F(D_c(t_c, S_{-c})) = q_c$. When $|\mathcal{M}| = 1$, for a given S_{-c}, there is a unique t_c that clears the market since a smaller cutoff will exceed the capacity while a larger one will leave a positive measure of the capacity empty. Define $T = \{T_c\}_{c \in \mathcal{C}}$, where $T_c(S)$ is the function that maps each S to the market-clearing cutoff t_c under S_{-c}. The result then follows from (i) showing the fixed points of T correspond to market-clearing cutoffs of stable matchings, (ii) establishing that T is monotone, (iii) applying Tarski's fixed point theorem to show that the set of market-clearing cutoffs is a lattice, and (iv) showing that there can only be one market-clearing cutoff as, if there were two, one would strictly exceed the capacities of at least one authority.

Suppose now that there are multiple groups, $|\mathcal{M}| > 1$. A significant complication arises: there is a potential continuum of cutoffs that would clear the market for authority c. A selection from this set is provided by the cutoffs induced by the optimal APM characterized by Theorem 1, $A^*_{m,c}(y_m, s) \equiv h_c^{-1}(h_c(s) + u'_{m,c}(y_m))$. We show that the APM cutoffs are

[21] Within-group fairness requires an authority to not reject an agent if it is admitting an agent from the same group with a lower score. Under our assumption that authorities prefer higher scores (h_c is strictly increasing), within-group fairness is satisfied if there is no blocking in discrete models.

unique among the market-clearing cutoffs in being compatible with stability. This is because, for any other t'_c, there is a set $\hat{\Theta}$ of agents (with positive measure) who have scores lower than the cutoff for their group and a set $\tilde{\Theta}$ of agents (with positive measure) who have scores higher than the cutoff for their group, but the authority is strictly better off by admitting $\hat{\Theta}$ and rejecting $\tilde{\Theta}$. We define $T_c(S)$ as the market-clearing cutoffs induced by the optimal APM, show that the fixed points of T_c correspond to market-clearing cutoffs of stable matchings, and then use this to establish the existence and uniqueness of the stable allocation.

This hints at a connection between the stable allocation and the allocation induced by all authorities pursuing the optimal APM, which we now make explicit. The demand set of c at μ, $D_c(\mu)$, is the set of agents who prefer c to their allocation under μ. A mechanism is *consistent with stability* if for all F with stable matching μ_F, it chooses $\mu_F(c)$ from $D_c(\mu_F)$. In other words, evaluated at the set of agents who demand an authority, this mechanism chooses the set of agents with which the authority is already matched. Moreover, we say that a mechanism ϕ is *equivalent* to ϕ' if it chooses the same agents under all full support measures. We now establish that single-authority-optimal APM, as characterized by Theorem 1, comprise the full set of mechanisms that are consistent with stability (up to equivalent mechanisms).

Proposition 4. *A mechanism is consistent with stability if and only if it is equivalent to the single-authority-optimal APM, A_c^*.*

Thus, not only is the optimal APM A^* inherent to the structure of stable allocations, but it also characterizes stability in this setting in the sense that any deviation from A^* would result in a violation of stability.

4.3 APM and DA Implement the Unique Stable Matching

We have now shown there is a unique stable matching and that any mechanism that is consistent with stability is equivalent to the optimal APM. These results suggest that authorities using their optimal APM can implement the unique stable matching. We now show that this is indeed the case: the standard Deferred Acceptance (DA) algorithm implements the unique stable matching when the authorities' choice rules follow the optimal APM.

First, we define the DA algorithm, following Abdulkadiroğlu, Che, and Yasuda (2015), which we describe below under our notational conventions for completeness. Each authority $c \in \mathcal{C}$ submits a choice function $Ch_c : 2^\Theta \to 2^\Theta$ such that $Ch_c(\Theta') \subseteq \Theta'$ and $F(Ch_c(\Theta')) \leq q_c$. To define the Deferred Acceptance mapping DA, first, consider a mapping $Q : \Theta \to \mathcal{R}$, where $Q(\theta)$ is an ordered list of authorities. The DA mapping $Q' = DA(Q)$ is determined as follows. Informally, every agent θ applies to her most preferred authority in $Q(\theta)$. Every authority

c tentatively accepts its applicants according to Ch_c. If all seats of c are assigned, it rejects the remaining agents. If an agent θ is rejected by c, $Q'(\theta)$ is obtained from $Q(\theta)$ by deleting c in $Q(\theta)$. If an agent θ is not rejected, then $Q'(\theta) = Q(\theta)$.

More formally, let $\mathcal{T}_c(Q) = \{\theta \in \Theta : \text{c is ranked first in } Q(\theta)\}$ be the set of agents that rank c as first choice. Each authority admits agents in $Ch_c(\mathcal{T}_c(Q))$ and rejects agents in $\mathcal{T}_c(Q) \setminus Ch_c(\mathcal{T}_c(Q))$. If $\theta \in \mathcal{T}_c(Q) \setminus Ch_c(\mathcal{T}_c(Q))$ for some c, then $Q'(\theta)$ is obtained from $Q(\theta)$ by deleting c from the top of $Q(\theta)$; otherwise $Q'(\theta) = Q(\theta)$. Repeated application of the DA mapping gives us the DA algorithm. That is, given the set of authorities C, the distribution of agents F and capacities q, let $Q^0(\theta) = \succ_\theta$ for all θ and $Q^t = DA(Q^{t-1})$ for $t \in \mathbb{N}$. By adapting Theorem 0 of Abdulkadiroğlu, Che, and Yasuda (2015), we obtain that this DA algorithm converges for every agent:

Lemma 1. *$DA^t(Q^0)$ converges pointwise.*

We let $Q^* : \Theta \to \mathcal{R}$ denote the limit of $DA^t(Q_0)$ as $t \to \infty$. Given $Q^*(\theta)$, we define $\mu^{DA}(\theta)$ as the $\succ_\theta$-maximal element of $Q^*(\theta)$. We say that μ^{DA} is the matching that is implemented by DA. We now show that if Ch_c is generated by the optimal APM A_c^* for all c, then μ^{DA} is the unique stable matching of this economy.

Theorem 4. *If all authorities use the optimal APM A_c^*, then DA implements the unique stable matching.*

This result demonstrates that Deferred Acceptance, a widely used and standard matching algorithm, can readily be combined with APM to compute and implement the unique stable matching in practice in a multi-authority setting. Moreover, as is well known (see Theorem 5.16 in Roth and Sotomayor, 1990), this establishes that the combination of DA with APM is strategy-proof for agents: truthfully reporting preferences is a dominant strategy.[22]

4.4 Extensions: Efficient Mechanisms and Dominance Under Decentralized Sequential Admissions

Efficient Mechanisms. In our main analysis, as is standard in market design, we considered the problem of implementing stable allocations. It is also natural to ask what kind of mechanisms might implement allocations that are *efficient* for the authorities. That is, what is optimal when a single, centralized authority wishes to design a mechanism that trades off

[22]Of course, it is not a dominant strategy for authorities to report truthfully (*i.e.*, use the optimal APM) under Deferred Acceptance. This has nothing to do with their using APM or restricting attention to DA: it is not possible to find a stable matching mechanism that makes truth-telling dominant for authorities (see Theorem 5.14 in Roth and Sotomayor, 1990).

match quality and diversity taking into account the effects of their mechanism on each authority? In Appendix G, we study authority-efficient mechanisms. Considering utilitarian efficiency over authorities as our welfare criterion, we show that the unique stable allocation is generally inefficient (Proposition 11). We propose an augmentation of an APM to solve this problem, an adaptive priority mechanism with quotas (APM-Q). When authorities evaluate the scores (but not necessarily diversity) in the same way, we show that an APM-Q implements efficient allocations (Theorem 6). The idea behind this hybrid mechanism is to construct a fictitious aggregate authority in our single object setting—one whose preferences aggregate individual authorities' score and diversity concerns—and then to use aggregate, market-level priorities with authority-specific quotas to implement efficient allocations.

Decentralized Admissions. In our main analysis of multiple authorities, we have considered a centralized allocation mechanism (Deferred Acceptance) and shown that this implements the unique stable matching. It is also interesting to consider what might happen as an equilibrium outcome between authorities in a decentralized setting. In Appendix F, we study such a decentralized setting in which the agents apply sequentially to the authorities, who then decide which agents to admit. We show that a mechanism implements a dominant strategy for an authority c if and only if it is the optimal APM A_c^* characterized by Theorem 1 (Theorem 5). Moreover, in any equilibrium in which authorities use the optimal APM, the equilibrium allocation is the unique stable matching of the economy (Proposition 10). Moreover, this is true whatever the order in which the authorities admit agents. These results imply that one could advise authorities to use APMs with confidence that outcomes will be stable and that they could do no better under any alternative mechanism in decentralized allocation contexts.

5 A Practical Method to Implement APM as a Market Design Solution

In this section, we demonstrate how to practically implement APM as a market design solution. First, we describe a general three-step roadmap for how APM can be implemented in practice along with the practical advantages of APM over alternatives. Second, as a proof-of-concept exercise, we use data from Chicago Public Schools to show that the gains from switching to APM are *necessarily* positive and may also be quantitatively significant.

5.1 A General Methodology to Implement APM in Practice

APM can be implemented in practice via a simple three-step approach.

Step I: Understand the Preferences of the Authorities. As with any market design solution, the designer must understand the preferences of the authorities. This can be achieved via standard stated- or revealed-preference approaches. Under a stated-preference approach, one would ask the authorities to compare pairs of allocations for a large set of allocations and use the corresponding revealed preferences to set-identify the utility functions that can represent these preferences (see Ben-Akiva, McFadden, and Train, 2019, for a review of such methods).[23] Under a revealed-preference approach, one would estimate preferences based on historical data, *e.g.,* by using the admissions rules adopted in previous years to structurally estimate preferences that are consistent with these rules. We will shortly adopt this approach using observational data from Chicago Public Schools. A practical advantage of APM over priority and quota design is the absence of a need to elicit authorities' beliefs.

Step II: Compute and Communicate the Optimal APM. Once preferences have been recovered, Theorem 1 provides an explicit formula for the optimal APM and its computation is immediate. The communication of the optimal APM to stakeholders is also simple and, in view of Algorithm 1, can be achieved via a tabular format in which the score boost that agents from various groups receive is described as a function of an agent's own rank within their group.[24]

We provide a simple example of how APM can be transparently communicated to stakeholders in Table 1. Suppose that there are four groups $\mathcal{M} = \{\text{Red}, \text{Green}, \text{Blue}, \text{Orange}\}$ and consider an authority $c = \text{Great School}$. After computing the optimal APM A_c^*, we can print the additive priority boost that the APM awards $h(A^*_{m,c}(x_{m,c})) - h(s) = u'_{m,c}(x_{m,c})$ as a function of $x_{m,c}$, which here can be interpreted as the agent's rank within their group.

This reveals a further practical advantage of APM over priority and quota policies: it is not necessary to treat different groups differently *ex ante* using an APM. This is because even if score boosts are *ex ante* identical, agents of different groups will not necessarily receive the same realized score boost *ex post*. In the example of Table 1, Blue and Orange agents receive the same score boosts as a function of the number of agents that are admitted from their groups. However, if Orange agents were to score worse *ex post* and fewer were to be admitted, then Orange agents would receive higher score boosts. This constitutes a significant benefit of APM over priority mechanisms, which must treat different groups differently *ex ante* to achieve a more equitable allocation of resources *ex post*.

[23]By using the axiomatization of Çelebi (2023) for when a choice rule is rationalized by a utility function that features separability between match quality and diversity, this stated-preference approach would also allow one to test our assumption of separable preferences by checking whether the stated preferences between allocations satisfy the acyclicity axiom. If acyclicity is violated, then one can advise an authority to use our

Table 1: A Simple Example of How to Communicate APM to Stakeholders

School: Great School								
Score Boost		Rank Within Group						
		0	5	10	15	20	25	...
Group	Blue	+10	+8	+6	+4	+2	+0	...
	Red	+5	+3	+1	+0	+0	+0	...
	Green	+0	+0	+0	+0	+0	+0	...
	Orange	+10	+8	+6	+4	+2	+0	...

Note: The table reports the score boost that an agent in each group receives at Great School depending on the agent's rank within their group at the school. For example, the tenth ranked Blue agent will receive six bonus points for admission.

Step III: Implement the Optimal APM Using the DA Algorithm. With the optimal APM computed and communicated, all that remains is to implement it. As Theorem 4 shows, running the agent-proposing Deferred Acceptance algorithm, in which all authorities tentatively accept and reject agents by running the optimal APM on the set of agents that apply, implements stable allocations.[25] A benefit of this DA implementation is that it is dominant for agents to truthfully report their preferences to the mechanism.

5.2 The Practical Advantages of APM Over *Ex Post* Adjustments

An alternative approach to achieve first-best optimality would be to replicate the optimal allocation implemented by APM using priority or quota policies that adjust *ex post*, *i.e.,* based on the realized distribution of scores. While this could, in theory, achieve optimality, we argue that this suffers from two major shortcomings relative to using the APM approach.

Legal Constraints and APM. In many contexts, *ad hoc* and *ex post* adjustments of priority and quota policies after observing the distribution of applicants' characteristics are illegal. For example, under Section 15 of the 2021 UK School Admissions Code (UK Department for Education, 2021) (bold in original):[26]

> "Admission authorities **must** set ('determine') admission arrangements annually. Where changes are proposed to admission arrangements, the admission authority

extension to APMs beyond separable preferences that we present in Appendix D.

[24] A mathematically equivalent communication strategy would be to express every agent's score boost as a function of the number of agents in their group that is admitted to the school.

[25] While the lack of authority-side strategy-proofness of DA (and any stable mechanism) does not pose an issue for preference *elicitation*, there is always the delicate issue of how you can encourage authorities to truthfully *report* their preferences to the mechanism. This is true for any kind of mechanism, *e.g.,* priorities and quotas, and not just APMs.

[26] Which acts as law pursuant to Section 92 of the UK School Standards and Framework Act of 1998.

> **must** first publicly consult on those arrangements [...] Consultation **must** be for a minimum of 6 weeks and **must** take place between **1 October** and **31 January** of the school year before those arrangements are to apply."

This law explicitly rules out the possibility of *ex post* adjustment of priorities, quotas, or any other admissions policy. Moreover, it is followed in practice. For example, the North Yorkshire Council (an admissions authority) published its proposed priority policy on January 2025 as part of its admissions arrangements for the 2026/2027 academic year (North Yorkshire Council, 2025). The UK is far from abnormal adopting such a law: laws in India (India's Right of Children to Free and Compulsory Education Act), South Africa (South Africa's National Education Policy Act), and New Zealand (the New Zealand Education and Training Act) explicitly require that allocation rules be based on clear and objective criteria that are communicated in advance to stakeholders such as students and parents.

Simplicity and Transparency of APM. The second shortcoming is the potentially challenging nature of communicating and comprehending *ex post* adjustments even if they are legal. This is because stakeholders (students and parents) would be forced to entertain the vast space of distributions and how they will be mapped into resulting priorities and quotas—a task that would be challenging even for economic analysts with access to years of historical data and state-of-the-art econometric techniques (see Pathak and Shi, 2021). From this perspective, *ex post* adjustments do not feature the transparency of communication for which market design solutions typically strive (see Sönmez, 2023, for a full discussion of the high value placed on the transparency of market design solutions).

5.3 A Proof-Of-Concept: Implementing This Roadmap for Chicago Public Schools

To show how this roadmap could be practically implemented, we now use Chicago Public Schools (CPS) data in a proof-of-concept exercise.

Institutional Context on Chicago Public Schools. Under current policy, CPS admits students to its selective exam schools based on two criteria. First, CPS ranks students according to a score that combines the results of a specialized entrance exam, prior standardized test scores, and grades in prior coursework. This composite score ranges from 0 to 900 and higher-scoring students are admitted before lower-scoring ones, reflecting CPS's desire to allocate seats in exam schools to the students with the best academic standing.[27]

[27] Ellison and Pathak (2021) argue that the absolute levels of these scores are meaningful, and "can roughly be thought of as corresponding to a student's national percentile [when renormalized to range from 0 to 100.]"

Second, CPS divides the census tracts in the city into four *tiers* based on socioeconomic characteristics.[28] Tier 1 tracts are the most disadvantaged, while Tier 4 tracts are the most advantaged. This is reflected in the composite scores of students from Tier 1, who represent 25% of the city's population but comprise relatively few of the high-scoring students (Ellison and Pathak, 2021). As a result, Tier 1 students would have a very small share in the city's top exam schools without affirmative action. To ensure more equal representation across socioeconomic status in these schools, between 2013 and 2017 CPS implemented a quota policy that reserves 17.5% of the seats for each tier, yielding a total of 70% reserve seats and 30% merit slots that are open to students from all tiers. CPS allocates the seats by first assigning the highest-scoring students (regardless of their tier) to the merit slots and then the highest-scoring students from each tier to the 17.5% reserve seats. Moreover, the size of reserves in CPS cannot be easily changed *ex post* and has a high procedural burden. The size of the reserves in CPS is specified in Board Rule 602.2 (Chicago Public Schools, 2022), and changing a rule requires a two-thirds majority of the CPS board after a mandatory 30-day public comment period, initiation of which itself requires a majority vote of the CPS board (Chicago Public Schools, 2024).

We focus on the most selective CPS school, Walter Payton College Preparatory High School (Payton), which has the highest cutoffs for each tier in each year in our data and would have very few tier 1 students without affirmative action.[29] Table 2 presents the cutoff scores of each tier (*i.e.,* the composite score of the last admitted student from each tier).

Our general analysis relies on two premises: authorities care about both match quality and diversity; and there exists uncertainty about the distribution of applicants' scores and characteristics. Both of these premises hold in the CPS system.

First, the cutoff students from less advantageous tiers face is lower than the cutoff for students from more advantageous tiers. Therefore, CPS has a *revealed* preference (and not merely a stated preference) for a diverse student body.

Second, cutoff scores vary across years. This implies that the distribution of applicant characteristics varies from year to year. Notably, the majority of this variation is not simply

[28]Concretely, 800 census tracts are divided into four tiers based on six characteristics of each census tract: (i) median family income, (ii) percentage of single-parent households, (iii) percentage of households where English is not the first language, (iv) percentage of homes occupied by the homeowner, (v) adult educational attainment, and (vi) average Illinois Standards Achievement Test scores for attendance-area schools. These characteristics are then combined to construct the socioeconomic score for the tract. Finally, the tracts are ranked according to socioeconomic scores and partitioned into 4 tiers with approximately the same number of school-age children. See Ellison and Pathak (2021) for a more detailed account of the CPS system.

[29]This approach follows the analysis in Ellison and Pathak (2021), who focus on the two most competitive schools, Northside College Preparatory High School (Northside) alongside Payton. In the years that we study, the cutoff scores for Northside are below some other schools frequently, which is why we restrict attention to Payton.

Table 2: Admissions Cutoff Scores for Payton

Cutoff Score	2013	2014	2015	2016	2017
Tier 1	801	838	784	769	771
Tier 2	845	840	831	826	846
Tier 3	871	883	877	853	875
Tier 4	892	896	891	890	894

Note: The table reports the score of the lowest-scoring student that was admitted to Payton in each of the four tiers and five years.

driven by level shifts in the scores that students obtain across years: variation that is not explained by common shocks to the level of cutoff scores is approximately as important as the easily observable variation in the level of cutoff scores.[30] Given this uncertainty and the fact that CPS uses a policy that processes reserves after open slots, we know by Theorem 2 that CPS' baseline policy cannot be rationalized as optimal (even if they are extremely risk-averse). Thus, without even estimating their preferences, we know that the gains from the optimal APM will be strictly positive. Nevertheless, it is always possible that the gains from APM could be small. To quantify the size of these gains and demonstrate what the optimal APM might look like, we now follow our three-step approach.

Step I: Understanding Preferences. We use a revealed-preference approach to calibrate preferences. We assume a parsimonious, parametric form for CPS's preferences in a similar fashion to Ellison and Pathak (2021):

$$\xi(\bar{s}, x; \beta, \gamma) = \bar{s} + \sum_{t=1}^{4} \beta |x_t - 0.25|^{\gamma} \tag{20}$$

where $\bar{s}$ is the average score of admitted students and $x_t \in [0, 1]$ is the proportion of tier t students. The parameters β and γ index the slope and curvature of utility in losses from unequal representation (deviations from admitting 25% from each tier) and are the two free parameters of our framework. Our goal in assuming these preferences is to arrive at some sense of the gains from APM while acknowledging that it is impossible to know the parametric class in which the authority's preferences lie. To probe robustness to this functional form assumption, in Appendix C.2 we consider two other parametric utility functions that: i) allow for asymmetric effects of under-representation and over-representation, and ii) only consider losses from under-represented tiers.

[30] Formally, defining $C_{m,t}$ as the cutoff for tier m in year t (as reported in Table 2) and demeaning this for each tier, we compute $\tilde{C}_{m,t} = C_{m,t} - \frac{1}{T}\sum_{t=1}^{T} C_{m,t}$. We then statistically decompose this into a year-specific component λ_t and an orthogonal tier-by-year residual $\epsilon_{m,t}$ that averages to zero within each year: $\tilde{C}_{m,t} = \lambda_t + \varepsilon_{m,t}$, $\lambda_t \perp \epsilon_{m,t}$, $\mathbb{E}_t[\epsilon_{m,t}] = 0$. We find that $\mathbb{V}[\hat{\epsilon}_t]/\mathbb{V}[\tilde{C}_{m,t}] = 0.54$. That is, 54% of the variation in cutoff-scores cannot be accounted for by aggregate shifts in the levels of scores over time.

Figure 4: The Estimated Optimal APM

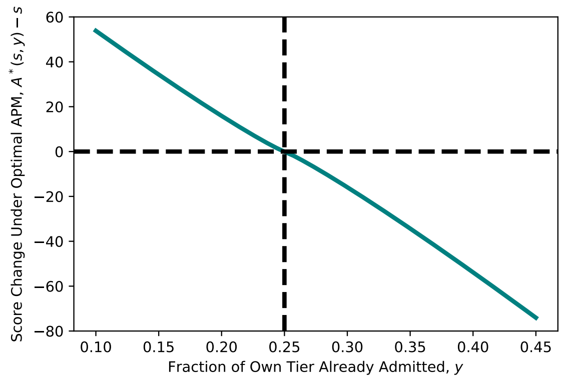

Note: This figure plots the change in a student's score when fraction y of students in their own tier has already been admitted under the estimated optimal APM, A^*. At $y = 0.25$ (the vertical dashed black line), the score is unchanged. For $y < 0.25$, students receive a score boost. For $y > 0.25$, students receive a score penalty. The range of the x-axis, $[0.1, 0.45]$, is chosen to cover the full range of fractions of admitted students under both the optimal APM and the CPS reserve policy from all tiers in all years of our sample (see Figure 5).

We estimate β and γ to best rationalize the pursued policy choice of 17.5% reserves for each tier as optimal within the class of all reserve policies. That is, using the empirical distribution of distributions of agent's scores and tiers across the five years of our data, we choose β and γ as those that minimize the sum of squared differences from zero of the authority's first-order conditions for optimal reserve choices. We believe this to be a reasonable approach as the size of the reserves is an important issue that is decided only after much deliberation.[31] This approach is also conservative as it estimates preferences to make the *status quo* reserve policy as good as possible. The full, formal approach is described in Appendix C.1. Performing this estimation yields estimated parameter values of $\beta^* = -209.5$ and $\gamma^* = 2.11$.

Step II: Computing and Communicating the Optimal APM. In Figure 4, we illustrate how the estimated optimal APM changes students' scores to arrive at their ultimate priorities. In accordance with the preferences we have assumed, students receive a score boost when their tier is underrepresented and a score penalty when their tier is overrepresented. As we found $\gamma^* = 2.11$, the estimated diversity preference is very close to quadratic. Thus, the optimal APM depicted in Figure 4 is very close to linear.

In Table 3, we demonstrate how the estimated optimal APM for CPS can be communi-

[31] These points are emphasized in Dur, Pathak, and Sönmez (2020): "This change was made at the urging of a Blue Ribbon Commission (BRC, 2011), which examined the racial makeup of schools under the 60% reservation compared to the old Chicago's old system of racial quotas. They advocated for the increase in tier reservations [to the current 70% level] on the basis it would be "improving the chances for students in neighborhoods with low performing schools, increasing diversity, and complementing the other variables."

Table 3: The Estimated Optimal APM in a Simple Tabular Format

Walter Payton College Preparatory High School									
Score Boost		Rank Within Group							
		10	15	20	25	30	35	40	45
Group	Tier 1	+128	+108	+90	+74	+58	+39	+20	+0
	Tier 2	+128	+108	+90	+74	+58	+39	+20	+0
	Tier 3	+128	+108	+90	+74	+58	+39	+20	+0
	Tier 4	+128	+108	+90	+74	+58	+39	+20	+0

Note: The table reports the score boost that a student in each group receives depending on their rank within their group, rounded to the nearest integer. As the proportion of admitted students from each group under the Optimal APM has a range in $[0.10, 0.45]$ (see Figure 5), the score boost is calibrated to be zero at the highest level of representation and set to be the maximum value of +128 for representation below 10% and zero above 45%.

cated to stakeholders using the tabular approach presented in Table 1. We restrict attention to the observed proportions of students from each tier $[0.10, 0.45]$ and recalibrate the score boost so that it is equal to zero at the highest level of representation. We believe this recalibration will be useful in communication as it is only required to keep track of bonuses instead of bonuses and penalties, which may look controversial to some stakeholders. The optimal APM gives the top-ranked students in each tier a score boost of 128 points out of a possible 900, equivalent to 14.2% of the total available points. This boost decreases to 0 almost linearly as the representation of a group approaches 45% of the student body. As the assumed preferences treat all groups in the same way, so does the optimal APM (the rows of Table 3 are identical). Therefore, the optimal APM treats groups in a symmetric manner: a group receives score boosts relative to another if and only if its members have lower scores.

Step III: Implementing the Optimal APM and Computing its Gains. We now use our estimated model to implement the optimal APM instead of the pursued reserve policy between 2013 and 2017 to quantify the welfare gains from using APM. To do this, we compare the empirical payoff under two mechanisms: the pursued quota policy, r^*, and the optimal APM from Theorem 1, A^*. The empirical payoff $\Xi(\phi, \hat{\Lambda}, \beta^*, \gamma^*)$ from Mechanism ϕ is the average estimated utility (*i.e.,* given β^* and γ^*) of the authority of the allocations that ϕ implements under the empirical distribution over distributions $\hat{\Lambda}$.

