Fair Dynamic Rationing

Vahideh Manshadi

Yale School of Management, New Haven, CT, vahideh.manshadi@yale.edu

Rad Niazadeh

University of Chicago Booth School of Business, Chicago, IL, rad.niazadeh@chicagobooth.edu

Scott Rodilitz

Yale School of Management, New Haven, CT, scott.rodilitz@yale.edu

We study the allocative challenges that governmental and nonprofit organizations face when tasked with equitable and efficient rationing of a social good among agents whose needs (demands) realize sequentially and are possibly correlated. As one example, early in the COVID-19 pandemic, the Federal Emergency Management Agency faced overwhelming, temporally scattered, a priori uncertain, and correlated demands for medical supplies from different states. To better achieve their dual aims of equity and efficiency in such contexts, social planners intend to maximize the minimum fill rate across agents, where each agent's fill rate must be irrevocably decided upon its arrival. For an arbitrarily correlated sequence of demands, we establish upper bounds on both the expected minimum fill rate (ex-post fairness) and the minimum expected fill rate (ex-ante fairness) achievable by any policy. Our bounds are parameterized by the number of agents and the expected demand-to-supply ratio, and they shed light on the limits of attaining equity in dynamic rationing. Further, we show that for any set of parameters, a simple adaptive policy of projected proportional allocation achieves the best possible fairness guarantee, ex post as well as ex ante. Our policy is transparent and easy to implement, as it does not rely on distributional information beyond the first conditional moments. Despite its simplicity, we demonstrate that this policy provides significant improvement over the class of non-adaptive target-fill-rate policies by characterizing the performance of the optimal such policy, which relies on full distributional knowledge. We obtain the performance guarantees of (i) our proposed adaptive policy by inductively designing lower-bound functions on its corresponding value-to-go, and (ii) the optimal target-fill-rate policy by establishing an intriguing connection to a monopoly-pricing optimization problem. Further, we extend our results to considering alternative objective functions and to rationing multiple types of resources. We complement our theoretical developments with a numerical study motivated by the rationing of COVID-19 medical supplies based on a projected-demand model used by the White House. In such a setting, our simple adaptive policy significantly outperforms its theoretical guarantee as well as the optimal target-fill-rate policy.

Key words: rationing, fair allocation, social goods, correlated demands, online resource allocation

1.   Introduction

In Spring 2020, with the COVID-19 pandemic surging across the US, states were relying on the Federal Emergency Management Agency (FEMA) to provide urgently needed medical equipment from the Strategic National Stockpile. Unequipped for such a widespread emergency, FEMA aimed to ration its limited supplies in order to address states' current needs while also retaining some of the stockpile in anticipation of future needs. However, the allocation decisions made by FEMA were inconsistent and lacked transparency, which frustrated state officials (Washington Post 2020a).[1] Because having access to medical equipment can be a matter of life or death for a COVID-19 patient, making allocation decisions which are efficient and equitable is of paramount importance (Emanuel et al. 2020). Achieving efficiency alone is easy: a first-come, first-serve policy allocates all of the supply to meet early-arriving needs. However, such a policy can be unfair to patients in states where needs materialize later.

The above is just one example of a fundamental sequential allocation problem that social planners face when aiming to allocate divisible goods as *efficiently* and *equitably* as possible to demanding agents that arrive over time.

1.1.   Overview of Contributions

In this paper, we take the first step toward theoretically studying the aforementioned class of problems. We develop a framework for fair dynamic rationing where agents' one-time needs (demands) for a divisible good realize sequentially and can be *arbitrarily* correlated. In particular, upon arrival of each agent's demand, the planner makes an irrevocable decision about their *fill rate* (FR). Toward jointly achieving efficiency and equity, the planner aims to *maximize the minimum* FR, either ex post or ex ante. To assess the performance of sequential allocation policies, we introduce measures of ex-post and ex-ante fairness guarantees. For this general setting:

 (i) We establish upper bounds on the ex-post and ex-ante fairness guarantees achievable by any policy. These bounds are parameterized by the *supply scarcity* (i.e., the expected demand-to-supply ratio) and the number of agents.

 (ii) Remarkably, we show that a simple, adaptive, and transparent policy called *projected proportional allocation* (PPA) simultaneously achieves our upper bounds on the ex-post and ex-ante fairness guarantees for any set of parameters.

(iii) We illustrate the power of adaptivity by characterizing the ex-post guarantee of the optimal target-fill-rate policy and showing that such a non-adaptive policy cannot achieve our upper bounds.

---

[1] "*We don't know how the federal government is making those decisions,*" said Casey Katims, the federal liaison for Washington state.

(iv) Finally, we demonstrate the effectiveness of our policy through an illustrative case study motivated by the allocation of COVID-19 medical supplies based on a model of demand which was used by the White House.

Introducing a framework for fair dynamic rationing: We study the allocation of a divisible good to agents arriving over time with varying levels of demand. We assume the demand sequence is drawn from an arbitrary but known joint distribution across all agents. To account for heterogeneity in the demand level of different agents, we set each agent's utility to be its FR. In our base model, we focus on the objective of maximizing the minimum FR across all agents. Such an objective— which is in the spirit of Rawlsian justice—maximizes the utility of the worst-off agent. As such, it takes fairness into consideration along with efficiency.[2] Due to the stochasticity of the demand sequence, we consider two versions of this objective function: the excepted minimum FR and the minimum expected FR (see eq. (ex-post) and eq. (ex-ante), respectively, as well as the subsequent discussion).

Like other online stochastic optimization problems, our sequential allocation problem can be formulated as a dynamic program (DP), and it similarly suffers from the curse of dimensionality as well as other practical limitations such as a lack of interpretability. Consequently, we aim to design sequential allocation policies that perform well while being practically appealing and computable in polynomial time. We assess the performance of a policy by computing its ex-post and ex-ante fairness guarantees for any given supply scarcity and number of agents. In defining our notions of such guarantees, we use the minimum FR achievable under deterministic demand as a benchmark (see eq. (1) and Definition 1 and their related discussion in Section 2) to separate the impact of demand stochasticity from the impact of supply scarcity. The ex-post (resp. ex-ante) fairness guarantee of a policy serves as a lower bound on the expected minimum (resp. minimum expected) FR that the policy achieves relative to our benchmark under all possible joint demand distributions.

Establishing upper bounds: In order to gain insight into the difficulty of achieving equity and efficiency in sequential allocation, we develop upper bounds on the achievable fairness guarantees of any online policy, even policies which cannot be computed in polynomial time. For intuition, consider the following example with two agents. The first agent has demand of B_1, where B_1 is a Bernoulli random variable with success probability $2/3$. The second agent has demand $B_1 \times B_2$, where B_2 is an independent Bernoulli random variable with success probability $1/2$. In other words, the demand sequence is equally likely to be $(0,0)$, $(1,0)$, or $(1,1)$. For such an instance, no sequential allocation policy can distinguish between the latter two scenarios after observing the first demand, which leads to a sub-optimal decision. Building on the above intuition, in Sections 3.1 and 3.5, we

---

[2] In Section 5, we generalize our objective function.

establish upper bounds on the ex-post and ex-ante guarantees of any policy (see Theorems 1 and 4).

As we later show, these bounds are indeed tight. Thus, conducting comparative statics with respect to the supply scarcity and the number of agents reveals several insights (see Figure 1): when demand is small relative to supply, the bounds on both fairness guarantees deteriorate with increased demand. However, in the over-demanded regime, the bounds are independent of the supply scarcity. Further, in both the under-demanded and over-demanded regimes, the ex-post fairness guarantee worsens with more agents. On the other hand, the ex-ante fairness guarantee is independent of the number of agents. This highlights the fundamental difference in our notions of fairness: the objective corresponding to ex-post fairness (the expected minimum FR) is concerned with fairness along all samples paths, whereas the objective corresponding to ex-ante fairness (the minimum expected FR) is only concerned with marginal fairness (see the related discussion in Section 2).

