Adaptive Priority Mechanisms in Discrete Economies

Oğuzhan Çelebi* Joel P. Flynn†
Michigan Yale

December 23, 2025

Abstract

In this companion note, we extend the results of Çelebi and Flynn (2025) 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.

*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

1 The Discrete Model

We begin by introducing the discrete analog of the model of Çelebi and Flynn (2025). We refer the reader to Çelebi and Flynn (2025) for additional details and interpretation and here focus on expositing the results.

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 $\xi : \mathbb{R}^{|\mathcal{M}|+1} \to \mathbb{R}$:

$$\xi\left(\bar{s}_h, x\right) \equiv \bar{s}_h + \sum_{m \in \mathcal{M}} u_m(x_m) \tag{1}$$

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

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 1 (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 $\mu(\theta') = 0$,

$$A_{m(\theta)}(x_{m(\theta)}(\mu, \omega), s(\theta)) \geq A_{m(\theta')}(x_{m(\theta')}(\mu, \omega) + 1, s(\theta')) \tag{2}$$

(ii) for all θ' with $m(\theta') = m(\theta)$ and $\mu(\theta') = 0$, $s(\theta) > s(\theta')$

2. *The resource is fully allocated:*

$$\sum_{m \in \mathcal{M}} x_m(\mu, \omega) = q \tag{3}$$

Definition 1 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.[1] 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 1. *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, $\tilde{\theta}_m$, 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 1. *$\tilde{\mu}$ is implemented under A^* in state ω and $\xi(\mu, \omega) = \xi(\tilde{\mu}, \omega)$.*

[1] 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.

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:

$$\begin{aligned} A^*_n(s(\hat{\theta}_n), x_n + 1) &\geq A^*_n(s(\tilde{\theta}'_n), x_n + 1) \geq A^*_n(s(\tilde{\theta}'_n), x'_n) \\ &\geq A^*_m(s(\hat{\theta}'_m), x'_m + 1) \geq A^*_m(s(\hat{\theta}'_m), x_m) \geq A^*_m(s(\tilde{\theta}_m), x_m) \end{aligned} \tag{4}$$

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 1 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 1. *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 2. *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

$$\begin{aligned} &h(s(\theta)) + u_{m(\theta)}(x_{m(\theta)}(\mu) + 1) - u_{m(\theta)}(x_{m(\theta)}(\mu)) > \\ &h(s(\theta')) + u_{m(\theta')}(x_{m(\theta')}(\mu)) - u_{m(\theta')}(x_{m(\theta')}(\mu) - 1) \end{aligned} \tag{5}$$

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. $\square$

Lemma 2 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 1, all matchings implemented by A^* give the authority the same payoff, proving the result. $\square$

Note that Lemma 1 and Theorem 1 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.

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 2 (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 $D : \mathcal{M}\cup\{R\} \to \{1, 2, \ldots, |\mathcal{M}|+1\}$ 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 3 (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 $m, n \in \mathcal{M}$:

$$h(1) + (u_n(1) - u_n(0)) > h(0) + (u_m(q) - u_m(q-1)) \tag{6}$$

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 2. *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 2. *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 1, $P(s, m) = s + u_m(1) - u_m(0)$ 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 3. *Either (i) there exists $t < q$, s_m, s_n such that*

$$u_m(t+1) - u_m(t) + h(s_m) = u_n(q-t) - u_n(q-t-1) + h(s_n) \tag{7}$$

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

$$u_m(t+1) - u_m(t) + h(1) < u_n(q-t) - u_n(q-t-1) + h(0) \tag{8}$$

$$u_m(t) - u_m(t-1) + h(0) > u_n(q-t+1) - u_n(q-t) + h(1) \tag{9}$$

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. $\square$

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 9, (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 8, (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 4. *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 7, 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 $(s_m - \epsilon, m)$. $\square$

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: $l \leq t$ or $l > t$. First, suppose that $l \leq t$. We have that:

$$u_m(l) - u_m(l-1) + h(s_m) > u_m(l+1) - u_m(l) + h(s_m) \geq u_n(q-l) - u_n(q-l+1) + h(s_n) \quad (10)$$

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 $l \leq t$. Thus, for sufficiently small $\epsilon > 0$, we have that:

$$u_m(l) - u_m(l-1) + h(s_m - \epsilon) > u_n(q-l) - u_n(q-l+1) + h(s_n) \quad (11)$$

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 $P(s_m - \epsilon, m) \geq P(s_n, n)$, which is a contradiction to Claim 4.