The empirical payoff under APM is 876.9, while it is 874.8 under the CPS reserve policy. Thus, the gains from APM are equivalent to increasing average scores by 2.1, holding diversity fixed. To benchmark the size of the gains, we require units in which they can be meaningfully expressed. To this end, we define the *loss from underrepresentation* as the payoff lost by CPS under its baseline policy from not admitting a fully balanced class, while holding fixed

Figure 5: Comparing Admissions under the Optimal APM and the CPS Policy

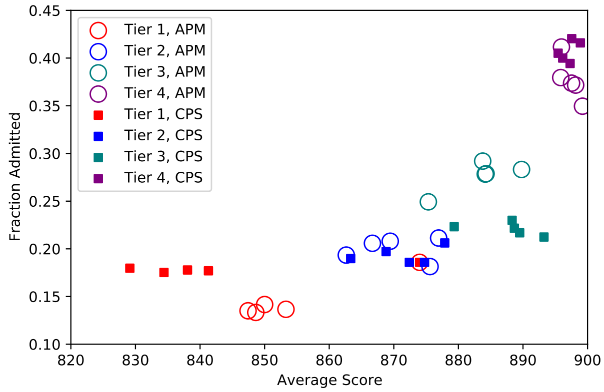

Note: Each point corresponds to one of the four tiers of students in one of the five years under either the optimal APM or the CPS policy. The x-axis corresponds to the average score of those admitted from that tier in that year under that policy. The y-axis corresponds to the fraction of admitted students from that tier in that year under that policy.

the average score of the class. This is equal to 5.6 points. Thus, the gains from APM are equal to 37.5% of the loss from underrepresentation incurred under the CPS policy.

To contextualize the magnitude of this improvement, we can compare the gains from the optimal APM to the gains from the 2012 CPS reform that gave rise to the CPS policy from 2013-2017 and increased the size of all reserves from 15% to 17.5%. Under the estimated preferences, the empirical payoff under the 15% reserve rule is 873.9, and so the gains from the reform are equivalent to increasing average scores by 0.9. Thus, the gains from switching to the optimal APM are 2.3 times larger than the gains from this recent reform.

These gains stem from the variation across years in the joint distribution of student scores and tiers. This can be seen in Table 2, which shows the variability in the scores of the marginally admitted students from tiers 1, 2, and 3. More systematically, we visualize the difference in outcomes under CPS' reserves and the optimal APM by plotting the average scores and fraction admitted for each tier for each year under both mechanisms in Figure 5. There are two main differences between the allocations. First, the APM allocates systematically fewer tier 1 and tier 4 students and more tier 3 students. These level effects are a consequence of the second difference: the APM admits a greater fraction of students from each tier (especially tiers 1 and 3) in the years in which that tier scores well. The fact that tier 3 students score well relative to their admissions level is what leads the authority to admit more tier 3 students and fewer tier 1 students. These positive selection and level effects generate the welfare gains.

Robustness of the Estimated Gains. While these estimates are intended to be taken as suggestive, In Appendix C.2 we nevertheless explore the robustness of our estimates to the three core assumptions of our analysis: (i) that CPS has correct beliefs over the possible distributions of student characteristics, (ii) that CPS has preferences that lie in the assumed parametric family, and (iii) that CPS optimizes the sizes of all four tiers. In all three cases, we continue to find substantial gains from switching to the optimal APM.

6 Conclusion

We study optimal mechanism design for authorities that wish to trade off match quality and diversity in the face of uncertainty about the distribution of agents' characteristics. We introduce Adaptive Priority Mechanisms (APM) and characterize an APM that is both optimal and can be specified solely in terms of the preferences of the authority. We study the priority and quota policies that are used in practice and show that they are optimal if and only if the authority is either risk-neutral or extremely risk-averse over diversity. Moreover, all authorities following the optimal APM implements the unique stable allocation when combined with the standard Deferred Acceptance algorithm.

Our analysis therefore proposes a new and practical allocation mechanism that provably improves upon existing alternatives and, we argue, can be transparently communicated to stakeholders and implemented using a leading existing approach (the Deferred Acceptance algorithm). Thus, we argue that APM could be a valuable tool for improving the design of real-world allocation mechanisms in many settings, including the allocation of seats at schools, places at universities, and medical resources to patients. Moreover, we have provided a concrete roadmap for how APM could be designed, communicated, and implemented in practice. Implementing this roadmap, our proof-of-concept analysis using Chicago Public Schools data suggests that the use of APM has the potential to yield considerable welfare gains over the *status quo*.

A Omitted Proofs of Main Results

A.1 Proof of Theorem 1

Proof. We characterize the optimal allocation for each $\omega \in \Omega$ and show that the claimed adaptive priority mechanism implements the same allocation. Fix an $\omega \in \Omega$ and suppress the dependence of F_ω and f_ω thereon, and define the utility index of a score as $\tilde{s} = h(s)$ with induced densities over $\tilde{s}$ given by $\tilde{f}_m$ for all $m \in \mathcal{M}$. Let the measure of agents from any group $m \in \mathcal{M}$ that is allocated the resource be $x_m \in [0, \bar{x}_m]$ where $\bar{x}_m = \int_{h(0)}^{h(1)} \tilde{f}_m(\tilde{s}) \mathrm{d}\tilde{s}$. Observe that, conditional on fixing the measures of agents from each group that are allocated the resource $x = \{x_m\}_{m\in\mathcal{M}}$, there is a unique optimal allocation (*i.e.*, ξ-maximal μ up to measure zero transformations). In particular, as g and h are continuous and strictly increasing, the optimal allocation conditional on x satisfies $\mu^*(\tilde{s}, m; x) = 1 \iff \tilde{s} \geq \underline{\tilde{s}}_m(x_m)$ for some thresholds $\{\underline{\tilde{s}}_m(x_m)\}_{m\in\mathcal{M}}$ that solve:

$$\int_{\underline{\tilde{s}}_m(x_m)}^{h(1)} \tilde{f}_m(\tilde{s})\mathrm{d}\tilde{s} = x_m \tag{21}$$

We can then express the problem of choosing the optimal $x = \{x_m\}_{m\in\mathcal{M}}$ as:

$$\max_{x_m\in[0,\bar{x}_m],\,\forall m\in\mathcal{M}} \sum_{m\in\mathcal{M}} \int_{\underline{\tilde{s}}_m(x_m)}^{h(1)} \tilde{s}\tilde{f}_m(\tilde{s})\mathrm{d}\tilde{s} + \sum_{m\in\mathcal{M}} u_m(x_m) \quad \text{s.t.} \sum_{m\in\mathcal{M}} x_m \leq q \tag{22}$$

where a solution exists by compactness of the constraint sets and continuity of the objective. We can derive necessary and sufficient conditions on the solution(s) to this problem by considering the Lagrangian:

$$\begin{aligned}\mathcal{L}(x,\lambda,\overline{\kappa},\underline{\kappa}) = &\sum_{m\in\mathcal{M}} \int_{\underline{\tilde{s}}_m(x_m)}^{h(1)} \tilde{s}\tilde{f}_m(\tilde{s})\mathrm{d}\tilde{s} + \sum_{m\in\mathcal{M}} u_m(x_m) \\ &+ \lambda\left(q - \sum_{m\in\mathcal{M}} x_m\right) + \sum_{m\in\mathcal{M}} \overline{\kappa}_m(\bar{x}_m - x_m) + \sum_{m\in\mathcal{M}} \underline{\kappa}_m x_m\end{aligned} \tag{23}$$

The first-order necessary conditions to this program are given by:

$$\frac{\partial\mathcal{L}}{\partial x_m} = -\underline{\tilde{s}}'_m(x_m)\underline{\tilde{s}}_m(x_m)\tilde{f}_m(\underline{\tilde{s}}_m(x_m)) + u'_m(x_m) - \lambda - \overline{\kappa}_m + \underline{\kappa}_m = 0 \tag{24}$$

$$\lambda\frac{\partial\mathcal{L}}{\partial\lambda} = \lambda\left(q - \sum_{m\in\mathcal{M}} x_m\right) = 0 \tag{25}$$

$$\overline{\kappa}_m\frac{\partial\mathcal{L}}{\partial\overline{\kappa}_m} = \overline{\kappa}_m(\bar{x}_m - x_m) = 0 \tag{26}$$

$$\underline{\kappa}_m\frac{\partial\mathcal{L}}{\partial\underline{\kappa}_m} = \underline{\kappa}_m x_m = 0 \tag{27}$$

for all $m \in \mathcal{M}$. By implicitly differentiating Equation 21, we obtain that:

$$-\underline{\tilde{s}}'_m(x_m)\tilde{f}_m(\underline{\tilde{s}}_m(x_m)) = 1 \tag{28}$$

Thus, we can simplify Equation 24 to:

$$\frac{\partial\mathcal{L}}{\partial x_m} = \underline{\tilde{s}}_m(x_m) + u'_m(x_m) - \lambda - \overline{\kappa}_m + \underline{\kappa}_m = 0 \tag{29}$$

Observe that all constraints are linear. Thus, if the objective function is strictly concave, the first-order conditions are also sufficient. Observe by Equation 28 that $\underline{\tilde{s}}_m(x_m)$ is a strictly decreasing function of x_m, and all cross-partial derivatives are zero. Therefore, the first summation is strictly concave. Moreover u'_m is a decreasing function of x_m by virtue of the assumption that u_m is concave for all $m \in \mathcal{M}$. Then the second summation is concave. Thus, the objective function is strictly concave and the optimal allocation is unique.

Thus, to verify that our claimed adaptive priority mechanism is a first-best mechanism, it suffices to show that the allocation it implements satisfies Equations 24 to 27. The adaptive priority mechanism $A_m(y_m, s) = h^{-1}\left(h(s) + u_m'(y_m)\right)$ in the transformed score space yields transformed scores $h\left(A_m(y_m, s)\right) = \tilde{s} + u_m'(y_m)$. Define x_m as the admitted measure of agents from group m under this mechanism. Agents in group $m \in \mathcal{M}$ are allocated the resource if and only if $\tilde{s} + u_m'(x_m) \geq s^C$ for some threshold s^C that solves:

$$\sum_{m \in \mathcal{M}} \int_{\max\{\min\{s^C - u_m'(x_m), h(1)\}, h(0)\}}^{h(1)} \tilde{f}_m(\tilde{s}) \mathrm{d}\tilde{s} = q \tag{30}$$

We can therefore partition $\mathcal{M}$ into three sets that are uniquely defined: (i) interior $\mathcal{M}_I = \{m \in \mathcal{M} | s^C - u_m'(x_m) \in (h(0), h(1))\}$; (ii) no allocation $\mathcal{M}_0 = \{m \in \mathcal{M} | s^C - u_m'(x_m) \geq h(1)\}$; (iii) full allocation $\mathcal{M}_1 = \{m \in \mathcal{M} | s^C - u_m'(x_m) \leq h(0)\}$. For all $m \in \mathcal{M}_0$, we implement $x_m = 0$. For all $m \in \mathcal{M}_1$, we implement $x_m = \bar{x}_m$. For all $m \in \mathcal{M}_I$, we implement $x_m \in (0, \bar{x}_m)$. For any $m \in \mathcal{M}_I$, the allocation threshold is $\underline{\tilde{s}}_m(x_m) = s^C - u_m'(x_m)$. For any $m \in \mathcal{M}_0$, the allocation threshold is $h(1)$. For any $m \in \mathcal{M}_1$, the allocation threshold is $h(0)$.

We now verify that this outcome satisfies the established necessary and sufficient conditions. For all $m \in \mathcal{M}_I$, by the complementary slackness conditions we have that $\underline{\kappa}_m = \overline{\kappa}_m = 0$. Substituting the above into Equation 24 for all $m \in \mathcal{M}_I$ we obtain that:

$$s^C - \lambda = 0 \tag{31}$$

which is satisfied for $\lambda = s^C$. As $q = \sum_{m \in \mathcal{M}} x_m$, the complementary slackness condition for λ is then satisfied. For all $m \in \mathcal{M}_0$, by complementary slackness we have that $\overline{\kappa}_m = 0$ and Equation 24 is satisfied by:

$$\underline{\kappa}_m = \lambda - h(1) - u_m'(0) \tag{32}$$

For all $m \in \mathcal{M}_I$, by complementary slackness we have that $\underline{\kappa}_m = 0$ and Equation 24 is satisfied by:

$$\overline{\kappa}_m = h(0) + u_m'(\bar{x}_m) - \lambda \tag{33}$$

This completes the proof of first-best optimality of A^*. Moreover, as the optimal allocation is unique for all ω, any allocation that differs from the allocation implemented by the optimal APM at any ω would not be first-best optimal. Therefore, any first-best-optimal mechanism must implement essentially the same allocation as A^*. $\square$

A.2 Proof of Theorem 2

Proof. First, we prove the if parts of the results. Part (i): When u_m is linear, u'_m is constant and the first-best optimal adaptive priority mechanism is a priority mechanism $P(s,m) = h^{-1}(h(s)+u'_m)$. Part (ii): When $\tilde{u}'_m(x_m) \geq k$ for $x_m \leq x_m^{\text{tar}}$ and $\tilde{u}'_m(x_m) = 0$ for $x_m > x_m^{\text{tar}}$ and $\sum_{m\in\mathcal{M}} x_m^{\text{tar}} < q$, observe that the optimal mechanism admits $x_m \geq x_m^{\text{tar}}$ for all $m \in \mathcal{M}$ in all states of the world, but conditional on $x_m \geq x_m^{\text{tar}}$ for all $m \in \mathcal{M}$ admits the highest-scoring set of agents. A quota $Q_m = x_m^{\text{tar}}$ and $Q_R = q - \sum_{m\in\mathcal{M}} x_m^{\text{tar}}$, with $D(R) = |\mathcal{M}|+1$ implements this allocation and is first-best optimal for any authority that is extremely risk-averse.

Second, we prove the only if parts of the results. Part (i): Assume the utility functions are not linear and let m denote a group where u'_m is not constant on $(0, q)$. We say that a state ω has full support if f_ω has full support. A state ω has *full support in m and n* if $f_\omega(\cdot, m) > 0$ and $f_\omega(\cdot, n) > 0$ for some m and n and positive measures of only m and n. Let ω be a state that has full support in m and n. Moreover, assume both groups have a measure q of agents. We first establish that in any optimal allocation, agents from both groups are allocated the resource.

Claim 1. *If preferences are non-trivial, then the optimal allocation has $x_n, x_m > 0$.*

Proof. Toward a contradiction, suppose without loss of generality that $x_n = 0$. This implies that $x_m = q$. By the necessary first-order condition from Theorem 1 (combing Equations 24 and 27), we have that:

$$u'_m(q) + h(0) = u'_n(0) + h(1) + \underline{\kappa}_n \geq u'_n(0) + h(1) \tag{34}$$

where the inequality follows as $\underline{\kappa}_n \geq 0$. Thus, we have that:

$$u'_m(q) - u'_n(0) \geq h(1) - h(0) > u'_m(q) - u'_n(0) \tag{35}$$

where the first inequality follows by rearranging Equation 34 and the second follows by the definition of non-triviality of preferences. This is a contradiction, thus $x_n, x_m > 0$ in any optimal allocation. □

We now establish an equation relating x_n and x_m that will be useful in the steps to come.

Claim 2. *Let ω have full support in m and n and let μ denote a cutoff matching with cutoffs s_m and s_n. Let x_m and x_n denote the measures of agents who are allocated the object at μ. μ is optimal if and only if $u'_m(x_m) + h(s_m) = u'_n(x_n) + h(s_n)$ and $x_n + x_m = q$.*

Proof. By the necessary and sufficient first-order conditions from Theorem 1, we again have that:

$$u'_m(x_m) + h(s_m) - \overline{\kappa}_m + \underline{\kappa}_m = u'_n(x_n) + h(s_n) - \overline{\kappa}_n + \underline{\kappa}_n \tag{36}$$

By Claim 1, we have $x_m, x_n > 0$. Thus, by the complementary slackness conditions (Equations 26 and 27), we have that $\overline{\kappa}_m = \underline{\kappa}_m = \overline{\kappa}_n = \underline{\kappa}_n = 0$. Thus, we obtain:

$$u'_m(x_m) + h(s_m) = u'_n(x_n) + h(s_n) \tag{37}$$

together with $x_n + x_m = q$, we have characterized the optimal allocation as claimed. $\square$

Continue to let x_m and x_n denote the measures of group m and n agents at the optimal allocation under ω, and s_m and s_n denote the cutoff scores for admission. There are now two cases to consider: (i) $u'_m(x_m)$ and $u'_n(x_n)$ are locally constant, (ii) $u'_m(x_m)$ or $u'_n(x_n)$ are not locally constant. If we are in case (i), we will construct an ω' with a unique optimal allocation x'_m and x'_n where $u'_m(x'_m)$ or $u'_n(x'_n)$ is not locally constant, and then show jointly how we arrive at a contradiction in both cases (i) and (ii).

To this end, suppose that we are in case (i). Let x^*_m and x^*_n denote the measures that are closest to x_m and x_n such that $u'_m(x_m)$ and $u'_n(x_n)$ are not locally constant, *i.e.*,:

$$x^*_k = \arg\min_{x'_k} \Bigg\{ |x_k - x'_k| \Big| u'_k(x_k) = u'_k(x'_k) \text{ and for all } \varepsilon > 0$$
$$\text{either } u'_k(x'_k - \varepsilon) > u'_k(x_k) \text{ or } u'_k(x'_k + \varepsilon) > u'_k(x_k) \Bigg\} \tag{38}$$

As u'_k is continuous, this minimum is attained and x^*_k is well-defined. Without loss of generality, assume $|x_m - x^*_m| \leq |x_n - x^*_n|$ and define both $\hat{x}_m = x^*_m$ and $\hat{x}_n = q - x^*_m$. We now construct a state ω' such that $\hat{x}$ is optimal:

Claim 3. *Define ω' where $F_m(1) - F_m(s_m) = \hat{x}_m$, $F_n(1) - F_n(s_n) = \hat{x}_n$ and ω' has full support in m and n. The allocation that admits the highest-scoring $\hat{x}_m$ group m agents and the highest-scoring $\hat{x}_n$ group n agents is the unique optimal allocation.*

Proof. By Claim 2, as $\hat{x}_m + \hat{x}_n = q$ by construction, $\hat{x}$ is optimal if and only if Equation 37 holds. To this end, observe that if we admit $\hat{x}$, then the cutoff scores are the same as under x as $F_m(1) - F_m(s_m) = \hat{x}_m$ and $F_n(1) - F_n(s_n) = \hat{x}_n$, by construction. Thus, we have that:

$$u'_m(\hat{x}_m) + h(s_m) = u'_m(x_m) + h(s_m) = u'_n(x_n) + h(s_n) = u'_n(\hat{x}_n) + h(s_n) \tag{39}$$

where the first equality holds by construction as $\hat{x}_m = x^*_m$ and $u'_m(x^*_m) = u'_m(x_m)$, the second equality holds by optimality of x, and the third equality holds as $|x_m - x^*_m| \leq |x_n - x^*_n|$, which implies that $u'_n(\hat{x}_n) = u'_n(x^*_n)$. Thus, Equation 37 holds, and $\hat{x}$ is optimal, as claimed. $\square$

Observe that this construction also applies trivially in case (ii) with $\hat{x}_m = x^*_m = x_m$ and $\hat{x}_n = x^*_n = x_n$. Thus, using this construction, we can now study cases (i) and (ii) together. In state ω', to implement this optimal allocation, we must have that $P(s, m) < P(s_n, n)$ for all but a measure zero set of s such that $s < s_m$. We will now construct another state ω'' such that any priority mechanism with this property is suboptimal.

First, suppose that $\hat{x}_m \leq x_m$ and fix some ε such that $\hat{x}_m - \varepsilon > 0$. Define $\tilde{s}_m$ as solving the following equation:

$$u'_m(\hat{x}_m - \varepsilon) + h(\tilde{s}_m) = u'_n(\hat{x}_n + \varepsilon) + h(s_n) \tag{40}$$

We then have that:

$$\tilde{s}_m = h^{-1}\left(h(s_n) + u'_n(\hat{x}_n + \varepsilon) - u'_m(\hat{x}_m - \varepsilon)\right) < h^{-1}\left(h(s_n) + u'_n(\hat{x}_n) - u'_m(\hat{x}_m)\right) = s_m \tag{41}$$

where the first equality rearranges Equation 40 and the second inequality uses the facts that $u'_n(\hat{x}_n) \geq u'_n(\hat{x}_n + \varepsilon)$ and $u'_m(\hat{x}_m) < u'_m(\hat{x}_m - \varepsilon)$. We now construct a state ω'' such that $(\hat{x}_m - \varepsilon, \hat{x}_n + \varepsilon)$ is optimal.

Claim 4. *Define ω'' where $1 - F_m(\tilde{s}_m) = \hat{x}_m - \varepsilon$, $1 - F_n(s_n) = \hat{x}_n + \varepsilon$ with full support in m and n. The allocation that admits the highest-scoring $(\hat{x}_m - \varepsilon, \hat{x}_n + \varepsilon)$ agents is the unique optimal allocation.*

Proof. Following the same steps as Claim 3, and the fact that Equation 40 holds by construction, we have that the claim holds. □

Observe that to implement this optimal allocation a priority mechanism must set $P(s, m) \geq P(s_n, n)$ for all but zero measure $s > \tilde{s}_m$. However, since $\tilde{s}_m < s_m$, this contradicts the optimality condition for state ω' that $P(s, m) < P(s_n, n)$ for all $s < s_m$. This is because for all but measure zero $s \in (\tilde{s}_m, s_m)$, which we have established is non-empty, we have that:

$$P(s, m) \geq P(s_n, n) > P(s, m) \tag{42}$$

which is a contradiction. To complete the proof, we need only consider the case that $\hat{x}_m > x_m$. In this case, we can apply essentially the same steps and the result follows. Concretely, instead increasing $\hat{x}_m$ by ε (that satisfies $\hat{x}_n - \varepsilon > 0$) and following the same steps yields the required contradiction.

We have now constructed three states $\omega, \omega', \omega''$ such that no priority mechanism can be optimal in each state when the authority is not risk-neural, completing the proof.

Part (ii): Assume that a quota policy is optimal, we now show that the authority's preferences must be extremely risk-averse. For each group $m \in \mathcal{M}$, let $c_m \in [0, 1]$ and $c_m \neq c_n$ if $m \neq n$. Let ω be such that the scores of agents from each group m are uniformly distributed between $[c_m, c_m + \epsilon]$, where ϵ is chosen to be small so that there is no overlap of these supports and each group has measure q agents. Let m_ω denote the group with the highest c_m at ω. Now, compute the optimal allocation at ω and denote the measure of admitted agents from each group at the optimal allocation by $\{x^*_m(\omega)\}_{m \in \mathcal{M}}$. We first show that under any optimal quota policy, the level of the quotas must be set equal to the optimal allocation for all but the highest-scoring group:

Claim 5. *If a quota Q attains the optimal allocation, then for each $m \neq m_\omega$, $Q_m = x^*_m(\omega)$.*

Proof. If $Q_m > x^*_m(\omega)$, then we admit $x_m \geq Q_m > x^*_m(\omega)$, which is suboptimal as there is a unique optimal allocation by Theorem 1. If $Q_m < x^*_m(\omega)$ and $m \neq m_\omega$, then $x_m = Q_m$ as $c_{m_\omega} > c_m + \varepsilon$ and no agent from group m can claim a merit slot. This is suboptimal. Thus, $Q_m = x^*_m(\omega)$ for all $m \neq m_\omega$. □

Next, create ω' by changing the highest-scoring group, *i.e.*, $m_\omega \neq m_{\omega'}$. Let $x^*_{m_\omega}(\omega')$ denote the measure of admitted agents from group m_ω under ω'. Applying Claim 5, If Q attains the optimal allocation, then it must be that $Q_{m_\omega} = x^*_{m_\omega}(\omega')$. Define Q^*_m by $Q^*_m = x^*_m$ for all $m \in \mathcal{M} \setminus \{m_\omega\}$ and $Q^*_{m_\omega} = x^*_{m_\omega}(\omega')$.

Now, we have proved that if Q is an optimal policy, then $Q_m = Q^*_m$ for all $m \in \mathcal{M} \setminus \{m_\omega\}$ and $Q_{m_\omega} = Q^*_{m_\omega}$. We now establish that merit slots must be processed after any positive measure quota slots if the merit slots are of positive measure:

Claim 6. *If there is a quota policy that attains the first-best, Q, then $Q_m = Q^*_m$ and either $\sum_{m\in\mathcal{M}} Q^*_m = q$, i.e., there are no merit slots (merit slot processing does not matter), or $\sum_{m\in\mathcal{M}} Q^*_m < q$ and merit slots are processed after any positive measure quota slots.*

Proof. We have already proved $Q_m = Q^*_m$. If $\sum_{m\in\mathcal{M}} Q^*_m = q$, there are no merit slots and any processing order yields the same result. If $\sum_{m\in\mathcal{M}} Q^*_m < q$, for a contradiction, assume merit slots are processed before quota slots for group m and $Q^*_m > 0$. There are two cases, $m \neq m_\omega$ and $m = m_\omega$. We start with the first case. Note that there is a cutoff s_m for group m with $s_m < c_m + \epsilon$ and all agents from group m who score above s_m are allocated the resource. Create ω'' by taking measure $x_m/2$ of these agents who are allocated the resource and give them scores above c_{m_ω} (the highest-scoring group at ω). The scores of the remaining $x_m/2$ agents are distributed uniformly at $[c_m, c_m + \epsilon]$.

We now observe that the optimal allocations at ω and ω'' are the same. This is because increasing the scores of already admitted agents does not change the preferences of the authority of whom to admit. Moreover, the optimal allocation at ω'' cannot be attained if the quota slots for group m are processed after the merit slots. This follows as, if merit slots are processed before quota slots for group m, a strictly positive measure of them would go to group m agents at ω'' since now they have a measure of agents with the highest scores, which violates optimality.

This proves the claim for $m \neq m_\omega$. To prove the result for $m = m_\omega$, replicate the above steps with ω' where m_ω is not the highest-scoring group. □

We now use these claims to establish that if quotas are first-best optimal, then (u, h) must agree with $(\tilde{u}, \tilde{h})$ on optimal allocations.