Achieving upper bounds: Since our upper bounds apply to all sequential policies including the optimal online policy (namely, the exponential-sized DP), it would be reasonable if no policy could achieve these upper bounds in polynomial-time. However, we show that not only are these upper bounds achievable, but they can be achieved by our PPA policy. To motivate our policy, let us consider a hypothetical situation where the demand sequence is known a priori. In that case, the optimal allocation under both objectives is to equalize the FRs and then maximize that FR (see eq. (1) and its related discussion). Alternatively, this can be written as a deterministic DP with a simple solution: at each time period, proportionally allocate the remaining supply based on the current demand and the total future demand (see Section 3.2 and Appendix A.3). When demand is stochastic, our PPA policy simply replaces all the future random demands by their projected values, namely, their conditional expectations (see eq. (4)).

In Sections 3.3 and 3.5, we analyze the ex-post and ex-ante fairness guarantees of the PPA policy and show that it achieves the best of both worlds: our lower bounds on the PPA policy's guarantees match the corresponding upper bounds for any supply scarcity and any number of agents. These two analyses rely on delicate inductive arguments. For ex-post fairness, we establish a lower bound on the value-to-go function of our PPA policy by analyzing the evolution of the minimum FR and progressively constructing a worst-case joint distribution for demand (see Lemma 1 and Figure 3). For ex-ante fairness, we demonstrate that the expected demand-to-supply ratio before the arrival of each agent is non-increasing when following the PPA policy, which enables us to bound the marginal expected FR for each agent (see Appendix A.6).

We highlight that beyond enjoying the best possible guarantees, our PPA policy is practically appealing: it is computationally efficient, interpretable, and transparent. In addition, it does not

require full distributional knowledge, as it only relies on the first conditional moments of the joint distribution for demand. This last property is significant because knowing the full distribution may not be feasible in many practical situations (we discuss one such example in Section 4). Further, policies which rely on detailed distributional knowledge can be prone to errors or perturbations (see Footnote 15 and Appendix A.2).

Establishing sub-optimality of target-fill-rate policies: In addition to showing that our PPA policy achieves the best possible guarantees, we extend our work to studying the subclass of target-fill-rate (TFR) policies. A TFR policy commits upfront to a fill rate τ, and upon arrival of each agent, it allocates a fraction τ of that agent's demand until it exhausts the supply. Our study of TFR policies is motivated by two reasons: (i) since such policies are transparent and easy-to-communicate, they are frequently used in practice, including at the outset of the COVID-19 pandemic when an initial formula allocated a fixed percentage of states' estimated needs (Washington Post 2020a), and (ii) TFR policies are a natural yet powerful class of non-adaptive policies (see Section 3.4). Consequently, comparing the performance of TFR policies with that of our adaptive PPA policy sheds light on the limitations of making non-adaptive decisions.

Intuitively, a TFR policy can perform poorly because it does not take advantage of information that reduces future uncertainty. For instance, consider a setting with two agents where the second agent's demand is perfectly correlated with the first agent's demand. A simple adaptive policy—such as our PPA policy—will perform optimally in such a setting because demand is deterministic upon the first agent's arrival. The PPA policy achieves such performance by crucially leveraging information about the second agent's demand when determining the first agent's fill rate. In contrast, a TFR policy targets the same fill rate regardless of the first agent's demand, and consequently cannot ensure that sufficient supply remains for the second agent. Based on this intuition, in Section 3.4, we provide a tight bound on the ex-post fairness guarantee of the optimal TFR policy (see Theorem 3), which can be considerably lower than the corresponding guarantee of our adaptive PPA policy (see Figure 4).

To characterize the ex-post fairness guarantee of the optimal TFR policy, we construct the worst-case total demand distribution against such a policy. In the proof, we establish a rather surprising connection to the literature on monopoly pricing and Bayesian mechanism design (see Hartline (2013) for more details on this literature). In particular, upon mapping the problem of finding the worst-case instance into the quantile space, our problem reduces to a constrained version of the (single-item) monopoly pricing problem (see Remark 2). We identify two key properties of the worst-case distribution in this constrained monopoly pricing problem, and by exploiting the connection to our original problem, we end up with the desired characterization of the worst-case total demand distribution against the optimal TFR policy. Due to this connection, our proof

technique and corresponding results can be of independent interest (e.g., see Alaei et al. (2019) for proof techniques and results in the same spirit).

Illustrative case study: To demonstrate the effectiveness of our policy, in Section 4 we conduct a numerical case study motivated by the allocative challenges that FEMA faced at the beginning of the COVID-19 pandemic (as discussed at the beginning of this section). Drawing upon a model cited by the White house at that time, we first highlight the sequential and heterogeneous nature of states' demands for medical supplies. Then, we augment that model by considering a multivariate Normal distribution for demand with various levels of correlation. Our simulation results illustrate the superior performance of our PPA policy compared to both its ex-post fairness guarantee and the optimal TFR policy. Further, the results suggest that our PPA policy performs nearly as well as a DP solution (which, as we discuss, suffers from many practical limitations).

Allocating medical supplies in a pandemic is just one motivating example of the challenges that arise when a governmental or nonprofit organization aims to ration supply among agents whose (a priori uncertain and correlated) needs realize sequentially. Other examples include the allocation of emergency aid when a natural disaster such as a hurricane or wildfire impacts multiple locations over time (Wang et al. 2019), as well as the distribution of food donations by mobile pantries that sequentially visit agencies (Lien et al. 2014).[3] Our proposed policy can effectively guide transparent allocation decisions in such contexts while also providing a guarantee on the fairness level of the process. Finally, as discussed in Section 5, our framework can be enriched to account for other practical considerations, such as (i) generalized objective functions that enable the social planner to balance equity and efficiency to varying degrees, and (ii) rationing multiple types of resources (see Corollaries 1, 2, and 3 in Section 5.2).

1.2. Related Work

We conclude this section by discussing how our work relates to and contributes to several streams of literature.

Fairness in static resource allocation: Considerations of fairness and its trade-off with efficiency have frequently arisen in the resource allocation literature in operations research and computer science.[4] We begin by discussing papers which study fairness in static (one-shot) allocation settings. The seminal work of Bertsimas et al. (2011) considers a general setting where a central

---

[3] As explained in detail in Lien et al. (2014), even though the daily demand for food donations from different agencies are not temporally scattered, they will only be observed by the operators upon their arrival at the sites.

[4] Other recent papers have focused on fairness in the contexts of pricing (Cohen et al. 2019), information acquisition (Cai et al. 2020), targeted interventions (Levi et al. 2019), service levels (Jiang et al. 2019), and online learning (Gupta and Kamble 2019). See also the work of Cayci et al. (2020) that considers fair resource allocation with online learning.

decision-maker allocates m divisible resources to n agents, each with a different utility function. Focusing on two commonly used notions of fairness in allocation, max-min and proportional fairness, the authors characterize the efficiency loss due to maximizing fairness (see also Bertsimas et al. 2012 and Bertsimas et al. 2013). If demand was deterministic in our setting, the optimal allocation would coincide with that of the max-min objective in Bertsimas et al. (2011). Namely, for both objectives, the optimal allocation consists of maximized equal FRs.

