Promoting External and Internal Equities Under Ex-Ante/Ex-Post Metrics in Online Resource Allocation

Karthik Abinav Sankararaman [1] **Aravind Srinivasan** [2] **Pan Xu** [3]

Abstract

This paper proposes two different models for equitable resource allocation in online settings. The first one is called *external* equity promotion, where sequentially arriving agents are heterogeneous in their external attributes, namely how many resources they demand, which are drawn from a probability distribution (accessible to the algorithm). The focus is then to devise an allocation policy such that every requester can get a fair share of resources *proportional to their demands*, regardless of their arrival time. The second is called *internal* equity promotion, where arriving requesters can be treated homogeneously in external attributes (demands) but are heterogeneous in internal traits such as demographics. In particular, each requester can be identified as belonging to one or several groups, and an allocation of resources is regarded as equitable when every group of requesters can receive a fair share of resources proportional to the percentage of that group in the whole population. For both models above, we consider as the benchmark a clairvoyant optimal solution that has the privilege to access all random demand realizations in advance. We consider two equity metrics, namely *ex-post* and *ex-ante*, and discuss the challenges under the two metrics in detail. Specifically, we present two linear program (LP)-based policies for external equity promotion under ex-ante with independent demands, each achieving an *optimal* CR of $1/2$ with respect to the benchmark LP. For internal equity promotion, we present optimal policies under both ex-ante and ex-post metrics.

[1]Meta AI, Menlo Park, USA [2]Department of Computer Science, University of Maryland, College Park, USA [3]Department of Computer Science, New Jersey Institute of Technology, Newark, USA. Correspondence to: Pan Xu <pxu@njit.edu>.

Proceedings of the 41^{st} International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

1. Introduction

We consider online resource allocation under a typical *non-profit* setting, where limited or even scarce resources are administered by a not-for-profit organization like a government, and our priority is *equity* such that every type of online agent could receive a fair share of the limited supplies. Examples include refugee resettlement (Ahani et al., 2021), allocation of public housing for lower-income families (Benabbou et al., 2018), distribution of emergency aid for natural disasters like wildfires (Wang et al., 2019), allocation of food donation from mobile pantries for needy families (Lien et al., 2014), and the distribution of urgent medical equipment to local hospitals and nursing homes during a pandemic (Manshadi et al., 2021).

In this work, we identify two different classes of models for equitable resource allocation in online settings. In the first class, which we refer to as *external* equity promotion, the arriving requesters can be heterogeneous in their external attributes, namely how much resource they demand, which is drawn from a probability distribution. The focus is then to devise an allocation policy such that every requester can get a fair share of resources *proportional to their demand*, regardless of their arrival time. Some representative works for this class investigate distributing medical suppliers (Manshadi et al., 2021) and allocating food donations to different agencies (Lien et al., 2014). By contrast, in the second class, the arriving requesters all demand exactly one unit of resource. We call this class *internal* equity promotion because the emphasis is on the differences between the demographics of requesters. The requesters can be identified as belonging to one or several groups, and the distribution of resources is regarded as equitable when every group of requesters can receive a fair share of resources proportional to the percentage of that group in the whole population. Internal equity has been studied in ride-hailing services (during peak hours especially), where we need to allocate (limited) available drivers to arriving riders such that riders of each group (based on gender, race, *etc.*) can receive a fair share of drivers (Ma et al., 2022; Nanda et al., 2020; Xu & Xu, 2020). The issue of internal equity has been raised again and well reported in the COVID-19 vaccine distribution. There is much news showing stark racial disparities existing in the early stages of the vaccine rollout across the country (Stolberg, 2021; Fitzsimmons, 2021; Board, 2021; Singer, 2021; Blackstock & Blackstock, 2021; Jones, 2021; Abcarian, 2021; Adamson, 2021). The article (Walker et al., 2021) published on March 5, 2021, in The New York Times, showed that "The vaccination rate for Black Americans is half that of white people, and the gap for Hispanic people is even larger, according to a New York Times analysis of state-reported race and ethnicity information."

In this work, we derive online allocation algorithms that optimize for equity, as well as compare their performance to the equity attainable by offline algorithms, on parsimonious models from both classes. For internal equity, we consider the simplified setting where every agent belongs to their own group. Surprisingly, even this simple problem of rationing a single resource among identical individuals is non-trivial, and we expect our techniques for solving it to be extendable to group settings in the future.

We now present models of external and internal equity promotion where equity can be defined with respect to two different metrics, sometimes called ex-post and ex-ante in the literature (Freeman et al., 2020; Aziz, 2020).

External Equity Promotion (EEP) under Ex-Post and Ex-Ante Metrics. Our setting is mainly inspired by the work (Manshadi et al., 2021). *There is an infinitely-divisible total supply, normalized to* 1*, which is allocated over* n *rounds*. During each round $i \in [n] := \{1, 2, \ldots, n\}$, an agent with random non-negative demand D_i is observed, and the n random demands $\mathcal{D} := (D_1, D_2, \ldots, D_n)$ can be either independent or correlated with each other. Let π denote an online allocation policy (or algorithm).[1] π has access to the distribution of $\mathcal{D}$ in advance, but it must make an irrevocable commitment $X_i \leq \min(D_i, R_i)$ to each realized demand D_i before seeing the realizations of demands D_j for $j > i$, where R_i refers to the resource remaining at the beginning of time i. We note that the allocated amounts $X_1, \ldots, X_n$ are random, depending on the realizations $D_1, \ldots, D_n$ and also potentially on the randomness in allocation policy π. We study the metrics of equity defined as follows.

Ex-Post Equity: $\sigma_P(\pi) = \mathbb{E}_{\mathcal{D},\pi}\big[\min_{i\in[n]} X_i/D_i\big]$;
Ex-Ante Equity: $\sigma_A(\pi) = \min_{i\in[n]} \mathbb{E}_{\mathcal{D},\pi}\big[X_i/D_i\big]$.

Remarks. (1) For any agent i with $D_i = 0$, we assume by default that any policy will make a commitment $X_i = 0$, with X_i/D_i understood to be 1. (2) The expectation in the two definitions above is taken over the randomness in both of the n random demands of $\{D_i | i \in [n]\}$ and the possible randomness in the policy π.

Internal Equity Promotion (IEP) under Ex-Post and Ex-Ante Metrics. Recall that internal equity has been an issue in applications such as ride-hailing services and distributing vaccines, where arriving agents each have a *uniform* and *unsplittable* demand while they are heterogeneous in their internal attributes like demographics. In this paper, we focus on individual-level equity.[2] Assume WLOG that each agent requests a unit resource (indivisible) and each represents a specific group different from others. Suppose we have a serving capacity of $b \in \mathbb{Z}^+$, *i.e.,* we can serve up to b agents, and we expect to see a random number N of arriving agents, where N follows a known Poisson distribution $\mathrm{Pois}(\lambda)$. Consider a given policy π. Similar to external equity, we have that (1) π has access to λ in advance and (2) upon the arrival of each agent, π should decide irrevocably if to serve her or not when the serving capacity remains (If yes, then the serving capacity should be decreased by 1 after serving the arriving agent). For each arriving agent $i \in [N]$, let $X_i = 1$ indicate i is served in π. We define the ex-post and ex-ante equity achieved by π as follows, respectively.

Ex-Post Equity: $\sigma_P(\pi) = \mathbb{E}_{N,\pi}\big[\min_{i\in[N]} X_i\big]$;
Ex-Ante Equity: $\sigma_A(\pi) = \mathbb{E}_N\Big[\min_{i\in[N]} \mathbb{E}_\pi\big[X_i\big]\Big]$.

Remarks. For the ex-post equity above, the expectation is taken over the randomness in the total number of arrivals (N) and that of π; as for the ex-ante equity, the outer expectation is over the total number of arrivals (N), while the inner is over the potential randomness of π.

Competitive Ratio (CR). Recall that all policies considered here are required to make an instant and irrevocable decision upon each arrival of online agents. When it comes to evaluating the performance of an online policy, a common benchmark, called clairvoyant optimal (denoted by OPT), enjoys the privilege that it can optimize its decision after observing all arriving agents. Consider a given metric of equity (either ex-post or ex-ante) and a given objective (maximization of either EEP or IEP). We say a policy π achieves a competitive ratio of $\alpha \in [0, 1]$ if $\sigma(\pi, \mathcal{I}) \geq \alpha \cdot \sigma(\mathrm{OPT}, \mathcal{I})$ for all possible instances $\mathcal{I}$, where $\sigma(\pi, \mathcal{I})$ and $\sigma(\mathrm{OPT}, \mathcal{I})$ denote the respective performance of π and OPT on the instance $\mathcal{I}$ under the pre-specified metric of equity.

In many cases, it is interesting enough to identify an optimal policy. For a given problem and a given metric of equity, we say a policy π is optimal if $\sigma(\pi, \mathcal{I}) \geq \sigma(\pi', \mathcal{I})$ for all possible instances $\mathcal{I}$ and all possible (online) policies π'.[3] *When it comes to policy design, we restrict our attention to policies that run efficiently with a running time polynomial in the input size*. Note that π is 1-competitive shows that

[1] In this paper, we use the two terms "algorithm" and "policy" interchangeably.

[2] See a detailed discussion on the challenge in promoting group-level internal equity in Appendix A.

[3] Here we just require π' should conform to the real-time decision-making requirement: There is no requirement on its time efficiency (*e.g.,* polynomial running time).

π is optimal but not vice versa: The former statement is stronger, indicating that π is not only optimal among all policies but also matches the clairvoyant optimal that can access well in advance all arriving agents together with their demands and demographics.

1.1. Ex-post vs. Ex-ante

In this section, we list a few facts regarding ex-post and ex-ante to expose the subtle differences between the two.

Observation 1. *For internal equity promotion (IEP) under ex-post, the greedy policy (*Greedy*), which will serve any arriving agent until capacity is exhausted, is* 1*-competitive.*

Proof of Observation 1. By definition of ex-post equity under IEP, no policy has any incentive to reserve any capacity for an arriving agent. We can verify that Greedy matches the clairvoyant optimal for any arrival sequence in the way that $\sigma_P(\mathsf{Greedy}) = \sigma_P(\mathrm{OPT}) = 1$ if $N \le b$ and $\sigma_P(\mathsf{Greedy}) = \sigma_P(\mathrm{OPT}) = 0$ if $N > b$, where N is the total number of arrivals and b is the serving capacity. □

Observation 2. *For both IEP and EEP, a clairvoyant optimal algorithm under ex-post can differ from that under ex-ante.*

Proof of Observation 2. Let OPT_P and OPT_A denote the respective clairvoyant optimal policy under ex-post and ex-ante. Consider an instance of IEP where $b = 1$ and $N = 2$ with probability 1. From Observation 1, we see $\mathrm{OPT}_P = \mathsf{Greedy}$ and $\sigma_P(\mathrm{OPT}_P) = 0$. However, OPT_A should serve agents 1 and 2 each with probability $1/2$, with $\sigma_A(\mathrm{OPT}_A) = 1/2$.