Claim 7. *The quota first policy with $Q_m = x_m^{tar}$ maximizes the utility with respect to $\tilde{u}, \tilde{h}$.*

Proof. This is clear as for $\tilde{u}, \tilde{h}$, diversity utility dominates until x_m^{tar} and has no effect after. □

This proves the result since if there exists a first-best optimal quota policy, then it is rationalized by $(\tilde{u}, \tilde{h})$ with $x_m^{tar} = Q_m^*$. Hence, if there is a first-best quota mechanism, the authority is extremely risk-averse. □

A.3 Proof of Theorem 3

Proof. We first prove the following lemma.

Lemma 2. *Any stable matching is a cutoff matching.*

Proof. Assume that μ is a stable matching. Let $S_{m,c} = \inf_\theta\{s_c(\theta) : m(\theta) = m, \mu(\theta) = c\}$. Since μ satisfies within-group fairness, for all m and $s' > S_{m,c}$, if $m(\theta) = m$ and $s_c(\theta) = s'$, $\mu(\theta) \succeq_\theta c$. Moreover, from part (iv) of the definition of matching, this extends to the case where $s' = S_{m,c}$. Concretely, suppose that $\mu(\theta) \neq c$, $c \succ_\theta \mu(\theta)$ and $s_c(\theta) = S_{m,c}$. Consider a sequence of types $\{\theta_k\}_{k\in\mathbb{N}}$ with common group m and scores $\{s_c(\theta_k)\}_{k\in\mathbb{N}}$ such that $s_c(\theta_k) > S_{m,c}$ for all $k \in \mathbb{N}$ and $s_k(\theta) \to S_{m,c}$. Define the set $\Theta^E = \{\theta \in \Theta : c \succ \mu(\theta)\}$, which must be open by part (iv) of the definition of a matching. We have that $\theta_k \notin \Theta^E$ for all $k \in \mathbb{N}$ but $\lim_{k\to\infty} \theta_k \in \Theta^E$, which contradicts that Θ^E is open. Thus, if μ is stable, then it is also a cutoff matching. □

Therefore, to characterize stable matchings, it is enough to characterize cutoffs that induce a stable matching, which we call *stable cutoffs*.

Definition 6. *A vector S is a market-clearing cutoff if it satisfies the following:*

1. *$D_c(S) \leq q_c$ for all c.*
2. *$D_c(S) = q_c$ if $S_{m,c} > 0$ for some $m \in \mathcal{M}$.*

Since an authority can admit different measures of agents from different groups, there is a continuum of cutoffs that clear the market given S_{-c}, as long as $\{(0, \dots, 0)\}$ is not the only market-clearing cutoff. Let $I(S_{-c})$ denote the set of market-clearing cutoffs. Let $I^*(S_{-c}) \subseteq I(S_{-c})$ denote the unique (by Lemma 3) cutoffs that implement the outcome under APM A_c^* when the authority faces the induced type measure over the set $\tilde{D}_c(S_{-c})$. Define the map $T_c : [0,1]^{|\mathcal{M}|\times|\mathcal{C}|} \to [0,1]^{|\mathcal{M}|}$ as $T_c(S) = I_c^*(S_{-c})$ with $T : [0,1]^{|\mathcal{M}|\times|\mathcal{C}|} \to [0,1]^{|\mathcal{M}|\times|\mathcal{C}|}$

given by $T = \{T_c\}_{c\in\mathcal{C}}$. We first show that the set of fixed points of T equals the set of stable cutoffs and that T is increasing.

Claim 8. *The set of fixed points of T equals the set of stable cutoffs.*

Proof. If S^*, with corresponding matching μ^* (by Lemma 2), is a fixed point of T, then each $c \neq c_0$ admits their most preferred measure q_c agents in $\tilde{D}_c(S^*_{-c})$ (by Theorem 1). Note that any $\hat{\Theta}$ that can block the matching must prefer c to their allocation at μ^* and therefore $\hat{\Theta} \subset \tilde{D}_c(S^*_{-c})$. Then there cannot be a $\hat{\Theta}$ that blocks μ^* at c since c already attains the first-best utility under $\tilde{D}_c(S^*_{-c})$ from the definition of $T_c(S)$ and Theorem 1.

Conversely, let S be any cut-off vector that supports a stable matching μ. Then for any $\hat{\Theta} \subseteq \tilde{D}_c(S_{-c}) \setminus \mu(c)$ and $\tilde{\Theta} \subseteq \mu(c)$ of equal measure, we must have $\mu(c) \succ_c \mu(c) \setminus \tilde{\Theta} \cup \hat{\Theta}$. Therefore, the optimal APM A^*_c applied to $\tilde{D}_c(S_{-c})$ under the distribution F chooses $\mu(c)$, yielding the cut-off vector $T_c(S) = S_c$ for all c. Thus any stable cut-off profile is a fixed point of the map T. □

Claim 9. *T is increasing.*

Proof. Fix an arbitrary $c \in \mathcal{C}$ and suppose that $S'_{-c} \geq S_{-c}$. Toward a contradiction suppose that there exists $m \in \mathcal{M}$ such that $t'_{c,m} = T_{c,m}(S') = I^*_c(S'_{-c}) < I^*_c(S_{-c}) = T_{c,m}(S) = t_{c,m}$, *i.e.,* the admissions threshold for group m at authority c goes down. Let f and f' be the induced joint densities of agents over scores at c and groups by the sets $\tilde{D}_c(S_{-c})$ and $\tilde{D}_c(S'_{-c})$, respectively. Let $\{x_{m,c}\}_{m\in\mathcal{M}}$ and $\{x'_{m,c}\}_{m\in\mathcal{M}}$ denote the measure agents who score above $t_{m,c}$ for their group (*i.e.,* admitted under A^*_c) under $\tilde{D}_c(S_{-c})$ and $\tilde{D}_c(S'_{-c})$, respectively. As $S'_{-c} \geq S_{-c}$, we have that $D^c(S_{-c}, S_c) \subseteq D^c(S'_{-c}, S_c)$ for all $S_c \in [0,1]^{|\mathcal{M}|}$. It follows that $f'(\theta_c) \geq f(\theta_c) > 0$ for all $\theta_c = (s_c, m_c) \in [0,1] \times \mathcal{M}$. As $t'_{c,m} < t_{c,m}$, f' has full support, and $f' \geq f$, we have that the measure of admitted group m agents increases $x'_{c,m} > x_{c,m}$. But as $\sum_{k\in\mathcal{M}} x'_k = \sum_{k\in\mathcal{M}} x_k = q$, we know that there exists an $m' \in \mathcal{M}$ such that $x'_{c,m'} < x_{c,m'}$. It follows that $t'_{c,m'} > t_{c,m'}$, otherwise, if $t'_{c,m'} \leq t_{c,m'}$, then $x'_{c,m'} \geq x_{c,m'}$. So we have:

$$\begin{aligned} h_c(t'_{c,m'}) + u'_{m',c}(x'_{c,m'}) &> h_c(t_{c,m'}) + u'_{m',c}(x_{c,m'}) \\ &\geq h_c(t_{c,m}) + u'_{m,c}(x_{c,m}) > h_c(t'_{c,m}) + u'_{m,c}(x'_{c,m}) \end{aligned} \tag{43}$$

where the first inequality follows by $t_{c,m'} < t'_{c,m'}$, $x_{c,m'} > x'_{c,m'}$, concavity of u_m and strictly increasing h_c. The second inequality follows by optimality. This is because the facts that $t_{c,m} > 0$ and $t'_{c,m'} < 1$ imply that $\overline{\kappa}_m = \underline{\kappa}_{m'} = 0$ and so Equation 29 implies that:

$$h_c(t_{c,m'}) + u'_{m',c}(x_{c,m'}) - \overline{\kappa}_{m'} = h_c(t_{c,m}) + u'_{m,c}(x_{c,m}) + \underline{\kappa}_m \tag{44}$$

with $\overline{\kappa}_{m'}, \underline{\kappa}_m \geq 0$. The final inequality follows as $t'_{c,m} < t_{c,m}$ and $x'_{c,m} > x_{c,m}$. But this contradicts the optimality condition for APM (Theorem 1), which implies that $T_c \neq I^*_c$, which is a contradiction. Hence, for all c and $m \in \mathcal{M}$, $T_{c,m}$ is an increasing function. $\square$

As $T : [0,1]^{|\mathcal{M}|\times|\mathcal{C}|} \to [0,1]^{|\mathcal{M}|\times|\mathcal{C}|}$ is monotone and $[0,1]^{|\mathcal{M}|\times|\mathcal{C}|}$ is a lattice under the elementwise order $\geq$, Tarski's fixed point theorem implies that the set of stable matching cutoffs is a non-empty lattice.

Finally, we use the fact that the set of stable cutoffs is a complete lattice to argue that there is a unique cutoff consistent with stability.

Claim 10. *The stable matching cutoffs are unique.*

Proof. Assume that there are multiple stable cutoffs. As the set of stable cutoffs is a lattice, there exists a largest (S^+) and smallest (S^-) stable cutoff vector, where $S^+ \geq S^-$, with strict inequality for some $m \in \mathcal{M}$, $c \in \mathcal{C}$ as $S^+ \neq S^-$. But then, as there is full support of agent types and authority c fills the capacity under stable cutoffs S^+, it must exceed its capacity under S^-, which is a contradiction. $\square$

The combination of Lemma 2 and Claim 10 completes the proof. $\square$

References

Abdulkadiroğlu, A., Y.-K. Che, and Y. Yasuda (2015): "Expanding 'choice' in school choice," *American Economic Journal: Microeconomics*, 7(1), 1–42.

Abdulkadiroğlu, A. (2005): "College Admission with Affirmative Action," *International Journal of Game Theory*, 33, 535–549.

Abdulkadiroğlu, A., and T. Sönmez (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Ahani, N., T. Andersson, A. Martinello, A. Teytelboym, and A. C. Trapp (2021): "Placement optimization in refugee resettlement," *Operations Research*, 69(5), 1468–1486.

Angrist, J. D., P. A. Pathak, and R. A. Zárate (2019): "Choice and consequence: Assessing mismatch at chicago exam schools," Discussion paper, NBER.

Athey, S., C. Avery, and P. Zemsky (2000): "Mentoring and diversity," *American Economic Review*, 90(4), 765–786.

Aygün, O., and I. Bó (2021): "College admission with multidimensional privileges: The Brazilian affirmative action case," *American Economic Journal: Microeconomics*, 13(3), 1–28.

Aygün, O., and B. Turhan (2017): "Large-scale affirmative action in school choice: Admissions to IITs in India," *American Economic Review*, 107(5), 210–13.

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

Azevedo, E. M., and J. D. Leshno (2016): "A supply and demand framework for two-sided matching markets," *Journal of Political Economy*, 124(5), 1235–1268.

Bansak, K., S. Lee, V. Manshadi, R. Niazadeh, and E. Paulson (2024): "Dynamic Matching with Post-Allocation Service and its Application to Refugee Resettlement," *Available at SSRN*.

Ben-Akiva, M., D. McFadden, and K. Train (2019): "Foundations of stated preference elicitation: Consumer behavior and choice-based conjoint analysis," *Foundations and Trends in Econometrics*, 10(1-2), 1–144.

Çelebi, O. (2023): "Diversity Preferences, Affirmative Action and Choice Rules," *arXiv preprint arXiv:2310.14442*.

Çelebi, O., and J. P. Flynn (2022): "Priority design in centralized matching markets," *The Review of Economic Studies*, 89(3), 1245–1277.

Chan, J., and E. Eyster (2003): "Does banning affirmative action lower college student quality?," *American Economic Review*, 93(3), 858–872.

Che, Y.-K., J. Kim, and F. Kojima (2019): "Stable matching in large economies," *Econometrica*, 87(1), 65–110.

Chicago Public Schools (2022): "Admissions Policy for Magnet, Selective Enrollment and Other GoCPS Schools and Programs," *Chicago Public Schools Policy*.

——— (2024): "Board Rule 1-2: Board Administrative Rules," *Chicago Public Schools Board Rule*.

Combe, J., U. M. Dur, O. Tercieux, C. Terrier, and M. U. Ünver (2022): *Market design for distributional objectives in (re) assignment: An application to improve the distribution of teachers in schools*. MIT Department of Economics.

Delacrétaz, D., S. D. Kominers, and A. Teytelboym (2023): "Matching mechanisms for refugee resettlement," *American Economic Review*, 113(10), 2689–2717.

Dessein, W., A. Frankel, and N. Kartik (2023): "Test-Optional Admissions," *arXiv preprint arXiv:2304.07551*.

Doğan, B. (2016): "Responsive affirmative action in school choice," *Journal of Economic Theory*, 165, 69–105.

Doğan, B., and A. Erdil (2022): "Widening Access in University Admissions," *Available at SSRN 4309391*.

Dur, U., S. D. Kominers, P. A. Pathak, and T. Sönmez (2018): "Reserve design: Unintended consequences and the demise of Boston's walk zones," *Journal of Political Economy*, 126(6), 2457–2479.

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

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

Ehlers, L., I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim (2014): “School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds,” *Journal of Economic Theory*, 153, 648–683.

Ellison, G., and P. A. Pathak (2021): “The Efficiency of Race-Neutral Alternatives to Race-Based Affirmative Action: Evidence from Chicago's Exam Schools,” *American Economic Review*, 111(3), 943–75.

Erdil, A., and T. Kumano (2019): “Efficiency and stability under substitutable priorities with ties,” *Journal of Economic Theory*, 184, 104950.

Goto, M., F. Kojima, R. Kurata, A. Tamura, and M. Yokoo (2017): “Designing matching mechanisms under general distributional constraints,” *American Economic Journal: Microeconomics*, 9(2), 226–62.

Guardian (2021): “Medical schools in England offering new students £10,000 to study elswehere,” *The Guardian*.

Hafalir, I. E., M. B. Yenmez, and M. A. Yildirim (2013): “Effective Affirmative Action in School Choice,” *Theoretical Economics*, 8, 325–363.

Imamura, K. (2020): “Meritocracy versus diversity,” Discussion paper, Boston College.

Kamada, Y., and F. Kojima (2017): “Stability concepts in matching under distributional constraints,” *Journal of Economic Theory*, 168, 107–142.

——— (2018): “Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition,” *Theoretical Economics*, 13(2), 761–793.

Kojima, F. (2012): “School Choice: Impossibilities for Affirmative Action,” *Games and Economic Behavior*, 75, 685–693.

Kominers, S. D., and T. Sönmez (2016): “Matching with slot-specific priorities: Theory,” *Theoretical Economics*, 11(2), 683–710.

Kornbluth, D., and A. Kushnir (2021): “Undergraduate Course Allocation through Competitive Markets,” *Available at SSRN 3901146*.

Mailath, G. J. (1987): “Incentive compatibility in signaling games with a continuum of types,” *Econometrica: Journal of the Econometric Society*, pp. 1349–1365.

North Yorkshire Council (2025): “School admission arrangements 2026 to 2027 - Executive report January 2025,” *North Yorkshire Council*.

Pathak, P. A., and P. Shi (2021): "How well do structural demand models work? Counterfactual predictions in school choice," *Journal of Econometrics*, 222(1), 161–195.

Pathak, P. A., T. Sönmez, M. U. Unver, and M. B. Yenmez (2021): "Leaving No Ethical Value Behind: Triage Protocol Design for Pandemic Rationing," Discussion paper.

Rostek, M. J., and J. H. Yoon (2023): "Equilibrium theory of financial markets: Recent developments," Working Paper 3710206, SSRN.

Roth, A. E. (1984): "Stability and polarization of interests in job matching," *Econometrica: Journal of the Econometric Society*, pp. 47–57.

——— (1985): "The college admissions problem is not equivalent to the marriage problem," *Journal of economic Theory*, 36(2), 277–288.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-sided matching: a study in game-theoretic modeling and analysis*. Econometric Society monographs, Cambridge.

Shi, P. (2022): "Optimal priority-based allocation mechanisms," *Management Science*, 68(1), 171–188.

Sönmez, T. (2023): "Minimalist market design: A framework for economists with policy aspirations," *arXiv preprint arXiv:2401.00307*.

Sönmez, T., and M. B. Yenmez (2022a): "Affirmative action in India via vertical, horizontal, and overlapping reservations," *Econometrica*, 90(3), 1143–1176.

——— (2022b): "Constitutional Implementation of Affirmative Action Policies in India," *arXiv preprint arXiv:2203.01483*.

Temnyalov, E. (2023): "An information theory of efficient differential treatment," *American Economic Journal: Microeconomics*, 15(1), 323–358.

Tuoi Tre News (2017): "Vietnamese college hopefuls miss dream school despite acing national exam," *Tuoi Tre News*, August 4, 2017.

Tversky, A. (1964): "Finite additive structures," *Michigan Mathematical Psychology Program*.

UK Department for Education (2021): "School Admissions Code," *UK Department for Education*.

Weitzman, M. L. (1974): "Prices vs. quantities," *The Review of Economic Studies*, 41(4), 477–491.

Westkamp, A. (2013): "An analysis of the German university admissions system," *Economic Theory*, 53(3), 561–589.

Online Appendix

Adaptive Priority Mechanisms

B Other Omitted Proofs

B.1 Proof of Proposition 1

Proof. Part (i): In state ω the payoff from admitting the highest-scoring minority students of measure $x(\omega)$ is:

$$q\omega + (1+\gamma-\omega)x(\omega) - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)x(\omega)^2 \tag{45}$$

Thus, the $x(\omega)$ that solves the FOC is given by:

$$x(\omega) = \frac{\kappa(1+\gamma-\omega)}{1+\kappa\gamma\beta} \tag{46}$$

Under our maintained assumptions, we have that:

$$x(\omega) = \frac{\kappa(1+\gamma-\omega)}{1+\kappa\gamma\beta} \leq \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega}-\underline{\omega}) < \min\{\kappa, q\} \tag{47}$$

and:

$$x(\omega) = \frac{\kappa(1+\gamma-\omega)}{1+\kappa\gamma\beta} \geq \frac{1+\gamma-\overline{\omega}}{\frac{1}{\kappa}+\gamma\beta} > \kappa(1-\underline{\omega}) \geq 0 \tag{48}$$

Thus, this level of minority admissions is feasible. Substituting, we have that:

$$V^* = q\mathbb{E}[\omega] + \frac{1}{2}\frac{\mathbb{E}[\kappa(1+\gamma-\omega)^2]}{1+\kappa\gamma\beta} \tag{49}$$

Consider now the APM $A(y) = \gamma(1-\beta y)$. Agents are allocated the resource if their modified scores exceed ω, with a uniform lottery over students with score exactly ω. Thus, in state ω, this policy admits measure $y(\omega)$ minorities that solve the fixed point equation:

$$y(\omega) = \min\left\{\kappa\int_0^1 \mathbb{I}[s + A(y(\omega)) \geq \omega]\mathrm{d}s, q\right\} = \min\{\kappa\left(1-\max\{\omega - A(y(\omega)), 0\}\right), q\} \tag{50}$$

Denote the RHS of this fixed point equation by the function $\mathrm{RHS}(y,\omega)$, which is continuous and decreasing in y. Moreover, $\mathrm{RHS}(0,\omega) = \min\{\kappa(1-\max\{\omega-\gamma,0\}), q\} > 0$ and $\mathrm{RHS}(\min\{\kappa,q\},\omega) < \min\{\kappa,q\}$. The second of these inequalities is true because the condition $\underline{\omega} > \gamma(1-\beta\min\{q,\kappa\})$ follows from our assumption that $\min\{\kappa,q\} > \frac{1+\gamma-\omega}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega}-\underline{\omega})$. Thus, there exists a unique $y(\omega)$ implemented by the APM. Moreover, let $y_A(\omega)$ denote the unique solution to the equation 50, which gives the measure of admitted minority students under APM A at state ω.

$$y_A(\omega) = \kappa(1-(\omega-\gamma(1-\beta y_A(\omega)))) = \frac{\kappa(1-\omega+\gamma)}{1+\kappa\gamma\beta} \tag{51}$$

Thus, A implements the optimal level of minority admissions characterized in equation 46 and $V_A = V^*$.

Part (ii): First, if we admit all minority students over some threshold $\hat{s}$, the total score of admitted minority students is $\kappa \int_{\hat{s}}^1 s \mathrm{d}s$. Moreover, when we admit measure x minority students where $x \leq \min\{\kappa, q\}$, this admissions threshold is defined by $x = \kappa \int_{\hat{s}}^1 \mathrm{d}s = \kappa(1-\hat{s})$. Thus, we have that $\hat{s} = 1 - \frac{x}{\kappa}$. Finally, the residual measure $q - x$ admitted majority students all score ω. Thus, the total score is given by $\bar{s} = q\omega + (1-\omega)x - \frac{1}{2\kappa}x^2$ for $0 \leq x \leq \min\{\kappa, q\}$. As both quota and priority policies always admit the highest-scoring minority students, the authority's utility is given by:

$$\mathcal{U} = q\mathbb{E}[\omega] + \mathbb{E}[(1+\gamma-\omega)x] - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)\mathbb{E}[x^2] \tag{52}$$

We now derive the admitted measure of minority students. In the absence of a priority or quota policy, $\alpha = 0$ or $Q = 0$, we have that $x = \kappa(1-\omega)$ measure minority students is admitted. Thus, under a quota policy Q, measure $x = \max\{Q, \kappa(1-\omega)\}$ minority students are admitted. Under a priority policy, the measure of admitted minority students is $x = \kappa \int_{\omega-\alpha}^1 \mathrm{d}x = \kappa(1+\alpha-\omega)$. In each case x is capped by $\min\{\kappa, q\}$ and floored by 0.

The expected utility function over quotas is given by one of four cases. First, $Q > \min\{\kappa, q\}$ and:

$$\mathcal{U}_Q(Q) = q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])\min\{\kappa, q\} - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)\min\{\kappa, q\}^2 \tag{53}$$

Second, $Q \in [\kappa(1-\underline{\omega}), \min\{\kappa, q\})$ and:[32]

$$\mathcal{U}_Q(Q) = q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])Q - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)Q^2 \tag{54}$$

Third, $Q \in (\kappa(1-\overline{\omega}), \kappa(1-\underline{\omega}))$ and:

$$\begin{aligned}\mathcal{U}_Q(Q) = q\mathbb{E}[\omega] &+ \int_{1-\frac{Q}{\kappa}}^{\overline{\omega}} \left((1+\gamma-\omega)Q - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)Q^2\right)\mathrm{d}\Lambda(\omega) \\ &+ \int_{\underline{\omega}}^{1-\frac{Q}{\kappa}} \left((1+\gamma-\omega)\kappa(1-\omega) - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)(\kappa(1-\omega))^2\right)\mathrm{d}\Lambda(\omega)\end{aligned} \tag{55}$$

Finally, $Q \leq \kappa(1-\overline{\omega})$ and:

$$\mathcal{U}_Q(Q) = q\mathbb{E}[\omega] + \mathbb{E}\left[(1+\gamma-\omega)\kappa(1-\omega)\right] - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)\mathbb{E}\left[(\kappa(1-\omega))^2\right] \tag{56}$$

We claim that the optimum lies in the second case. See that in case two the strict maximum is attained at $Q^* = \frac{1+\gamma-\mathbb{E}[\omega]}{\frac{1}{\kappa}+\gamma\beta} \in (\kappa(1-\underline{\omega}), \min\{\kappa, q\})$, by our assumptions that $\min\{\kappa, q\} > \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega} - \underline{\omega})$ and $\kappa(1-\underline{\omega}) < \frac{1+\gamma-\overline{\omega}}{\frac{1}{\kappa}+\gamma\beta}$. Moreover, in case three, the first derivative of the

[32] By our maintained assumptions we have that this interval has non-empty interior.

payoff is given by:

$$\mathcal{U}'_Q(Q) = \int_{1-\frac{Q}{\kappa}}^{\overline{\omega}} \left((1+\gamma-\omega) - \left(\frac{1}{\kappa} + \gamma\beta\right) Q \right) \mathrm{d}\Lambda(\omega) \tag{57}$$

Thus, checking that the sign of this is positive amounts to verifying that for all $Q \in (\kappa(1-\overline{\omega}), \kappa(1-\underline{\omega}))$, we have that:

$$Q < \frac{1+\gamma-\mathbb{E}[\omega|\omega \geq 1-\frac{Q}{\kappa}]}{\frac{1}{\kappa}+\gamma\beta} \tag{58}$$

As the RHS is an increasing function of Q, it suffices to show that:

$$\kappa(1-\underline{\omega}) < \frac{1+\gamma-\overline{\omega}}{\frac{1}{\kappa}+\gamma\beta} \tag{59}$$

which we have assumed. Moreover, the expected utility in the first case equals $\mathcal{U}_Q(\kappa(1-\underline{\omega}))$, thus is lower than the optimum of the second case. The expected utility in the fourth case equals $\mathcal{U}_Q(\kappa(1-\overline{\omega}))$, thus is lower than the optimum of the third case. We therefore have that:

$$V_Q = q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])Q^* - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right) Q^{*2} \tag{60}$$

We now turn to characterizing the value of priorities. There are three cases to consider. First, when $\kappa(1+\alpha-\overline{\omega}) \geq \min\{\kappa, q\}$ we have that $x = \min\{\kappa, q\}$ and:

$$\mathcal{U}_P(\alpha) = q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])\min\{\kappa,q\} - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\min\{\kappa,q\}^2 \tag{61}$$

Second, when $\kappa(1+\alpha-\underline{\omega}) \geq \min\{\kappa, q\} \geq \kappa(1+\alpha-\overline{\omega})$ we have that:

$$\begin{aligned}\mathcal{U}_P(\alpha) = q\mathbb{E}[\omega] &+ \int_{\underline{\omega}}^{1+\alpha-\min\{\frac{q}{\kappa},1\}} \left((1+\gamma-\omega)\min\{\kappa,q\} - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\min\{\kappa,q\}^2 \right) \mathrm{d}\Lambda(\omega) \\ &+ \int_{1+\alpha-\min\{\frac{q}{\kappa},1\}}^{\overline{\omega}} \left((1+\gamma-\omega)\kappa(1+\alpha-\omega) - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)[\kappa(1+\alpha-\omega)]^2 \right) \mathrm{d}\Lambda(\omega)\end{aligned} \tag{62}$$

Finally, when $\min\{\kappa, q\} \geq \kappa(1+\alpha-\underline{\omega})$, we have that:

$$\mathcal{U}_P(\alpha) = q\mathbb{E}[\omega] + \mathbb{E}[(1+\gamma-\omega)\kappa(1+\alpha-\omega)] - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\mathbb{E}[(\kappa(1+\alpha-\omega))^2] \tag{63}$$

We claim that the optimum under our assumptions lies only in the third case. First, we argue that there is a unique local maximum in the third case. Second, we show the value in the second case is decreasing in α. By continuity, the unique optimum then lies in the third case.