Focusing on indivisible goods, Donahue and Kleinberg (2020) considers the trade-off between fairness and utilization when demand is distributed across different agents. A priori, only demand distributions are known. However, after a one-shot allocation decision, all demand values realize. The fairness notion considered in this line of work is in the same spirit of our notion of ex-ante fairness: they require that an individual's chance of receiving the resource should not significantly depend on the group to which the individual belongs. Similarly, by maximizing the minimum expected FR, we aim to reduce the impact of an agent's place in the sequence of arrivals. Sharing similar motivation to our paper, Grigoryan (2020) and Pathak et al. (2020) consider equitable COVID-19 vaccine allocation. However, the settings (e.g., offline and deterministic), models, and techniques in both papers differ drastically from those in this work.

Also falling within the category of static allocation of indivisible goods, a stream of papers in computer science considers allocation problems when agents' valuations are deterministically known. For deterministic algorithms, recent research has centered on the existence of allocations which satisfy certain fairness properties, such as envy-freeness up to any good (see, e.g., Chaudhury et al. (2020) and references therein). For randomized algorithms, the closest to our work is the recent work of Freeman et al. (2020), which uses notions of ex-post and ex-ante fairness and explores whether both can be achieved simultaneously. They develop a randomized algorithm that is approximately fair ex post and precisely fair ex ante. We ask a similar question, albeit in a dynamic divisible-good setting with random and correlated demand, and we affirmatively answer it: our PPA policy exactly achieves the best possible fairness guarantee ex post as well as ex ante (see Theorems 2 and 4).

Fairness in dynamic resource allocation: We now turn our attention to papers that consider fairness in dynamic (online) allocation settings. In terms of modeling, closest to our work are Lien et al. (2014) and Sinclair et al. (2020). Motivated by the distribution of food donations by mobile pantries,[5] Lien et al. (2014) introduced the problem of sequential resource allocation which coincides with our base model and the ex-post fairness objective function, in that it aims to maximize the

---

[5] For other examples of work in this application area, see Solak et al. (2014), Orgut et al. (2018), and Eisenhandler and Tzur (2019).

expected minimum FR (although it only studies the special case of independent demands). The recent work of Sinclair et al. (2020) considers a similar model; however, it focuses on a multi-criteria objective which is based on an allocation's distance from the optimal offline Nash Social Welfare solution. We note that their notion of fairness is also different in nature from ours.[6]

The algorithmic aspects of both Lien et al. (2014) and Sinclair et al. (2020) consist of designing novel heuristics and numerically evaluating them against a relevant benchmark (the intractable DP solution and the Nash Social Welfare solution, respectively). On the other hand, we take a theoretical approach and analyze fairness guarantees for the policies we design. Further, we provide upper bounds on the performance of any policy (including the DP solution), which serves a dual purpose: (i) it establishes that our policy is the best possible one if we aim to achieve both ex-ante and ex-post fairness guarantees, and (ii) it highlights the fundamental limits of achieving equity in a dynamic setting.

In settings with multiple types of resources, Azar et al. (2010) and Bateni et al. (2016) study online versions of Fisher markets and develop policies with fairness guarantees under two different arrival models. The former assumes an adversarial model whereas the latter considers demand that belongs to a general class of stochastic processes.[7] There are fundamental differences between our work and the aforementioned papers. Just to name one, the settings of Azar et al. (2010) and Bateni et al. (2016) are motivated by online advertising, where demanding agents (advertisers) are offline and items (impressions) arrive in an online fashion. Demanding agents have a large budget compared to the price of each arriving item, and they derive item-specific utilities. Consequently, the fairness notion is concerned with the total utility of each agent, which is a function of *all* items allocated to it during the horizon. In contrast from such a setting, demanding agents in our work arrive in an online fashion while the supply side is offline, and each demanding agent receives a *single* allocation. The recent works of Ma and Xu (2020) and Nanda et al. (2020) are closer to our setting in that the demanding agents arrive online; however, they differ in several aspects: (i) the underlying arrival process is known i.i.d. where arriving demand belongs to various groups, (ii) they focus on group-level fairness, and (iii) they consider a matching setting, i.e., allocating indivisible goods.

The objectives of ex-post and ex-ante fairness which we study in our problem bear some resemblance to the objective in the online contention resolution scheme (OCRS) problem, although the

---

[6] In Sinclair et al. (2020), their notion of fairness is with respect to the absolute allocation, i.e., if possible, agents' allocation should be equalized regardless of differences in their needs. In contrast, we aim for an allocation which is proportional to need.

[7] We remark that papers considering general convex objective functions, such as Agrawal and Devanur (2014) and Balseiro et al. (2020), admit many common fairness objective functions as special cases. See also Mehta (2012) for more details.

two problems are not directly comparable. The OCRS is basically a rounding algorithm that aims to uniformly preserve the marginals induced by a fractional solution while obtaining feasibility of the final allocation. This technique has found application in many settings such as Bayesian online selection, oblivious posted pricing mechanisms, and stochastic probing models (see, e.g., Alaei 2014, Feldman et al. 2016, and Lee and Singla 2018). The OCRS problem diverges from ours because that setting focuses on designing randomized policies for allocating indivisible goods, while our focus is on divisible goods (consequently, restricting to deterministic policies is without loss).

Dynamic allocation of social goods: On a broader level, our paper is related to the literature on dynamic allocation of social goods and services, such as public housing, donated organs, and emergency care. Examples of centralized allocation policies include Kaplan (1984), Ashlagi et al. (2013), Agarwal et al. (2019), and Ashlagi et al. (2019); examples of decentralized mechanisms are Leshno (2017), Anunrojwong et al. (2020), and Arnosti and Shi (2020). For the most part, the aforementioned papers focus on the analysis of social welfare in steady-state models where both demand and supply dynamically arrive. We complement this literature by focusing on equitable allocation in a non-stationary framework where a fixed amount of supply must be rationed across demand that arrives over time.

Online resource allocation: From a technical point of view, our work is related to the rich literature on online resource allocation and prophet inequalities, which started from the seminal work of Krengel and Sucheston (1978) and Samuel-Cahn et al. (1984). For an informative survey, we refer the interested reader to Lucier (2017). We highlight that in terms of modeling demand, our work departs from the prevailing approaches in this literature, namely adversarial, i.i.d., or random permutation arrival models. In our work, we assume that the sequence of demands can be arbitrarily correlated and the joint distribution is known in advance. In terms of modeling demand, our work is closest to a few papers that consider prophet inequalities with correlated demand (Rinott and Samuel-Cahn 1992, Truong and Wang 2019, Immorlica et al. 2020). However, the nature of the online decisions is different; in our model, a fraction of a divisible good is allocated to each arriving demand, whereas in prophet inequality settings, an indivisible good is allocated to a single agent.

Finally, our PPA policy relies on re-optimizing the FR by replacing all future random demands by their expected values. As such, it is related to the stream of papers in revenue management and dynamic programming that theoretically analyze the performance of such heuristics. Great examples of work in this direction include Ciocan and Farias (2012), Jasin and Kumar (2012), Balseiro and Brown (2019), and Calmon et al. (2020).

2. Preliminaries

Problem setup: Consider a planner that is using a sequential allocation policy—also referred to as an online policy—to allocate a divisible resource of supply s among n agents. Without loss of generality, we normalize the total supply so that $s = 1$. Agents arrive sequentially over time periods $1, 2, \ldots, n$, and we index agents according to the period in which they arrive. Once agent i arrives, their demand $d_i \in \mathbb{R}_{\geq 0}$ is realized and observed by the planner. Based on the observed demand d_i and the history up to time period i, the sequential policy makes an irrevocable decision by allocating an amount x_i of the resource to this agent. The allocated amount x_i cannot exceed the agent's realized demand d_i nor can it exceed the remaining supply before agent i's arrival, which we denote by s_i. Thus, x_i is a feasible allocation if $x_i \in [0, \min\{s_i, d_i\}]$. Given the feasible allocation x_i and the demand d_i, agent i's *fill rate* (FR) is defined as $\frac{x_i}{d_i}$.[8] After allocating x_i to agent i, the remaining supply before the arrival of agent $i + 1$ is $s_{i+1} = s_i - x_i$.