For EEP, it is slightly tricky to see that OPT_P can differ from OPT_A. Consider an instance of EEP where $n = 3$, $D_1 = \mathrm{Ber}(1)$, $D_2 = \mathrm{Ber}(1)$, and $D_3 = \mathrm{Ber}(1/2)$, where $\mathrm{Ber}(p)$ denotes a Bernoulli random variable of mean p. We can verify that OPT_A will completely ignore agent $i = 3$ by committing $X_1 = X_2 = 1/2$ to agents $i = 1, 2$. The resulting ex-ante equity is $\sigma_A(\mathrm{OPT}_A) = 1/2$. Note that by definition, the ex-ante equity on agent $i = 3$ in OPT_A should be $1/2 \cdot 1 + 1/2 \cdot 0 = 1/2$. In contrast, OPT_P will commit $X_1 = X_2 = X_3 = 1/3$ when $D_3 = 1$ and $X_1 = X_2 = 1/2$ when $D_3 = 0$. The resulting ex-post equity should be $\sigma_P(\mathrm{OPT}_P) = 1/2 \cdot 1/3 + 1/2 \cdot 1/2 = 5/12$. □

Observation 3. *For external equity promotion (EEP) under ex-post with correlated demands, no policy can achieve a competitive ratio (CR) better than $O(1/\ln n)$.*[4]

[4]Note that in the case of correlated demands, it may take exponential space (in terms of n) to represent the input distribution of n demands. The negative results shown in Observations 3 and 4 hold even if the algorithm has access to this full (possibly exponentially-large) representation.

Example 1. *Consider an instance of EEP due to (Manshadi et al., 2021) as follows. We have n arriving agents and a unit supply. The distribution of $\mathcal{D} = (D_i)$ is specified as $\mathcal{D} = \mathbf{D}_k := (1, \ldots, 1, 0, \ldots, 0)$ with probability $1/n$ for each $k \in [n]$, where $\mathbf{D}_k$ consists of k consecutive ones followed by $n - k$ zeros.* ■

Proof of Observation 3. Focus on the instance as stated in Example 1. We show that **Claim (1)**: for any policy π, it achieves an ex-post equity $\sigma_P(\pi) \le 1/n$; and **Claim (2)**: for the clairvoyant optimal (OPT), it achieves an ex-post equity $\sigma_P(\mathrm{OPT}) \sim \ln n/n$.

We show **Claim (1)** first. Consider any policy π and let $\alpha_i \in [0, 1]$ be the total commitment of resources to agent i when $D_i = 1$. Under the metric of ex-post, we can assume WLOG that $\alpha_1 \ge \alpha_2 \ge \cdots \ge \alpha_n$ (it gains nothing to make a larger commitment than previous ones). Observe that π achieves an ex-post equity equal to $\sigma_P(\pi) = \sum_{i=1}^n \alpha_i/n$. Since the total supply is 1, we have $\sum_{i=1}^n \alpha_i \le 1$, which suggests that $\sigma_P(\pi) \le 1/n$. Now we show **Claim (2)**. Recall that the clairvoyant optimal (OPT) is allowed to optimize its decision after observing the full realization of $\mathcal{D}$. Thus, OPT will choose to allocate an amount of $1/k$ resources to each agent with no-zero demand when $\mathcal{D} = \mathbf{D}_k$. As a result, $\sigma_P(\mathrm{OPT}) = \frac{1}{n}\sum_{k=1}^n 1/k \sim \ln n/n$. By definition of competitive ratio, we establish that any policy will achieve a competitive ratio no more than $O(1/\ln n)$ for EEP under ex-post. □

Observation 4. *For external equity promotion (EEP) under ex-ante with correlated demands, no policy can achieve a competitive ratio (CR) better than $O(1/\ln n)$.*

Example 2. *Consider an instance of EEP as follows. We have n arriving agents and a unit supply. The distribution of $\mathcal{D} = (D_i)$ is specified as $\mathcal{D} = \mathbf{D}'_k := (1, 1, \ldots, 1, \infty, \ldots, \infty)$ with probability $1/n$ for each $k \in [n]$, where $\mathbf{D}'_k$ consists of k consecutive ones followed by $n - k$ infinites.* ■

Proof of Observation 4. We show that on Example 2, **Claim 1**: a clairvoyant optimal (OPT) achieves an ex-ante equity *equal* to $1/n$; and **Claim 2**: any policy π achieves an ex-ante equity at most $O(1/(n \ln n))$.

We show **Claim 1** first. Observe that $\sigma_A(\mathrm{OPT}) \le 1/n$ since for agent $i = n$, the marginal demand distribution has $D_n = 1$ with probability $1/n$ and $D_n = \infty$ with probability $1 - 1/n$. Second, $\sigma_A(\mathrm{OPT}) \ge 1/n$ since there exists at least an (offline) policy ALG that achieves an ex-ante equity equal to $1/n$. ALG operates as follows: For each realization $\mathcal{D} = \mathbf{D}'_k$, ALG will commit $X_i = 1$ for $i = k$ and $X_i = 0$ for all $i \ne k$. In this case, we can verify that $\mathbb{E}[X_i/D_i] = 1/n$ for all $i \in [n]$, and thus, $\sigma_A(\mathrm{ALG}) = 1/n$.

We show **Claim 2**. Consider a given policy π. Let $\alpha_i \in [0,1]$ be the fraction of resource committed to agent i when π observes the realization of $D_i = 1$. Thus, $\mathbb{E}[X_i/D_i] = \alpha_i \cdot (1-(i-1)/n)$. An optimal policy π can be obtained by solving the program below,

$$\max \left[\min_{i \in [n]} \Big(1 - \frac{i-1}{n}\Big)\alpha_i \right] : \sum_{i=1}^{n} \alpha_i = 1, \alpha_i \in [0,1], \forall i.$$

We can verify that the program above has an optimal value of $1/(n\sum_{k=1}^{n} 1/k) \sim 1/(n \ln n)$, suggesting $\sigma_A(\pi) \leq 1/(n \ln n)$. Thus, by definition, we conclude that any policy π achieves a CR of no more than $O(1/\ln n)$ for EEP under ex-ante with correlated demands. $\square$

Example 2 suggests another nuance between ex-post and ex-ante.

Observation 5. *Any policy π achieves an ex-ante equity as least as large as the ex-post equity on any instance $\mathcal{I}$ of EEP or IEP. However, this does not necessarily mean π achieves a competitive ratio under ex-ante as least as large as that under ex-post when restricted on any instance $\mathcal{I}$.*

Proof of Observation 5. Note that the function $f : \mathbb{R}^n \to \mathbb{R}$ defined as $f(x_1, \ldots, x_n) = \min_{1 \leq i \leq n} x_i$ is concave over $\mathbb{R}^n$. Thus, $\mathbb{E}[f(\mathbf{X})] \leq f(\mathbb{E}[\mathbf{X}])$ for any random vector-valued variable $\mathbf{X} \in \mathbb{R}^n$. This suggests that $\sigma_P(\pi) \leq \sigma_A(\pi)$ for all possible instances of EEP and IEP. The second part of the claim can be seen in Example 2 as follows: Consider a policy π^* that will always commit $1/n$ resource to all agent i with $D_i = 1$. We can verify that $\sigma_P(\pi^*) = \sigma_P(\mathrm{OPT}) = 1/n^2$. Thus, π^* is 1-competitive (and surely optimal) on Example 2 under ex-post. However, as Observation 4 indicates, any policy achieves a CR of no more than $O(1/\ln n)$ on Example 2 under ex-ante. $\square$

In this paper, we assume that (1) for EEP (under either ex-ante or ex-post), we focus on *independent demands* only. This is necessary since, otherwise, no constant competitive ratio can be obtained, as suggested by Observations 3 and 4 and (2) for IEP, we focus on the metric of *ex-ante* only since Greedy is 1-competitive (optimal) for ex-post due to Observation 1.

1.2. Main Results

Throughout this paper, we focus on efficiently-computable policies only. Specifically, all policies presented in this paper are computable within polynomial time in terms of the input size.

Theorem 1. *[Section 3] There exists an LP-based policy with attenuations that achieves an optimal competitive ratio (CR) of $1/2$ for EEP under ex-ante.*

Remarks on Theorem 1. (i) In Appendix B, we offer another policy (DTH) achieving the same optimal CR of $1/2$ for EEP under ex-ante. It features a double-layer thresholding system that designates two specific thresholds, namely (p_i, β_i), for each arriving agent i, and these serve different purposes. The first one, $p_i \in [0,1]$, is obtained by solving an alternative benchmark LP (6) and ensures that agent i only gets *served anything at all* if $F_i(D_i) \leq p_i$, where F_i is the CDF of the demand distribution of i. The second one, β_i, ensures that agent i is allocated at most β_i of the resource under any circumstances. DTH distinguishes itself from single-layer threshold-based frameworks widely used in online resource allocation (Alaei et al., 2012; 2013; Manshadi et al., 2021). (ii) The optimalities of the $1/2$-competitiveness for ATT and DTH are with respect to benchmark LP (1) and LP (6), respectively, and the two LPs are equivalent to each other. We are not yet certain if the optimality holds unconditionally, *i.e.,* when directly compared against a clairvoyant optimal. That being said, we manage to get an unconditional upper bound of $1/(2-2/\mathrm{e}) \approx 0.791$; see Section 3.1.

We also consider two special cases: (1) EEP under ex-ante when every demand takes large non-zero values (*i.e.,* at least one) only, and (2) EEP under ex-post when every demand takes small values only. For each case, we offer an optimal (or nearly optimal) policy. *We stress that both optimality results are unconditional,* i.e., *when evaluated directly against a clairvoyant optimal policy*.

Theorem 2. *[Appendix D] There exists an optimal policy for EEP under ex-ante when every agent's demand takes values of either zero or at least one (possibly non-integers).*

Theorem 3. *[Appendix E] There exists a $(1-\epsilon)$-competitive policy for EEP under ex-post when every agent's demand is upper bounded by $K \cdot \epsilon^2/\ln(1/\epsilon))$, where $K > 0$ is an absolute constant.*

Remarks on Theorems 2 and 3. (i) The idea behind Theorem 2 is inspired by the work (Papadimitriou et al., 2021), which introduced the problem of designing approximation algorithms for online matching using an online optimal as the benchmark. We successfully construct an LP that serves as an upper bound for any online optimal and then design an online policy whose performance can match the LP value. (ii) The optimality of the result in Theorem 2) and the near optimality of that in Theorem 3 in CR are both *unconditional*, *i.e.,* when evaluated against a clairvoyant optimal.

Theorem 4. *[Section 4] There is an optimal policy for IEP under ex-ante.*

Remarks on Theorem 4. (1) The optimality is independent of any benchmark, *i.e.,* it is directly compared against a clairvoyant optimal. (2) The result can be generalized to *all* unimodal distributions, including Poisson and binomial distributions, where there is one single peak in the probability

Table 1. Summary of results related to **External Equity Promotion** (EEP) obtained in the paper. "Indep./Corre." in the second row represents "independent" and "correlated" demands, respectively. In the first row, we present only the upper bounds of competitiveness for all online policies under ex-post and ex-ante, respectively. The term "optimal*" at the lower-right corner represents optimality with respect to the benchmark LP (1).

	Ex-Post	Ex-Ante
Corre. Demands	$\sigma_P(\mathrm{ALG}) = O(1/\ln n), \forall\, \mathrm{ALG}$ (**Observation 3**, Section 1.1)	$\sigma_A(\mathrm{ALG}) = O(1/\ln n), \forall\, \mathrm{ALG}$ (**Observation 4**, Section 1.1)
Indep. Demands	$\sigma_P(\tilde{\pi}) = 1 - \epsilon$, for small demands. (**Theorem 3**, Appendix E)	**Algorithm 4** is optimal for large demands. (**Theorem 2**, Appendix D)
	Not considered (for general cases).	**Algorithm 1 (**ATT**)** is optimal* (Section 3). **Algorithm 3 (**DTH**)** is optimal* (Appendix B).