First, it is helpful to write $\bar{x}(\alpha) = \kappa(1+\alpha-\mathbb{E}[\omega])$ and $\varepsilon = \kappa(\mathbb{E}[\omega]-\omega)$. The value in the third case can then be re-expressed as:

$$\begin{aligned}\mathcal{U}_P(\alpha) &= q\mathbb{E}[\omega] + \mathbb{E}[(1+\gamma-\omega)\left(\bar{x}(\alpha)+\varepsilon\right)] - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\mathbb{E}[(\bar{x}(\alpha)+\varepsilon)^2] \\ &= q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])\bar{x}(\alpha) - \mathbb{E}[\omega\varepsilon] - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\bar{x}(\alpha)^2 - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\mathbb{E}[\varepsilon^2]\end{aligned} \tag{64}$$

Finally, we have that $\mathbb{E}[\varepsilon^2] = \kappa^2\text{Var}[\omega]$ and $\mathbb{E}[\omega\varepsilon] = \text{Cov}[\omega,\varepsilon] = -\kappa\text{Var}[\omega]$. Thus:

$$\mathcal{U}_P(\alpha) = q\mathbb{E}[\omega] + (1+\gamma-\mathbb{E}[\omega])\bar{x}(\alpha) - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\bar{x}(\alpha)^2 + \frac{\kappa}{2}\left(1-\kappa\gamma\beta\right)\text{Var}[\omega] \tag{65}$$

We then see that the optimal α^* in this range sets $\bar{x}(\alpha^*) = Q^* < \min\{\kappa, q\}$. It remains only to check that this optimal α^* indeed lies within this case, or equivalently that $\kappa(1+\alpha^*-\underline{\omega}) \leq \min\{\kappa, q\}$. To this end, see that $\kappa(1+\alpha^*-\mathbb{E}[\omega]) = Q^*$, and:

$$\begin{aligned}\kappa(1+\alpha^*-\underline{\omega}) &= Q^* + \kappa(\mathbb{E}[\omega]-\underline{\omega}) \leq Q^* + \kappa(\overline{\omega}-\underline{\omega}) \\ &\leq \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega}-\underline{\omega}) < \min\{\kappa, q\}\end{aligned} \tag{66}$$

where the final inequality follows by our assumption that $\min\{\kappa, q\} > \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega}-\underline{\omega})$.

Second, in the second case we have that the first derivative of the payoff in α is given by:

$$\begin{aligned}\mathcal{U}_P'(\alpha) &= \int_{1+\alpha-\min\{\frac{q}{\kappa},1\}}^{\overline{\omega}} \frac{d}{d\alpha}\left((1+\gamma-\omega)\kappa(1+\alpha-\omega) - \frac{1}{2}\left(\frac{1}{\kappa}+\gamma\beta\right)\left[\kappa(1+\alpha-\omega)\right]^2\right)\mathrm{d}\Lambda(\omega) \\ &= \kappa\int_{1+\alpha-\min\{\frac{q}{\kappa},1\}}^{\overline{\omega}}\left((1+\gamma-\omega) - \left(\frac{1}{\kappa}+\gamma\beta\right)\left(\bar{x}(\alpha)+\varepsilon(\omega)\right)\right)\mathrm{d}\Lambda(\omega)\end{aligned} \tag{67}$$

Checking that the sign of this is negative for all α such that $\kappa(1+\alpha-\underline{\omega}) \geq \min\{\kappa, q\} \geq \kappa(1+\alpha-\overline{\omega})$ then amounts to checking that:

$$\bar{x}(\alpha) > \frac{1+\gamma-\mathbb{E}[\omega|\omega \geq 1+\alpha-\min\{\frac{q}{\kappa},1\}]}{\frac{1}{\kappa}+\gamma\beta} - \mathbb{E}\left[\varepsilon(\omega)|\omega \geq 1+\alpha-\min\{\frac{q}{\kappa},1\}\right] \tag{68}$$

for all $\bar{x}(\alpha) \in [\min\{\kappa, q\} - \kappa(\mathbb{E}[\omega]-\underline{\omega}), \min\{\kappa, q\} - \kappa(\mathbb{E}[\omega]-\overline{\omega})]$. So it suffices to check that the minimal possible value of the LHS exceeds the maximal possible value of the RHS. A sufficient condition for this is that:

$$\min\{\kappa, q\} - \kappa(\mathbb{E}[\omega]-\underline{\omega}) > \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} - \kappa(\mathbb{E}[\omega]-\overline{\omega}) \tag{69}$$

Which holds as we assumed that $\min\{\kappa, q\} > \frac{1+\gamma-\underline{\omega}}{\frac{1}{\kappa}+\gamma\beta} + \kappa(\overline{\omega}-\underline{\omega})$. Substituting the optimal priority policy $\bar{x}(\alpha) = Q^*$ in equation 63, we obtain

$$V_P = q\mathbb{E}[\omega] + (1 + \gamma - \mathbb{E}[\omega])Q^* - \frac{1}{2}\left(\frac{1}{\kappa} + \gamma\beta\right)Q^{*2} + \frac{\kappa}{2}(1 - \kappa\gamma\beta)\,\text{Var}[\omega] \tag{70}$$

We have now established that:

$$\Delta = V_P - V_Q = \frac{\kappa}{2}(1 - \kappa\gamma\beta)\,\text{Var}[\omega] \tag{71}$$

Part (iii): We have V^*, V_Q, V_P. Thus, we can compute the loss from restricting to quota policies:

$$\mathcal{L}_Q = \frac{1}{2}\frac{\kappa\text{Var}[\omega]}{1 + \kappa\gamma\beta} \tag{72}$$

To find the loss from restricting to priority policies, we compute:

$$\mathcal{L}_P = \mathcal{L}_Q - \Delta = \frac{1}{2}(\kappa\gamma\beta)^2\frac{\kappa\text{Var}[\omega]}{1 + \kappa\gamma\beta} \tag{73}$$

Enveloping over these losses yields the claimed formula. □

B.2 Proof of Proposition 2

Proof. Adapting Definition 5 to the single-object setting, we say that a matching μ admits a cutoff structure if there exists $S(\omega) = \{S_m(\omega)\}_{m\in\mathcal{M}}$ such that $\mu(s, m; \omega) = 1$ if $s > S_m(\omega)$ and $\mu(s, m; \omega) = 0$ if $s < S_m(\omega)$. A mechanism admits a cutoff structure if it admits a cutoff structure at every ω. We will first prove that any monotone APM admits a cutoff structure.

Lemma 3. *A monotone APM admits a cutoff structure.*

Proof. For a contradiction, assume it does not. Then there exists ω and matching μ implemented by the monotone APM such that for some $m \in \mathcal{M}$, $s > s'$, $\mu(s, m; \omega) = 0$ but $\mu(s', m; \omega) = 1$. Let x_m denote the measure of group m agents allocated the resource at μ. Since A is a monotone APM and $s > s'$, we have that $A(x_m, s) > A(x_m, s')$, which contradicts that A implements μ. □

We now use Lemma 3 to show that a monotone APM implements an essentially unique allocation. Assume for a contradiction that A is monotone and implements two different allocations in some state $\omega \in \Omega$, μ and μ', such that for some set of agents $\hat{\Theta}$ with positive measure (*i.e.*, $F_\omega(\hat{\Theta}) > 0$), we have $\mu(\theta; \omega) = 1$ and $\mu'(\theta; \omega) = 0$ for all $\theta \in \hat{\Theta}$. Let x_m and x'_m denote the measure of type m agents assigned the resource at μ and μ'. For notational simplicity, we suppress the dependence on ω for the remainder of the proof. Define:

$$\begin{aligned} s_m &= \inf\{s : \mu(s, m) = 1\} \\ s'_m &= \inf\{s : \mu'(s, m) = 1\} \end{aligned} \tag{74}$$

As the matchings μ and μ' differ by a positive measure and both are cut-off matchings that are implemented by an APM, there exists m and n such that $x_m > x'_m$ and $x'_n > x_n$. Since $x_m > x'_m$ and both matchings are cut-off matchings, it follows that $s'_m > s_m$. Likewise, as $x'_n > x_n$ and both matchings are cut-off matchings, $s_n > s'_n$. Note that these imply that, there exists $\hat{s}_m \in (s_m, s'_m)$ and $\hat{s}_n \in (s'_n, s_n)$ such that $\mu(\hat{s}_m, m) = 1$, $\mu'(\hat{s}_m, m) = 0$, $\mu(\hat{s}_n, n) = 0$ and $\mu'(\hat{s}_n, n) = 1$. Thus, the following inequalities hold:

$$A_m(\hat{s}_m, x_m) > A_n(\hat{s}_n, x_n) \geq A_n(\hat{s}_n, x'_n) > A_m(\hat{s}_m, x'_m) \geq A_m(\hat{s}_m, x_m) \tag{75}$$

where the first inequality follows as $\mu(\hat{s}_m, m) = 1$ while $\mu(\hat{s}_n, n) = 0$, the second inequality follows as $x_n < x'_n$ and A is monotone, the third inequality follows as $\mu'(\hat{s}_n, n) = 1$ and $\mu'(\hat{s}_m, m) = 0$ and the fourth inequality follows as $x'_m < x_m$ and A is monotone. This equation yields $A_m(\hat{s}_m, x_m) > A_m(\hat{s}_m, x_m)$, which is a contradiction, proving that all allocations implemented by A are essentially the same. □

B.3 Proof of Proposition 3

Proof. Let x^*_m denote the measure of group m agents in the optimal allocation, with $x^* = \{x^*_m\}_{m \in \mathcal{M}}$. A priority policy $P(s, m) = h^{-1}(h(s) + u'_m(x^*_m)) = A_m(x^*_m, s)$ implements the same allocation as the optimal adaptive priority mechanism and by Theorem 1, is optimal. A quota mechanism with (Q, D) where $Q_m = x^*_m$ implements x^* for all D, and is therefore optimal. □

B.4 Proof of Proposition 4

Proof. If ϕ is equivalent to A^*_c, Claim 8 implies that ϕ is consistent with stability.

We prove consistency with stability implies that ϕ is equivalent to A^*_c by the contrapositive. To this end, suppose that ϕ is not equivalent to A^*_c. It follows that there exists a full-support density $\{\tilde{f}(s_c, m)\}_{s_c \in [0,1], m \in \mathcal{M}}$ such that ϕ yields a different allocation than A^*_c under $\tilde{f}$. The rest of the proof constructs a full-support measure F with unique stable matching μ_F such that $\tilde{f}$ is the induced density of scores and groups of the agents who demand authority c at μ_F. Given such an F, we will have that ϕ cannot be consistent with stability as it yields a different allocation than A^*_c, which itself yields $\mu_F(c)$, the set of agents c is matched to in the unique stable matching.

We first define some notation. Given a density f, for any set of types $\check{\Theta} \subseteq \Theta$, we define the marginal density of agents with score $s_c \in [0, 1]$ at authority c in group $m \in \mathcal{M}$ as:

$$f_{\text{marg}(\check{\Theta})}(s_c, m) = \int_{\check{\Theta}} \mathbb{I}[s_c(\theta) = s_c, m(\theta) = m] \, \mathrm{d}F(\theta) \tag{76}$$

To construct such an F, we proceed in three steps. First, take a full-support density f^0 that satisfies the following two conditions: i) Define $\hat{S}_c \in [0, 1]^{|\mathcal{M}|}$ as the cutoff vector that obtains by applying A^*_c to $\tilde{f}$.[33] We assume that f^0 is such that authority c's cutoff vector that is consistent with the unique stable matching, μ_{F_0}, coincides with $\hat{S}_c$; ii) for all $m \in \mathcal{M}$

[33] Which exists as any monotone APM admits a cutoff structure (Lemma 2) and the optimal APM is monotone (Theorem 1).

and $s_c < \hat{S}_{m,c}$, $f^0_{\text{marg}(\Theta)}(s,m) < \tilde{f}(s,m)$; and iii) all authorities have strictly positive cutoffs for all groups at the unique stable matching.

Second, transform f^0 into a new density f^1 that differs from f^0 on the set of types that is matched with c under μ_{F_0}, which we call Θ_c.[34] We define the scaling factor $\iota^1(s_c,m)$ as:

$$\iota^1(s_c,m) = \frac{\tilde{f}(s_c,m)}{f^0_{\text{marg}(\Theta_c)}(s_c,m)} \tag{77}$$

Moreover, we define:

$$f^1(\theta) = \begin{cases} f^0(\theta)\iota^1(s_c(\theta),m(\theta)) & \text{if} \quad \theta \in \Theta_c, \\ f^0(\theta) & \text{otherwise.} \end{cases} \tag{78}$$

This changes the scores of the types who are allocated to c under μ_{F^0} but does not change their total measure, their composition, or their scores at any other authority. Thus, the unique stable matching under f^1, μ_{F^1}, coincides with μ_{F^0}. Moreover, by assumption i) of step 1, we have that $f^1_{\text{marg}(\Theta_c)}(s_c,m) = \tilde{f}(s_c,m)$ for all m and $s_c \geq \hat{S}_{m,c}$.

Third, transform f^1 into a new density f^2 that differs on the set of unmatched agents under f^0 (and also therefore f^1 by step 2), $\tilde{\Theta}$, and define the set of types who strictly prefer c to their assignment under μ_{F^0} (and also therefore μ_{F^1} by step 2), $\hat{\Theta}_c$.[35] We define a new scaling factor $\iota^2(s_c,m)$ as:

$$\iota^2(s_c,m) = \frac{\tilde{f}(s_c,m) - f_{\text{marg}(\hat{\Theta}_c)}(s_c,m)}{f_{\text{marg}(\tilde{\Theta})}(s_c,m)} \tag{79}$$

which is strictly positive by assumption ii) of step 1. We then define f^2 as:

$$f^2(\theta) = \begin{cases} f^1(\theta)(1+\iota^2(s_c(\theta),m(\theta))) & \text{if} \quad \theta \in \tilde{\Theta}, \\ f^1(\theta) & \text{otherwise.} \end{cases} \tag{80}$$

By construction, $f^2_{\text{marg}(\hat{\Theta}_c)}(s_c,m) = \tilde{f}(s_c,m)$ for all m and $s_c < \hat{S}_{m,c}$. Moreover, $\mu_{F^2} = \mu_{F^1} = \mu_{F^0}$ as all $\theta \in \tilde{\Theta}$ remain unmatched.

We have now constructed a full-support density f^2 with unique stable matching μ_{F^2} (by Theorem 3) such that the density over $D_c(\mu_{F^2})$ coincides with $\tilde{f}$. Moreover, by Claim 8, A^*_c selects $\mu_{F^2}(c)$ from $D_c(\mu_{F^2})$. As ϕ selects a different allocation from $D_c(\mu_{F^2})$ (as it has density of types $\tilde{f}$), it is inconsistent with stability. $\square$

[34]Formally, $\Theta_c = \{\theta : \theta \in D_c(\mu_{F^0}), s_c(\theta) \geq \hat{S}_{m(\theta),c)}\}$.

[35]Formally, $\tilde{\Theta} = \{\theta : \theta \in D_c(\mu_{F^1}), s_c(\theta) < \hat{S}_{m(\theta),c}, s_{c'}(\theta) < S^{\mu_{F1}}_{m(\theta),c'}$ for all $c' \neq c\}$, where $S^{\mu_{F1}}_{m,c'}$ denotes the group m cutoff at authority c' at the stable matching μ_{F^1}, which is strictly positive by assumption iii) of step 2. Moreover, $\hat{\Theta}_c = \{\theta : \theta \in D_c(\mu_{F^1}), s_c(\theta) < \hat{S}_{m(\theta),c}\}$.

B.5 Proof of Lemma 1

Proof. Fix an arbitrary type $\theta \in \Theta$. Suppose that $DA^t(Q^0)(\theta)$ does not converge to some $r \in \mathcal{R}$. As $\mathcal{R}$ is a finite set, this implies that there exist three time indices $t_0, t_1, t_2 \in \mathbb{N}$ such that $t_0 < t_1 < t_2$ with the property that $DA^{t_0}(Q_0)(\theta) = DA^{t_2}(Q_0)(\theta) \neq DA^{t_1}(Q_0)(\theta)$. Moreover, we know that $DA^{t_1}(Q_0)$ must delete at least the top-ranked element of $DA^{t_0}(Q_0)(\theta)$. As DA only deletes elements, it cannot be true that $DA^{t_0}(Q_0)(\theta) = DA^{t_2}(Q_0)(\theta)$. □

B.6 Proof of Theorem 4

Let $T^t_{m,c}$ denote the step t cut-offs for group m agents at authority c at iteration t of the mapping T (see the proof of Theorem 3 for the formal definition) starting with the market-clearing cutoff being 0 at all authorities. Formally, $T^t(S) = T \circ T^{t-1}(S)$ with initial condition $T^1 = T(0)$. Similarly, let $H^t_{m,c}$ denote the step t cut-offs for group m agents at authority c at step t of the Deferred Acceptance procedure when Ch_c is the optimal APM for all c. Formally, $H^t_{m,c} = \inf_{\theta \in \{\tilde{\theta} \in Ch_c(\mathcal{T}_c(Q^t)) : m(\tilde{\theta}) = m\}} s_c(\theta)$.

It is immediate that $T^1_{m,c} = H^1_{m,c}$ from the definitions and this forms the base case for our proof by induction that $H^t \leq T^t$ for all $t \in \mathbb{N}$. To this end, we now prove the inductive step:

Claim 11. *Suppose that $H^{t'}_{m,c} \leq T^{t'}_{m,c}$ for all m and c for all $t' \leq t$. Then $H^{t+1}_{m,c} \leq T^{t+1}_{m,c}$.*

Proof. Let $\Theta^{t+1}_{m,c}$ denote the set of agents admitted to c at iteration $t+1$ of the mapping T. Formally

$$\Theta^{t+1}_{m,c} = \{\theta \in D_c(T^t) \text{ and } s_c(\theta) \geq T^{t+1}_{m(\theta),c}\} \tag{81}$$

We now show that no authority receives an application under DA from an agent that is preferable to the agents they have admitted at step $t+1$ of the iterative algorithm T.

Claim 12. *If θ applies to c at step $t+1$ of DA and $s_c(\theta) \geq T^{t+1}_{m(\theta),c}$, then $\theta \in \Theta^{t+1}_{m,c}$.*

Proof. Suppose for a contradiction that θ applies to c at step $t+1$ of DA and $s_c(\theta) \geq T^{t+1}_{m(\theta),c}$ but $\theta \notin \Theta^{t+1}_{m,c}$. By definition, we have that $\theta \notin D_c(T^t)$. Let c' denote the authority such that $\theta \in D_{c'}(T^t)$. By the definition of D, we then have that $s_{c'}(\theta) \geq T^t_{m,c'}$. As T is monotone, $s_c(\theta) \geq T^{t+1}_{m(\theta),c} \geq T^t_{m(\theta),c}$. Thus, both $s_c(\theta) \geq T^t_{m(\theta),c}$ and $s_{c'}(\theta) \geq T^t_{m,c'}$ and both c and c' are attainable for θ at iteration t of T. As θ chooses c' over c, we know that $c' \succ_\theta c$. By the inductive hypothesis, we have that for all $t' \leq t$, $H^{t'}_{m(\theta),c'} \leq T^t_{m(\theta),c'} \leq s_{c'}(\theta)$, and so θ is not rejected from c' in any previous step of DA. Thus, c' is not deleted from the preference list of θ and θ applies to c' or a more preferred authority at step $t+1$. Thus, θ does not apply to c at step $t+1$ of DA, which is a contradiction. □

Now suppose for a contradiction $H^{t+1}_{m,c} > T^{t+1}_{m,c}$ for some m and c. By monotonicity of T and the inductive hypothesis, we have that $H^{t+1}_{m,c} > T^{t+1}_{m,c} \geq T^t_{m,c} \geq H^t_{m,c}$ and $H^{t+1}_{m,c} > 0$. As F has full support, the set of group m agents that rank c first and have scores in $(H^t_{m,c}, H^{t+1}_{m,c})$ has strictly positive measure. Moreover, all such agents are tentatively admitted at c at step t and rejected from c at step $t+1$. As some positive measure of agents are rejected at step $t+1$, c must fill its capacity at step $t+1$.

Let $\hat{x}^{t+1}_{m,c}$ denote the measure of group m agents admitted to c at step $t+1$ of DA and $\tilde{x}^{t+1}_{m,c}$ denote the measure of agents in $\tilde{D}_c(T^t_{-c})$ with $s_c(\theta) \geq T^{t+1}_{m,c}$. Combining the contradictive hypothesis that $H^{t+1}_{m,c} > T^{t+1}_{m,c}$ with Claim 12, we obtain that $\hat{x}^{t+1}_{m,c} < \tilde{x}^{t+1}_{m,c}$, where the strict inequality follows from the full-support assumption. As c fills its capacity at step $t+1$ of DA, there must be m' such that $\hat{x}^{t+1}_{m',c} > \tilde{x}^{t+1}_{m',c}$.

Moreover, Claim 12 and $\hat{x}^{t+1}_{m',c} > \tilde{x}^{t+1}_{m',c}$ imply that there exists θ' admitted to c at step $t+1$ of DA such that $m(\theta') = m'$ and $s_c(\theta') < T^{t+1}_{m',c}$. Further, as $H^{t+1}_{m,c} > H^t_{m,c}$ and $H^{t+1}_{m,c} > T^{t+1}_{m,c}$, there exists θ that applies to and is rejected by c at step $t+1$ of DA such that $m(\theta) = m$ and $s_c(\theta) \geq T^{t+1}_{m,c}$. Then we must have

$$A^*_{m',c}(\tilde{x}^{t+1}_{m',c}, T^{t+1}_{m',c}) \geq A^*_{m',c}(\hat{x}^{t+1}_{m',c}, s_c(\theta')) > A^*_{m,c}(\hat{x}^{t+1}_{m,c}, s_c(\theta)) \geq A^*_{m,c}(\tilde{x}^{t+1}_{m,c}, T^{t+1}_{m,c}) \tag{82}$$

where the first and third inequalities holds by monotonicity of A^* and second holds as θ' is accepted while θ is rejected at step $t+1$ of DA where c was using the optimal APM. Thus, by full support, there exists a strictly positive measure of types with $m(\theta) = m'$ who have adaptive priorities $A^*_{m',c}(\tilde{x}^{t+1}_{m'c}, s(\theta)) \in (A^*_{m,c}(\tilde{x}^{t+1}_{m,c}, T^{t+1}_{m,c}), A^*_{m',c}(\tilde{x}^{t+1}_{m',c}, T^{t+1}_{m',c}))$ and who are not assigned to c at step $t+1$ of the iteration of T. This is a contradiction as these agents must be admitted to c under the optimal APM. □

Having shown that $H^t \leq T^t$ for all $t \in \mathbb{N}$, we now show that H^t is increasing:

Claim 13. *$H^t_{m,c}$ is increasing for all m and c.*

Proof. If c does not fill its capacity at time t, then as F has full support, $H^t_{m,c} = H^{t-1}_{m,c} = \ldots = H^1_{m,c} = 0$ and so $H^t_{m,c} \geq H^{t-1}_{m,c}$. Suppose now that c fills its capacity at time t.

Let $x^t_{m,c}$ denote the measure of agents from group m that are accepted by c at step t of DA. First, if $x^t_{m,c} = x^{t+1}_{m,c}$ for all m, then as all accepted agents reapply in $t+1$ and under the optimal APM, Ch_c chooses higher scoring agents before lower scoring ones, $H^{t+1}_{m,c} \geq H^t_{m,c}$.

Second, suppose that $x^t_{m',c} \neq x^{t+1}_{m',c}$ for some m'. As c fills its capacity, there exists m such that $x^{t+1}_{m,c} < x^t_{m,c}$. Moreover, as all agents admitted in step t reapply in step $t+1$ of DA, we have that $H^{t+1}_{m,c} > H^t_{m,c}$. Suppose for a contradiction there is an m'' such that $H^{t+1}_{m'',c} < H^t_{m'',c}$. As all agents admitted in step t reapply in step $t+1$, we have $x^{t+1}_{m'',c} \geq x^t_{m'',c}$. However, this contradicts that Ch_c is the optimal APM as

$$A^*_{m,c}(x^{t+1}_{m,c}, H^{t+1}_{m,c}) > A^*_{m,c}(x^t_{m,c}, H^t_{m,c}) = A^*_{m'',c}(x^t_{m'',c}, H^t_{m'',c}) > A^*_{m'',c}(x^{t+1}_{m'',c}, H^{t+1}_{m'',c}) \tag{83}$$

where the first and third inequalities hold as A^* is a monotone APM and the second equality holds the adaptive priority of the cutoff agent from each group must be equal by the definition of an adaptive priority mechanism. □

Let $T^* = \lim_{t\to\infty} T^t$ denote the cut-offs in the unique stable matching, μ. Claims 11 and 13 imply that $H^t \to H^{DA}$ with $H^{DA} \leq T^*$. Moreover, as $H^t \leq T^*$ for all t, no agent is rejected from their match under μ at any step of DA, which implies that all agents are matched to a weakly more preferred authority under the DA outcome μ^{DA} compared to μ. Let $x^{DA}_{m,c}$ denote the measure of group m agents matched to c at μ^{DA}.

Claim 14. $\mu^{DA} = \mu$.

Proof. Suppose that some θ is rejected from a school c at any step $t' < \hat{t}$. Then $A^*_{m,c}(x^t_{m,c}, H^t_{m,c})$ is a weakly increasing function of t for all $t \geq \hat{t}$. This is because at each step, all accepted agents reapply (meaning that there exists at least one group m with $x^{t+1}_{m,c} \leq x^t_{m,c}$) and $H^t_{m,c}$ is increasing. Moreover, due to the full support assumption, if an agent of group m is rejected from c at any step $t' < \hat{t}$, for all $m' \in \mathcal{M}$ and $t \geq \hat{t}$, we have

$$A^*_{m',c}(x^t_{m',c}, H^t_{m',c}) \geq A^*_{m,c}(x^t_{m,c}, H^t_{m,c}) \tag{84}$$

with equality whenever $H^t_{m',c} > 0$.