To model the uncertainty about future demands, we consider a Bayesian setting where the d_i's are stochastic and arbitrarily correlated such that $\vec{d} = (d_1, d_2, \ldots, d_n)$ is drawn from a joint distribution $\vec{\mathcal{F}} \in \Delta\left(\mathbb{R}_{\geq 0}^n\right)$ *known* by the planner. We further denote the *supply scarcity* (i.e., the expected demand-to-supply ratio) by $\mu \triangleq \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\sum_{i \in [n]} d_i\right]$,[9] which is equivalent to the total expected demand since we normalize the total supply to be 1. For simplicity of presenting our results, we consider joint distributions that assign non-zero probability to at least one sample path of demands with $d_n \neq 0$. Equivalently, we assume d_n is not deterministically equal to zero.[10]

As detailed earlier, our setup is motivated by the distributional operations of a governmental or nonprofit organization. Consequently, we focus on an egalitarian planer that intends to balance the equity and efficiency of the allocation. To this end, the planner's objective is to maximize the minimum achieved FR among the agents, i.e., $\min_{i \in [n]} \frac{x_i}{d_i}$, given the uncertainty in the demands. Maximizing such an objective has its roots in the classic literature on welfare economics (e.g., Arrow 1963) and has been studied more recently in similar contexts in operations research (e.g., Lien et al. 2014). It provides equity through its focus on the worst FR across all agents—in contrast to the sum of FRs—and provides efficiency by aiming to maximize this FR—in contrast to allocating an equally minimal amount of the resource to all agents.[11]

Objectives & fairness guarantees: Since demands are a priori uncertain in the setup described above, the planner should consider appropriate metrics to aggregate over uncertain outcomes.

---

[8] If $x_i = d_i = 0$, we set the FR to 1 as a convention.

[9] For any $a \in \mathbb{N}$, we use $[a]$ to refer to the set $\{1, 2, \ldots, a\}$.

[10] This assumption is without loss of generality, as one can alternatively re-define n to be the smallest index such that $d_{n'}$ is deterministically equal to zero for $n' > n$.

[11] We consider a broader class of objectives that subsumes the minimum FR in Section 5.2.1.

We now formally define the planner's objectives by considering two different metrics: the ex-post minimum FR and the ex-ante minimum FR. For any sequential allocation policy π, the ex-post minimum FR of policy π is its expected minimum FR, i.e.,

$$W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi) \triangleq \mathbf{E}_{\vec{d} \sim \mathcal{F}} \left[\min_{i \in [n]} \left\{ \frac{x_i}{d_i} \right\} \right] , \tag{ex-post}$$

where $\vec{x} = (x_1, x_2, \ldots, x_n)$ is the sequence of allocations generated by π. On the other hand, the ex-ante minimum FR of policy π is its minimum expected FR, i.e.,

$$W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi) \triangleq \min_{i \in [n]} \left\{ \mathbf{E}_{\vec{d} \sim \mathcal{F}} \left[\frac{x_i}{d_i} \right] \right\} . \tag{ex-ante}$$

For a randomized policy π, we abuse notation and again use $W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)$ and $W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi)$ to denote the expectation of the above two quantities over the policy π's internal randomness.[12]

These two objectives represent two different notions of fairness: eq. (ex-post) aims for equity in outcomes, whereas eq. (ex-ante) aims for equity in *expected* outcomes. We largely focus on the ex-post minimum FR for two main reasons. First, when allocating supplies in response to a rare event like a pandemic or natural disaster, agents only observe one realized outcome. Because the ex-ante minimum FR is only concerned with marginal fairness, it can have unfair outcomes for *every* sample path, i.e., every realized demand sequence. In contrast, the ex-post minimum FR considers each full sample path; every sample path with positive probability which results in an unfair outcome reduces $W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)$. Second, by Jensen's inequality, the ex-post minimum FR serves as a lower bound on the ex-ante minimum FR, i.e., for any policy π,

$$W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi) \geq W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi).$$

However, for a fixed ex-post minimum FR, achieving a higher ex-ante minimum FR is desirable because it reduces systematic biases against a particular agent, e.g., the last-arriving agent. In the extreme case where $W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi) = W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)$, one particular agent receives the smallest FR, *regardless* of the sample path. On the other hand, $W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi) > W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)$ implies that the worst-off agent varies across different sample paths.

Having defined our notions of fairness, we first observe that if the sequence of demand is deterministic, then the policy that maximizes *both* the ex-post and the ex-ante minimum FR is simply equalizing all FRs. Namely,

$$\overline{W} \triangleq \max_{\pi} \left\{ W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi) \right\} = \max_{\pi} \left\{ W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi) \right\} = \frac{x_1}{d_1} = \frac{x_2}{d_2} = \cdots = \frac{x_n}{d_n} = \min \left\{ 1, \frac{1}{\mu} \right\}. \tag{1}$$

---

[12] In principle, we allow randomization of our policies in this paper; however, as will be clear later, all of our proposed policies are deterministic and no randomization is needed to obtain our targeted performance guarantees.

The above observation highlights that when total demand exceeds supply, even without stochasticity, we cannot guarantee a minimum FR better than $1/\mu$, simply due to the scarcity of supply. However, if the sequence of demands is stochastic and correlated, a minimum FR of $1/\mu$ may not be achievable. Consequently, we evaluate policies based on how they perform relative to $\overline{W}$. For a policy π and a joint demand distribution $\vec{\mathcal{F}}$, we say that the policy achieves *ex-post fairness* (resp. *ex-ante fairness*) of $W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)/\overline{W}$ (resp. $W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi)/\overline{W}$). We aim to design a policy with guarantees on both ex-post and ex-ante fairness that hold universally for all joint demand distributions $\vec{\mathcal{F}}$ with n agents and supply scarcity μ. We refer to the universal lower bounds of a policy π as its *fairness guarantees*, which we formally define below.

DEFINITION 1 (**Ex-post/Ex-ante Fairness Guarantee**). A sequential allocation policy π achieves an *ex-post fairness guarantee* (resp. *ex-ante fairness guarantee*) of $\kappa_{\mathrm{P}}(\mu, n)$ (resp. $\kappa_{\mathrm{A}}(\mu, n)$), if for all $n \in \mathbb{N}$ and $\mu \in \mathbb{R}_{\geq 0}$,

$$\inf_{\vec{\mathcal{F}} \in \Delta\left(\mathbb{R}_{\geq 0}^n ; \mu\right)} \frac{W_{\mathrm{P}}^{\vec{\mathcal{F}}}(\pi)}{\overline{W}} \geq \kappa_{\mathrm{P}}(\mu, n) \quad \left(\text{resp.} \quad \inf_{\vec{\mathcal{F}} \in \Delta\left(\mathbb{R}_{\geq 0}^n ; \mu\right)} \frac{W_{\mathrm{A}}^{\vec{\mathcal{F}}}(\pi)}{\overline{W}} \geq \kappa_{\mathrm{A}}(\mu, n) \right),$$

where $\Delta\left(\mathbb{R}_{\geq 0}^n ; \mu\right)$ denotes the domain of joint demand distributions with n agents and supply scarcity μ.