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

$$u_m(t+1) - u_m(t) + h(s_m) = u_n(q-t) - u_n(q-t-1) + h(s_n) \quad (12)$$

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

$$u_m(l) - u_m(l-1) + h(s_m) \leq u_n(q-l+1) - u_n(q-l) + h(s_n) \quad (13)$$

From concavity of u_n and u_m:

$$u_m(l+1) - u_m(l) + h(s_m) < u_n(q-l) - u_n(q-l-1) + h(s_n) \quad (14)$$

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

$$u_m(l+1) - u_m(l) + h(s_m + \epsilon) < u_n(q-l) - u_n(q-l-1) + h(s_n) \quad (15)$$

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 4.

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

Proof. Assume for a contradiction this is not the case. Then $x_m^* < x_m(\omega)$ and $x_m^* < x_m^{tar}$. Since $\sum_{m\in\mathcal{M}} x_m^{tar} \leq q$ and $x_m^* < x_m^{tar}$, there exists $n \in \mathcal{M}$ such that $x_n^* > x_n^{tar}$. 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:

$$h(s_m) + u_m(x_m^* + 1) - u_m(x_m^*) > h(s_n) - u_n(x_n^*) + u_m(x_n^* - 1) \tag{16}$$

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

Claim 6. *If the authority is extremely risk-averse,* $x_m^* > x_m^{tar}$ *and* $x_n^* > x_n^{tar}$*,* $\mu^*(s, m) = 0$ *and* $\mu^*(s', n) = 1$*, then* $s' > s$*.*

Proof. Assume for a contradiction that $s > s'$.[2] The difference in the utility of the authority when allocating the resource to (s, m) rather than (s', n) is given by

$$h(s) + u_m(x_m^* + 1) - u_m(x_m^*) - (h(s') - u_n(x_n^*) + u_m(x_n^* - 1)) = h(s) - h(s') > 0 \tag{17}$$

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_m^{tar} 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 $Q_m = x_m^{\text{tar}}$ 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 7. *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:*

$$u_m(t) - u_m(t - 1) + h(0) \geq u_n(l + 1) - u_n(l) + h(1) \tag{18}$$

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:

$$u_m(t) - u_m(t - 1) + h(0) \geq u_n(l + 1) - u_n(l) + h(1 - \epsilon_1) \tag{19}$$

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

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

Claim 8. *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 $i \in \{1, \dots, q\}$ (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 $j \in \{1, \dots, Q_m\}$ (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

$$u_m(Q_m + 1) - u_m(Q_m) + h(\hat{s}/2) \geq u_n(Q_n + 1) - u_n(Q_n) + h(\hat{s} - \epsilon_{j'}) \tag{20}$$

Let s^* be the score of the lowest-scoring group n agent (*i.e.*, $s^* = \min_{i \in \{1,\dots,q\}} \hat{s} - \epsilon_i$). 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

$$u_m(Q_m + 1) - u_m(Q_m) + h(\hat{s}/2 + \hat{\epsilon}) \leq u_n(Q_n + 1) - u_n(Q_n) + h(\hat{s} - \epsilon_{j'}) \tag{21}$$

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 20, proving the claim. □

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

Claim 9. *If merit slots are processed last, then for all $l \geq Q_m$ and $j \geq Q_n$*

$$u_m(l + 1) - u_m(l) = u_n(j + 1) - u_n(j) \tag{22}$$

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

$$\delta = (u_m(l + 1) - u_m(l)) - (u_n(j + 1) - u_n(j)) \tag{23}$$

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. □

Taken together, claims 7 and 9 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 7 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 9. This proves the only if part of (3), which finishes the proof of the result. □

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:

$$\xi_c(\bar{s}_{h_c}, x_c) = \bar{s}_{h_c} + \sum_{m \in \mathcal{M}} u_{m,c}(x_{m,c}) \tag{24}$$

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 1. *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. $\mu'(c) = \mu(c) \setminus \theta' \cup \theta$ 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 3 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 3. *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:

$$\begin{aligned} &A_{m(\theta)}(n_{m(\theta)}, s(\theta)) > A_{m(\hat{\theta})}(n_{m(\hat{\theta})}, s(\hat{\theta})) \geq A_{m(\tilde{\theta})}(n'_{m(\tilde{\theta})}, s(\tilde{\theta})) \\ &> A_{m(\theta)}(n'_{m(\theta)}, s(\theta)) \geq A_{m(\theta)}(n_{m(\theta)}, s(\theta)) \end{aligned} \tag{25}$$

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. □

5 Dominance of APM in Sequential Discrete Economies

We finally extend the dominance of APM to discrete economies in which 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 4. *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 1 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

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

ΩÇelebi and Flynn

Çelebi, O., and J. P. Flynn (2025): “Adaptive Priority Mechanisms,” Working Paper 34035, National Bureau of Economic Research.