Suppose toward a contradiction, $\mu^{DA} \neq \mu$, which implies that μ^{DA} is not stable. Then there exists a school c, and a positive measure set of agents $\hat{\Theta} \in D_c(\mu^{DA})$ such that $\hat{\Theta} \notin \mu^{DA}(c)$ and $\hat{\Theta}$ is selected by the optimal APM from the demand set $D_c(\mu^{DA})$ under F. As a result, there exists types θ and θ' such that $\theta \in D_c(\mu^{DA})$, $\mu^{DA}(\theta) \neq c$, $m(\theta) = m$, $\mu^{DA}(\theta') = c$, $m(\theta') = m'$ and

$$A^*_{m,c}(x^{DA}_{m,c}, s_c(\theta)) > A^*_{m',c}(x^{DA}_{m',c}, s_c(\theta')) \tag{85}$$

as otherwise, $\mu^{DA}(c)$ would have been selected from $D_c(\mu^{DA})$. Moreover, as $\theta \in D_c(\mu^{DA})$ and $\mu^{DA}(\theta) \neq c$, θ must be rejected from c at some earlier step t. Then

$$A^*_{m',c}(x^{DA}_{m',c}, s_c(\theta')) \geq A^*_{m',c}(x^{DA}_{m',c}, H^{DA}_{m',c}) \geq A^*_{m,c}(x^{DA}_{m,c}, H^{DA}_{m,c}) > A^*_{m,c}(x^{DA}_{m,c}, s_c(\theta)) \tag{86}$$

where the first inequality holds as $\theta' \in \mu^{DA}(c)$, the second holds as θ is previously rejected by c (by Equation 84), and the third holds as $\theta \in D_c(\mu^{DA})$ and $\theta \notin \mu^{DA}_c$ (which implies $H^{DA}_{m,c} > s_c(\theta)$). Equations 85 and 86 contradict each other, and therefore $\mu^{DA} = \mu$. $\square$

C Details and Robustness of the Quantitative Exercise

In this appendix, we provide omitted details on our estimation strategy and the three robustness exercises we highlighted in the main text.

C.1 Estimation Details

In this appendix, we describe in full detail how we estimate the preference parameters (β, γ). We index reserve mechanisms by the reserve sizes of the four socioeconomic tiers $r = (r_1, r_2, r_3, r_4)$. We let $\bar{s}(r, y)$ and $x(r, y)$ denote the average scores and tier percentages that would be obtained in year y, with distribution F_y, under reserve policy r. The payoff of the policymaker under reserve policy r is given by $\Xi(r, \Lambda; \beta, \gamma)$, as per Equation 7:

$$\Xi(r, \Lambda; \beta, \gamma) = \mathbb{E}_\Lambda[\xi(\bar{s}(r, y), x(r, y); \beta, \gamma)] \tag{87}$$

Define the expected marginal benefit of increasing reserve i and decreasing reserve j as:

$$G_{ij}(r, \Lambda; \beta, \gamma) = \frac{\partial}{\partial r_i}\Xi(r, \Lambda; \beta, \gamma) - \frac{\partial}{\partial r_j}\Xi(r, \Lambda; \beta, \gamma) \tag{88}$$

Any (interior) optimal reserve policy r^* must equate the expected marginal benefit of increasing reserve i and decreasing reserve j at r^* to zero for all (i, j) pairs, *i.e.,* $G_{ij}(r^*, \Lambda; \beta, \gamma) = 0$ for all $\{i, j\} \subset \{1, 2, 3, 4\}$ such that $j > i$. These six first-order conditions yield six moments.

We take empirical analogs of the theoretical moments and estimate preference parameters by minimizing the sum of squared deviations of these moments from zero. We take CPS' pursued reserve policy from 2013 to 2017 as optimal, $\hat{r}^* = (0.175, 0.175, 0.175, 0.175)$. We estimate the empirical joint distribution of students' scores and tiers in CPS in each year $\hat{F}_y$ for $y \in \{2013, 2014, 2015, 2016, 2017\}$. We take the distribution over distributions by setting $\hat{\Lambda}$ as a distribution that places equal probability on each of the measured distributions

$$\mathbb{P}_{\hat{\Lambda}}[F = \tilde{F}] = \begin{cases} \frac{1}{5}, & \text{if } \tilde{F} = \hat{F}_y \text{ for } y \in \{2013, 2014, 2015, 2016, 2017\} \\ 0, & \text{otherwise.} \end{cases} \tag{89}$$

We plug these sample estimates into the theoretical moment functions. This yields six empirical moment functions, $G_{ij}(\hat{r}^*, \hat{\Lambda}; \beta, \gamma)$, that depend only on the preference parameters. Motivated by the theoretical necessity of $G_{ij}(r^*, \Lambda; \beta, \gamma) = 0$, we estimate the preference parameters by minimizing the sum of squared deviations of the empirical moments from zero:

$$(\beta^*, \gamma^*) \in \arg\min_{\beta, \gamma} \sum_{i=1}^{4} \sum_{j>i} G_{ij}(\hat{r}^*, \hat{\Lambda}; \beta, \gamma)^2 \tag{90}$$

C.2 Robustness

In this appendix, we document three robustness exercises that probe how the gains from switching to the optimal APM in CPS change for: (i) alternative estimates of the distribution of distributions; (ii) alternative preferences specifications; (iii) an alternative estimation assumption that CPS must keep the same size of reserves for each tier.

Robustness to the Distribution. Our baseline analysis took the beliefs of CPS to be the true empirical distribution of student distributions over years. To test robustness to this assumption, we take $\hat{\Lambda}$ as a Dirac distribution on the realized distribution for each of the five years of our data and re-estimate the preference parameters. In Figure 6, we plot the difference in welfare under the optimal APM and CPS reserve policies over the full range of these re-estimated parameters (*i.e.,* we take the minimum and maximum of the estimated parameters across years as the ranges for the axes). We find that the gains from APM range from 2.0 to 3.5, while our baseline estimate was 2.1. Thus, the point estimate of our welfare gains from APM appears to be conservative by this metric.

Estimation with Homogeneous Reserves. As we have motivated, in this section we estimate an alternative model, where CPS chooses a single reserve size, r, instead of separate reserve sizes for all tiers. Formally, we replace the vector of reserve sizes of the four socioeconomic tiers, $r = (r_1, r_2, r_3, r_4)$ by $r = (r, r, r, r)$. In this setting, we define the marginal benefit of increasing reserve size as

$$G(r, \Lambda; \beta, \gamma) = \frac{\partial}{\partial r} \Xi(r, \Lambda; \beta, \gamma) \tag{91}$$

Figure 6: Robustness of the Gains from APM to the Empirical Distribution of Distributions

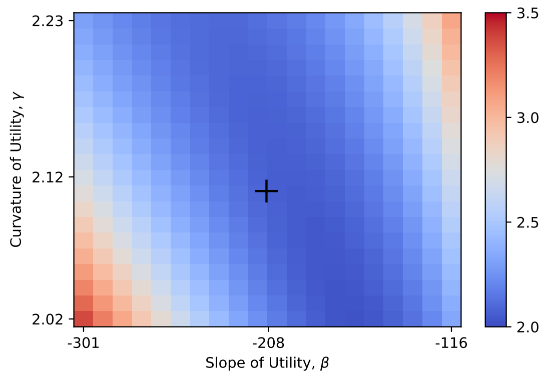

Note: This chart plots the difference in empirical payoffs from the optimal APM and CPS reserve policy under alternative parameter values, with the shaded colors corresponding to the numerical value of the gains from APM, ranging from 2.0 to 3.5. The black '+' indicates our baseline parameter values. The ranges for the axes are obtained by estimating β and γ separately for each year of our data and separately taking the minimum and maximum estimated values of each set of estimated parameters.

Figure 7: Estimated Slope of Utility Under Homogeneous Reserves

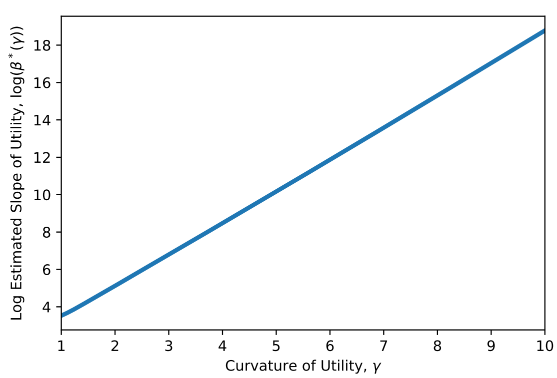

Note: This graph plots the estimated logarithm of the slope of utility $\log \beta^*(\gamma)$ in the homogeneous reserve case as we vary the curvature of utility $\gamma \in [1, 10]$.

Figure 8: Payoff Gains from APM Under Homogeneous Reserves

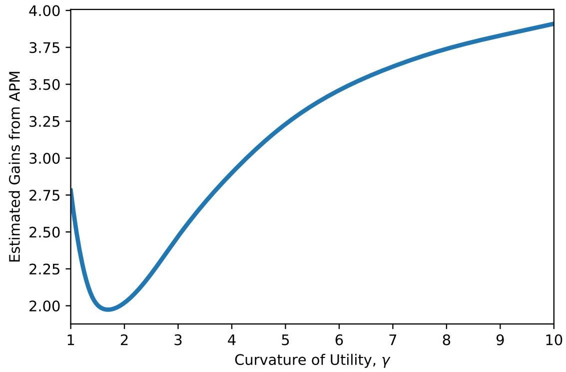

Note: This graph plots the estimated difference in payoffs in the homogeneous reserves case between the optimal APM and the CPS policy as we vary the curvature of utility $\gamma \in [1, 10]$.

As in the general model, any (interior) reserve policy r^* must satisfy $G(\hat{r}^*, \hat{\Lambda}; \beta, \gamma) = 0$. This first-order condition yields one moment, and so we can estimate one parameter. To this end, we fix γ, and for each $\gamma \in [1, 10]$, and we estimate $\beta^*(\gamma)$ as the exact solution to the following empirical moment condition:

$$G(\hat{r}^*, \hat{\Lambda}; \beta^*(\gamma), \gamma) = 0 \tag{92}$$

Figure 7 plots the logarithm of the estimated $\beta^*(\gamma)$. The estimated $\beta^*(\gamma)$ is increasing in γ. As the loss term $|x_t - 0.25|$ is in $(0, 1)$, $\beta^*(\gamma)$ is increasing and convex in γ, where $\beta^*(1) = 34$ and $\beta^*(10) = 1.436 \times 10^8$. In Figure 8, we plot the gains as a function of γ, which shows that even though the estimated value for β moves quite a lot, the empirical gains range from 2 to 4 points. This also shows that the estimated gain from APM of 2.1 under our benchmark specification is close to the lower bound of the estimated gains under the alternative specification with homogeneous reserves.

Gains from APM Under Different Utility Functions. In this section, as we have motivated, we estimate alternative objective functions to investigate the robustness of our findings. First, we analyze a setting that includes a loss term only for underrepresented tiers (and does not penalize overrepresentation of any tier). To this end, we replace the term $|0.25 - x_t|$ with $\min\{0, (0.25 - x_t)\}$ and perform the same estimation with the following parametric utility function:

$$\xi(\bar{s}, x; \beta, \gamma) = \bar{s} + \beta \sum_{t=1}^{4} (\min\{0, (0.25 - x_t)\})^{\gamma} \tag{93}$$

The estimated parameter values are $\beta^* = -52058$ and $\gamma^* = 3.87467$. We compute the difference between the empirical payoffs under APM and the CPS reserve policy to be 0.262, which is significantly lower than our estimate of 2.1. However, the reason for this is that the diversity domain is estimated to be less important under this specification, and the diversity loss under the CPS policy is 2.71. Thus, improvements from APM correspond to 9.6% of the loss from underrepresentation, which is attenuated relative to our baseline specification, but remains non-negligible.

Second, we allow CPS to care differentially about underrepresentation and overrepresentation by considering a utility function with separate coefficients for underrepresented and overrepresented tiers. To this end, we define the following loss function:

$$f(x_t, \beta_l, \beta_h, \gamma) = \begin{cases} \beta_l (0.25 - x_t)^\gamma & \text{if } x_t \leq 0.25 \\ \beta_h (x_t - 0.25)^\gamma & \text{if } x_t > 0.25 \end{cases} \tag{94}$$

where β_l indexes the loss from underrepresentation of a tier, while β_h indexes the loss from overrepresentation. We then perform the same estimation with:

$$\xi(\bar{s}, x; \beta, \gamma) = \bar{s} + \sum_{t=1}^{4} f(x_t, \beta_l, \beta_h, \gamma) \tag{95}$$

This yields the following estimated values: $\beta_l^* = -1362270$, $\beta_h^* = -12278$, $\gamma^* = 5.28021$. We compute the difference between the empirical payoffs under APM and the CPS reserve policy to be 0.195 and the loss from underrepresentation under the CPS policy to be 2.24. Thus, we conclude that improvement from APM corresponds to 8.7% of loss from underrepresentation under the CPS policy, which is similar to what we obtain under the specification in which there is no loss from overrepresentation.

D Extension to More General Authority Preferences

In this Appendix, we first relax the assumptions underlying Equation 6 to allow for non-separable diversity preferences, non-concave diversity preferences, and non-separable score and diversity preferences. We show how these changes in assumptions lead to certain modified APM becoming first-best optimal.

D.1 Non-Separable Diversity Preferences

First, we relax Equation 6 and instead suppose that the authority's preferences satisfy the following assumption:

Assumption 1. *The authority's utility function can be represented as:*

$$\xi\left(\bar{s}_h, x\right) \equiv g\left(\bar{s}_h + u(x)\right) \tag{96}$$

for some continuous, strictly increasing function $g : \mathbb{R} \to \mathbb{R}$ and a concave, partially differentiable u in each argument.

In this environment, we define a *non-separable APM* $\tilde{A} = \{\tilde{A}_m\}_{m \in \mathcal{M}}$ where $\tilde{A}_m : \mathbb{R}^{|\mathcal{M}|} \times [0,1] \to \mathbb{R}$. This implements allocation μ in state ω as per Definition 2 (under the modification of point 1 in Definition 2 to allow A_m to depend on x rather than just x_m).

We generalize Theorem 1 to show that the following non-separable APM uniquely implements the first-best optimal allocation:

Proposition 5. *The non-separable APM* $\tilde{A}^*_m(y,s) \equiv h^{-1}(h(s) + u^{(m)}(y))$ *and uniquely implements the first-best optimal allocation.*[36]

Proof. Follow every step in the proof of Theorem 1 with $\sum_{m\in\mathcal{M}} u_m(x_m)$ replaced by $u(x)$ and $u'_m(x_m)$ replaced by $u^{(m)}(x)$. $\square$

Thus, allowing for non-separable diversity preferences does not substantially change the analysis of adaptive priority mechanisms. One must simply adapt the APM to be non-separable to allow cross-group diversity concerns to shape the marginal benefits of admitting agents from various groups. The main difference is that this a non-separable APM does not necessarily allow the simple implementation of Algorithm 1. This is because, in the presence of cross-group adaptive priorities, it is no longer enough to rank agents within their own group. A small adaptation to this algorithm that dynamically admits agents, starting from the highest-scoring agents in each group, would naturally implement the unique first-best optimal allocation.

D.2 Non-Separable Score and Diversity Preferences

Second, we relax Equation 6 and instead suppose that the authority's preferences are represented by:

Assumption 2. *The authority's Bernoulli utility function can be represented as:*

$$\xi\left(\bar{s}_h, x\right) \tag{97}$$

where ξ *is strictly increasing, continuously differentiable, concave, and satisfies the Inada conditions* $\lim_{x_m\to 0} \xi_{x_m}(\bar{s}_h, x) = \infty$.

We define an aggregate-score-dependent APM $\hat{A} = \{\hat{A}_m\}_{m\in\mathcal{M}}$ where $\hat{A}_m : \mathbb{R}^{|\mathcal{M}|} \times [0,1] \times \mathbb{R} \to \mathbb{R}$. This implements allocation μ in state ω as per Definition 2 (where point 1 in Definition 2 is modified to allow A_m to depend on both x and $\bar{s}_h$).

In this more general setting, we now find an aggregate-score-dependent APM that implements the optimal allocation.

Proposition 6. *The following aggregate-score-dependent APM implements any first-best optimal allocation:*

$$A_m(y, s, \bar{s}_h) \equiv h^{-1}\left(h(s) + \frac{\xi_{x_m}\left(\bar{s}_h, y\right)}{\xi_{\bar{s}_h}\left(\bar{s}_h, y\right)}\right) \tag{98}$$

Proof. Follow Theorem 1 with the following substitution: replace $\sum_{m\in\mathcal{M}} \int_{\underline{\tilde{s}}_m(x_m)}^{h(1)} \tilde{s}\tilde{f}_m(\tilde{s})\mathrm{d}\tilde{s} + \sum_{m\in\mathcal{M}} u_m(x_m)$ with $\xi\left(\bar{s}_h(y,\omega), x\right)$ where $\bar{s}_h(y,\omega) = \sum_{m\in\mathcal{M}} \int_{\underline{\tilde{s}}_m(x_m)}^{h(1)} \tilde{s}\tilde{f}_{m,\omega}(\tilde{s})\mathrm{d}\tilde{s}$. By the Inada assumption, any optimal allocation is interior and so $\bar{\kappa}_m = \underline{\kappa}_m = 0$. Our previous arguments no longer establish strict concavity of the problem. The existence of an optimal allocation follows from continuity of $\xi(\bar{s}_h(x,\omega), x)$ in x for all $\omega \in \Omega$. The interiority of any optimal x implies that all optimal allocations satisfy the set of $|\mathcal{M}|$ first-order conditions:

$$\lambda = \underline{\tilde{s}}_m(x_m)\xi_{\bar{s}_h}(\bar{s}_h, x) + \xi_{x_m}(\bar{s}_h, x) \tag{99}$$

[36] Where we define $u^{(m)}(y) = \frac{\partial}{\partial y_m} u(y)$.

Under the claimed aggregate-score-dependent APM, all agents are admitted whose transformed scores exceed some threshold s^C. Thus, the marginally admitted students in each group m are those with transformed scores:

$$\underline{\tilde{s}}_m(x_m)\xi_{\bar{s}_h}(\bar{s}_h, x) + \xi_{x_m}(\bar{s}_h, x) = s^C \tag{100}$$

which satisfies the necessary conditions with $\lambda = s^C$. Thus, any optimal allocation is implemented by the claimed aggregate-score-dependent APM. □

D.3 Non-Concave Preferences

Third, we relax Equation 6 and instead suppose that u_m are not necessarily concave.

Proposition 7. *If μ is an optimal allocation, then μ is implemented by A^*.*

Proof. Without concavity, the optimal allocation characterized in the proof of Theorem 1 is no longer unique. However, the Lagrangian conditions we have derived are still necessary for any optimal allocation $x = \{x_m\}_{m\in\mathcal{M}}$. Thus, any optimal allocation is implemented by A^*. □

This result shows that any optimal allocation is implemented by the optimal APM. However, when $\{u_m\}_{m\in\mathcal{M}}$ are not concave, A^* is not necessarily monotone. Therefore, A^* does not necessarily implement a unique allocation. Indeed, it is possible that A^* implements suboptimal allocations, as it will implement any locally optimal allocation. Therefore, a mechanism defined by an arbitrary selection from the allocations implemented by A^* would not be first-best optimal. However, A^* may still help decision-making in this setting as it implements any optimal allocation.

There are two substantial differences in this optimal policy from our baseline APM. First, the policy depends on the aggregate score index. Second, without assumptions on the shape of the distribution of agents, there is no guarantee that this policy is monotone and thus no guarantee that it implements a unique policy. Taken together, this shows that the core economics of APM extend to non-separable preferences. However, some of the practical advantages of APM—not needing to condition on aggregate scores and a guarantee of unique implementation of the first-best—do not necessarily carry over without further structure.

E Noisy and Endogenous Scores

In this appendix, we study two extensions of our baseline model. In the first, authorities care about ability but can only see scores, which are endogenously generated by the agents by expending costly effort to generate them. In the second, authorities continue to care about ability but can only see scores, which are now a noisy measure thereof. In both instances, we show that the optimal APM from Theorem 1 continues to be optimal. However, in the case with noisy scores, the authority must deconvolute agents' scores to compute revised scores before applying APM.

E.1 Scores Determined by Ability and Effort

We now consider a setting in which the scores observed by the authority are the outcome of costly efforts undertaken by students. Each agent's type denotes their ability and group $\theta = (a, m) \in \Theta \equiv [0, 1] \times \mathcal{M}$. As in the main text, f_ω is the density of types over Θ in the state of the world ω and has full support. Each agent chooses their score s and pays cost $c(s, a)$, which denotes the cost of effort that needs to be exerted by an agent with ability a to obtain score s. We assume that c is strictly convex in s, twice continuously differentiable and satisfies strictly decreasing differences in s and a, *i.e.,* it is less costly for higher ability agents to obtain a given score s. We also assume that $c(0, a) = 0$ and $c(s, a)$ is symmetric in s around 0. A standard example that satisfies all of these assumptions is the cost function $c(s, a) = \left(1 - \frac{a}{2}\right) s^2$. If the agent is admitted, they receive utility of one from admission and receive zero admission utility otherwise.

In this setting, a strategy profile is a tuple (τ, ϕ) comprising: (i) effort choice functions for each group: $\tau = \{\tau_m\}_{m \in \mathcal{M}}$ with $\tau_m : [0, 1] \to \mathbb{R}$ representing the effort choice of a group m agent with ability $a \in [0, 1]$, (ii) a mechanism $\phi : \Omega \to \mathcal{U}$ mapping states to allocations chosen by the authority. A strategy profile (τ, ϕ) is *separating* if τ_m is invertible for all $m \in \mathcal{M}$. If a strategy profile is separating, with some abuse of notation we can write the allocation under mechanism ϕ in state ω directly as a function of groups m and the authority's perceptions of ability given observed scores $\hat{a} = \tau_m^{-1}(s)$, which we denote by $\phi_\omega(\tau_m^{-1}(s), m)$. Given a separating strategy profile, an agent in group m with ability a who chooses score s (and the principal therefore perceives to have ability $\hat{a} = \tau_m^{-1}(s)$) has expected utility that can be written in the following way:

$$U^m(a, \tau_m^{-1}(s), s; \phi) = \int_\Omega \phi_\omega(\tau_m^{-1}(s), m)\, \mathrm{d}\Lambda(\omega) - c(s, a) \tag{101}$$

We say that a separating strategy profile is a first-best optimal separating equilibrium if: (i) agents find it optimal to choose their prescribed report

$$\tau_m(a) \in \arg\max_{s \in \mathbb{R}} U^m(a, \tau_m^{-1}(s), s) \tag{102}$$

and (ii) ϕ is first-best optimal.

We now show that such a separating equilibrium exists under a richness condition on the environment. Let $\mu^*_\omega(a, m)$ denote the optimal allocation (which is implemented by the optimal APM that uses a when abilities are observable). We assume that:

Assumption 3. *The admissions probability of an agent with ability a in group m under the optimal allocation $H^m(a) = \int_\Omega \mu^*_\omega(a, m) \mathrm{d}\Lambda(\omega)$ is a strictly increasing and C^2 function for all $m \in \mathcal{M}$.*

Assumption 3 imposes two key conditions on the optimal allocation. The first is that for every $a \in (0, 1)$ and $m \in \mathcal{M}$, the probability that (a, m) is allocated the object lies strictly between zero and one. That is, there is positive probability that (a, m) is admitted and positive probability that it is not. This condition can be satisfied regardless of the specific preferences of the authority, provided those preferences satisfy our general assumptions.[37] The second condition is that the admission probability varies smoothly with a, ruling out the possibility that small changes in ability lead to discontinuous jumps in the probability of admission under the optimal allocation.

Proposition 8. *There exists a first-best optimal separating equilibrium. In any such equilibrium, the authority's mechanism implements essentially the same allocations as the optimal APM A^*.*

Proof. Suppose that the authority chooses a first best-optimal mechanism ϕ^*. That is, in all $\omega \in \Omega$, the authority implements the allocation μ^*_ω. Suppose also that agents in group m have an invertible effort choice function and define the perceived ability by the authority when they see a score s as $\hat{a} = \tau_m^{-1}(s)$. We now show that $U^m(a, \hat{a}, s; \phi^*)$ satisfies a set of conditions such that agents of all abilities in group m find it optimal to separate. In the statement and the proof of the following lemma, we suppress the superscript and dependence on ϕ^* for notational simplicity.

Lemma 4. *The function $U(a, \hat{a}, s)$ satisfies the following conditions:*

1. *$U(a, \hat{a}, s)$ is C^2 on $[0,1]^2 \times \mathbb{R}$.*
2. *U_2 is strictly positive.*
3. *U_{13} is strictly positive.*
4. *For each $a \in [0,1]$, the equation $U_3(a, a, s) = 0$ has a unique solution in s, denoted $\rho(a)$, which maximizes $U(a, a, s)$. Moreover, $U_{33}(a, a, \rho(a)) < 0$.*
5. *There exists $k > 0$ such that for all $(a, s) \in [0,1] \times \mathbb{R}$, if $U_{33}(a, a, s) \geq 0$, then $|U_3(a, a, s)| > k$.*
6. *The function $\frac{U_3(a,\hat{a},s)}{U_2(a,\hat{a},s)}$ is strictly monotonic in a for all $\hat{a} \in [0,1]$, $s \in \mathbb{R}$.*

Proof. We verify each condition in turn.

1. The first-order partial derivatives are:

$$U_1 = -c_a(s, a), \quad U_2 = H'(\hat{a}), \quad U_3 = -c_s(s, a).$$

The second-order partial derivatives are:

$$\begin{aligned} &U_{11} = -c_{aa}(s, a), \qquad &&U_{12} = 0, \qquad &&U_{13} = -c_{as}(s, a), \\ &U_{22} = H''(\hat{a}), &&U_{23} = 0, &&U_{33} = -c_{ss}(s, a). \end{aligned}$$

All these derivatives exist and are continuous by assumption, so U is C^2.
2. By Assumption 3, we have $H'(\hat{a}) > 0$.
3. Since $U_{13} = -c_{as}(s, a)$ and $c(s, a)$ satisfies strictly decreasing differences, $U_{13} > 0$.

[37] To see this, consider any (a, m). If the measure of group m agents with abilities higher than a exceeds the capacity q, then a is not admitted in the optimal allocation. Conversely, if the total measure of agents from other groups, along with group m agents with scores higher than a, is strictly less than capacity q, then (a, m) is always admitted, regardless of the authority's preferences.