Table 2. Summary of results related to **Internal Equity Promotion** (IEP) obtained in the paper. "Indep./Corre." on the second row represents "independent" and "correlated" demands, respectively. Note that in IEP, each arriving agent is assumed to request one unit (*deterministic and unsplittable*) resource; thus, "Indep./Corre." does not apply here. In the third row, the term "$N \sim \mathrm{Pois}(\lambda)$" means that the total number of arriving agents, denoted by N, follows a Poisson distribution with parameter $\lambda > 0$, where λ is accessible to the algorithm. Note that Greedy is one-competitive under ex-post, implying it is optimal among all online policies. However, we only claim that RD (Algorithm 2) is optimal among all online policies under ex-ante; we have not yet evaluated its competitiveness.

	Ex-Post	Ex-Ante
Indep./Corre. Demands	N/A	N/A
$N \sim \mathrm{Pois}(\lambda)$	Greedy is 1-competitive (**Observation 1**, Section 1.1)	Algorithm 2 is optimal (**Theorem 4**, Section 4)

density function of the random number of total arrivals. (3) Theorem 4 establishes the existence of a policy that is optimal among all possible online policies for Internal Equity Promotion (IEP) under the ex-ante metric. However, we have not yet assessed its competitiveness by comparing it against the clairvoyant optimal policy, which is a promising avenue for future research.

1.3. Other Related Work

For external equity promotion, to the best of our knowledge, the closest work to this paper is due to (Manshadi et al., 2021), which has focused on external equity promotion under ex-post. The key differences from here are that they assume demands can be arbitrarily correlated, and they choose the worst-case performance, *i.e.,* the least ex-post equity possibly achieved over all possible demand distributions. They then compare the performance of each policy against a distribution-free benchmark, which is $\min(1, 1/\mathbb{E}[\sum_{i\in[n]} D_i])$. In contrast, we focus on the external equity promotion under ex-ante and choose as a benchmark the performance of a clairvoyant optimal. Roughly speaking, the study of (Manshadi et al., 2021) aims to devise a policy to maximize its worst-case performance (defined over all possible demand distributions), while we try to find a policy whose gap to a clairvoyant optimal over all possible instances is minimized.

As for internal equity promotion, there are several works that have considered a more general setting than ours. Generally, they assume that there are multiple supplying agents and use a bipartite graph to model the network of static supplying agents and dynamic demanding agents (S. Sankar et al., 2021; Ma et al., 2022; Nanda et al., 2020). They examine equity under the concept of *group-level fairness*. However, none of them has established an optimal policy as we do here.

There is a substantial body of research on the fair allocation of online resources (Banerjee et al., 2023; Zhou et al., 2023; Hosseini et al., 2023; Freeman et al., 2020; Aziz, 2020; Gkatzelis et al., 2021), which have explored settings quite distinct from ours. In these studies, they assume the existence of sets of agents and items (whether divisible or indivisible), with each agent having a particular valuation over the items. Their primary goal is to develop allocation policies that optimize objectives like envy-freeness and proportionality, which differ significantly from our focus. Another line of studies has considered resource allocation in an online setting like ours. However, they either consider adversarial arrival setting, or they focus on regret bounds compared against the optimal (Gkatzelis et al., 2021; Sinha et al., 2024; Cayci et al., 2020; Balseiro et al., 2021).

2. Benchmark Linear Program (LP)

For EEP under ex-ante, recall that we assume (1) we have a unit normalized supply and (2) the demands of the n arriving agents $\{D_i | i \in [n]\}$ are independent of each other. For ease of exposition, we consider a discrete case, where there is a joint support $\{d_j | 1 \leq j \leq m\}$ with $0 \leq d_1 < d_2 < \cdots < d_m$ such that $\Pr[D_i = d_j] = p_{ij}$ for all $i \in [n], j \in [m]$. We use i to index both the arriving time and the arriving agent. Consider a given clairvoyant optimal OPT. Let Z_{ij} be the (random) amount of resources committed to agent i conditioning on $D_i = d_j$ for $i \in [n], j \in [m]$. Thus, $z_{ij} := \mathbb{E}[Z_{ij}/d_j | D_i = d_j]$ is the expected Filling Rate (FR) for agent i when $D_i = d_j$. By definition, we assume $Z_{ij} = 0$ and $z_{ij} = 1$ if $d_j = 0$. Consider the LP below.

$$\max \min_{i \in [n]} \Big(\sum_{j \in [m]} p_{ij} \cdot z_{ij} \Big) \tag{1}$$

$$\sum_{i \in [n]} \sum_{j \in [m]} p_{ij} \cdot z_{ij} \cdot d_j \leq 1 \tag{2}$$

$$z_{ij} \cdot d_j \leq 1 \qquad \forall i \in [n], j \in [m] \tag{3}$$

$$0 \leq z_{ij} \leq 1 \qquad \forall i \in [n], j \in [m] \tag{4}$$

Throughout this paper, we refer to the LP above as LP (1) (with Constraints (2)-(4)).

Lemma 1. *The optimal value of* LP (1) *is a valid upper bound of a clairvoyant optimal for EEP under ex-ante.*

Proof. By definitions of $\{Z_{ij}\}$ and $\{z_{ij}\}$, we see that the ex-ante equity on agent i in OPT should be $\sum_{j \in [m]} p_{ij} \mathbb{E}[Z_{ij}/d_j] = \sum_{j \in [m]} p_{ij} \cdot z_{ij}$. Thus, the objective of LP (1) encodes the exact ex-ante equity achieved in OPT. To prove Lemma 1, it suffices to show that the feasibility of all constraints for $\{z_{ij}\}$. As for Constraint (2): Observe that $\sum_{i \in [n]} \sum_{j \in [m]} p_{ij} \cdot z_{ij} \cdot d_j$ denotes the expected total resources committed to all arriving agents, which should be no larger than the total supply of 1; Constraint (3) follows from $Z_{ij} \leq 1$, and thus, $z_{ij} \cdot d_j = \mathbb{E}[Z_{ij}] \leq 1$; the last constraint holds since z_{ij} represents the expected filling rate for agent i when $D_i = d_j$. □

Remarks on LP (1). First, LP (1) is a valid benchmark for External Equity Promotion (EEP) under ex-ante, regardless of whether the distributions of demands $\{D_i | i \in [n]\}$ are independent or correlated. This can be seen from the above proof of Lemma 1. Based on the current definition of z_{ij}, which represents the expected filling rate when $D_i = d_j$, we can verify that $\mathbf{z} = (z_{ij})$ continue to satisfy all the constraints in LP (1) even when demands are correlated. This leads to the fact that LP (1) remains a valid benchmark for EEP under ex-ante for correlate demands. Second, the equivalence between LP (1) and LP (6) (in Appendix B.1) holds for both independent and correlated demands. Thus, both LPs are valid benchmarks for EEP under ex-ante even for correlate demands.

3. An LP-Based Policy with Attenuations

In the following, we present a policy with simulation-based attenuations. The simulation-based attenuation has been used previously in several stochastic optimization problems, *e.g.,* stochastic knapsack (Ma, 2018), stochastic matching (Adamczyk et al., 2015; Brubach et al., 2020), and matching policy design in ride-hailing platforms (Feng et al., 2019; Dickerson et al., 2020). The main idea is as follows. We have a random event A and a target value $\gamma \in (0, 1)$. Suppose we can show that $\Pr[A] \geq \gamma$, although we are unaware of the exact value of $\Pr[A]$. By applying Monte-Carlo simulation, we can first get a sharp estimate of $\Pr[A]$, and then apply attenuations to make A occur with a probability almost equal to γ.

Let $\{z_{ij}\}$ be an optimal solution to LP (1). The overall picture of our algorithm is as follows. Let $R_i \in [0, 1]$ be the remaining supply at (the beginning of) time i with $R_1 = 1$. We aim to achieve an expected filling rate *equal* to $z_{ij}/2$ when $D_i = d_j$ for all i, j. This will lead to a $1/2$-competitive policy by linearity of expectation.[5] We will show that before attenuation, the expected filling rate $\mathbb{E}[\min(d_j, R_i)]/d_j \geq z_{ij}/2$, and that we can make this an equality by adding proper attenuations. The full details of our algorithm (ATT) are as follows.

Lemma 2. ATT *is valid with β_{ij} as specified in Step 6.*

Proof. We prove by induction over time $i \in [n]$. Consider the base case when $i = 1$ with $R_i = 1$. For a given j with $d_j \geq 1$, we have $\mathbb{E}[\min(1, R_1/d_j)] = 1/d_j \leq 1$. Thus,

$$\beta_{ij} = (z_{ij}/2)/\mathbb{E}[\min(1, R_i/d_j)] = z_{ij} \cdot d_j/2 \leq 1/2,$$

which is due to the feasibility of $\{z_{ij}\}$ in Constraint (3) of LP-(1). For a given j' with $d_{j'} < 1$, we see that $\beta_{i,j'} \leq 1/2$ by definition.

[5] We technically achieve a competitive ratio of $1/2 - 1/n^c$ for any desired constant $c > 0$, since the Monte-Carlo simulation has some (controllable) error. For convenience, we omit the lower-order term for clarity.

Algorithm 1 An LP-based policy for external equity promotion under ex-ante (ATT).

1: **Offline Phase**:
 ▷ *The offline phase takes as input the distributions of* $\{D_i | i \in [n]\}$*, and output* $\{\beta_{ij}\}$*, where* $\beta_{ij} \in [0, 1]$ *denotes the attenuation factor applied to agent* i *when* $D_i = d_j$*.*◁
2: Solve LP (1) and let $\{z_{ij}\}$ be an optimal solution.
3: *Initialization*: When $i = 1$, set $\beta_{ij} = (z_{ij}/2)/\mathbb{E}[\min(1, R_i/d_j)]$ for all $j \in [m]$ with $R_i = 1$.
4: **for** $i = 2, \cdots, n$ **do**
5: Applying Monte-Carlo method to simulate Step 10 to Step 13 for all the rounds $i' = 1, 2, \cdots, i-1$ of Online Phase, we can get a sharp estimate of $\mathbb{E}[\min(1, R_i/d_j)]$ for all $j \in [m]$, where $R_i \in [0, 1]$ denotes the (random) remaining supply at the beginning of time i.
6: Set $\beta_{ij} = (z_{ij}/2)/\mathbb{E}[\min(1, R_i/d_j)]$ for all $j \in [m]$. ▷ *It is valid due to Lemma 2.*◁
7: **end for**
8: **Online Phase**:
9: **for** $i = 1, \ldots, n$ **do**
10: Let $R_i \in [0, 1]$ be the remaining supply at i.
11: **end for**
12: **if** Agent i arrives with $D_i = d_j$ **then**
13: With probability β_{ij}, we commit an amount of $\min(d_j, R_i)$ resources to i; with probability $1 - \beta_{ij}$, commit none.
14: **end if**

Now, consider a given $i > 1$ and assume that $0 \leq \beta_{i',j} \leq 1$ for all $1 \leq i' < i$ and $j \in [m]$. For each agent $1 \leq i' < i$, let $X_{i'}$ be the expected amount of resources committed to i' in ATT. By the nature of ATT, we see that

$$\begin{aligned}
\mathbb{E}[X_{i'}] &= \sum_{j \in [m]} p_{i',j} \cdot \beta_{i',j} \cdot \mathbb{E}[\min(R_{i'}, d_j)] \\
&= \sum_{j \in [m]} p_{i',j} \cdot \frac{z_{i',j}/2}{\mathbb{E}[\min(1, R_{i'}/d_j)]} \cdot \mathbb{E}[\min(R_{i'}, d_j)] \\
&= \sum_{j \in [m]} p_{i',j} \cdot (z_{i',j}/2) \cdot d_j.
\end{aligned}$$

Thus, the expected value of the remaining supply at the beginning of i should satisfy

$$\begin{aligned}
\mathbb{E}[R_i] &= 1 - \sum_{i' < i} \mathbb{E}[X_{i'}] \\
&= 1 - \sum_{i' < i} \sum_{j \in [m]} p_{i',j} \cdot (z_{i',j}/2) \cdot d_j \geq 1/2.
\end{aligned}$$

Consider a given j with $d_j \geq 1$. We see that $\beta_{ij} = (z_{ij}/2)/\mathbb{E}[\min(1, R_i/d_j)] = z_{ij}/(2\mathbb{E}[R_i]) \leq z_{ij} \leq 1$. Similarly, for a given j' with $d_{j'} < 1$, we have $\beta_{i,j'} \leq \beta_{ij} \leq 1$. □

Now we have all the ingredients to prove Theorem 1.