Our goals are (i) to understand the limits of achieving fairness in sequential allocation by computing upper bounds on the achievable guarantees, and (ii) to obtain tight lower bounds by designing policies with strong ex-post guarantees as well as ex-ante guarantees. We show in Section 3 that no gap exists between the achievable upper and lower bounds under both ex-post and ex-ante notions. More specifically, we show how to obtain *exactly matching* upper and lower bounds for both notions of fairness using a *single* adaptive policy.

3. Optimal Bounds on Fairness Guarantees

In this section, we present our main results for the setting introduced in Section 2. First, we focus on ex-post fairness in Section 3.1 and establish parameterized upper bounds on the ex-post fairness guarantee achievable by any sequential allocation policy—whether adaptive or non-adaptive, computationally efficient (i.e., with polynomial running time) or not. Then, somewhat surprisingly, we show that such upper bounds can be achieved by our policy, which is introduced and analyzed in Sections 3.2 and 3.3. Next, to illustrate the power of our simple adaptive algorithm, in Section 3.4 we characterize the ex-post fairness guarantee of the best policy which non-adaptively aims for a particular target fill rate, and we show that our policy performs favorably compared to such a policy. Finally, in Section 3.5 we turn our attention to the notion of ex-ante fairness, and we show that our policy also achieves the best possible ex-ante fairness guarantee.
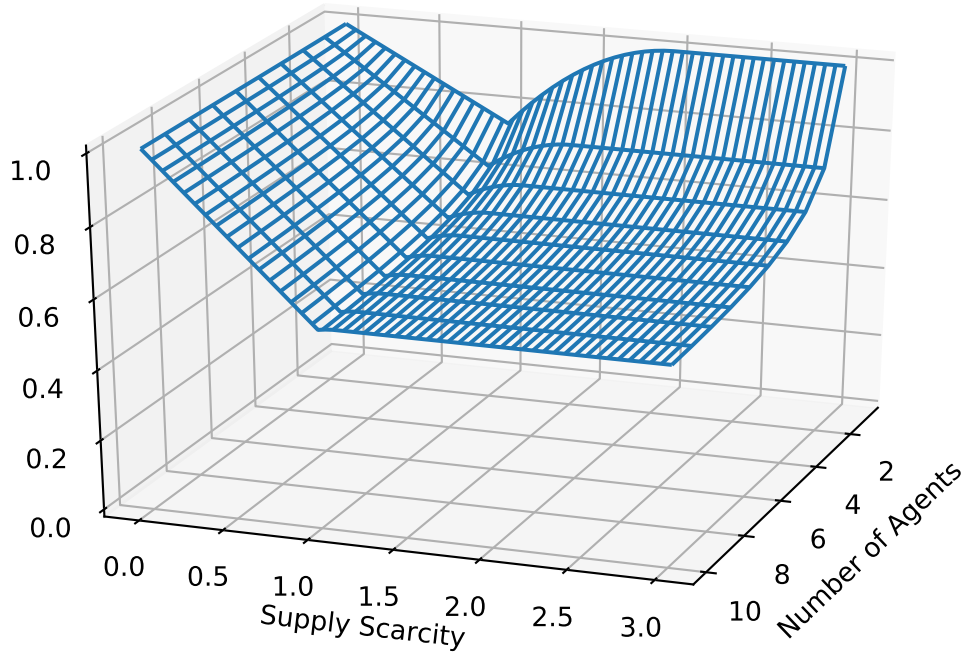
Figure 1   The upper bound on the ex-post fairness guarantee of any policy, as a function of the supply scarcity
and the number of agents.

3.1.   Upper Bound on Ex-post Fairness Guarantee

We begin this section by establishing a fundamental limit on ex-post fairness for any allocation policy when faced with stochastic and sequential demands. The main result of this subsection is the following theorem:

THEOREM 1 (**Upper Bound on Ex-post Fairness Guarantee**). *Given a fixed number of agents $n \in \mathbb{N}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, no sequential allocation policy obtains an ex-post fairness guarantee (see Definition 1) greater than $\kappa_{\mathrm{P}}(\mu, n)$, defined as*

$$\kappa_{\mathrm{P}}(\mu, n) \triangleq \begin{cases} 1 - \left(\frac{n}{2(n+1)}\right)\mu, & \mu \in [0, 1) \\ \mu - \left(\frac{n}{2(n+1)}\right)\mu^2, & \mu \in [1, \frac{n+1}{n}) \\ \frac{n+1}{2n}, & \mu \in [\frac{n+1}{n}, +\infty) \end{cases} \tag{2}$$

See Figure 1 for an illustration of this upper bound as a function of the supply scarcity μ and the number of agents n. Per Definition 1, the ex-post fairness guarantee is relative to the achievable minimum FR when demands are deterministic, namely $\overline{W} = \min\{1, 1/\mu\}$. Consequently, this upper bound provides insight into the unavoidable loss in efficiency and equity when demands are a priori uncertain and realize sequentially. In particular, we remark that the achievable fairness guarantee crucially depends on the supply scarcity. In the regime where $\mu < 1 + \frac{1}{n}$, which we refer to as

the *under-demanded regime*, $\kappa_{\mathrm{P}}(\mu, n)$ initially worsens as μ increases before hitting its minimum (for any fixed n) when expected demand equals supply, i.e., at $\mu = 1$. This suggests that the stochastic nature of demand is most harmful when expected demand exactly equals supply. On the other hand, in the *over-demanded regime* where $\mu \geq 1 + \frac{1}{n}$, the achievable fairness guarantee is independent of μ. Given that we are usually in the over-demanded regime in our motivating applications, Theorem 1 ensures that supply scarcity does not contribute to the loss in fairness due to uncertain, correlated, and sequential demand. Fixing μ, the upper bound always decreases with n, implying that achieving fairness can be more challenging for a larger population of agents with stochastic demands, even if the total expected demand of the population remains the same. Finally, we highlight that the bound is always at least $1/2$ regardless of the supply scarcity and the number of agents, and it attains its minimum when $\mu = 1$ and $n \to +\infty$.

The proof of Theorem 1 relies on establishing two hard instances with similar structures, one for $\mu < 1 + 1/n$ and one for $\mu \geq 1 + 1/n$. The details of the proof are presented in Appendix A.1. Here, we present the instance for the over-demanded regime along with a sketch of our analysis. In this instance, there are n possible equally-likely scenarios, i.e., scenario σ happens with probability $1/n$ for $\sigma \in [n]$. In scenario σ, the first σ agents have equal demand of $\frac{2\mu}{n+1}$ and the rest have no demand. We illustrate this instance in Figure 2(a).[13]

First, note that the total supply scarcity for the above hard instance is μ (as shown in Appendix A.1). Next, consider any sequential policy that faces a non-zero demand from agent i. The policy cannot distinguish among possible scenarios $i, i+1, \ldots, n$. Consequently, its allocation decision for agent i will be independent of the scenario. In light of this observation, any policy can be sufficiently described by a set of (possibly random) allocations with expected values $\vec{y} = (y_1, y_2, \ldots, y_n)$, such that if agent i has non-zero demand, then they receive an expected allocation y_i. Given $\vec{y}$, the minimum FR for scenario σ is

$$r_\sigma \triangleq \frac{(n+1)}{2\mu} \mathbf{E}_\pi[\min\{x_1, x_2, \ldots, x_\sigma\}] \leq \frac{(n+1)}{2\mu} \ \min\{y_1, y_2, \ldots, y_\sigma\} \leq \frac{(n+1)}{2\mu} y_\sigma \ , \tag{3}$$

where the first inequality is due to the expectation of a minimum being less than the minimum over expectations (Jensen's inequality).