4. The equation $U_3(a, a, s) = 0$ is equivalent to $c_s(s, a) = 0$, which has a unique solution at $s = 0$ under our assumptions. Moreover, $U_{33}(a, a, \rho(a)) = -c_{ss}(0, a) < 0$ by strict convexity of c.
5. Since $c(s, a)$ is strictly convex, $U_{33} = -c_{ss}(s, a) < 0$ for all (a, s). Thus $U_{33} \geq 0$ never holds, and the implication holds vacuously.
6. Since $U_2 = H'(\hat{a})$ does not depend on a, and $U_3 = -c_s(s, a)$ is strictly increasing in a (by strictly decreasing differences), the ratio $\frac{U_3}{U_2}$ is strictly increasing in a. □

We now use the arguments from Mailath (1987) to establish separation. Consider the differential equation:

$$\frac{\mathrm{d}\tau}{\mathrm{d}a} = \frac{-U_2(a, a, \tau)}{U_3(a, a, \tau)}, \quad \tau(0) = 0 \tag{103}$$

Under the conditions on U established by Lemma 4 above, Proposition 5 and Theorem 2 in Mailath (1987) establish that this differential equation has a unique solution $\tilde{\tau}$ that is moreover strictly monotone and continuous.[38] Given this solution and the properties on U established by Lemma 4 above, Theorem 3 in Mailath (1987) establishes that $\tilde{\tau}$ satisfies the incentive compatibility condition of Equation 102.

To complete the proof, suppose that all groups can exert effort to determine their scores, and that agents in each group m use a one-to-one strategy $\tilde{\tau}_m$ characterized above. Under the optimal APM, for each group m, the actions of agents from other groups affect their admission probabilities, and therefore their payoffs, only through their inferred type. Since $\tilde{\tau}_m$ is one-to-one for all m, the function $U^m(a, \hat{a}, s; \phi^*)$ is the same (i) when the actions of other groups are observable, and (ii) when all groups use one-to-one strategies $\tilde{\tau}_m$. This establishes the first part of the result. The second part of the result follows immediately from Theorem 1. □

Thus, the presence of endogenously chosen efforts does not affect the optimality of APM.

E.2 Scores as Noisy Signals of Unobserved Ability

We will now consider a setting in which the scores observed by the authority are noisy measures of the innate ability of the agents and the authority cares about innate abilities instead of observed scores. Formally, each agent's type denotes their ability and group $\theta = (a, m) \in \Theta \equiv [0, 1] \times \mathcal{M}$. As in the main text, f_ω is the density of types over Θ in the state of the world ω and has full support. We also assume that h is the identity function to simplify the notation. We define the aggregate ability (instead of score) index of any allocation as:

$$\bar{a}(\mu, \omega) = \int_\Theta \mu(a, m) a \mathrm{d}F_\omega(a, m) \tag{104}$$

Our assumptions on the preferences of the authority remain the same, with $\bar{a}$ replacing $\bar{s}_h$.

Differently from the main analysis, the authority does not directly observe the abilities, but observes noisy measures thereof: $s = a + \epsilon$, where $\epsilon \sim N(0, \sigma^2)$ and σ is known to the

[38] This Proposition holds so long as $U_2 = H'$ is bounded above, which it is in our context as the domain of H' is compact and H' is continuous by Assumption 3.

authority.[39] Thus, at each $\omega \in \Omega$, thedistribution of scores for group m is given by

$$g_\omega(s|m) = \int_0^1 f_\omega(a|m)t\,(s-a)\,\mathrm{d}a \tag{105}$$

where t is the density of a normal random variable with mean zero and variance σ^2 and, with some abuse of notation, $f_\omega(\cdot)$ is the marginal distribution of ability. It turns out that the critical object is the expectation of ability conditional on score, which depends on the observed joint distribution of scores g and groups, and is denoted by $\mathbb{E}_g[a|s,m]$ (this expression is computed explicitly in the proof of Proposition 9). To address this, we extend APM with a two step procedure: 1) Apply a de-noising procedure in which they take the observed score distribution g and apply standard Fourier de-convolution to compute conditional expectations of students abilities given their scores and their groups $\hat{a}(s,m)$, 2) Implement the standard optimal APM that takes as input the de-noised scores. As the ability is not observable, the definitions of allocation μ, mechanism ϕ and optimality are the same, but we refer to them as *feasible* allocations and *optimal* mechanisms.[40]

Proposition 9. *The following APM is monotone and optimal.*

$$A_m(y_m, s) = \hat{a}(s,m) + u'_m(y_m) \tag{106}$$

where the authority de-noises scores to compute $\hat{a}(s,m) = \mathbb{E}_g[a|s,m]$.

Proof. We first show that under our assumptions on the ability distribution f_ω and error distribution t, the realized ability distribution f_ω and $\mathbb{E}_\omega[a|s,m]$ can be computed by the authority without observing f_ω. To see how, note that taking the Fourier transform of the both sides of Equation 105, we obtain $\hat{g}_\omega(\xi|m) = \hat{f}_\omega(\xi|m)\hat{t}(\xi)$. As $\hat{t}(\xi) = e^{-\frac{\sigma^2\xi^2}{2}} > 0$ everywhere, $\hat{f}_\omega(\xi|m) = \frac{\hat{g}_\omega(\xi|m)}{\hat{t}(\xi)}$. As f_ω is continuous (and therefore integrable), it can be identified uniquely by inverting $\hat{f}_\omega(\xi)$, which we just constructed as a function of two objects observed by the authority, g and t. Thus, for the remainder of the proof, we take expectations with respect to ω rather than g, as the two are equivalent.

The authority's posterior distribution of ability after seeing the score s of a group m agent at state ω is

$$\pi_\omega(a|s,m) = \frac{f_\omega(a|m)t(s-a)}{g_\omega(s|m)} \tag{107}$$

And the conditional expectation of ability given score at state ω is

[39] We assume normally distributed noise for analytical convenience, particularly to guarantee identifiability of the latent ability distribution via Fourier deconvolution. The analysis can be extended to other distributions that are integrable and have a Fourier transform that is strictly positive everywhere on $\mathbb{R}$. These conditions hold, for example, for any symmetric, log-concave, unimodal distribution with full support (e.g., Laplace, logistic). Extensions to such distributions follow from the same deconvolution and monotonicity arguments used here.

[40] This is because any mechanism that can implement allocations that can admit students based on differences in unobserved ability conditional on scores is infeasible when ability is not observable. We use the word optimal instead of first-best-optimal in this section as infeasible mechanisms may yield higher utility than feasible mechanisms.

$$\mathbb{E}_\omega[a|s,m] = \int_0^1 a\pi_\omega(a|s,m)\mathrm{d}a = \frac{\int_0^1 af_\omega(a|m)t(s-a)\mathrm{d}a}{\int_0^1 f_\omega(a|m)t(s-a)\mathrm{d}a} \tag{108}$$

As $t'(s-a) = -\frac{s-a}{\sigma^2}t(s-a)$, we have the following:

$$\begin{aligned}\text{sign}\left(\frac{\mathrm{d}}{\mathrm{d}s}\mathbb{E}_\omega[a|s,m]\right) = \text{sign}\Bigg(&\int_0^1 af_\omega(a|m)t(s-a)\,da \cdot \int_0^1 (s-a)f_\omega(a|m)t(s-a)\,\mathrm{d}a \\ &- \int_0^1 a(s-a)f_\omega(a|m)t(s-a)\,da \cdot \int_0^1 f_\omega(a|m)t(s-a)\,\mathrm{d}a\Bigg)\end{aligned} \tag{109}$$

Defining $w_{s,m}(a) \equiv f_\omega(a)t(s-a)$ and normalizing to obtain the distribution

$$\mathrm{d}G^\omega_{s,m}(a) = \frac{w_{s,m}(a)\,\mathrm{d}a}{\int_0^1 w_{s,m}(u)\ \mathrm{d}u} \tag{110}$$

We can re-express the sign of the slope of the conditional expectation by dividing every term by the normalizing constant that defines $G_{s,m}$. Doing this, we obtain that:

$$\begin{aligned}\text{sign}\left(\frac{\mathrm{d}}{\mathrm{d}s}\mathbb{E}_\omega[a|s,m]\right) = \text{sign}\Bigg(&\left(\int_0^1 a\ \mathrm{d}G^\omega_{s,m}(a)\right)\left(\int_0^1 (s-a)\ \mathrm{d}G^\omega_{s,m}(a)\right) \\ &- \int_0^1 a(s-a)\ \mathrm{d}G^\omega_{s,m}(a)\Bigg).\end{aligned} \tag{111}$$

As a is strictly increasing in a, $s-a$ is strictly decreasing in a and both functions are not constant almost everywhere, by Chebyshev's Algebraic Inequality, sign $\left(\frac{d}{ds}\mathbb{E}_\omega[a \mid s]\right) > 0$ (See Theorem 1 in Wagener (2006) for a statement of the inequality). The optimality of $A_m(y,s) = \hat{a}(s,m) + u'_m(y_m)$ then follows from Theorem 1, replacing s in our general model with $\hat{a}(s,m) = \mathbb{E}_g[a|s]$. The monotonicity of APM follows from the previous arguments which established that $\hat{a}^{(s)}(s,m) > 0$. □

This establishes that, after scores have been deconvoluted, the standard APM remains optimal. However, the optimal deconvolution requires the authority to condition on the observed joint distribution of scores and groups. This raises potentially steeper practical challenges for communicating and implementing APM. Of course, in practice, one could commit to a de-noising scheme that does not adapt to the state. This will lose the optimality guarantee of APM but is likely to feature simpler practical implementation.

F Decentralized Admissions and APM Dominance

In this Appendix, we study which mechanisms are optimal under decentralized admissions. We consider a setting in which the agents apply sequentially to the authorities, who then decide which agents to admit. We index the stage of the game by $t \in \mathcal{T} = \{1, \ldots, |\mathcal{C}| - 1\}$.

Each stage corresponds to a (non-dummy) authority $I(t)$, where $I : \mathcal{T} \to \mathcal{T}$. At each stage t, any unmatched agents choose whether to apply to authority $I(t)$. Given the set of applicants, authority $I(t)$ chooses to admit a subset of these agents, who are then matched to the authority. Given this, histories are indexed by the path of the measure of agents who have not yet matched, $h^{t-1} = (F, F_1, \ldots, F_{t-1}) \in \mathcal{H}^{t-1}$. Given each history h^{t-1} and set of applicants $\Theta_c^A \subseteq \Theta$, a strategy for an authority returns a set of agents $\Theta_c^G \subseteq \Theta$ whom they will admit such that $\Theta_c^G \subseteq \Theta_c^A$ and $F_t(\Theta_c^G) \leq q_c$ for each time at which they could move $t \in \mathcal{T}$, $a_{c,t} : \mathcal{H}^{t-1} \times \mathcal{P}(\Theta) \to \mathcal{P}(\Theta)$, where $\mathcal{P}(\Theta)$ is the power set over Θ.[41] A strategy for an agent returns a choice of whether to apply to authorities at each history and time $t \in \mathcal{T}$ for all agent types $\theta \in \Theta$, $\sigma_{\theta,t} : \mathcal{H}^{t-1} \to [0, 1]$.

Within this context, our notion of equilibrium is that of subgame perfect equilibrium:

Definition 7 (Equilibrium)**.** *A strategy profile* $\Sigma = \{\{a_{c,t}\}_{c \in \bar{\mathcal{C}}}, \{\sigma_{\theta,t}\}_{\theta \in \Theta}\}_{t \in \mathcal{T}}$ *is a subgame perfect equilibrium if* $a_{c,t}$ *maximizes authority utility given* Σ *for all* $c \in \bar{\mathcal{C}}$ *and* $t \in \mathcal{T}$ *and* $\sigma_{\theta,t}$ *is maximal according to agent preferences for all* $\theta \in \Theta$ *and* $t \in \mathcal{T}$*.*

We moreover say that a strategy $a_{\tilde{c},t}$ for an authority $\tilde{c}$ at time t is *dominant* if it maximizes authority utility regardless of the strategies of all other authorities and agents, $\{\{a_{c,t}\}_{c \in \bar{\mathcal{C}} \setminus \{\tilde{c}\}}, \{\sigma_{\theta,t}\}_{\theta \in \Theta}\}_{t \in \mathcal{T}}$, and the order in which authorities admit agents, I. Moreover, an equilibrium Σ is in *dominant strategies* if $a_{c,t}$ is dominant for all $c \in \bar{\mathcal{C}}$ and $t \in \mathcal{T}$. We denote the unique probabilistic allocation of agents to authorities induced by Σ as $\mu_\Sigma : \Theta \to \Delta(\mathcal{C})$. A probabilistic allocation μ_Σ is deterministic if $\mu_\Sigma(\theta)$ is a Dirac measure on some authority $c \in \mathcal{C}$ for all $\theta \in \Theta$. A deterministic allocation μ_Σ corresponds to a matching μ if $\mu_\Sigma(\theta)$ is a Dirac measure on $\mu(\theta)$ for all $\theta \in \Theta$.

We now establish that the single-authority optimal APM characterizes dominance.

Theorem 5. *A mechanism implements a dominant strategy for an authority if and only if it implements essentially the same allocations as* A_c^**.*

Proof. We prove that APM A_c^* implements a dominant strategy by backward induction. Consider the terminal time $t = |\mathcal{C}| - 1$. Some measure of agents λ applies to the authority. Regardless of the measure λ, by Theorem 1 we have that the APM A_c^* is first-best optimal (to see this more concretely, simply index λ by an arbitrary $\omega \in \Omega$ and apply Theorem 1). Thus, A_c^* is dominant. Moreover, from Theorem 1, any strategy that differs from A_c^* on a strictly positive measure set cannot be optimal. Thus any dominant strategy implements essentially the same allocation as A_c^*. Consider now any time $t < |\mathcal{C}| - 1$, precisely the same argument applies and A_c^* is (essentially uniquely) dominant. $\square$

The intuition behind this result is that each authority takes as given the set of agents that will accept it. Thus, given this measure of agents, they can do no better than to employ the same APM that a single authority would, which is A_c^* by Theorem 1.

Theorem 5 provides a powerful rationale for focusing on APMs in decentralized markets at both positive and normative levels. Normatively, this result allows an analyst to advise an authority regarding how it should conduct its admissions. This is important because

[41] Formally, so that $F_t(\Theta_c^G)$ is well defined, we require that authorities' strategies be measurable in the Borel sigma algebra over Θ.

any policy that does not coincide with the APM we derive — such as the popular priority and quota mechanisms outside of the cases delimited by Theorem 2 — will disadvantage an authority. Positively, this result allows a sharp prediction that the equilibrium matching between agents and authorities will be the unique stable matching (as per Theorem 3):

Proposition 10. *For all equilibria Σ^* where authorities use A^*, the allocation μ_{Σ^*} is deterministic and corresponds to the unique stable matching of this economy.*

Proof. We first prove the following claim.

Claim 15. *μ_{Σ^*} is (almost surely) a deterministic allocation that corresponds to a cutoff matching μ^*.*

Proof. Since there is a continuum of agents, under any Σ^*, with probability 1, any authority c faces a given set of agents who apply Θ_c^{A,Σ^*} with induced measure $\lambda_c^{\Sigma^*}$. As c uses APM A_c^*, with probability 1, any agent θ is admitted to an authority if and only if $s_c(\theta) \geq S_{m,c}^{\Sigma^*}$, where $S_{m,c}^{\Sigma^*}$ denotes the cutoffs when APM A_c^* is applied to agent measure $\lambda_c^{\Sigma^*}$. Since the agents have strict preferences, in any equilibrium, each agent applies to the $\succeq_\theta$ −maximal authority in $\{c : s_c(\theta) \geq S_{m,c}^{\Sigma^*}\}$, and is admitted, which establishes that μ_{Σ^*} is (almost surely) deterministic allocation that corresponds to a cutoff matching with cutoffs $S_{m,c}^{\Sigma^*}$. □

We now establish that μ_{Σ^*} is the unique stable matching of the economy.

Claim 16. *μ^* is the unique stable matching of this economy.*

Proof. For a contradiction, assume μ_{Σ^*} is not stable. Let S denote the unique cutoffs associated with μ_{Σ^*}. Since μ_{Σ^*} is not stable, by Claim 8, S is not a fixed point of T. Let $t_c = T_c(S)$. Since S is not a fixed point of T, there exists $m \in \mathcal{M}$ and $c \in \mathcal{C}$ such that $t_{m,c} \neq S_{m,c}$. Moreover, let $\{x_{m,c}^t\}_{m\in\mathcal{M}}$ and $\{x_{m,c}^s\}_{m\in\mathcal{M}}$ denote the measure of agents in $\tilde{D}_c(S_{-c})$ who are above the admission thresholds for authority c under t_c and S_c. As in Claim 9, note that if there exists m, c such that $t_{m,c} > S_{m,c}$, then from full support, we have that $x_{m,c}^s > x_{m,c}^t$. Since the authority fills its capacity in both cases, there must exists m' such that $x_{m',c}^t > x_{m',c}^s$ which is only possible if $S_{m,c} > t_{m,c}$. By an identical argument, if there is m, c such that $t_{m,c} < S_{m,c}$, then there exists m' $S_{m',c} < t_{m',c}$. Therefore, whenever $t_{m,c} \neq S_{m,c}$, there exists c and m, m' such that $t_{m,c} > S_{m,c}$ and $S_{m',c} > t_{m',c}$. But now we have shown the following:

$$h_c(S_{c,m'}) + u_{m'}(x_{c,m'}^s) > h_c(t_{c,m'}) + u_{m'}(x_{c,m'}^t) \geq h_c(t_{c,m}) + u_m(x_{c,m}^t) > h_c(S_{c,m}) + u_m(x_{c,m}^s)$$

where the first inequality follows by $t_{c,m'} < S_{c,m'}$, $x_{c,m'}^s < x_{c,m'}^t$, and concavity of u_m. The second inequality follows by optimality. This is because the facts that $t_{c,m} > 0$ and $t'_{c,m'} < 1$ imply that the Lagrange multipliers in the proof of Theorem 1 $\overline{\kappa}_m = \underline{\kappa}_{m'} = 0$. The final inequality follows since $t_{m,c} > S_{m,c}$ and $x_{m,c}^t < x_{m,c}^s$. However, this is a contradiction since $h_c(S_{c,m'}) + u_{m'}(x_{c,m'}^s) > h_c(S_{c,m}) + u_m(x_{c,m}^s)$ implies that there exists $\varepsilon > 0$, an agent θ with score $s_c(\theta) = S_{c,m'} - \epsilon$ and type $m(\theta) = m$ has higher score under A^* than the agent θ' with score $s_c(\theta') = S_{c,m}$ and type $m(\theta') = m$. Since θ' is admitted to c, θ would be if it applied to c. Moreover, from full support, there is such θ whose top choice is c and the strategy of this agent is not a best response, which is a contradiction. □

The combination of Claims 15 and 16 completes the proof. □

The intuition is that if an equilibrium matching under A^* was not the unique stable matching, then it must be that some agents are applying suboptimally and failing to select the most preferred authority that they can attend, which contradicts that the outcome is consistent with equilibrium. Proposition 10 also shows that the allocation implemented in the dominant strategy equilibrium is the same for all possible orderings of authorities, I.

G Efficient Mechanisms with Multiple Authorities

We have characterized stability and the decentralized outcome (and shown that they are the same), but two natural questions remain. First, is the stable outcome efficient for the authorities?[42] Second, if not, what kind of centralized solution can remedy any inefficiency? We show that the stable outcome is generally inefficient and that a modified, centralized APM mechanism restores efficiency.

G.1 Inefficiency of the Decentralized Outcome

The notion of efficiency that we will consider is utilitarian efficiency over authorities. A mechanism in the multi-authority setting is a function $\phi : \Omega \to \mathcal{U}$, where $\mathcal{U}$ is the set of matchings (which, by definition, encodes the feasibility requirement imposed in the single authority setting). We define the total authority value Ξ_T of a mechanism ϕ under distribution $\Lambda \in \Delta(\Omega)$ as the total expected utility of the allocations induced by that mechanism:

$$\Xi_T(\phi, \Lambda) = \sum_{c \in \bar{\mathcal{C}}} \Xi_c(\phi, \Lambda) \tag{112}$$

A mechanism is efficient if it maximizes total authority value for all possible distributions:

Definition 8 (Efficiency)**.** *A mechanism ϕ^* is efficient if:*

$$\Xi_T(\phi^*, \Lambda) = \sup_{\phi} \Xi_T(\phi, \Lambda) \tag{113}$$

for all $\Lambda \in \Delta(\Omega)$.

For this section, so that scores are directly comparable across authorities and allocations are interior, we impose the following assumption:

Assumption 4. *Scores and preferences are such that $s_c(\theta) = s_{c'}(\theta)$, $h_c = h$ and $g_c = Id$, where Id is the identity function, for all $c, c' \in \bar{\mathcal{C}}$ and $\theta \in \Theta$. Moreover, $\lim_{x \to +0} u'_{m,c}(x) = \infty$, $u_{m,c}$ is strictly concave for all $m \in \mathcal{M}$ and $c \in \bar{\mathcal{C}}$, and for all $\theta \in \Theta$, c_0 is less preferred than c for all $c \in \bar{\mathcal{C}}$.*

Assumption 4 makes scores a common numeraire across authorities and is akin to the standard quasi-linearity assumption in mechanism design. For example, it may be suitable in settings where the score is derived from a common index of academic attainment, such as in

[42]We have shown that the decentralized outcome corresponds to the unique stable matching. Naturally, stable allocations need not be efficient for the agents. Here, we will see if they are efficient for the authorities.

Chicago Public Schools. This assumption does not impose that all authorities have common marginal rates of substitution between scores and diversity, as they are allowed unrestricted heterogeneity in diversity preferences. We add the Inada condition for analytical tractability. We argue that it is also reasonable to assume that failing to admit any agents from a given group is intolerable for authorities. We add that the outside option is ranked lower than all authorities so that any agent is willing to be assigned to any of the authorities. When authorities control highly desirable resources, such as elite school or university seats or essential medical supplies, we argue that this is a reasonable assumption.

With the efficiency benchmark defined, we now demonstrate that the decentralized equilibrium outcome can fail to be efficient. We prove this result via an explicit example with two authorities, c and c' of capacity $\frac{1}{4}$, and two groups of agents, m and m' of measure $\frac{1}{2}$. All agents in group m prefer c' to c and all agents in group m' prefer c to c'. Authority c values admitting group m agents more on the margin than group m' agents, and authority c' values admitting group m' agents more on the margin than group m agents. Using the optimal APMs, both authorities admit more agents of the group whose admissions they value relatively less than the efficient benchmark. The intuition for this is that both authorities "steal" the high-scoring agents of the group whom they relatively less value from the other authority, an externality that they do not internalize.

Proposition 11 (Equilibrium Inefficiency)**.** *All authorities using the privately optimal APMs $\{A_c^*\}_{c\in\mathcal{C}}$ is not necessarily efficient.*

Proof. We prove the result by explicitly constructing an economy in which the optimal APMs lead to inefficiency. There are two authorities, c and c', both with capacity $1/2$ and two groups of agents, m and m'. Both agent groups have a measure of 1 and their scores are uniformly distributed in $[1/2, 1]$. Authorities' utility functions are given by

$$\xi_c\left(\bar{s}_h, x\right) \equiv \bar{s}_h + \frac{1}{4}\sqrt{x_m} + \frac{1}{8}\sqrt{x_{m'}} \tag{114}$$

$$\xi_{c'}\left(\bar{s}_h, x\right) \equiv \bar{s}_h + \frac{1}{4}\sqrt{x_{m'}} + \frac{1}{8}\sqrt{x_m} \tag{115}$$

with $h(x) \equiv x$ while all agents of type m prefer authority c' to c while all agents of type m' prefer authority c to c'.[43]

We will now derive the stable outcome of this economy, which is (up to measure zero transformation) the unique outcome implemented when the authorities use the optimal APM. Let x_m^c denote the measure of type m agents at authority c. First, note that higher-scoring agents from the same group go to the more preferred authority. To see why this is true, note that if $m(\theta) = m(\theta') = m$, $s(\theta) > s(\theta')$ and $\mu(\theta) = c$ while $\mu(\theta') = c'$, c and θ would violate within group fairness since θ has higher priority at c than θ' regardless of the allocation. As a result, in any stable allocation μ, the highest-scoring $x_{m'}^c$ type m' agents are assigned to c and the next highest-scoring $x_{m'}^{c'}$ agents are assigned to authority c', while rest

[43]This assumption on the preferences and the distribution of scores violate our full support assumption, but adding an arbitrarily small full support density to all types makes arbitrarily small changes in the utility under the stable matching and optimal allocation but complicates the calculation, so we omit it for expositional clarity.

of the type m' agents are not assigned to any authority. The allocation for type m agents is analogous. Moreover, since $q = 1/2$ for both authorities, $x^{c'}_{m'}$=1/2-$x^{c}_{m'}$ and x^{c}_{m}=1/2-$x^{c'}_{m}$ and the allocation is completely determined by the measures $x^{c}_{m'}$ and $x^{c'}_{m}$.