Proof. For each agent $i \in [n]$, let X_{ij} be the amount of resources committed to i conditioning on $D_i = d_j$ in ATT. By the nature of ATT, we have

$$\begin{aligned}
\mathbb{E}\Big[X_{ij} \mid D_i = d_j\Big] &= \beta_{ij} \cdot \mathbb{E}[\min(d_j, R_i)] \\
&= \frac{z_{ij}/2}{\mathbb{E}[\min(1, R_i/d_j)]} \cdot \mathbb{E}[\min(d_j, R_i)] = z_{ij}/2 \cdot d_j.
\end{aligned}$$

Therefore, the expected filling rate of agent i in ATT, denoted by FR_i, should be

$$\mathsf{FR}_i = \sum_{j \in [m]} p_{ij} \cdot \mathbb{E}\big[X_{ij}/d_j \mid D_i = d_j\big] = \sum_{j \in [m]} p_{ij} \cdot z_{ij}/2.$$

Thus, the ex-ante equity achieved by ATT should have

$$\begin{aligned}
\sigma_A(\mathsf{ATT}) &= \min_{i \in [n]} \mathsf{FR}_i = \min_{i \in [n]} \Big(\sum_{j \in [m]} p_{ij} \cdot z_{ij}/2 \Big) \\
&= \frac{1}{2} \Big(\min_{i \in [n]} \sum_{j \in [m]} p_{ij} \cdot z_{ij} \Big) \geq \frac{\mathrm{OPT}}{2},
\end{aligned}$$

where the last inequality is due to Lemma 1 and where OPT denotes the ex-ante equity achieved by a clairvoyant optimal policy. □

3.1. Upper Bound of $1/2$ Relative to LP (1)

We now show that without any assumptions, it is not possible to be more than 1/2-competitive relative to LP (1), as demonstrated by the following example.

Example 3. *Consider an instance of EEP with IID demands for agents* $i = 1, \ldots, n$ *that satisfy*

$$D_i = \begin{cases} 1, & w.p.\ 1/n; \\ \infty, & w.p.\ 1 - 1/n. \end{cases}$$

We can verify that the LP can achieve an ex-ante equity of $1/n$. However, now consider any online algorithm with ex-ante equity $\sigma_A(\text{ALG}) \in [0, 1/n]$. This means $\mathbb{E}[X_i/D_i] \geq \sigma_A(\text{ALG})$ for all i. Since $\mathbb{E}[X_i/D_i] = \frac{1}{n} \cdot \mathbb{E}[X_i|D_i = 1] + (1 - \frac{1}{n}) \cdot 0$ for all i, this implies that $\mathbb{E}[X_i|D_i = 1] \geq n\sigma_A(\text{ALG})$ for all $i < n$. Hence, $\sum_{i<n} \mathbb{E}[X_i] \geq (n-1)\frac{1}{n}n\sigma_A(\text{ALG}) = (n-1)\sigma_A(\text{ALG})$ and we have that $\mathbb{E}[R_n] \leq 1 - (n-1)\sigma_A(\text{ALG})$. Now, we have

$$\begin{aligned}\sigma_A(\text{ALG}) &\leq \mathbb{E}[X_n/D_n] = \frac{1}{n} \cdot \mathbb{E}[X_n|D_n = 1] \\ &\leq \frac{1}{n} \cdot \mathbb{E}[R_n] \leq \frac{1}{n}(1 - (n-1)\sigma_A(\text{ALG}))\end{aligned}$$

which implies that $n\sigma_A(\text{ALG}) \leq 1 - (n-1)\sigma_A(\text{ALG})$. In other words, $\sigma_A(\text{ALG}) \leq \frac{1}{2n-1}$. Therefore,

$$\frac{\sigma_A(\text{ALG})}{\sigma_A(\text{LP})} \leq \frac{1/(2n-1)}{1/n} = \frac{n}{2n-1}.$$

Remarks. As a byproduct, Example 3 suggests an *unconditional* upper bound of $1/(2 - 2/\mathsf{e}) \approx 0.791$ for EEP under ex-ante. Note that for Example 3, the performance of any *clairvoyant* optimal (OPT) is no more than $(1/n)(1 - 1/\mathsf{e})$. This can be seen as follows: By symmetry, we assert that OPT would distribute the demand 1 among all arriving demands with realizations of 1. Consider a specific i with $D_i = 1$, which occurs with probability $1/n$. The expected allocation given $D_i = 1$ should be $\mathbb{E}[X_i \mid D_i = 1] = \mathbb{E}[1/(1 + Z)] = 1 - 1/\mathsf{e}$, where $Z \sim \text{Pois}(1 - 1/n)$ denotes the number of realizations of 1 among the remaining agents. Thus, no online policy can achieve a CR better than $(1/(2n-1))/((1 - 1/\mathsf{e}) \cdot (1/n)) = 1/(2 - 2/\mathsf{e})$ when $n \to \infty$.

4. Internal Equity Promotion under Ex-Ante

We start with the case when the probability density function of the random number of total arrivals (N) has a single peak, which is known as a unimodal distribution.[6] Examples include Poisson and binomial distributions. Consider a given (online) policy π. Recall that (1) π has access to the distribution of N but has no idea when the arrival process stops and (2) upon the arrival of one agent, π should decide immediately and irrevocably if to serve her in case there is at least one serving capacity. These two properties suggest that π just needs to *non-adaptively* select a sequence of values $\mathbf{p}_N = \{p_{N,k} | k = 1, 2, 3, \ldots\}$, where $p_{N,k} \in [0, 1]$ denotes the probability that π should accept the kth arriving agent in case capacity remains. In the following, we try to figure out how to compute an optimal sequence $\mathbf{p}_N$ to maximize the resulting ex-ante equity. Throughout this section, we refer to the kth arriving agent as k for simplicity.

[6]More precisely, we assume there is one single set S consisting of either one single or multiple consecutive integer values on which the probability density function gets locally maximized.

Assume that π is parameterized by an integer $M \gg 1$ such that π will ignore all arrivals of k with $k > M$. For each $k \in [M] = \{1, 2, \ldots, M\}$, let $c_k = \Pr[N = k]$ and α_k be the probability that k (*i.e.,* the kth agent) will be accepted in π *unconditionally*: Note that α_k includes the probability that π accepts k and that π has at least one remaining capacity for the kth arrival, while α_k excludes the probability that k arrives. Let $c_0 = \Pr[N = 0]$. By definition of the ex-ante equity, we have

$$\sigma_A(\pi) = \mathbb{E}_N\left[\min_{i \in [N]} \mathbb{E}_\pi[X_i]\right] = c_0 + \sum_{k=1}^{M} c_k \cdot \min_{1 \leq i \leq k} \alpha_i.$$

Recall that $b \in \mathbb{Z}^+$ denotes the total serving capacity. Consider the below program:

$$\max_{\boldsymbol{\alpha}} F(\boldsymbol{\alpha}) := c_0 + \sum_{k=1}^{M} c_k \cdot \min_{1 \leq i \leq k} \alpha_i, \tag{5}$$
$$\text{subject to} \sum_{k=1}^{M} \alpha_k \leq b; \alpha_k \in [0, 1], \forall k \in [M].$$

Lemma 3. *There exists an optimal solution $\boldsymbol{\alpha} = (\alpha_k)$ to Program (5) in which $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_M$.*

Proof. Suppose there exists an optimal solution $\boldsymbol{\alpha} = (\alpha_k)$ in which there exist two indices $\ell > m$ such that $\alpha_\ell < \alpha_m$. Observe that we can always update the current solution $\boldsymbol{\alpha}$ to another $\boldsymbol{\alpha}'$ by decreasing α_m to α_ℓ. We can verify that $\boldsymbol{\alpha}'$ is feasible and the objective value remains the same. By applying this argument repeatedly, we get the claim. □

For each given positive integer ℓ, let $g(\ell) := \Pr[1 \leq N \leq \ell]/\ell$. Let $\ell^* = \text{argmax}_{\ell=1,2,\ldots} g(\ell)$. Under the current unimodal assumption, the probability density function of N has a single peak, and so does the function g. Since the function g gets maximized at $\ell = \ell^*$, we claim that $g(\ell)$ will first increase when $1 \leq \ell \leq \ell^*$ and then decrease afterwards at $\ell > \ell^*$.

Lemma 4. *Consider a given integer b and $M \gg 1$. There exists an optimal solution $\boldsymbol{\alpha} = (\alpha_k)$ to Program (5) such that (1) If $b \leq \ell^*$, then $\alpha_i = b/\ell^*$ for all $1 \leq i \leq \ell^*$ and $\alpha_i = 0$ for all $i > \ell^*$; (2) If $b > \ell^*$, then $\alpha_i = 1$ for all $1 \leq i \leq b$ and $\alpha_i = 0$ for all $i > b$.*

Proof. From Lemma 3, we can assume WLOG that there is an optimal solution $\boldsymbol{\alpha} = \{\alpha_i\}$ to Program (5) with $\alpha_i \geq \alpha_{i+1}$ for all $i \geq 1$.

Focus on the first case $b \leq \ell^*$. Suppose there exists a $k < \ell^*$ such that $\alpha_1 = \alpha_2 = \cdots = \alpha_k > \alpha_{k+1}$. We can

Algorithm 2 An optimal Randomized-Rounding-based policy for IEP under ex-ante (RD).