In order to establish our upper bound, we set up a factor-revealing linear program as presented in Figure 2(b). The LP maximizes the expected minimum FR subject to three sets of natural constraints that must hold for any sequential policy:

- The minimum FR in scenario σ cannot exceed the FR for agent σ, as shown in eq. (3).

---

[13] We remark that similar settings can occur in practice. As one example, consider the challenge of allocating limited disaster-relief supplies to towns damaged by a hurricane which may continue on its destructive path or may veer back out to sea.
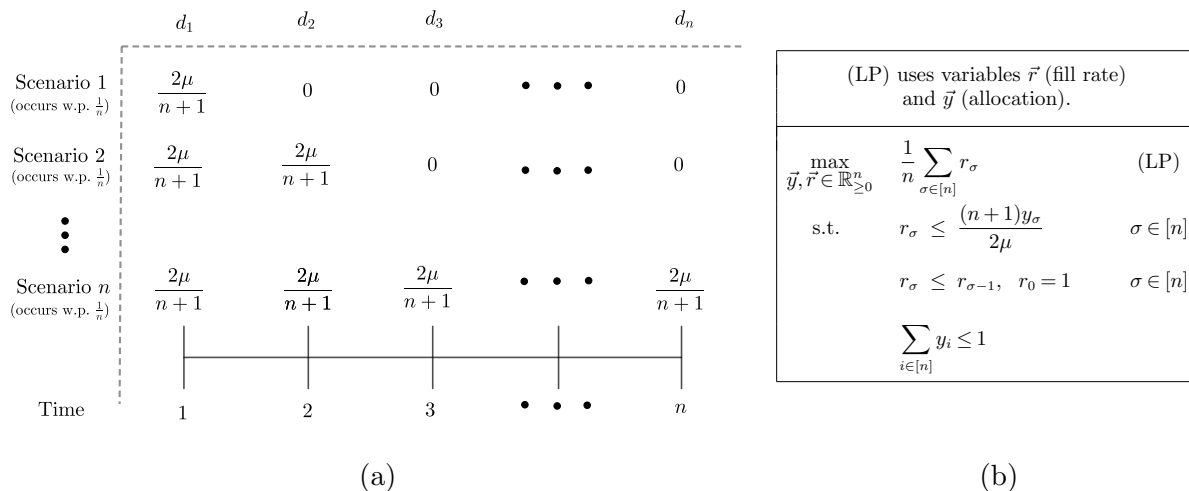
Figure 2   **(a) The instance and (b) the factor-revealing LP which establish an upper bound of $\kappa_{\mathrm{p}}(\mu, n)$ for the over-demanded regime (i.e., when $\mu \geq 1 + \frac{1}{n}$).**

- The minimum FR in scenario σ is at most the minimum FR in scenario $\sigma - 1$.
- The total amount of expected allocations cannot exceed the available supply of 1.

In Appendix A.1, we provide an upper bound on the optimal value of this LP by presenting a feasible solution to its dual. To complete the proof of Theorem 1, we must scale by $\overline{W} = \min\{1, 1/\mu\}$ to translate this upper bound on the expected minimum FR into an upper bound on ex-post fairness (see Definition 1).

Having provided the proof sketch of Theorem 1, we finish this section by noting that it is not clear whether the bound in eq. (2) can be achieved, even by the optimal online policy which can be found via a DP. Furthermore, there are significant limitations and drawbacks to a DP approach for maximizing the expected minimum FR in this setting. First, (i) the state space of such a DP is exponentially large for correlated demands, which makes the DP intractable.[14] In addition, (ii) solving the DP requires full distributional knowledge, and (iii) the DP decisions may lack transparency and interpretability, which are highly desirable properties in our motivating applications.[15]

Remarkably, in the following subsection, we design a simple adaptive policy that not only achieves the best possible ex-post fairness guarantee of $\kappa_{\mathrm{p}}(\mu, n)$, but also offers several corresponding advantages over a DP solution: (i) it can be computed efficiently, (ii) it only requires knowledge of

---

[14] While the exponential-sized state space does not formally prove the computational hardness of finding the optimal online policy, it suggests that the naive DP approach will be computationally intractable; additional evidence suggesting computational hardness is due to Papadimitriou and Tsitsiklis (1987), which shows that the closely related general problem of finding the optimal policy in partially observable MDPs is PSPACE-hard.

[15] Of course, the DP solution for ex-post fairness can lead to significantly sub-optimal ex-ante fairness, and it can also be sensitive to small perturbations in the demand distributions. We provide a simple example to illustrate both of those drawbacks in Appendix A.2.

the conditional first moments of agents' demands, and (iii) its decisions can be clearly explained. Additionally, as shown in Section 3.5, it simultaneously attains the best-possible ex-ante fairness guarantee.

3.2. Projected Proportional Allocation Policy

We introduce our policy, referred to as the *projected proportional allocation* (PPA) policy, through the following simple intuition. Consider a planner that (magically) has access to all the demand realizations $\vec{d}$. As already discussed in Section 2, to maximize the minimum FR when the demand realizations are known a priori, the planner should equalize the FR of all agents by allocating $x_i^* = \min\left(d_i, \frac{d_i}{\sum_{j\in[n]} d_j}\right)$ to each agent i. If $\sum_{j\in[n]} d_j$ is at most the initial supply (which we normalize to 1), then each agent i obtains a full allocation of $x_i = d_i$ in such a solution. This results in the maximum equal FR of 1. Otherwise, all the agents will have an equal FR of $1/\sum_{j\in[n]} d_j$, which is $1/\mu$ when each demand is equal to its expected value.

This solution can alternatively be obtained by solving a DP that returns allocations $x_n^*, x_{n-1}^*, \ldots, x_1^*$ maximizing the minimum FR. By a simple induction argument, given the remaining supply s_i at period i, this DP maintains the following invariant at each period i (refer to Appendix A.3 for details):

$$x_i^* = \min\left\{d_i, s_i \frac{d_i}{d_i + \sum_{j\in[i+1:n]} d_j}\right\} = \min\left\{d_i, s_i \frac{d_i}{d_i + [\text{total future demand}]}\right\}.^{16} \tag{4}$$

Notably, the above invariant suggests a *sequential* implementation of the optimal solution at each period i that only uses the knowledge of d_i (i.e., the current demand at period i) and $\sum_{j\in[i+1:n]} d_j$ (i.e., the total future demand from period $i+1$ to n). Now consider a setting with incomplete information, namely, with only knowledge of the current sample path of the observed demands up to period i, which we denote by $\vec{d}_{[1:i]} \triangleq (d_1, d_2, \ldots, d_i)$. Our PPA policy implements a version of the above policy by replacing the exact realization of total future demand with the conditional first moment of this random variable given the current sample path. More precisely:

- Given the remaining supply s_i, the PPA policy allocates an amount

$$x_i = \min\left\{d_i, s_i \frac{d_i}{d_i + \mu_{i+1}}\right\} \qquad \text{(PPA's \textsc{update rule})}$$

of the (divisible) resource to agent i upon their arrival, where

$$\mu_{i+1} \triangleq \mathbf{E}_{\vec{d}\sim\vec{\mathcal{F}}}\left[\sum_{j\in[i+1:n]} d_j \,\Big|\, \vec{d}_{[1:i]}\right]$$

---

[16] For any $a, b \in \mathbb{N}$ we use $[a:b]$ to refer to the set $\{a, a+1, \ldots, b\}$ if $a \leq b$ (and the empty set otherwise).