Next, note that at μ, the adaptive priority of the lowest-scoring type m and m' agents must be equal at both authorities. To see why this is true, take authority c without loss of generality. Let $s^{c}_{m'} = 1 - x^{c}_{m'}$ and $s^{c}_{m} = 1 - x^{c}_{m'} - x^{c}_{m}$ denote the scores of the lowest-scoring type m and m' agents and A_m denote the optimal APM. For a contradiction, assume $A_m(x^c_m, s^c_m) > A_{m'}(x^c_{m'}, s^c_{m'})$. Since agents of type m' with scores lower than s^c_m are unassigned at μ, for small enough ϵ, a type m agent with score $s^c_m - \epsilon$ and authority c blocks the matching. Similarly, assume $A_m(x^c_m, s^c_m) < A_{m'}(x^c_{m'}, s^c_{m'})$. Since agents of type m' with scores lower than $s^{c'}_m$ are assigned to authority c or unmatched at μ, a type m' agent with score $s^c_{m'} - \epsilon$ and authority c blocks the matching μ. Thus, the following equations must be satisfied:

$$A_m(x^c_m, s^c_m) = A_{m'}(x^c_{m'}, s^c_{m'}) \text{ and } A_m(x^{c'}_m, s^{c'}_m) = A_{m'}(x^{c'}_{m'}, s^{c'}_{m'}) \tag{116}$$

As the optimal APM in this setting is given by:

$$A^*_{\hat{m},\hat{c}}(y_{\hat{m}}, s) \equiv s + u'_{\hat{m},\hat{c}}(y_{\hat{m}}) \tag{117}$$

for all $\hat{m} \in \{m, m'\}$ and $\hat{c} \in \{c, c'\}$, we have that:

$$1 - x^c_{m'} + \frac{1}{8}\frac{1}{\sqrt{x^c_{m'}}} = 1 - x^c_{m'} - x^c_m + \frac{1}{4}\frac{1}{\sqrt{1/2 - x^c_{m'}}} \tag{118}$$

and:

$$1 - x^{c'}_m + \frac{1}{8}\frac{1}{\sqrt{x^{c'}_m}} = 1 - x^{c'}_m - x^{c'}_{m'} + \frac{1}{4}\frac{1}{\sqrt{1/2 - x^{c'}_m}} \tag{119}$$

These equations are identical up to relabelling and so $x^c_{m'} = x^{c'}_{m'} = x^*$ for some x^*. Thus, we need to find the solution to the following single equation to characterize the allocation:

$$1 - x^* + \frac{1}{8}\frac{1}{\sqrt{x^*}} = \frac{1}{2} + \frac{1}{4}\frac{1}{\sqrt{1/2 - x^*}} \tag{120}$$

Observe that this equation can be rewritten as the fixed point equation:

$$x^* = \frac{1}{2} + \frac{1}{8}\frac{1}{\sqrt{x^*}} - \frac{1}{4}\frac{1}{\sqrt{1/2 - x^*}} \tag{121}$$

We observe that the RHS satisfies the following properties: (i) $\lim_{x^* \to 0} \text{RHS}(x^*) = \infty$, (ii) $\lim_{x^* \to \frac{1}{2}} \text{RHS}(x^*) = -\infty$, and (iii) $\text{RHS}'(x^*) < 0$ for all $x^* \in (0, \frac{1}{2})$. Thus, there exists a unique solution. Moreover, we can guess-and-verify that this solution is $x^* = \frac{1}{4}$.

In summary, if both authorities use the optimal APM, then the outcome is

$$\mu(\theta) = \begin{cases} c & \text{if } m(\theta) = m, s(\theta) \in [1/2, 3/4) \text{ or } m(\theta) = m', s(\theta) \in [3/4, 1] \\ c' & \text{if } m(\theta) = m', s(\theta) \in [1/2, 3/4) \text{ or } m(\theta) = m, s(\theta) \in [3/4, 1] \\ \theta & \text{otherwise} \end{cases} \tag{122}$$

In this outcome, both authorities have an average score of 3/4 and admit measure 1/4 agents from both groups, giving them a utility of 15/16. Thus, total utilitarian welfare is 15/8 under the decentralized outcome.

We now show that this does not attain the efficient benchmark. A necessary condition for the (utilitarian) efficient outcome is that for c:

$$\frac{1}{4}\frac{1}{\sqrt{x_m^c}} = \frac{1}{8}\frac{1}{\sqrt{1/2 - x_m^c}} \tag{123}$$

and for c':

$$\frac{1}{4}\frac{1}{\sqrt{x_{m'}^{c'}}} = \frac{1}{8}\frac{1}{\sqrt{1/2 - x_{m'}^{c'}}} \tag{124}$$

This implies that $x_m^c = x_{m'}^{c'} = 4/10$ and $x_{m'}^c = x_m^{c'} = 1/10$. In this case, the same set of agents is admitted overall, so the score contribution to utility remains 3/4 on average across the authorities. Total utilitarian welfare is now:

$$3/2 + 1/2 \times \sqrt{4/10} + 1/4 \times \sqrt{1/10} \approx 1.895 > 1.875 = 15/16 \tag{125}$$

Completing the proof. $\square$

G.2 An Efficient Centralized Mechanism

The inefficiency of each authority using a decentralized APM stems from the implicit incompleteness of markets: if we added the ability for authorities to pay each other for agents, then they would have willingness-to-pay to do so at the equilibrium allocation. A centralized mechanism can remedy this issue by ensuring the cross-sectional allocation of agents to authorities is optimal.

We propose the following augmentation of an APM to solve this problem, an *adaptive priority mechanism with quotas* (APM-Q). The idea behind this hybrid mechanism is to use aggregate, market-level priorities with authority-specific quotas. To this end, an APM-Q comprises an aggregate non-separable APM $\tilde{A} = \{\tilde{A}_m\}_{m\in\mathcal{M}}$ with $\tilde{A}_m : \mathbb{R}^{|\mathcal{M}|} \times [0,1] \to \mathbb{R}$ and a profile of quota functions $Q = \{Q_{m,c}\}_{m,c\in\mathcal{M}}$ with with $Q_{m,c} : \mathbb{R}^{|\mathcal{M}|} \to \mathbb{R}_+$. Intuitively, the aggregate APM pins down the aggregate measures of allocations of each group to *any* authority $\{x_m\}_{m\in\mathcal{M}}$, where $x_m = \sum_{c\in\bar{\mathcal{C}}} x_{m,c}$. The non-separability of this APM simply means that the measures of all groups matter for the adaptive priority of any agent. Given the aggregate measure of allocation for group m, the quota function for authority c assigns $Q_{m,c}(\{x_m\}_{m\in\mathcal{M}})$ agents of type m to authority c.

Definition 9 (Adaptive Priority Mechanism with Quotas). *An adaptive priority mechanism with quotas $(\tilde{A}, Q)$ comprises a non-separable APM $\tilde{A}$ and a quota function Q. An APM-Q* *implements allocation μ if the following are satisfied:*

1. *Aggregate allocations are in order or priorities: $\mu(\theta) \in \bar{\mathcal{C}}$ if and only if for all θ' with $\mu(\theta') = c_0$, we have that:*

$$\tilde{A}_{m(\theta)}(\{x_m(\mu)\}_{m\in\mathcal{M}}, s(\theta)) > \tilde{A}_{m(\theta')}(\{x_m(\mu)\}_{m\in\mathcal{M}}, s(\theta')) \tag{126}$$

2. *Authority-level allocations are given by the corresponding quota functions:*

$$x_{m,c}(\mu) = Q_{m,c}(\{x_m(\mu)\}_{m\in\mathcal{M}}) \tag{127}$$

3. *The resources are fully allocated:*

$$\sum_{m\in\mathcal{M}} x_{m,c}(\mu) = q_c \tag{128}$$

By appropriate choice of the APM and quota functions, we can derive an APM-Q that is efficient. To this end, define the optimally-allocated aggregate utility from diversity:

$$\begin{aligned} \tilde{u}(\{x_m\}_{m\in\mathcal{M}}) &= \max_{\{x_{m,c}\}_{c\in\mathcal{C}}} \sum_{c\in\mathcal{C}} \sum_{m\in\mathcal{M}} u_{m,c}(x_{m,c}) \\ \text{s.t.} &\sum_{c\in\mathcal{C}} x_{m,c} \leq x_m, \ \sum_{m\in\mathcal{M}} x_{m,c} \leq q_c, \, \forall m \in \mathcal{M}, c \in \mathcal{C} \end{aligned} \tag{129}$$

Moreover, define the marginal value of aggregate group m admissions $\tilde{u}^{(m)}(y) = \frac{\partial}{\partial y_m}\tilde{u}(y)$ and the marginal value of authority capacity $\tilde{u}_{q_c}(y) = \frac{\partial}{\partial q_c}\tilde{u}(y)$. Using these marginal values, we can design an efficient APM-Q that combines market-level APMs with authority-level quotas:

Theorem 6 (Efficient APM-Q)**.** *Every allocation induced by the following APM-Q $(\tilde{A}, Q)$ is efficient:*

1. *The non-separable APM is given by $\tilde{A}_m(y, s) = h^{-1}\left(h(s) + \tilde{u}^{(m)}(y)\right)$*
2. *The quota functions are given by $Q_{m,c}(y) = \left(u'_{m,c}\right)^{-1}\left(\tilde{u}^{(m)}(y) + \tilde{u}_{q_c}(y)\right)$*

Proof. First, we define a fictitious *composite authority* with utility function defined over vectors of total scores $\bar{s}_h = \{\bar{s}_h^c\}_{c\in\mathcal{C}}$, and aggregate allocation to each group $x = \{x_m\}_{m\in\mathcal{M}}$. To do this, we define:

$$\begin{aligned} \tilde{u}(\{x_m\}_{m\in\mathcal{M}}) &= \max_{\{x_{m,c}\}_{c\in\mathcal{C}}} \sum_{c\in\mathcal{C}} \sum_{m\in\mathcal{M}} u_{m,c}(x_{m,c}) \\ \text{s.t.} &\sum_{c\in\mathcal{C}} x_{m,c} \leq x_m, \ \sum_{m\in\mathcal{M}} x_{m,c} \leq q_c, \, \forall m \in \mathcal{M}, c \in \mathcal{C} \end{aligned} \tag{130}$$

and $\tilde{\bar{s}}_h = \sum_{c\in\mathcal{C}} \bar{s}_h^c$. We write the utility function of this composite authority as

$$\tilde{\xi}\left(\tilde{\bar{s}}_h, x\right) = \tilde{\bar{s}}_h + \tilde{u}(x) \tag{131}$$

We first establish that $\tilde{u}$ satisfies the properties necessary to invoke Proposition 5, which establishes the optimality of the claimed APM for the fictitious authority.

Claim 17. *The function $\tilde{u}$ is concave and partially differentiable in each argument.*

Proof. First, we establish concavity. That is, for all $\lambda \in [0,1]$ and $x, x' \in \mathbb{R}_+^{|\mathcal{M}|}$, we have that $\tilde{u}(\lambda x' + (1-\lambda)x) \geq \lambda \tilde{u}(x') + (1-\lambda)\tilde{u}(x)$. Let $\{x^*_{m,c}\}_{m\in\mathcal{M},c\in\mathcal{C}}$ and $\{x^{*'}_{m,c}\}_{m\in\mathcal{M},c\in\mathcal{C}}$ correspond to optimal values under x and x'. Under $\tilde{x} = \lambda x' + (1-\lambda)x$, we have that $\lambda x^{*'}_{m,c} + (1-\lambda)x^*_{m,c}$ is feasible for all $m \in \mathcal{M}$ and $c \in \mathcal{C}$. Thus, we have that:

$$\begin{aligned}
\tilde{u}(\tilde{x}) &\geq \sum_{m\in\mathcal{M}} \sum_{c\in\mathcal{C}} u_{m,c}(\lambda x^{*'}_{m,c} + (1-\lambda)x^*_{m,c}) \\
&\geq \sum_{m\in\mathcal{M}} \sum_{c\in\mathcal{C}} \lambda u_{m,c}(x^{*'}_{m,c}) + (1-\lambda)u_{m,c}(x^*_{m,c}) \\
&= \lambda \tilde{u}(x') + (1-\lambda)\tilde{u}(x)
\end{aligned} \tag{132}$$

where the second inequality is by concavity of $u_{m,c}$ for all $m \in \mathcal{M}, c \in \mathcal{C}$.

Second, we establish partial differentiability in each argument. That is, for all $x \in \mathbb{R}_{++}^{|\mathcal{M}|}$, $\frac{\partial}{\partial x_m}\tilde{u}(x) = \tilde{u}^{(m)}(x)$ exists. This follows by Corollary 5 in Milgrom and Segal (2002). Concretely, the domain of optimization can be taken to be a compact and convex subset of a normed vector space – a sufficiently large cube in $\mathbb{R}_+^{|\mathcal{M}|\times|\mathcal{C}|}$ equipped with the Euclidean norm, for example. The objective function does not depend on x, and constraints are linear in x (and therefore both continuous and continuously differentiable). Moreover, as $x \gg 0$, there exists a $\{x_{m,c}\}$ that satisfies all constraints with strict inequality. □

It follows that the objective function of the composite authority satisfies Assumption 1, and so Proposition 5 implies that the non-separable APM $\tilde{A}_m(y,s) = h^{-1}\left(h(s) + \tilde{u}^{(m)}(y)\right)$ uniquely implements the first-best optimal allocation for the composite authority.

It remains to establish that the quota functions implement the optimal allocation $\{x_{m,c}\}$ conditional on $\{x_m\}$. Let λ_m be the Lagrange multiplier on the x_m constraint, γ_c be the Lagrange multiplier on the q_c constraint and $\underline{\kappa}_{m,c}$ be the Lagrange multiplier on the positivity constraint. Under our maintained Inada condition, we have that $\underline{\kappa}_{m,c} = 0$. Moreover, by Corollary 5 in Milgrom and Segal (2002), we have that $\tilde{u}^{(m)}(x) = \lambda_m$, $\tilde{u}_{q_c}(x) = \gamma_c$, and $u'_{m,c}(x^*_{m,c}) = \lambda_m + \gamma_c - \underline{\kappa}_{m,c}$. Hence, we obtain that:

$$x^*_{m,c} = \left(u'_{m,c}\right)^{-1}\left(\tilde{u}^{(m)}(x) + \tilde{u}_{q_c}(x)\right) \tag{133}$$

Thus, the following profile of quota functions implements the optimal cross-sectional allocation:

$$Q_{m,c}(x) = \left(u'_{m,c}\right)^{-1}\left(\tilde{u}^{(m)}(x) + \tilde{u}_{q_c}(x)\right) \tag{134}$$

Completing the proof. □

The proof of this result constructs a fictitious aggregate authority in our single object setting. The claimed APM is optimal for this aggregate authority by a non-separable adaptation of Theorem 1. The substantial step in this proof establishes that $\tilde{u}$ is increasing, concave, and differentiable by employing the restrictions provided by Assumption 4. Then, given the allocation induced by this APM, we construct the quota function to optimally allocate the level of aggregate admissions induced by the APM.

Intuitively, this mechanism remedies inefficiency by "completing markets." There is a common "market price" for each group given by $\mathcal{P}_m = \tilde{u}^{(m)}(x)$ and an authority-level "shadow price of admissions" $\mathcal{P}_c = \tilde{u}_{q_c}(x)$. Authorities are allocated agents so that the marginal benefit of additional agents equals the sum of the market price and shadow price of admissions $u'_{m,c}(x_{m,c}) = \mathcal{P}_m + \mathcal{P}_c$. Hence, through the completion of markets, a centralized planner can allocate agents efficiently and internalize the externalities that prevented efficiency under the decentralized outcome. Notice that this market involves relatively few prices as it involves only $|\mathcal{M}|+|\mathcal{C}|$ shadow prices rather than the full set of $|\mathcal{M}|\times|\mathcal{C}|$ marginal values.

H Extension of the Main Results to Discrete Economies

In this Appendix, we extend the results in the main text to discrete economies and thereby establish that the core of our analysis generalizes from the continuum framework. Concretely, we show that appropriate analogs of Theorems 1, 2, and 4 carry over to discrete economies. Together, these establish the optimality of APM and characterize the (sub)-optimality of priorities and quotas. However, as discrete economies do not necessarily admit a unique stable matching (as is well known), the first part of Theorem 3 does not hold (uniqueness). This notwithstanding, we demonstrate that agent-proposing DA, when combined with the optimal APM, implements the agent-optimal stable allocation.

H.1 Primitives

An authority has q resources to allocate. At each state ω, the economy the authority faces corresponds to agents $\Theta^\omega = \{\theta_1, \ldots, \theta_{N(\omega)}\}$ where $q \leq |N(\omega)|$. As in the continuum case, $\theta \in [0,1] \times \mathcal{M}$ denotes the type of an agent who has score s and belongs to group m. We denote the score and group of any type θ by $s(\theta)$ and $m(\theta)$, respectively. For simplicity, we assume that no two agents have the same score at any ω, formally, if $\{\theta, \theta'\} \subseteq \Theta^\omega$, then $s(\theta) \neq s(\theta')$.

An allocation $\mu : \Theta \to \{0,1\}$ specifies for any type $\theta \in \Theta$ whether they are assigned to the resource. The set of possible allocations is $\mathcal{U}$ and Ω is the set of all possible economies. An allocation is feasible if it allocates no more than measure q of the resource. A mechanism is a function $\phi : \Omega \to \mathcal{U}$ that returns a feasible allocation for any possible Θ^ω.

The authority believes ω has distribution $\Lambda \in \Delta(\Omega)$. $x(\mu,\omega) = \{x_m(\mu,\omega)\}_{m\in\mathcal{M}}$ denotes the number of agents of each group allocated the resource at matching μ, while $\bar{s}_h(\mu,\omega) = \sum_{\theta\in\Theta^\omega} \mu(\theta)h(s(\theta))$ denotes the utility the authority derives from scores at μ. The preferences of the authority are given by $\xi : \mathbb{R}^{|\mathcal{M}|+1} \to \mathbb{R}$:

$$\xi\left(\bar{s}_h, x\right) \equiv \bar{s}_h + \sum_{m\in\mathcal{M}} u_m(x_m) \tag{135}$$

where h is continuous and strictly increasing and u_m is concave for all $m \in \mathcal{M}$.

H.2 Optimal Mechanisms in Discrete Economies

We adapt our definition of the Adaptive Priority Mechanisms to the discrete setting. An *adaptive priority policy* $A = \{A_m\}_{m \in \mathcal{M}}$, where $A_m : \mathbb{R} \times [0,1] \to \mathbb{R}$. The adaptive priority policy assigns priority $A_m(y_m, s)$ to an agent with score s in group m when y_m of agents of the same group is allocated the object. Given an adaptive priority policy, an APM implements allocations in the following way:

Definition 10 (Adaptive Priority Mechanism)**.** *An adaptive priority mechanism, induced by an adaptive priority A, implements an allocation μ in state ω if the following are satisfied:*

1. *Allocations are in order of priorities: $\mu(\theta) = 1$ if and only if*
   (i) for all θ' with $m(\theta') \neq m(\theta)$ and $\mu(\theta') = 0$,
   $$A_{m(\theta)}(x_{m(\theta)}(\mu, \omega), s(\theta)) \geq A_{m(\theta')}(x_{m(\theta')}(\mu, \omega) + 1, s(\theta')) \tag{136}$$
   (ii) for all θ' with $m(\theta') = m(\theta)$ and $\mu(\theta') = 0$, $s(\theta) > s(\theta')$
2. *The resource is fully allocated:*
   $$\sum_{m \in \mathcal{M}} x_m(\mu, \omega) = q \tag{137}$$

Definition 10 makes two modifications relative to the continuum model. First, the measures of agents from each group are replaced by the number of agents from each group. Second, when $m(\theta) \neq m(\theta')$, the adaptive priority of θ' is now evaluated in the case where an extra agent from $m(\theta')$ is assigned the resource.[44] Unlike the continuum case, it is possible for a monotone APM to implement two different allocations, since it can assign the same priority to two different agents, which could happen only for a zero-measure set of agents in the continuum model.

Define $A^*_m(y_m, s) \equiv h(s) + u_m(y_m) - u_m(y_m - 1)$, which will turn out to be the optimal APM. We first show that A^* is monotone, and all allocations that A^* implements give the authority the same utility.

Lemma 5. *A^* is monotone. Moreover, if A^* implements μ and $\mu' \neq \mu$ in state ω, then $\xi(\mu, \omega) = \xi(\mu', \omega)$.*

Proof. Monotonicity is immediate from the definition of A^* and concavity of u_m. Assume that A^* implements two different allocations, μ and μ' at ω. Let x_l and x'_l denote the number of group $l \in \mathcal{M}$ agents assigned the resource at μ and μ'. Since A^* is monotone and $\mu \neq \mu'$, there are m and n such that $x_m > x'_m$ and $x'_n > x_n$. Let $\tilde{\theta}_l$ and $\tilde{\theta}'_l$ denote the lowest-scoring type l agent assigned the resource at μ and μ', respectively. Similarly, let $\hat{\theta}_l$ and $\hat{\theta}'_l$ denote the highest-scoring type l agents who is not assigned the resource at μ and μ', respectively. Let $\tilde{\mu}$ denote the matching given by: $\tilde{\mu}(\theta) = \mu(\theta)$ if $\theta \notin \{\tilde{\theta}_m, \hat{\theta}'_n\}$, $\tilde{\mu}(\tilde{\theta}_m) = 0$ while $\mu(\hat{\theta}_n) = 1$. $\tilde{\mu}$ starts with μ, takes the resource away from the lowest-scoring group m agent who has it,

[44]This was not the case in the continuum model since all types of agents have measure 0 and therefore replacing θ with θ' has no effect the evaluation of diversity.

$\tilde{\theta}_m$, and allocates it to the highest-scoring group n agent who does not have it, $\hat{\theta}_n$. Note that since A^* is monotone, from $x_m > x'_m$ and $x'_n > x_n$, under μ', $\hat{\theta}_n$ is already allocated the resource while $\tilde{\theta}_m$ is not.

Claim 18. *$\tilde{\mu}$ is implemented under A^* in state ω and $\xi(\mu, \omega) = \xi(\tilde{\mu}, \omega)$.*

Proof. Since A^* implements μ and $\mu(\hat{\theta}_n) = 0$, we have that $A^*_m(s(\tilde{\theta}_m), x_m) \geq A^*_n(s(\hat{\theta}_n), x_n + 1)$. Conversely, since A^* also implements μ' and $\mu'(\hat{\theta}'_m) = 0$, we have that $A^*_n(s(\tilde{\theta}'_n), x'_n) \geq A^*_m(s(\hat{\theta}'_m), x'_m + 1)$. Moreover, since $x_m > x'_m$ and $x'_n > x_n$, we have that $s(\hat{\theta}'_m) \geq s(\tilde{\theta}_m)$ and $s(\hat{\theta}_n) \geq s(\tilde{\theta}'_n)$. From this, it follows that:

$$\begin{aligned} A^*_n(s(\hat{\theta}_n), x_n + 1) &\geq A^*_n(s(\tilde{\theta}'_n), x_n + 1) \geq A^*_n(s(\tilde{\theta}'_n), x'_n) \\ &\geq A^*_m(s(\hat{\theta}'_m), x'_m + 1) \geq A^*_m(s(\hat{\theta}'_m), x_m) \geq A^*_m(s(\tilde{\theta}_m), x_m) \end{aligned} \tag{138}$$

where the first inequality holds as $s(\hat{\theta}_n) \geq s(\tilde{\theta}'_n)$, the second inequality holds as $x'_n > x_n$ (which implies $x'_n \geq x_n + 1$) and A^*_n is decreasing in its second argument, the third inequality holds as A^* also implements μ' (as stated above), the fourth inequality holds as $x'_m < x_m$ (which implies $x'_m + 1 \leq x_m$) and A^*_n is decreasing in its second argument, and the fifth inequality holds as $s(\hat{\theta}'_m) \geq s(\tilde{\theta}_m)$. Thus, $A^*_m(s(\tilde{\theta}_m), x_m) \leq A^*_n(s(\hat{\theta}_n), x_n + 1)$. This shows that $A^*_m(s(\tilde{\theta}_m), x_m) = A^*_n(s(\hat{\theta}_n), x_n + 1)$, which implies that $\tilde{\mu}$ is implemented under A^* and $\xi(\mu, \omega) = \xi(\tilde{\mu}, \omega)$. □

Note that Claim 18 shows that starting from a matching μ which is implemented by A^*, taking away the object from a particular agent who does not have it in μ' and allocating it to a particular agent who has it in μ', we arrive at another matching $\tilde{\mu}$ that is implemented under A^* and gives the authority the same payoff. Therefore, starting from any μ that is implemented by A^* and repeating this construction (by replacing μ at step i with $\tilde{\mu}$ at step $i-1$) where at each step we take the resource from an agent who is not allocated the resource at μ' and assign it to an agent who is, in finitely many steps we arrive at μ'. Since the payoff stays the same at each step, μ' gives the authority the same payoff as μ. □

Theorem 7. *If μ is implemented by A^*, then μ is an optimal matching.*

Proof. First, note that an optimal matching exists since the economy (and therefore the set of matchings) is finite. We first show the following lemma.

Lemma 6. *If μ is not implemented by A^*, then there exists μ' that gives the authority a strictly higher payoff.*

Proof. If μ is not implemented by A^*, then there exists θ and θ' such that $\mu(\theta) = 0$, $\mu(\theta') = 1$ and either $m(\theta) = m(\theta')$ and $s(\theta) > s(\theta')$ or $m(\theta) \neq m(\theta')$ and

$$\begin{gathered} h(s(\theta)) + u_{m(\theta)}(x_{m(\theta)}(\mu) + 1) - u_{m(\theta)}(x_{m(\theta)}(\mu)) > \\ h(s(\theta')) + u_{m(\theta')}(x_{m(\theta')}(\mu)) - u_{m(\theta')}(x_{m(\theta')}(\mu) - 1) \end{gathered} \tag{139}$$

However, in both cases, a μ' that allocates the resource to θ instead of θ' (while not changing any other agent's matching) strictly improves the utility of the authority. □

Lemma 6 proves that the optimal matching cannot be a matching that is not implemented by A^*. Since the optimal matching exists, then it is implemented by A^*. From Lemma 5, all matchings implemented by A^* give the authority the same payoff, proving the result. $\square$

Note that Lemma 5 and Theorem 7 imply that any mechanism that is defined by an arbitrary singleton selection from the set of matchings that A^* implements would achieve an optimal matching under any ω and therefore would be first-best optimal.

H.3 Suboptimality of Priorities and Quotas in Discrete Economies

Now, we define Priority and Quota Mechanisms in the discrete model and extend our (sub)optimality results to discrete economies.

A *priority policy* $P : \Theta \to [0, 1]$ awards an agent of type $\theta \in \Theta$ a priority $P(\theta)$.

Definition 11 (Priority Mechanisms)**.** *A priority mechanism, induced by a priority policy P, allocates the resource in order of priorities until measure q has been allocated, with ties broken uniformly and at random.*

A *quota policy* is given by (Q, D), where $Q = \{Q_m\}_{m\in\mathcal{M}}$ and $D : \mathcal{M}\cup\{R\} \to \{1, 2, \dots, |\mathcal{M}|+1\}$ is a bijection. The vector Q reserves Q_m objects for agents in group m, with residual capacity $Q_R = q - \sum_{m\in\mathcal{M}} Q_m$ open to agents of all types. The bijection D (often called the precedence order) determines the order in which the groups are processed.