1: **Offline Phase**:
$\triangleright$ *The offline phase takes as input $\boldsymbol{\alpha}^*(N,b) \in [0,1]^B$, and output a random binary vector $\mathbf{A} = \{0,1\}^B$, where $B = \max(b, \ell^*)$, $\ell^* = \text{argmax}_{\ell=1,2,\ldots} \Pr[1 \le N \le \ell]/\ell$, and $\boldsymbol{\alpha}^*(N,b)$ is the truncated vector specified in Lemma 4 with only non-zero entries retained.*$\triangleleft$
2: Apply dependent rounding (Gandhi et al., 2006) to $\boldsymbol{\alpha}^*(N,b) \in [0,1]^B$; let $\mathbf{A} = \{0,1\}^B$ be the random vector output.
3: **Online Phase**:
4: For each $k = 1, 2, \ldots$, accept the kth arrival if $\mathbf{A}[k] = 1$ and $k \le B$; reject it otherwise.

twist the current solution $\boldsymbol{\alpha}$ to $\boldsymbol{\alpha}'$ as follows: decrease every α_i with $i \le k$ by ϵ/k and increase α_{k+1} by ϵ, where $\epsilon = (\alpha_k - \alpha_{k+1})/(1+1/k)$. We can verify that after the twisting, (1) the resulting solution $\boldsymbol{\alpha}'$ keeps being feasible with $\alpha_1 = \cdots = \alpha_k = \alpha_{k+1}$ and the total sum remains invariant; (2) the change in the objective value due to the twisting should be $F(\boldsymbol{\alpha}') - F(\boldsymbol{\alpha}) = \epsilon\big(c_{k+1} - \sum_{1 \le i \le k} c_i/k\big) \ge 0$. This is due to the fact that $g(\ell)$ increases when $1 \le \ell \le \ell^*$, which implies that $c_{k+1} \ge \sum_{1 \le i \le k} c_i/k$ for all $1 \le k < \ell^*$. Apply this argument repeatedly, we end up with another optimal solution satisfying $\alpha_1 = \cdots = \alpha_{\ell^*}$. Now assume K is the largest integer with $\ell^* < K \le M$ such that $\alpha_K > 0$. Now we can apply the following twisting: decrease α_K to 0 and increase each α_i with $1 \le i \le \ell^*$ by α_K/ℓ^*. Note that we can always make it since $b \le \ell^*$. We can verify that (1) the solution after the twist keeps being feasible and (2) the change in the objective value (the objective value after the twisting minus before) should be $\Delta = \alpha_K(\sum_{1 \le i \le \ell^*} c_i/\ell^* - c_K) \ge 0$. This is due to the fact that $g(\ell)$ decreases when $\ell > \ell^*$, which implies that $\sum_{1 \le i \le \ell^*} c_i/\ell^* \ge c_K$ for all $K > \ell^*$. Repeating these procedures, we end up with an optimal solution such that $\alpha_\ell = 0$ for all $\ell > \ell^*$. Thus, we claim that there exists an optimal solution that satisfies $\alpha_i = b/\ell^* \le 1$ for all $1 \le i \le \ell^*$ and all the rest are zeros. We can apply a similar analysis to prove the second case when $b > \ell^*$. $\square$

Lemma 4 offers an exact characterization of an optimal solution to Program (5). For a given random variable N and b, let $\boldsymbol{\alpha}^*(N,b)$ be the solution as specified in Lemma 4. For the second case when $b > \ell^*$, there is a trivial policy π whose performance matches $\boldsymbol{\alpha}^*(N,b)$: π simply accepts all the first b arrivals and ignores all the rest. As for the first case when $b \le \ell^*$, it is slightly tricky to design a policy such that its performance matches that specified in $\boldsymbol{\alpha}^*(N,b)$, *i.e.*, the first ℓ^* arrivals each will be accepted with a probability equal to b/ℓ^*. Let $B = \max(b, \ell^*)$. Observe that for both cases, $\boldsymbol{\alpha}^*(N,b)$ will have its first B entries being non-zero and all the rest being zero.

We present a randomized policy (RD) in Algorithm 2, whose performance matches exactly $\boldsymbol{\alpha}^*(N,b)$ via the dependent-rounding technique (Gandhi et al., 2006). For notation convenience, we use $\boldsymbol{\alpha}^*(N,b)$ to denote the truncated version as well, which includes only the first B non-zero entries.

Lemma 5. *The performance of* RD *matches exactly that specified in $\boldsymbol{\alpha}^*(N,b)$.*

Proof. Consider the case when $b > \ell^*$. We see that both $\mathbf{A}$ and $\boldsymbol{\alpha}^*(N,b)$ will have all B entries equal to 1. This suggests that RD will accept the first b arrivals, each with probability one. As for the case $b \le \ell^*$, by the property of dependent rounding, we see that $\Pr[\mathbf{A}[k] = 1] = b/\ell^*$ for each $1 \le k \le \ell^*$, which suggests that kth arriving agent will be accepted with a probability equal to b/ℓ^*. $\square$

5. Conclusions and Future work

In this paper, we considered EEP and IEP under two different equity metrics, namely ex-post and ex-ante. For EEP under ex-ante with independent demands, we presented two LP-based policies, each achieving a competitive ratio of at least $1/2$, which is optimal against the current benchmark LP. For IEP under ex-ante and ex-post, we presented an optimal policy for each case. Our work opens a few research directions. The first one is EEP under ex-post with independent demands. Can we get a constant-competitive policy similar to that under ex-ante? One direct challenge is to design an appropriate convex program such that (1) it is polynomial-time computable and (2) it can offer a valid upper bound on the performance of a clairvoyant optimal policy. Another question is whether we can beat $1/2$ under the natural setting when each demand takes value in $[0, 1]$. Perhaps we need to tighten the current competitive-ratio analyses and/or design a stronger benchmark LP with additional constraints.

Acknowledgments

Aravind Srinivasan was supported in part by NSF grant CCF-1918749. Pan Xu was partially supported by NSF CRII Award IIS-1948157. The authors gratefully acknowledge the valuable comments provided by the anonymous reviewers. Pan Xu would also like to thank Will Ma for the insightful discussions that contributed to this paper.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Abcarian, R. Column: The answer to equitable distribution of covid-19 vaccines is hiding in plain sight. https://www.latimes.com/opinion/story/2021-02-24/vaccine-equity-community-health-clinic-role, 2021. Accessed: 2021-03-02.

Adamczyk, M., Grandoni, F., and Mukherjee, J. Improved approximation algorithms for stochastic matching. In *Algorithms-ESA 2015: 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pp. 1–12. Springer, 2015.

Adamson, P. Vaccine equity: If we build it, will they come? https://www.statnews.com/2021/02/24/vaccine-equity-if-we-build-it-will-they-come/, 2021. Accessed: 2021-03-20.

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

Alaei, S., Hajiaghayi, M., and Liaghat, V. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 18–35. ACM, 2012.

Alaei, S., Hajiaghayi, M., and Liaghat, V. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pp. 11–25. Springer, 2013.

Aziz, H. Simultaneously achieving ex-ante and ex-post fairness. In *International Conference on Web and Internet Economics*, pp. 341–355. Springer, 2020.

Balseiro, S., Lu, H., and Mirrokni, V. Regularized online allocation problems: Fairness and beyond. In *International Conference on Machine Learning*, pp. 630–639. PMLR, 2021.

Banerjee, S., Gkatzelis, V., Hossain, S., Jin, B., Micha, E., and Shah, N. Proportionally fair online allocation of public goods with predictions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pp. 20–28. International Joint Conferences on Artificial Intelligence Organization, 8 2023.

Benabbou, N., Chakraborty, M., Ho, X.-V., Sliwinski, J., and Zick, Y. Diversity constraints in public housing allocation. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 973–981. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

Blackstock, U. and Blackstock, O. White americans are being vaccinated at higher rates than black americans. such inequity cannot stand. https://www.washingtonpost.com/opinions/2021/02/01/racial-inequality-covid-vaccine/, 2021. Accessed: 2021-03-16.

Board, T. E. How new york's vaccine program missed black and hispanic residents. https://www.nytimes.com/2021/02/01/opinion/new-york-covid-vaccine.html, 2021. Accessed: 2021-03-15.

Brubach, B., Sankararaman, K. A., Srinivasan, A., and Xu, P. Attenuate locally, win globally: An attenuation-based framework for online stochastic matching with timeouts. *Algorithmica*, 82(1):64–87, 2020.

Cayci, S., Gupta, S., and Eryilmaz, A. Group-fair online allocation in continuous time. *Advances in Neural Information Processing Systems*, 33:13750–13761, 2020.

Dickerson, J., Sankararaman, K., Srinivasan, A., and Xu, P. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. *ACM Transactions on Economics and Computation*, 2020.

Feng, Y., Niazadeh, R., and Saberi, A. Linear programming based online policies for real-time assortment of reusable resources. *Chicago Booth Research Paper*, (20-25), 2019.

Fitzsimmons, E. G. Black and latino new yorkers trail white residents in vaccine rollout. https://www.nytimes.com/2021/01/31/nyregion/nyc-covid-vaccine-race.html, 2021. Accessed: 2021-03-05.

Freeman, R., Shah, N., and Vaish, R. Best of both worlds: ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 21–22, 2020.

Gandhi, R., Khuller, S., Parthasarathy, S., and Srinivasan, A. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)*, 2006.

Gkatzelis, V., Psomas, A., and Tan, X. Fair and efficient online allocations with normalized valuations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 5440–5447, 2021.

Hosseini, H., Huang, Z., Igarashi, A., and Shah, N. Class fairness in online matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5673–5680, 2023.

Jones, G. A. How we're working toward racial equity in distributing coronavirus vaccines in d.c. `https://www.washingtonpost.com/opinions/local-opinions/how-were-working-toward-racial-equity-in-distributing-coronavirus-vaccines-in-dc/2021/02/25/13fed062-75e4-11eb-9537-496158cc5fd9_story.html`, 2021. Accessed: 2021-03-20.

Lien, R. W., Iravani, S. M., and Smilowitz, K. R. Sequential resource allocation for nonprofit operations. *Operations Research*, 62(2):301–317, 2014.

Ma, W. Improvements and generalizations of stochastic knapsack and markovian bandits approximation algorithms. *Mathematics of Operations Research*, 43(3):789–812, 2018.

Ma, W., Xu, P., and Xu, Y. Group-level fairness maximization in online bipartite matching. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, pp. 1687–1689, 2022.

Manshadi, V., Niazadeh, R., and Rodilitz, S. Fair dynamic rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 694–695, 2021.

Nanda, V., Xu, P., Sankararaman, K. A., Dickerson, J., and Srinivasan, A. Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 2210–2217, 2020.

Papadimitriou, C., Pollner, T., Saberi, A., and Wajc, D. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 763–764, 2021.

S. Sankar, G., Louis, A., Nasre, M., and Nimbhorkar, P. Matchings with group fairness constraints: Online and offline algorithms. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pp. 377–383, 8 2021.

Singer, N. Where do vaccine doses go, and who gets them? the algorithms decide. `https://www.nytimes.com/2021/02/07/technology/vaccine-algorithms.html`, 2021. Accessed: 2021-03-15.

Sinha, A., Joshi, A., Bhattacharjee, R., Musco, C., and Hajiesmaili, M. No-regret algorithms for fair resource allocation. *Advances in Neural Information Processing Systems*, 36, 2024.

Stolberg, S. G. As biden pushes for racial equity in vaccination, data is lagging. `https://www.nytimes.com/2021/02/09/us/politics/biden-vaccination-race-data.html`, 2021. Accessed: 2021-03-02.

Walker, A. S., Singhvi, A., Holder, J., Gebeloff, R., and Avila, Y. Pandemic's racial disparities persist in vaccine rollout. `https://www.nytimes.com/interactive/2021/03/05/us/vaccine-racial-disparities.html`, 2021. Accessed: 2021-03-15.

Wang, Y., Bier, V. M., and Sun, B. Measuring and achieving equity in multiperiod emergency material allocation. *Risk Analysis*, 39(11):2408–2426, 2019.

Xu, Y. and Xu, P. Trade the system efficiency for the income equality of drivers in rideshare. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 4199–4205, 2020.

Zhou, S., Bai, R., and Wu, X. Multi-agent online scheduling: MMS allocations for indivisible items. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 42506–42516, 2023.

A. Challenges in Generalizing from Individual to Group-Level for Internal Equity Promotion

Focus on the Internal Equity Promotion (IEP). Let us first discuss how to generalize the current ex-post equity metric from individual to group-level. Surprisingly, there seem to be numerous ways to define group-level equity for IEP.