Note that the conditional expected future demand μ_{i+1} given all previously-realized demands $\vec{d}_{[1:i]}$ is a function of $\vec{d}_{[1:i]}$; however, for ease of notation, we use μ_{i+1} without any input arguments.

We remark that the PPA policy is simple, computationally efficient, and solely uses first-moment knowledge about the future demands. Further, because the allocation decisions of the PPA policy depend smoothly on the first moment of future demand, these decisions are robust to small changes in the scale of any marginal distribution. Yet, as we show in Sections 3.3 and 3.5, this simple policy remarkably achieves the best possible guarantee for both notions of fairness (ex-post and ex-ante), even though these two notions are quantitatively different whenever $n > 1$.

We conclude this subsection by highlighting an important technical property of the PPA policy.

REMARK 1. The PPA policy can only run out of supply at the end of period i if $\mu_{i+1} = 0$, or equivalently, only if all future demands $\vec{d}_{[i+1:n]}$ are deterministically equal to zero, conditional on the current realized sample path of demands $\vec{d}_{[1:i]}$. This property holds simply because

$$s_{i+1} = s_i - x_i \geq s_i \frac{\mu_{i+1}}{d_i + \mu_{i+1}}.$$

3.3. Ex-post Fairness of PPA Policy

In this section, we analyze the ex-post fairness guarantee of our PPA policy. In the following theorem, we show that this simple policy indeed achieves the best possible ex-post fairness guarantee.

THEOREM 2 (**Ex-post Fairness Guarantee of PPA Policy**). *Given a fixed number of agents $n \in \mathbb{N}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, the PPA policy achieves an ex-post fairness guarantee (see Definition 1) of at least $\kappa_{\mathrm{P}}(\mu, n)$ (defined in eq. (2)).*

3.3.1. Proof of Theorem 2 In order to prove the above theorem, we would have liked to analyze the evolution of the minimum FR, which we denote with f_i at the end of period $i - 1$, i.e., $f_1 = 1$, $f_i = \min\{f_{i-1}, \frac{x_{i-1}}{d_{i-1}}\}$ for $i \in [2 : n+1]$. Instead, we consider the evolution of a closely related stochastic process, which makes the analysis simpler. We define this surrogate stochastic process as follows:

$$\beta_1 \triangleq \min\left\{1, \frac{n+1}{n\mu}\right\} \qquad \beta_i \triangleq \min\left\{\beta_{i-1}, \frac{x_{i-1}}{d_{i-1}}\right\}, \quad i \in [2 : n+1]. \tag{5}$$

First, we note that $\beta_i = \min\{f_i, \frac{n+1}{n\mu}\}$, $i \in [n+1]$. Next, recall that s_i denotes the remaining supply after agent $i - 1$ arrives and receives an allocation. We observe that under the PPA policy, s_i evolves according to

$$s_1 = 1 \qquad s_i = s_{i-1} - \min\left\{d_{i-1}, \frac{d_{i-1}}{d_{i-1} + \mu_i} s_{i-1}\right\}, \quad i \in [2 : n]. \tag{6}$$

With the above observations, the main step of the proof is carefully analyzing the evolution of (β_k, s_k) under the PPA policy, which enables us to lower bound the final expected minimum FR in the following lemma.

LEMMA 1 (**Lower Bound on Expected Minimum FR**). *Under the PPA policy, for all $i \in [n+1]$ and any subsequence of demand realizations $\vec{d}_{[1:i-1]}$,*

$$\mathbf{E}_{\vec{d} \sim \mathcal{F}}\left[f_{n+1} \mid \vec{d}_{[1:i-1]}\right] \geq \beta_i \left(1 - \frac{n+1-i}{2(n+2-i)} \frac{\mu_i}{s_i} \beta_i\right) \tag{7}$$

where β_i is defined in eq. (5).[17]

Since the objective of our dynamic decision-making problem has no per-stage rewards and consists only of a terminal reward (i.e., the minimum FR), Lemma 1 can be thought of as establishing a lower bound on the value-to-go function of the PPA policy. Before providing the proof for this key lemma, we lay out the two remaining steps that finish the proof of Theorem 2: (i) plugging $i = 1$ into inequality (7) to obtain a lower bound on $\mathbf{E}_{\vec{d} \sim \mathcal{F}}[f_{n+1}]$, and (ii) scaling the obtained lower bound result by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$, which provides an ex-post fairness guarantee (see Definition 1).

Proof of Lemma 1: We will show that inequality (7) holds via backwards induction. The base case of $i = n+1$ is trivial as it follows from the observation we made earlier: $\beta_i = \min\{f_i, \frac{n+1}{n\mu}\}$.

Now let us consider $i = k < n+1$. Instead of proving inequality (7), we prove a stronger result:

$$\mathbf{E}_{\vec{d} \sim \mathcal{F}}\left[f_{n+1} \mid \vec{d}_{[1:k]}\right] \geq \beta_k \left(1 - \frac{n+1-k}{2(n+2-k)} \frac{d_k + \mu_{k+1}}{s_k} \beta_k\right). \tag{8}$$

Establishing inequality (8) means that the inequality in (7) holds for any realization of agent k's demand. Consequently, it will hold when we take an expectation over agent k's demand. In order to prove inequality (8), we consider two different cases that can arise depending on the remaining supply s_k, agent k's demand d_k, and the future expected demand μ_{k+1}. In the following, we introduce and analyze these cases separately.

(i) *Sufficient supply* $(s_k \geq \beta_k(d_k + \mu_{k+1}))$: Recall that according to the PPA policy, $x_k = \min\{d_k, \frac{d_k}{d_k + \mu_{k+1}} s_k\}$. Therefore, in this case, either $x_k = d_k$, i.e., the PPA policy meets the entire demand, or $x_k = \frac{d_k}{d_k + \mu_{k+1}} s_k \geq \beta_k d_k$, i.e., the PPA policy attains an FR of at least β_k. According to the dynamics specified in (5) and (6), this implies

$$\beta_{k+1} = \beta_k \quad \text{and} \quad \frac{\mu_{k+1}}{s_{k+1}} = \frac{\mu_{k+1}}{s_k - \min\{d_k, \frac{d_k}{d_k + \mu_{k+1}} s_k\}} \leq \frac{\mu_{k+1}}{s_k - \frac{d_k}{d_k + \mu_{k+1}} s_k} = \frac{d_k + \mu_{k+1}}{s_k}.$$

Using our inductive hypothesis when $i = k+1$,

$$\mathbf{E}_{\vec{d} \sim \mathcal{F}}\left[f_{n+1} \mid \vec{d}_{[1:k]}\right] \geq \beta_k \left(1 - \frac{n-k}{2(n+1-k)} \frac{d_k + \mu_{k+1}}{s_k} \beta_k\right) \triangleq \text{RHS}^{(1)}. \tag{9}$$

The lower bound given by $\text{RHS}^{(1)}$ is a linear function of $d_k + \mu_{k+1}$, as illustrated by the dotted red lines in all panels of Figure 3 (in the regime where $d_k + \mu_{k+1} \in [0, s_k/\beta_k]$). This linear

---

[17] If $s_i = 0$, then by Remark 1, we must also have $\mu_i = 0$. In such cases, we take the convention that $\frac{\mu_i}{s_i} = 0$.

function has a non-positive slope and an intercept of β_k. We can further lower bound this function for any $d_k + \mu_{k+1} \in [0, s_k/\beta_k]$ by another linear function with the same intercept of β_k and a smaller (more negative) slope. In particular, since $\frac{n-k}{n+1-k} \leq \frac{n+1-k}{n+2-k}$, we have:

$$\text{RHS}^{(1)} \geq \beta_k \left(1 - \frac{n+1-k}{2(n+2-k)} \frac{d_k + \mu_{k+1}}{s_k} \beta_k \right), \tag{10}$$

which proves inequality (8) in the sufficient supply case (see the blue lines in all panels of Figure 3).