Definition 12 (Quota Mechanisms)**.** *A quota mechanism, induced by a quota policy (Q, D), proceeds by allocating $Q_{D^{-1}(k)}$ objects to agents from group $D^{-1}(k)$ (if there are sufficient agents from this group) to the resource in ascending order of k, and in descending order of score within each k. If there are insufficiently many agents of any group to fill the quota, the residual capacity is allocated to a final round in which all agents are eligible.*

We also extend the definitions of risk-neutrality and high risk aversion to the discrete setting. Authority preferences are *non-trivial* if for all $m, n \in \mathcal{M}$:

$$h(1) + (u_n(1) - u_n(0)) > h(0) + (u_m(q) - u_m(q - 1)) \tag{140}$$

The authority is *risk-neutral* if for all $m \in \mathcal{M}$, $u_m(x) = c_m x$ for some $c_m \geq 0$ and all $x \in \{0, 1, \dots, q\}$. Define $\tilde{u}$ and $\tilde{h}$ as follows: there exists x_m^{tar} such that $\tilde{u}_m(x_m+1)-\tilde{u}_m(x_m) = 0$ for all $x_m \geq x_m^{\text{tar}}$ and $\tilde{u}_m(x_m+1)-\tilde{u}_m(x_m) \geq h(1)-h(0)$ for $x_m < x_m^{\text{tar}}$ and where $\sum_{m\in\mathcal{M}} x_m^{\text{tar}} \leq q$. Let $\tilde{\xi}$ denote the preferences of the authority under $\tilde{u}$ and $\tilde{h}$. The authority with preferences ξ is *extremely risk-averse* if the set of optimal allocations under ξ and $\tilde{\xi}$ coincide for all ω.

Theorem 8. *The following statements are true:*

1. *If there is no uncertainty, then there exist first-best priority and quota mechanisms.*
2. *Suppose that the authority has non-trivial preferences. There exists a first-best priority mechanism if and only if the authority is risk-neutral. This mechanism is given by $P(s, m) = s + u_m(1) - u_m(0)$.*
3. *Suppose that the authority has non-trivial preferences. There exists a first-best quota mechanism if and only if the authority is extremely risk-averse. This mechanism is given by $Q_m = x_m^{tar}$ and $D(R) = |\mathcal{M}| + 1$.*

Proof. Part (1):

Claim 19. *Let μ denote an optimal allocation at ω. Then μ is a cutoff matching.*

Proof. If μ is not a cutoff matching, then there exists (s, m) and (s', m) where $\mu(s, m) = 1$, $\mu(s', m) = 0$ and $s' > s$. Define μ' by setting: $\mu'(s, m) = 0$, $\mu(s', m) = 1$ and $\mu(\tilde{s}, \tilde{m}) = \mu(\tilde{s}, \tilde{m})$ for all $(\tilde{s}, \tilde{m})$ such that $(\tilde{s}, \tilde{m}) \notin \{(s, m), (s', m)\}$. Observe that, $\xi(\mu', \omega) - \xi(\mu, \omega) = s' - s > 0$. Therefore, μ is not an optimal allocation, which is a contradiction. □

Let μ denote an optimal allocation under ω, $\{\hat{s}_m(\mu, \omega)\}_{m \in \mathcal{M}}$ denote the cutoff scores at μ and s^* denote an arbitrary number. Any priority policy that assigns $P(\hat{s}_m(\omega), m) = s^*$ for all $m \in \mathcal{M}$ and is strictly increasing in the first argument allocates the resource to any agent who has a higher score than the cutoff for their group and implements the optimal allocation.

Let x_m denote the number of group m agents who are allocated the resource at an optimal allocation under ω. Then a quota policy that sets $Q_m = x_m$ allocates the resource to any agent who has a higher score than the cutoff for their group and implements the optimal allocation.

Part (2): The if part of the result follows from observing the priority policy $P(s, m) = s + u_m(1) - u_m(0)$ is equivalent to the optimal APM A^* under risk neutrality since $u_m(1) - u_m(0) = u_m(y_m + 1) - u_m(y_m)$ for all m, y_m. Thus, by Theorem 7, $P(s, m) = s + u_m(1) - u_m(0)$ is first-best optimal.

To prove the only if part, assume risk neutrality does not hold and let m denote a group such that u_m does not satisfy risk neutrality. For a contradiction, assume that P is an optimal priority policy. First, we observe that $P(s, m)$ must be strictly increasing in s for all m. To see why, assume $P(s, m) = P(s', m)$ where $s > s'$ and just consider an ω where there are $q - 1$ group m agents with scores strictly higher than s, and no other agents. Clearly, the optimal allocation would be to allocate the resource to all agents but (s', m), while P allocates the resource to (s', m) with at least probability $1/2$.

Second, let m denote a group such that u_m does not satisfy risk neutrality. Take another arbitrary group n. We have the following:

Claim 20. *Either (i) there exists $t < q$, s_m, s_n such that*

$$u_m(t + 1) - u_m(t) + h(s_m) = u_n(q - t) - u_n(q - t - 1) + h(s_n) \tag{141}$$

or (ii) there exists $t < q$ such that

$$u_m(t + 1) - u_m(t) + h(1) < u_n(q - t) - u_n(q - t - 1) + h(0) \tag{142}$$

$$u_m(t) - u_m(t - 1) + h(0) > u_n(q - t + 1) - u_n(q - t) + h(1) \tag{143}$$

Proof. From non-triviality, we know that $u_m(1) - u_m(0) + h(1) > u_n(q) - u_n(q - 1) + h(0)$ and $u_n(1) - u_n(0) + h(1) > u_m(q) - u_m(q - 1) + h(0)$. The result then follows from the fact that h is continuous and strictly increasing and u_m and u_n are concave. □

We first prove the result under case (ii). Fix two agents with scores $s_m \in (0, 1)$, who belong to group m and $s_n \in (0, 1)$, who belong to group n. Assume that there are $t - 1$

group m agents and $q-t$ group n agents with higher scores than $\max\{s_n, s_m\}$, so a total of t group m agents and $q-t+1$ group n agents. Note that in this case, only one agent will not be allocated the resource in the optimal allocation, and that would be either (s_m, m) or (s_n, n). From equation 143, (s_m, m) is more preferred than (s_n, n) and therefore it must be that $P(s_n, n) < P(s_m, m)$, as otherwise P would not be optimal. Next, assume that there are t group m agents and $q-t-1$ group n agents with higher scores than $\max\{s_n, s_m\}$. From equation 142, (s_n, n) is more preferred than (s_m, m) and therefore it must be that $P(s_m, m) < P(s_n, n)$, which is a contradiction.

We now prove the result under case (i).

Claim 21. *In case (i), any optimal priority policy P must satisfy $P(s_m+\epsilon, m) > P(s_n, n)$ for all $\epsilon > 0$ and $P(s_m-\epsilon, m) < P(s_n, n)$ for all $\epsilon > 0$*

Proof. From Equation 141, we see that when there are t group m agents and $q-t-1$ group n agents with higher scores, $(s_m+\epsilon, m)$ is strictly preferred to (s_n, n), which is strictly preferred to $(s_m-\epsilon, m)$. □

Since u_m is not linear, there exists an l such that $u_m(l+1)-u_m(l) < u_m(l)-u_m(l-1)$. There are two possibilities: $l \le t$ or $l > t$. First, suppose that $l \le t$. We have that:

$$u_m(l)-u_m(l-1)+h(s_m) > u_m(l+1)-u_m(l)+h(s_m) \ge u_n(q-l)-u_n(q-l+1)+h(s_n) \quad (144)$$

where the first inequality follows from $u_m(l+1)-u_m(l) < u_m(l)-u_m(l-1)$ and the second inequality follows as $u_m(t+1)-u_m(t)+h(s_m) = u_n(q-t)-u_n(q-t-1)+h(s_n)$, u_m and u_n are concave, and $l \le t$. Thus, for sufficiently small $\epsilon > 0$, we have that:

$$u_m(l)-u_m(l-1)+h(s_m-\epsilon) > u_n(q-l)-u_n(q-l+1)+h(s_n) \quad (145)$$

Given this inequality, we see that when there are $l-1$ group m agents and $q-l$ group n agents with higher scores, $(s_m-\epsilon, m)$ is strictly preferred to (s_n, n). Thus, to implement the optimal allocation, it must be that $P(s_m-\epsilon, m) \ge P(s_n, n)$, which is a contradiction to Claim 21.

Second, suppose that $l > t$. We know that:

$$u_m(t+1)-u_m(t)+h(s_m) = u_n(q-t)-u_n(q-t-1)+h(s_n) \quad (146)$$

As $l > t$, from concavity of u_m and u_n,

$$u_m(l)-u_m(l-1)+h(s_m) \le u_n(q-l+1)-u_n(q-l)+h(s_n) \quad (147)$$

From concavity of u_n and u_m:

$$u_m(l+1)-u_m(l)+h(s_m) < u_n(q-l)-u_n(q-l-1)+h(s_n) \quad (148)$$

Thus, for sufficiently small $\epsilon > 0$, we have that:

$$u_m(l+1)-u_m(l)+h(s_m+\epsilon) < u_n(q-l)-u_n(q-l-1)+h(s_n) \quad (149)$$

Given this inequality, we see that when there are l group m agents and $q-l-1$ group n agents with higher scores, (s_n, n) is strictly preferred to $(s_m+\epsilon, m)$. Thus, to implement the optimal allocation, it must be that $P(s_m+\epsilon, m) \leq P(s_n, n)$, which is a contradiction to Claim 21.

Part (3): To prove the if part, fix an ω and let μ^* denote the optimal allocation under ω. Let x^*_m denote the number of group m agents allocated the resource at μ^* and $x_m(\omega)$ denote the total number of group m agents under ω.

Claim 22. *If the authority is extremely risk-averse, then* $x^*_m \geq \min\{x_m(\omega), x^{tar}_m\}$

Proof. Assume for a contradiction this is not the case. Then $x^*_m < x_m(\omega)$ and $x^*_m < x^{tar}_m$. Since $\sum_{m\in\mathcal{M}} x^{tar}_m \leq q$ and $x^*_m < x^{tar}_m$, there exists $n \in \mathcal{M}$ such that $x^*_n > x^{tar}_n$. Let s_n denote the score of the lowest-scoring group n agent who is allocated the resource, and let s_m denote the score of any group m agent who is not allocated the resource, which exists as $x^*_m < x_m(\omega)$. Since the authority is extremely risk-averse, we have the following:

$$h(s_m) + u_m(x^*_m+1) - u_m(x^*_m) > h(s_n) - u_n(x^*_n) + u_m(x^*_n - 1) \tag{150}$$

However, this contradicts the optimality of μ^* and proves the claim. □

Claim 23. *If the authority is extremely risk-averse,* $x^*_m > x^{tar}_m$ *and* $x^*_n > x^{tar}_n$*,* $\mu^*(s,m) = 0$ *and* $\mu^*(s',n) = 1$*, then* $s' > s$*.*

Proof. Assume for a contradiction that $s > s'$.[45] The difference in the utility of the authority when allocating the resource to (s, m) rather than (s', n) is given by

$$h(s) + u_m(x^*_m+1) - u_m(x^*_m) - (h(s') - u_n(x^*_n) + u_m(x^*_n - 1)) = h(s) - h(s') > 0 \tag{151}$$

which is a contradiction to optimality of μ^*. □

The previous two claims show that under any ω, the optimal allocation admits (i) the highest-scoring x^{tar}_m agents from each group (provided that they exist) and (ii) highest-scoring agents who are not in (i), until the capacity is exhausted. Clearly, the quota policy $Q_m = x^{\text{tar}}_m$ and $D(R) = |\mathcal{M}| + 1$ implements this outcome at every ω.

To prove the only if part, assume that $\{Q_m\}_{m\in\mathcal{M}}$ is part of an optimal quota policy.

Claim 24. *For and each* $m, n \in \mathcal{M}$ *and any* t, l *such that* $t \leq Q_m$*,* $Q_m > 0$ *and* $l \geq Q_n$*, we have that:*

$$u_m(t) - u_m(t-1) + h(0) \geq u_n(l+1) - u_n(l) + h(1) \tag{152}$$

Proof. Assume that at ω, there are t group m agents, one of which one has score 0 and $l+1$ group n agents with scores higher than $1-\epsilon_1$ and q agents from other groups who have scores higher than $1-\epsilon_2$, where $\epsilon_1 > \epsilon_2 > 0$. As $t \leq Q_m$ and $Q_n < l+1$, t group m agents and $Q_n < l+1$ group n agents are admitted under Q. Since Q is optimal for all ϵ_1, we must have that:

$$u_m(t) - u_m(t-1) + h(0) \geq u_n(l+1) - u_n(l) + h(1-\epsilon_1) \tag{153}$$

The statement then follows from continuity of h by taking the limit $\epsilon_1 \to 0$. □

[45] Remember that $s' = s$ was ruled out by assumption.

Claim 25. *Merit slots are processed last at the optimal quota policy.*

Proof. For a contradiction, assume there is a merit slot that is processed before a quota slot. Let l denote the last merit slot that precedes a quota slot. Let m denote a group that has a quota slot after l. We consider a state in which: (i) there are q group n agents with scores $\hat{s} - \epsilon_i$, where $\epsilon_i > 0$ for all $i \in \{1, \ldots, q\}$ (let $\underline{\hat{s}}$ denote the score of the highest-scoring agent from this group), (ii) there are Q_m group m agents with scores $\hat{s} + \epsilon_j$ for $j \in \{1, \ldots, Q_m\}$ (let $\overline{\hat{s}}$ denote the score of lowest-scoring agent from this group) and one with score $\hat{s}/2$, and (iii) q agents from other groups with scores in $(\underline{\hat{s}}, \overline{\hat{s}})$. A group m agent with score $\hat{s} + \epsilon_k$ for some k is matched to l, thus $(\hat{s}/2, m)$ is matched to a later quota slot, while some agents with type $(\hat{s} - \epsilon_j, n)$ are rejected for some j. Let $\hat{s} - \epsilon_{j'}$ be the score of the highest-scoring such agent. From the optimality of the quota policy we have that

$$u_m(Q_m + 1) - u_m(Q_m) + h(\hat{s}/2) \geq u_n(Q_n + 1) - u_n(Q_n) + h(\hat{s} - \epsilon_{j'}) \tag{154}$$

Let s^* be the score of the lowest-scoring group n agent (*i.e.*, $s^* = \min_{i \in \{1,\ldots,q\}} \hat{s} - \epsilon_i$). Next, consider the modified version of the above state, all group n agents are the same, but all of the other Q_m group m agents as well as q agents from other groups now have scores in $(s^* - \hat{\epsilon}, s^*)$ and the group m agent who had a score of $\hat{s}/2$ now has a score of $\hat{s}/2 + \hat{\epsilon}$ for $\hat{\epsilon} > 0$. Note that now the group n agent with score $\hat{s} - \epsilon_{j'}$ is allocated the slot l or an earlier slot, while the agent $(\hat{s}/2 + \hat{\epsilon}, m)$ is not allocated to any slot. Thus

$$u_m(Q_m + 1) - u_m(Q_m) + h(\hat{s}/2 + \hat{\epsilon}) \leq u_n(Q_n + 1) - u_n(Q_n) + h(\hat{s} - \epsilon_{j'}) \tag{155}$$

which, since h is strictly increasing, implies that $u_m(Q_m + 1) - u_m(Q_m) + h(\hat{s}/2) < u_n(Q_n + 1) - u_n(Q_n) + h(\hat{s} - \epsilon_{j'})$. This contradicts Equation 154, proving the claim. $\square$

Given the previous two claims, the following claim proves the result.

Claim 26. *If merit slots are processed last, then for all $l \geq Q_m$ and $j \geq Q_n$*

$$u_m(l + 1) - u_m(l) = u_n(j + 1) - u_n(j) \tag{156}$$

Proof. Assume for a contradiction this does not hold. Without loss of generality, assume $u_m(l + 1) - u_m(l) > u_n(j + 1) - u_n(j)$ and define δ as

$$\delta = (u_m(l + 1) - u_m(l)) - (u_n(j + 1) - u_n(j)) \tag{157}$$

Consider a state with $q - 1$ agents with scores higher than s^*, of which exactly Q_m are group m agents and Q_n are group n agents. Moreover, there is one more group m agent with score $s' < s^*$ (denote this agent by θ_m) and one more group n agent with score $s'' \in (s', s^*)$ where $h(s'') - h(s') < \delta$ (denote this agent by θ_n). Note that all agents apart from θ_m and θ_n are allocated the resource before the final merit slot. Moreover, since θ_n has a higher score, she obtains the final merit slot. However, this is a contradiction to the optimality of Q as $h(s'') - h(s') < \delta$ and allocating that resource to θ_m gives the authority higher utility. This proves the claim. $\square$

Taken together, claims 24 and 26 prove that a fictitious authority that is extremely risk-averse with $x_m^{\text{tar}} = Q_m$ agrees with the authority on the optimal allocation, for all ω. To see this, observe that claim 24 implies that diversity preferences dominate any concern for scores when a group is allocated less than Q_m. Moreover, conditional on being allocated at least Q_m, it is as if there is no residual diversity preference, by claim 26. This proves the only if part of (3), which finishes the proof of the result. $\square$

H.4 Characterization of Stable Allocations in Discrete Economies

In this section, we extend our discrete model to the multiple authority case. Of course, in discrete models, there can be multiple stable matchings. This notwithstanding, we show that agent-proposing DA, when combined with the optimal APM, implements the agent-optimal stable allocation.

Let Θ_0 denote the set of agents. $\mathcal{C} = \{c_0, c_1, \ldots, c_{|\mathcal{C}|-1}\}$ denotes the set of authorities. q_c denotes the capacity of authority c and $q_{c_0} \geq |\Theta_0|$. $\theta = (s, m, \succ) \in [0,1]^{|\mathcal{C}|} \times \mathcal{M} \times \mathcal{R} = \Theta$, where $\mathcal{R}$ is set of all complete, transitive, and strict preference relations over $\mathcal{C}$ such that c_0 is less preferred than all $c \in \mathcal{C}$. For each type θ, $s_c(\theta)$ denotes the score of θ at authority c and $m(\theta)$ denotes the group of θ.

A matching in this environment is a function $\mu : \mathcal{C} \cup \Theta \to 2^{\Theta} \cup \mathcal{C}$ where $\mu(\theta) \in \mathcal{C}$ is the authority any type θ is assigned and $\mu(c) \subseteq \Theta$ is the set of agents assigned to authority c, which satisfies $|\mu(c)| \leq q_c$ for all c. $x_c(\mu) = \{x_{m,c}(\mu)\}_{m \in \mathcal{M}}$ denotes the number of agents of each group assigned to authority c at μ while $\bar{s}_{h_c}(\mu) = \sum_{\theta \in \mu(c)} h(s(\theta))$ denotes the score utility the authority derives from μ. The preferences of the authority are given by:

$$\xi_c(\bar{s}_{h_c}, x_c) = \bar{s}_{h_c} + \sum_{m \in \mathcal{M}} u_{m,c}(x_{m,c}) \tag{158}$$

where h_c is continuous and strictly increasing and $u_{m,c} : \mathbb{R} \to \mathbb{R}$ is concave for all $m \in \mathcal{M}$ and $c \in \bar{\mathcal{C}}$.

In Theorem 3, we showed that in the continuum model, there is stable matching and this matching is a cutoff matching. It is well known that in discrete models there may be multiple stable matchings, so the first part of the result does not hold. However, we can extend the second part of Theorem 3. Recall that in the discrete setting, a matching μ is stable if there are no blocking pairs, that is, there does not exist an agent θ and an authority $c \in \mathcal{C}$ (which includes the dummmy authority) such that $c \succ_\theta \mu(\theta)$ and either (i) c does not fill its capacity or (ii) there exists $\theta' \in c$ such that $\xi_c(\bar{s}_{h_c}(\mu'), x_c(\mu')) > \xi_c(\bar{s}_{h_c}(\mu), x_c(\mu))$, where $\mu'(c) = \mu(c) \setminus \theta' \cup \theta$.

Proposition 12. *If μ is a stable matching, then it is a cutoff matching.*

Proof. If μ is not a cutoff matching, then there exist $\theta = (s, m)$, $\theta' = (s', m)$ and $c \in \overline{\mathcal{C}}$ such that $\mu(\theta') = c$, $c >_\theta \mu(\theta)$ and $s > s'$. Define μ' as follows. $\mu'(c) = \mu(c) \backslash \theta' \cup \theta$ and $\mu'(c') = \mu(c')$ for all $c' \neq c$. As $s > s'$ and h_c is strictly increasing, $\xi_c(\bar{s}_{h_c}(\mu'), x_c(\mu')) > \xi_c(\bar{s}_{h_c}(\mu), x_c(\mu))$, which contradicts the stability of μ. $\square$

In discrete markets, there may be multiple stable matchings. We now show that when all authorities use the optimal APM, Agent-Proposing Deferred Acceptance implements the agent-optimal stable matching, in other words, Theorem 4 holds if we replace the unique stable matching with agent optimal stable matching. For this result, we also assume that the authorities preferences are strict in the sense that given the finite economy Θ and two different matchings μ and μ' such that $\mu(c) \neq \mu(c')$, $\xi_c(\bar{s}_{h_c}(\mu), x_c(\mu)) \neq \xi_c(\bar{s}_{h_c}(\mu'), x_c(\mu'))$.

Theorem 9. *Agent-Proposing Deferred Acceptance implements the agent-optimal stable matching when all authorities use the optimal APM.*

Proof. We first define substitutable preferences in this setting, following Definition 6.2 in Roth and Sotomayor (1990). A choice rule satisfies substitutes if an agent θ is chosen from a set of agents Θ and $\theta' \neq \theta$, then θ must be chosen from $\Theta' \equiv \Theta \setminus \theta'$.

We first show that under our assumptions on the preferences, the optimal APM for an authority satisfies this property. Suppose that θ is chosen from some Θ under the optimal APM for some authority. For a contradiction, suppose that θ is not chosen from Θ'. As the optimal APM admits higher-scoring agents before lower-scoring ones, this means that there are strictly fewer $m(\theta)$ agents chosen from Θ'. Let n_m and n'_m denote the number of m group m agents chosen from Θ and Θ'. Then we have $n_{m(\theta)} > n'_{m(\theta)}$. As θ is not chosen from Θ', there must be a group m' such that more m' agents are chosen from Θ' compared to Θ. Let $\hat{\theta}$ denote the highest scoring m' agent that is not chosen from Θ and $\tilde{\theta}$ denote the lowest scoring m' agent chosen from Θ'. Note that $s(\hat{\theta}) \geq s(\tilde{\theta})$, which yields the following inequalities:

$$\begin{aligned} &A_{m(\theta)}(n_{m(\theta)}, s(\theta)) > A_{m(\hat{\theta})}(n_{m(\hat{\theta})}, s(\hat{\theta})) \geq A_{m(\tilde{\theta})}(n'_{m(\tilde{\theta})}, s(\tilde{\theta})) \\ &> A_{m(\theta)}(n'_{m(\theta)}, s(\theta)) \geq A_{m(\theta)}(n_{m(\theta)}, s(\theta)) \end{aligned} \tag{159}$$

where the first inequality holds as θ is chosen from Θ while $\hat{\theta}$ was not chosen (and the preferences were assumed to be strict), the second inequality holds as $n'_{m(\tilde{\theta})} > n'_{m(\hat{\theta})}$ and $s(\hat{\theta}) \geq s(\tilde{\theta})$, the third inequality holds as $\tilde{\theta}$ is chosen from Θ' while θ was not chosen (and the preferences were assumed to be strict), and the fourth inequality holds as $n_{m(\theta)} > n'_{m(\theta)}$, which is a contradiction—as the first and final terms are identical.

Given this, the theorem follows from Theorem 6.8 in Roth and Sotomayor (1990): when authorities have substitutable preferences (and preferences are strict), the agent-proposing Deferred Acceptance algorithm produces the agent-optimal stable matching. ∎

H.5 Dominance of APM in Sequential Discrete Economies

We finally extend the dominance of APM to discrete economies. As in Appendix F, agents apply to the authorities sequentially, who decide which agents to admit. We index the stage of the game by $t \in \mathcal{T} = \{1, \ldots, |\mathcal{C}| - 1\}$. Each stage corresponds to an authority $I(t)$, where $I : \mathcal{T} \to \mathcal{T}$. At each stage t, any unmatched agents choose whether to apply to authority $I(t)$. Given the set of applicants, authority $I(t)$ chooses to admit a subset of these agents. Given this, histories are indexed by the path of the remaining of agents who have not yet matched, $h^{t-1} = (\Theta_0, \Theta_1, \ldots, \Theta_{t-1}) \in \mathcal{H}^{t-1}$. Given each history h^{t-1} and set of applicants $\Theta_c^A \subseteq \Theta$, a strategy for an authority returns a set of agents $\Theta_c^G \subseteq \Theta$ whom they will admit such that $\Theta_c^G \subseteq \Theta_c^A$ and $|\Theta_c^G| \leq q_c$ for each time at which they could move $t \in \mathcal{T}$,

$a_{c,t} : \mathcal{H}^{t-1} \times \mathcal{P}(\Theta) \to \mathcal{P}(\Theta)$, where $\mathcal{P}(\Theta)$ is the power set over Θ. A strategy for an agent returns a choice of whether to apply to authorities at each history and time for all agent types $\theta \in \Theta$, $\sigma_{\theta,t} : \mathcal{H}^{t-1} \to [0,1]$. We moreover say that a strategy $a_{\tilde{c},t}$ for an authority $\tilde{c}$ at time t is *dominant* if it maximizes authority utility regardless of $\{\{a_{c,t}\}_{c \in \mathcal{C}/\{\tilde{c}\}}, \{\sigma_{\theta,t}\}_{\theta \in \Theta}\}_{t \in \mathcal{T}}$ and I.

Theorem 10. *The APM A_c^* is a dominant strategy for all authorities.*

Proof. We prove that APM A_c^* implements a dominant strategy for all authorities in all stages by backward induction. Consider the terminal time $t = |\mathcal{C}| - 1$. Some set of agents $\hat{\Theta} \subseteq \Theta$ applies to the authority. Regardless of $\hat{\Theta}$, by Theorem 7 we have that the set of agents chosen under any selection from APM A_c^* is first-best optimal. Thus, A_c^* is dominant. Consider now any time $t < |\mathcal{C}| - 1$, precisely the same argument applies and A_c^* is dominant. $\square$

References

Mailath, G. J. (1987): "Incentive compatibility in signaling games with a continuum of types," *Econometrica: Journal of the Econometric Society*, pp. 1349–1365.

Milgrom, P., and I. Segal (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70(2), 583–601.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-sided matching: a study in game-theoretic modeling and analysis*. Econometric Society monographs, Cambridge.

Wagener, A. (2006): "Chebyshev's algebraic inequality and comparative statics under uncertainty," *Mathematical Social Sciences*, 52(2), 217–221.