Suppose we have a collection $\mathcal{G} = \{g\}$ of groups, where each group g represents a specific protected class of people (*e.g.,* female Latino). For each $g \in \mathcal{G}$, we assume the total number of arrivals from it, denoted by $\mathcal{N}_g$, follows a Poisson distribution of mean $\lambda_g > 0$, *i.e.,* $\mathcal{N}_g \sim \text{Pois}(\lambda_g)$, and $\{\mathcal{N}_g | g \in \mathcal{G}\}$ are independent from each other. For any online policy π, we assume it has access to $\mathcal{G}$ and $\boldsymbol{\lambda} := \{\lambda_g | g \in \mathcal{G}\}$. For each group g, let X_g denote the total number of agents that belong to g and get served by π. Set $\boldsymbol{\mathcal{N}} = (\mathcal{N}_g)_g$, which denotes the vector of the numbers of arrivals among all groups. Below are a few sample definitions for the group-level equity metric:

$$\sigma_1(\pi) = \mathbb{E}_{\boldsymbol{\mathcal{N}},\pi}\Big[\min_{g\in\mathcal{G}} \frac{X_g}{\mathcal{N}_g}\Big],\ \sigma_2(\pi) = \mathbb{E}_{\boldsymbol{\mathcal{N}},\pi}\Big[\min_{g\in\mathcal{G}} \frac{X_g}{\lambda_g}\Big]$$

$$\sigma_3(\pi) = \mathbb{E}_{\boldsymbol{\mathcal{N}}}\Big[\min_{g\in\mathcal{G}} \mathbb{E}_\pi\Big[\frac{X_g}{\mathcal{N}_g}\Big]\Big],\ \sigma_4(\pi) = \mathbb{E}_{\boldsymbol{\mathcal{N}}}\Big[\min_{g\in\mathcal{G}} \frac{\mathbb{E}_\pi[X_g]}{\mathcal{N}_g}\Big],\ \sigma_5(\pi) = \mathbb{E}_{\boldsymbol{\mathcal{N}}}\Big[\min_{g\in\mathcal{G}} \mathbb{E}_\pi\Big[\frac{X_g}{\lambda_g}\Big]\Big],$$

$$\sigma_6(\pi) = \min_{g\in\mathcal{G}} \mathbb{E}_{\boldsymbol{\mathcal{N}},\pi}\Big[\frac{X_g}{\mathcal{N}_g}\Big],\ \ \sigma_7(\pi) = \min_{g\in\mathcal{G}} \mathbb{E}_{\boldsymbol{\mathcal{N}},\pi}\Big[\frac{X_g}{\lambda_g}\Big].$$

Remarks on the above definitions. (1) There are no concerns for the cases where $\mathcal{N}_g$ appears in the denominator (e.g., σ_1, σ_3, σ_4, and σ_6): we are sure that $X_g = 0$ when $\mathcal{N}_g = 0$, and we can simply ignore this by assuming $0/0 = 1$. We list the above equity metrics in descending order of robustness, from high to low. (2) When each arriving agent represents a specific distinct group, we see that σ_1 is reduced to the ex-post equity metric defined in the paper (since we can simply ignore all those groups with $\mathcal{N}_g = 0$).

In the following, we take σ_1 for example and use it to demonstrate that our result in Observation 1 does not hold for IEP under the group-level equity metric as defined in σ_1. Specifically, we prove that Greedy*, which is to serve whatever arriving agent until capacity is exhausted, is not one-competitive any more*. Consider an input instance with $b = 2$ (serving capacity), and $\mathcal{G} = \{g_1, g_2\}$ with $\lambda_1 = \lambda_2 = 1$. Recall that an online policy is one-competitive iff its performance matches that of OPT (a clairvoyant optimal policy). Thus, to show Greedy is not one-competitive, it suffices to identify an arriving sequence of online agents, on which Greedy achieves an equity strictly worse (lower) than OPT.

Recall that $\mathcal{N}_1$ and $\mathcal{N}_2$ denote the numbers of arrivals of agents from g_1 and g_2, respectively. Consider the scenario when $\mathcal{N}_1 = 2$ and $\mathcal{N}_2 = 1$. For this case, we see that: (1) OPT can achieve an equity of $1/2$ by serving one agent each from g_1 and g_2, respectively. (2) Since Greedy is an online policy, the corresponding equity is determined by the arrival sequence of the three agents. We can verify that Greedy achieves an equity of $1/2$ if the unique agent of g_1, denoted by a, arrives either first or second among the three, which occurs with probability $2/3$; and Greedy achieves an equity of 0 if a arrives last. As a result, conditioning on $\mathcal{N}_1 = 2$ and $\mathcal{N}_2 = 1$, we see Greedy accomplishes an equity of $1/2 \cdot 2/3 + 0 \cdot 1/3 = 1/3$. Thus, we claim that Greedy is surely not one-competitive.

B. A Double-Thresholding LP-Based Policy for EEP under Ex-Ante

B.1. An Alternative LP Equivalent to LP (1)

We note the following alternate formulation of the LP relaxation. Let F_i denote the CDF of the random demand D_i. Then LP (1) can be equivalently written as

$$\max\ \gamma \tag{6}$$

$$\text{s.t.} \sum_{i=1}^{n} \int_0^{p_i} \min\{F_i^{-1}(q), 1\} dq \leq 1 \tag{7}$$

$$\int_0^{p_i} \frac{\min\{F_i^{-1}(q), 1\}}{F_i^{-1}(q)} dq = \gamma, \forall i = 1, \ldots, n \tag{8}$$

$$\gamma, p_1, \ldots, p_n \in [0, 1] \tag{9}$$

The equivalence to LP (1) can be explained as follows. In the original LP, to make $\sum_j p_{ij} z_{ij}$ equal to a target γ for an agent i, it is optimal sorting in increasing order of d_j, and sequentially raise each z_{ij} to the maximum level, *i.e.,* $z_{ij} = \min\{1/d_j, 1\}$, before proceeding to raise the next z_{ij}. In Appendix C, we present a toy example demonstrating the equivalence of two benchmark LPs (1) and (6).

B.2. A Double-Thresholding Policy: Statement and Analysis

Let $p_1, \ldots, p_n$ refer to a fixed optimal LP solution, which defines our algorithm.

Definition 1. *Define $\tilde{D}_i = \min\{D_i, 1\}\mathbb{1}(F_i(D_i) \le p_i)$, which should be interpreted as the effective demand of agent i that is served by the LP. By LP constraint* (7)*, note that we have $\sum_{i=1}^n \mathbb{E}[\tilde{D}_i] \le 1$.*

Our algorithm is based on the following observation. It suffices to focus on the effective demand of each agent i, *i.e.,* set an upper threshold of $F_i^{-1}(p_i)$ which D_i must fall below in order for agent i to be served at all (with tie-breaking defined for discrete distributions by having the agent draw a uniform quantile in [0,1]). To achieve a competitive ratio of β, we claim that it suffices to show $\mathbb{E}[\frac{X_i}{\tilde{D}_i}|F_i(D_i) \le p_i] \ge \beta$ for all i, recalling that X_i is the random variable for the amount of resource rationed to agent i. Moreover, if $\mathbb{E}[\frac{X_i}{\tilde{D}_i}|F_i(D_i) \le p_i] = \beta$, *i.e.,* we don't overserve agent i, then we show that the expected resource consumed by agent i is at most $\beta \cdot \mathbb{E}[\tilde{D}_i]$ (which will be useful for leaving enough for future agents). These two facts are formalized in the following proposition, whose proof boils down to showing that two different correlations both work in our favor.

Proposition 1. *Suppose $\mathbb{E}\Big[\frac{X_i}{\tilde{D}_i} \mid F_i(D_i) \le p_i\Big] = \beta$ for an agent i. Then, we have (1) $\mathbb{E}[\frac{X_i}{D_i}] \ge \beta\gamma$, and (2) $\mathbb{E}[X_i] \le \beta \cdot \mathbb{E}[\tilde{D}_i]$.*

Proof. We prove both statements conditioned on a particular value of R_i. By independence, this does not affect the realization of D_i.

To prove the first statement, we write

$$\begin{aligned}\mathbb{E}\left[\frac{X_i}{D_i}\right] &= p_i \cdot \mathbb{E}\Big[\frac{X_i}{\tilde{D}_i} \cdot \frac{\tilde{D}_i}{D_i}\Big| F_i(D_i) \le p_i\Big] = p_i \cdot \mathbb{E}\Big[\frac{\min\{R_i, \tilde{D}_i, \beta_i\}}{\min\{D_i, 1\}} \cdot \min\{1, \frac{1}{D_i}\}\Big| F_i(D_i) \le p_i\Big] \\ &\ge p_i \cdot \mathbb{E}\Big[\frac{\min\{R_i, \tilde{D}_i, \beta_i\}}{\min\{D_i, 1\}}\Big| F_i(D_i) \le p_i\Big]\mathbb{E}\Big[\min\{1, \frac{1}{D_i}\}\Big| F_i(D_i) \le p_i\Big] = \beta \cdot \gamma.\end{aligned}$$

Note that $\frac{\min\{R_i, \tilde{D}_i, \beta_i\}}{\min\{D_i, 1\}}$ is a decreasing function of D_i, because it equals 1 if $D_i \le \min\{R_i, \beta_i\}$, equals $\min\{R_i, \beta_i\}/D_i$ if $D_i \ge \min\{R_i, \beta_i\}$ but $D_i \le 1$, and equals constant $\min\{R_i, \beta_i\}$ if $D_i \ge 1$. $\min\{1, \frac{1}{D_i}\}$ is also a decreasing function of D_i, which is why we get the inequality above. Finally, the last equality holds because $p_i \cdot \mathbb{E}[\min\{1, \frac{1}{D_i}\}|F_i(D_i) \le p_i] = \mathbb{E}[\frac{\tilde{D}_i}{D_i}] = \gamma$ by LP constraint (8).

To prove the second statement, we write

$$\begin{aligned}\mathbb{E}[X_i] &= p_i \cdot \mathbb{E}\Big[\frac{X_i}{\tilde{D}_i} \cdot \tilde{D}_i\Big| F_i(D_i) \le p_i\Big] = p_i \cdot \mathbb{E}\Big[\min\{\frac{R_i}{\tilde{D}_i}, 1, \frac{\beta_i}{\tilde{D}_i}\} \cdot \tilde{D}_i\Big| F_i(D_i) \le p_i\Big] \\ &\le p_i \cdot \mathbb{E}\Big[\min\{\frac{R_i}{\tilde{D}_i}, 1, \frac{\beta_i}{\tilde{D}_i}\}\Big| F_i(D_i) \le p_i\Big] \cdot \mathbb{E}\Big[\tilde{D}_i\Big| F_i(D_i) \le p_i\Big] = \beta \cdot \mathbb{E}[\tilde{D}_i].\end{aligned}$$

The inequality uses the fact that conditional on $F_i(D_i) \le p_i$, we have $\frac{X_i}{\tilde{D}_i} = \min\{\frac{R_i}{\tilde{D}_i}, 1, \frac{\beta_i}{\tilde{D}_i}\}$ which is decreasing in $\tilde{D}_i$. The final equality uses the fact that $\mathbb{E}[\tilde{D}_i] = p_i \cdot \mathbb{E}[\tilde{D}_i|F_i(D_i) \le p_i]$. □

Note that since γ is an upper bound on the fairness achievable by any clairvoyant (see Lemma 1), the first statement in Proposition 1 suggests the following method for defining a β-competitive algorithm. We first set an upper threshold $\beta_1 \in [0, 1]$ so that $\mathbb{E}[\frac{X_1}{\tilde{D}_1}|F_1(D_1) \le p_1] = \beta$, where β is the fixed competitive ratio we are targeting and $X_1 = \min\{\tilde{D}_1, \beta_1\}$ is the amount serviced to agent 1 (recall that $\tilde{D}_i = 0$ if $F_i(D_i) > p_i$). Next, with R_2 denoting the remaining resource for agent 2, we show that $\mathbb{E}[\frac{X_2}{\tilde{D}_2}|F_2(D_2) \le p_2] \ge \beta$ if $X_2 = \min\{R_2, \tilde{D}_2, \beta_2\}$ with $\beta_2 = 1$. Therefore, it is possible to reduce β_2 so that $\mathbb{E}[\frac{X_2}{\tilde{D}_2}|F_2(D_2) \le p_2] = \beta$. We repeat this process for $i = 1, \ldots, n$, which is formalized in Algorithm 3.

Algorithm 3 A double-thresholding LP-based policy for EEP under ex-ante (DTH).

1: **Offline Phase**:
 ▷ *The offline phase takes as input the distributions of $D_1, \dots, D_n$ and output upper thresholds $\beta_1, \dots, \beta_n$.* ◁
2: *Initialization*: Set $\beta_1 = \cdots = \beta_n = 1$.
3: **for** $i = 1, \cdots, n$ **do**
4: Apply the Monte-Carlo method to simulate **Online Phase** up to the end of round i.
5: Update β_i to be a value in [0,1] so that $\mathbb{E}\left[\frac{X_i}{\tilde{D}_i}\middle| F_i(D_i) \le p_i\right] = \beta$, where $\tilde{D}_i = \min\{D_i, 1\}\mathbb{1}(F_i(D_i) \le p_i)$.
 ▷ *We prove this is always possible in Theorem 5.* ◁
6: **end for**
7: **Online Phase**:
8: **for** $i = 1, \dots, n$ **do**
9: Commit $X_i = \min(R_i, \tilde{D}_i, \beta_i)$ of the resource to agent i.
10: **end for**

Theorem 5. DTH *(Algorithm 3) is valid when $\beta = 1/2$,* i.e., *it is possible to set $\beta_1, \dots, \beta_n \in [0, 1]$ so that* (5) *holds for all i. By Proposition 1,* DTH *is 1/2-competitive.*

Proof. We prove Theorem 5 by induction. Consider the arrival of an agent i and suppose that $\mathbb{E}[\frac{X_{i'}}{\tilde{D}_{i'}}|F_{i'}(D_{i'}) \le p_{i'}] = \beta$, which implies $\mathbb{E}[X_{i'}] \le \beta \cdot \mathbb{E}[\tilde{D}_{i'}]$ for all previous agents $i' < i$, by the second statement of Proposition 1. We show that

$$\mathbb{E}\Big[\frac{\min\{R_i, \tilde{D}_i\}}{\tilde{D}_i}\Big| F_i(D_i) \le p_i\Big] \ge \beta, \tag{10}$$

i.e., agent i can be satisfactorily served as long as we are willing to devote the maximum resource to i whenever $F_i(D_i) \le p_i$. If this is true, then it will be possible to reduce the service so that $\mathbb{E}[\frac{X_i}{\tilde{D}_i}|F_i(D_i) \le p_i] = \beta$, completing the induction and implying that $\min_i \mathbb{E}[\frac{X_i}{D_i}] \ge \beta\gamma$ via the first statement of Proposition 1.

To prove (10), note that R_i is independent of $\tilde{D}_i$, and the function is decreasing in R_i. Moreover, it is also concave in R_i; hence, by Jensen's inequality, its expectation is minimized when all mass is placed on 0 or 1. The expectation of R_i is

$$1 - \sum_{i'<i} \mathbb{E}[X_{i'}] \ge 1 - \beta \sum_{i'<i} \mathbb{E}[\tilde{D}_{i'}] \ge 1 - \beta,$$

hence, the LHS of (10) is lower-bounded by its value when R_i is a random variable that equals 1 w.p. $1-\beta$ and 0 otherwise. Since $\tilde{D}_i \le 1$ by definition, we get

$$\mathbb{E}\Big[\frac{\min\{R_i, \tilde{D}_i\}}{\tilde{D}_i} \;\Big|\; F_i(D_i) \le p_i\Big] \ge 1(1-\beta) + 0(\beta) = 1 - \beta.$$

Therefore, setting $\beta = 1/2$ guarantees that (10) is satisfied for any agent i, completing the proof. □

C. Equivalence of the Two Benchmark LPs (1) and (6)

Lemma 6. *The two benchmark linear programs, LPs* (1) *and* (6)*, are equivalent to each other.*

Proof. Let us start with LP (1), and suppose $\mathbf{z}^* = \{z_{ij}|i \in [n], j \in [m]\}$ be an optimal solution to LP (1) with an optimal value of γ^*. For each agent $i \in [n]$, let $R_i(\mathbf{z}^*)$ denote the expected total amount of resources committed to i and $\gamma_i(\mathbf{z}^*)$ the resulting expected filling ratio with respect to the strategy of $\mathbf{z}^*$. Thus, we have

$$\gamma_i(\mathbf{z}^*) = \sum_{j\in[m]} p_{ij} \cdot z^*_{ij}, \quad R_i(\mathbf{z}^*) = \sum_{j\in[m]} p_{ij} \cdot z^*_{ij} \cdot d_j.$$

By the nature of LP (1), we can assume WLOG that $\gamma_i(\mathbf{z}^*) = \gamma^*$ for all $i \in I$. Moreover, we can safely assume that any clairvoyant optimal policy (OPT) achieves the expected filling ratio of $\gamma_i(\mathbf{z}^*)$ for agent i always by using up all the budget of $R_i(\mathbf{z}^*)$ filling demands in an increasing order of their sizes. Specifically, that means for each agent i, there exists a unique index $j_i(\mathbf{z}^*)$ such that $z^*_{ij} = \min(1, 1/d_j)$ for all $j \le j_i(\mathbf{z}^*)$ and $z^*_{ij} = 0$ for all $j > j_i(\mathbf{z}^*)$.

For each agent i, set $p_i(\mathbf{z}^*) := \sum_{1\le j\le j_i(\mathbf{z}^*)} p_{ij}$. In the following, we show the feasibility of $\{p_i(\mathbf{z}^*)|i\in[n]\}$ and γ^* to LP (5).

$$\begin{aligned}\int_0^{p_i(\mathbf{z}^*)} \min\left(F_i^{-1}(q),1\right)\mathsf{d}q &= \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot\min(d_j,1) = \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot d_j\cdot\min(1,1/d_j)\\ &= \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot d_j\cdot z^*_{ij} = \sum_{j=1}^{m} p_{ij}\cdot d_j\cdot z^*_{ij} = R_i(\mathbf{z}^*).\end{aligned}$$

Thus,

$$\sum_{i\in[n]}\int_0^{p_i(\mathbf{z}^*)} \min\left(F_i^{-1}(q),1\right)\mathsf{d}q = \sum_{i\in[n]} R_i(\mathbf{z}^*)\le 1, \tag{11}$$

where the inequality above is due to the feasibility of $\mathbf{z}^*$ to LP (1). Meanwhile, for each agent $i\in[n]$,

$$\begin{aligned}\int_0^{p_i(\mathbf{z}^*)} \frac{\min\left(F_i^{-1}(q),1\right)}{F_i^{-1}(q)}\mathsf{d}q &= \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot\frac{\min(d_j,1)}{d_j} = \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot\min(1,1/d_j)\\ &= \sum_{j=1}^{j_i(\mathbf{z}^*)} p_{ij}\cdot z^*_{ij} = \sum_{j=1}^{m} p_{ij}\cdot z^*_{ij} = \gamma_i(\mathbf{z}^*) = \gamma^*.\end{aligned}$$

Note that $p_i(\mathbf{z}^*) := \sum_{1\le j\le j_i(\mathbf{z}^*)} p_{ij}\in[0,1]$ and $\gamma^*\in[0,1]$. As a result, we establish the feasibility of $\{p_i(\mathbf{z}^*)|i\in[n]\}$ and γ^* to LP (5), which implies that LP (5) has an optimal value as least as large as LP (1).

Now, we show the opposite direction. Let $\mathbf{p}^* = (p^*_i)$ be an optimal solution to LP (5) with an optimal value of γ^*. For each $i\in[n]$, let $\tilde{j}_i(\mathbf{p}^*)$ be the index of demand such that $F_i^{-1}(p^*_i) = d_{\tilde{j}_i(\mathbf{p}^*)}$. For each $i\in[n]$ and $j\in[m]$, set $z_{ij}(\mathbf{p}^*) = \min(1,1/d_j)$ if $j\le\tilde{j}_i(\mathbf{p}^*)$ and $z_{ij}(\mathbf{p}^*)=0$ if $j>\tilde{j}_i(\mathbf{p}^*)$. Thus, $z_{ij}(\mathbf{p}^*)\in[0,1]$ and $z_{ij}(\mathbf{p}^*)\cdot d_j\in[0,1]$ for any i and j. Moreover,

$$\begin{aligned}\sum_{i\in[n]}\sum_{j\in[m]} p_{ij}\cdot z_{ij}(\mathbf{p}^*)\cdot d_j &= \sum_{i\in[n]}\sum_{1\le j\le\tilde{j}_i(\mathbf{p}^*)} p_{ij}\cdot\min(1,1/d_j)\cdot d_j = \sum_{i\in[n]}\sum_{1\le j\le\tilde{j}_i(\mathbf{p}^*)} p_{ij}\cdot\min(1,d_j)\\ &= \sum_{i\in[n]}\int_0^{p^*_i}\min\left(F_i^{-1}(q),1\right)\mathsf{d}q\le 1,\end{aligned}$$

where the last inequality follows from the feasibility of $\mathbf{p}^*$ to LP (5). Therefore, we establish the feasibility of $\mathbf{z}(\mathbf{p}^*) := \{z_{ij}(\mathbf{p}^*)|i\in[n],j\in[m]\}$ to LP (1). Note that for each given $i\in[n]$,

$$\begin{aligned}\sum_{j\in[m]} p_{ij}\cdot z_{ij}(\mathbf{p}^*) &= \sum_{j\le\tilde{j}_i(\mathbf{p}^*)} p_{ij}\cdot\min(1,1/d_j) = \sum_{j\le\tilde{j}_i(\mathbf{p}^*)} p_{ij}\cdot\frac{\min(1,d_j)}{d_j}\\ &= \int_0^{p^*_i}\frac{\min\left(F_i^{-1}(q),1\right)}{F_i^{-1}(q)}\mathsf{d}q = \gamma^*.\end{aligned}$$

This means that the objective of LP (1) is equal to γ^* on the feasible solution of $\mathbf{z}(\mathbf{p}^*)$. As a result, we claim that LP (1) has an optimal value as least as large as LP (5). □

D. An Optimal Policy for EEP with Ex-Ante under the Large-Demand Assumption

Consider a special case when every D_i with $1\le i\le n$ takes values either 0 or at least 1. WLOG assume that we have a joint support $0=d_1<d_2=1<\cdots<d_m$. Recall that $p_{ij}=\Pr[D_i=d_j]$ for all $i\in[n], j\in[m]$, and $\{D_i\}$ are all independent of each other. We aim to construct a strengthened LP exclusively for an optimal (online) policy (π^*) instead of

a clairvoyant optimal (OPT). For each given $i \in [n], j \in [m]$, let Z_{ij} be the (random) amount of resources committed by π^* to agent i when $D_i = d_j$, and $z_{ij} = \mathbb{E}[Z_{ij}/d_j | D_i = d_j]$ be the expected filling rate for agent i with $D_i = d_j$.

$$\max \min_{i \in [n]} \Big(\sum_{j \in [m]} p_{ij} \cdot z_{ij} \Big) \tag{12}$$

$$1 - \sum_{i' < i} \sum_{j \in [m]} p_{i',j} \cdot z_{i',j} \cdot d_j \geq z_{ij} \cdot d_j \quad \forall i \in [n], j \in [m] \tag{13}$$

$$0 \leq z_{ij} \leq 1 \qquad i \in [n], j \in [m] \tag{14}$$

Remarks on LP (12). Though LP (12) shares similar structures with LP (1), the two serve fundamentally different purposes. As shown in Lemma 7, LP (12) is proposed to upper bound the performance of an online optimal policy, which is subject to the real-time decision-making requirement. This contrasts with that LP (1) is designed to upper bound the performance of a clairvoyant optimal, which has the privilege to access full realizations of random demands before actions.

Lemma 7. *The optimal value of* LP (12) *is a valid upper bound of an (online) optimal policy for EEP under ex-ante.*

Proof. Similar to the proof of Lemma 1, we can verify that the objective of LP-(12) encodes the ex-ante equity achieved by an optimal policy π^*. It will suffice to show the validity of Constraint (13). Consider a given time $i \in [n]$ and let R_i be the remaining supply at (the beginning of) time i. Observe that

$$\mathbb{E}[R_i] = \mathbb{E}\Big[1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot Z_{i',j}\Big] = 1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot \mathbb{E}[Z_{i',j}] = 1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot z_{i',j} \cdot d_j.$$

Observe that $R_i \geq Z_{ij}$ for all $i \in [n]$ and $j \in [m]$, which leads to the fact that

$$\mathbb{E}[R_i] \geq \mathbb{E}[Z_{ij}] \Rightarrow 1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot z_{i',j} \cdot d_j \geq z_{ij} \cdot d_j, \quad \forall i \in [n], j \in [m].$$

This establishes Constraint (13). □

Let $\{z_{ij}\}$ be an optimal solution to LP (12). In the following, we present an LP-based algorithm with simulation-based attenuations (ATT-L) such that it achieves an expected filling rate exactly equal to z_{ij} for each agent $i \in [n]$ when $D_i = d_j$ for each $j \in [m]$.

Algorithm 4 An optimal LP-based policy for EEP under ex-ante with the large-demand assumption (ATT-L).

1: **Offline Phase**:
 ▷ *The offline phase will take as input the distributions of $\{D_i | i \in [n]\}$, and output $\{\beta_{ij}\}$, where $\beta_{ij} \in [0,1]$ denotes the attenuation factor applied to agent i when $D_i = d_j$.* ◁
2: Solve LP (12) and let $\{z_{ij}\}$ be an optimal solution.
3: *Initialization*: When $i = 1$, set $\beta_{ij} = z_{ij} \cdot d_j$ for all $j \in [m]$.
4: **for** $i = 2, \cdots, n$ **do**
5: Applying Monte-Carlo method to simulate Step 10 to Step 13 for all the rounds $i' = 1, 2, \cdots, i-1$ of Online Phase, we can get a sharp estimate of $\mathbb{E}[R_i]$, where $R_i \in [0,1]$ denotes the (random) remaining supply at the beginning of time i.
6: Set $\beta_{ij} = (z_{ij} \cdot d_j)/\mathbb{E}[R_i]$ for all $j \in [m]$.
7: **end for**
8: **Online Phase**:
9: **for** $i = 1, \ldots, n$ **do**
10: Let $R_i \in [0,1]$ be the remaining supply at (the beginning of) i.
11: **if** Agent i arrives with $D_i = d_j$ **then**
12: With probability β_{ij}, we commit an amount of R_i resources to i; with probability $1 - \beta_{ij}$, commit none.
13: **end if**
14: **end for**

Remarks. The last line (13) in ATT-L is always valid since (1) If $d_j = 0$, then we have $\beta_{ij} = 0$ and we will commit none with probability $1 - \beta_{ij} = 1$; If $d_j > 0$, we have $d_j \geq 1 \geq R_i$ that suggests the feasibility of the commitment of R_i for $D_i = d_j$.

Now we prove the main Theorem 2 by showing ATT-L is a *feasible* and *optimal* policy.

Proof. Let X_{ij} be the amount of resources committed to agent i in ATT-L conditioning on $D_i = d_j$ with $i \in [n], j \in [m]$. We show by induction over $i \in [n]$ that for each $i \in [n], j \in [m]$ (1) $\beta_{ij} \leq 1$; (2) the expected filling rate $\mathsf{FR}_{ij} := \mathbb{E}[X_{ij}/d_j] = z_{ij}$.

Consider the base case when $i = 1$. We have $\beta_{ij} = z_{ij} \cdot d_j \leq 1$ due to Constraint (13) for $i = 1$. In this case, we have $R_i = 1$ and thus, $\mathsf{FR}_{ij} = \mathbb{E}[X_{ij}/d_j] = \beta_{ij} \cdot R_i/d_j = z_{ij}$.

Now consider a given $i > 1$ and assume that $\beta_{i',j} \leq 1$ and $\mathsf{FR}_{i',j} = \mathbb{E}[X_{i',j}/d_j] = z_{i',j}$ for all $i' < i$ and $j \in [m]$. Let R_i be the remaining supply at the beginning of i. We have that

$$\begin{aligned}
\mathbb{E}[R_i] &= \mathbb{E}\Big[1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot X_{i',j}\Big] \\
&= 1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot \mathbb{E}[X_{i',j}] \qquad (15) \\
&= 1 - \sum_{1 \leq i' < i} \sum_{j \in [m]} p_{i',j} \cdot z_{i',j} \cdot d_j \;\; \text{(By assumption over all } i' < i.\text{)} \\
&\geq z_{ij} \cdot d_j. \;\; \text{(Due to Constraint (13).)} \qquad (16)
\end{aligned}$$

Therefore, we claim that $\beta_{ij} \leq 1$ for the case of i and all $j \in [m]$. In this case,

$$\mathsf{FR}_{ij} = \mathbb{E}[X_{ij}/d_j] = \beta_{ij} \cdot \mathbb{E}[R_i/d_j] = \frac{z_{ij} \cdot d_j}{\mathbb{E}[R_i]} \cdot \mathbb{E}[R_i/d_j] = z_{ij}.$$

We complete the induction part. Observe the ex-ante equity achieved by ATT-L should be

$$\sigma_A(\mathsf{ATT\text{-}L}) = \min_{i \in [n]} \Big(\sum_{j \in [m]} p_{ij} \cdot \mathsf{FR}_{ij} \Big) = \min_{i \in [n]} \Big(\sum_{j \in [m]} p_{ij} \cdot z_{ij} \Big) \geq \sigma_A(\pi^*),$$

where the last inequality is due to Lemma 7 and where π^* denotes an optimal policy. □

E. An Almost-Optimal Policy for EEP with Ex-Post under the Small-Demand Assumption

Consider EEP and the ex-post setting. We now present a simple condition under which there is an efficient online algorithm $\tilde{\pi}$ whose competitive ratio is at least $(1 - \epsilon)$ times that of any online algorithm, where $\epsilon > 0$ is a given (small) parameter that is bounded away from (and smaller than) 1. This sufficient condition is, for a certain absolute constant $K > 0$, that with probability one, every D_i is at most $K\epsilon^2/\ln(1/\epsilon)$ for all $1 \leq i \leq n$. We will not try to optimize K or the other constants in our proof; the proof below follows by taking K sufficiently small, but positive.

Denote $S = \sum_i D_i$, and let M denote $K\epsilon^2/\ln(1/\epsilon)$. Then, the following are simple consequences of the standard Chernoff-Hoeffding bounds:

$$\begin{aligned}
\forall \lambda \in [0,1], \quad \Pr[S \geq \mu(1+\lambda)] &\leq e^{-\mu\lambda^2/(3M)}; \qquad (17) \\
\forall \lambda \geq 1, \quad \Pr[S \geq \mu(1+\lambda)] &\leq e^{-\mu\lambda \ln(1+\lambda)/(4M)}; \qquad (18) \\
\forall \lambda \in [0,1], \quad \Pr[S \leq \mu(1-\lambda)] &\leq e^{-\mu\lambda^2/(2M)}. \qquad (19)
\end{aligned}$$

These three bounds easily follow by noting that each $D'_i := D_i/M$ lies in $[0,1]$, and by applying the standard Chernoff-Hoeffding bounds to $\sum_i D'_i$. The key parameter to note is the small value M in the denominators of the exponents in all three of these bounds, which leads to strong tail bounds.

Let $\mu = \sum_i \mathbb{E}[D_i]$ denote the total expected arrival of demand and let $\delta := \epsilon/10$. We next present our algorithm $\tilde{\pi}$ and show its near-optimality; we consider two cases based on the value of μ.

Case I: $\mu \leq 1-\delta$**.** Here, $\tilde{\pi}$ simply keeps setting $X_i = d_i$ in response to the i^{th} demand d_i if the remaining budget R_i is at least d_i, and sets $X_i = 0$ otherwise. Note that if the total demand S is at most 1, we achieve a competitive ratio of 1. Thus,

$$\sigma_P(\tilde{\pi}) \geq 1 \cdot \Pr[S \leq 1] = 1 - \Pr[S > 1] \geq 1 - \epsilon,$$

where the final inequality follows from (17) and (18), using our assumption that $\mu \leq 1-\delta$ and by taking the constant $K > 0$ small enough. Since the competitive ratio of any other algorithm is trivially at most 1, we are done in this case.

Case II: $\mu > 1-\delta$**.** Here, $\tilde{\pi}$ keeps setting $X_i = \frac{(1-\delta)}{\mu} \cdot d_i$ in response to the i^{th} demand d_i if the remaining budget R_i is at least $\frac{(1-\delta)}{\mu} \cdot d_i$, and sets $X_i = 0$ otherwise. Note that if the total demand S is at most $\mu/(1-\delta)$, we achieve a competitive ratio of $\frac{(1-\delta)}{\mu}$. Thus, we obtain, in our Case II:

$$\sigma_P(\tilde{\pi}) \geq \frac{(1-\delta)}{\mu} \cdot \Pr[S \leq \mu/(1-\delta)] = \frac{(1-\delta)}{\mu} \cdot (1 - \Pr[S > \mu/(1-\delta)]) \geq \frac{(1-\delta)}{\mu} \cdot (1-\delta)\,, \tag{20}$$

where (20) follows from (17), using our assumption that $\mu > 1-\delta$ and again by taking $K > 0$ small enough.

We now show that no other online algorithm π' can do much better. Splitting based on the value of S, and specifically letting $\mathcal{A}$ denote the event "$S > \mu(1-\delta/2)$," we get

$$\begin{aligned} \sigma_P(\pi') &\leq \Pr[\overline{\mathcal{A}}] \cdot 1 + \Pr[\mathcal{A}] \cdot \mathbb{E}\left[\min_i (X_i/D_i) \;\middle|\; \mathcal{A}\right] \leq \Pr[\overline{\mathcal{A}}] + \mathbb{E}\left[\min_i (X_i/D_i) \;\middle|\; \mathcal{A}\right] \\ &\leq \frac{\delta}{2\mu} + \mathbb{E}\left[\min_i (X_i/D_i) \;\middle|\; \mathcal{A}\right], \end{aligned} \tag{21}$$

where the upper-bound on $\Pr[\overline{\mathcal{A}}]$ used in (21) follows from (19) by taking K small enough again, and via our assumption that $\mu > 1-\delta$.

We now upper-bound $\mathbb{E}\left[\min_i (X_i/D_i) \;\middle|\; \mathcal{A}\right]$. Since

$$\min_i (X_i/D_i) \leq \frac{\sum_i X_i}{\sum_i D_i} = \frac{\sum_i X_i}{S} \leq \frac{1}{S},$$

we get

$$\mathbb{E}\left[\min_i (X_i/D_i) \;\middle|\; \mathcal{A}\right] \leq \mathbb{E}\left[(1/S) \;\middle|\; \mathcal{A}\right] \leq \frac{1}{\mu(1-\delta/2)}. \tag{22}$$

Substituting (22) into (21) and comparing with (20), we see that $\sigma_P(\tilde{\pi}) \geq (1-\epsilon) \cdot \sigma_P(\pi')$ for any other π', in Case II as well.