(ii) *Insufficient supply $(s_k < \beta_k(d_k + \mu_{k+1}))$*: In this case, the allocation of the PPA policy is $x_k = \frac{d_k}{d_k + \mu_{k+1}} s_k$, which results in an FR less than β_k, i.e., $\frac{x_k}{d_k} = \frac{s_k}{d_k + \mu_{k+1}} < \beta_k$. According to the dynamics specified in (5) and (6), this implies

$$\beta_{k+1} = \frac{s_k}{d_k + \mu_{k+1}} \quad \text{and} \quad \frac{\mu_{k+1}}{s_{k+1}} = \frac{\mu_{k+1}}{s_k - \frac{d_k}{d_k + \mu_{k+1}} s_k} = \frac{d_k + \mu_{k+1}}{s_k}.$$

Using our inductive hypothesis when $i = k+1$,

$$\mathbf{E}_{\vec{d} \sim \mathcal{F}} \left[f_{n+1} \mid \vec{d}_{[1:k]} \right] \geq \frac{s_k}{d_k + \mu_{k+1}} \left(1 - \frac{n-k}{2(n+1-k)} \right) \triangleq \text{RHS}^{(2)}. \tag{11}$$

The lower bound given by $\text{RHS}^{(2)}$ is a convex homographic function of $d_k + \mu_{k+1}$, as illustrated by the dashed red lines in all panels of Figure 3 (in the regime where $d_k + \mu_{k+1} \in [s_k/\beta_k, +\infty)$). To further lower bound this function by a linear function, note that for any variable z the following inequality holds:

$$\left(\frac{n+2-k}{2(n+1-k)} \right) \frac{s_k}{z} - \beta_k \left(1 - \frac{n+1-k}{2(n+2-k)} \frac{\beta_k}{s_k} z \right) = \frac{(n+1-k)\beta_k^2}{2(n+2-k)s_k z} \left(\frac{(n+2-k)s_k}{(n+1-k)\beta_k} - z \right)^2 \geq 0 \ .$$

The proof of the above inequality is purely algebraic and we omit it for brevity. Substituting $z = d_k + \mu_{k+1}$ in this inequality, we have:

$$\text{RHS}^{(2)} \geq \beta_k \left(1 - \frac{n+1-k}{2(n+2-k)} \frac{d_k + \mu_{k+1}}{s_k} \beta_k \right), \tag{12}$$

which proves inequality (8) in the insufficient supply case (again, see the blue lines in all panels of Figure 3).

Combining the above cases proves inequality (8) everywhere, which immediately implies the inductive hypothesis, i.e., inequality (7), for $i = k$, thus finishing the proof of the lemma. $\square$

3.4. Simple Non-adaptive Policies and Ex-post Fairness

As discussed in the previous sections, our PPA policy is *adaptive*, that is, the FR for agent i (and its corresponding allocation decision) can depend not only on the observed demand d_i but also on the exact sample path up to time i as well as the remaining supply s_i. In contrast to an adaptive
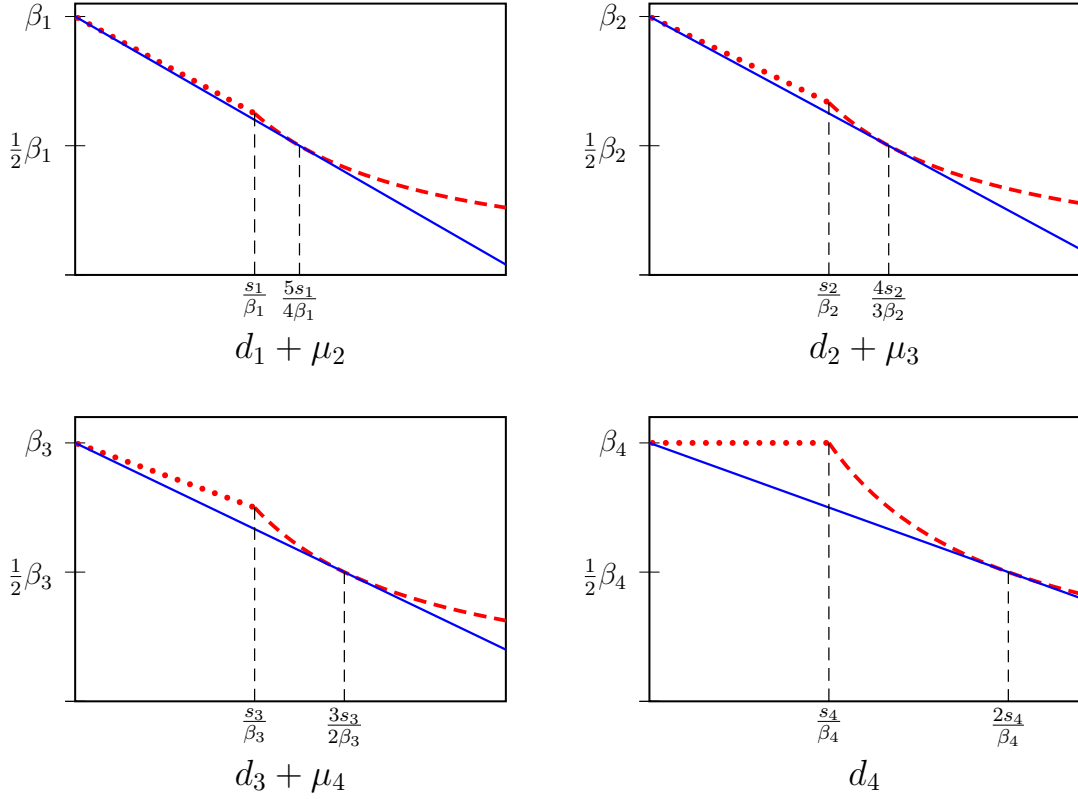
Figure 3 Lower bounds on the expected minimum FR given by eq. (8) (blue solid lines), eq. (9) (red dotted lines), and eq. (11) (red dashed lines) when $n = 4$ for $k \in [4]$.

policy, a *non-adaptive* policy commits to a sequence of feasible allocation maps $\{x_i(d_i)\}_{i \in [n]}$ upfront, where $x_i(d_i) \in [0, d_i]$ is the allocation decision for agent i when agent i has demand d_i.[18] If the non-adaptive policy's allocation decision $x_i(d_i)$ exceeds the remaining supply s_i, then agent i instead receives the entire remaining supply.

For settings that we consider, adaptivity can indeed help with improving the expected minimum FR of a policy. As an example, compare running our PPA policy versus the best non-adaptive policy on an instance with three agents. In this instance, the demands $\vec{d} = (d_1, d_2, d_3)$ follow one of the two possible sample paths $(\epsilon_1, 1, 1)$ or $(\epsilon_2, 1, 0)$ with equal probabilities $1/2$, where $\epsilon_1, \epsilon_2 \geq 0$ and $\epsilon_1 \neq \epsilon_2$. After agent 1's demand is realized, the PPA policy knows exactly which sample path is happening. By calculating the exact total demand of agents 2 and 3, it obtains the optimal expected minimum FR of $\frac{1}{2} \times 1 + \frac{1}{2} \times \frac{1}{2} = 3/4$ for small ϵ_1, ϵ_2. However, a non-adaptive policy cannot distinguish between the two possible sample paths after agent 1's demand is realized. Therefore,

---

[18] For ease of presentation, we focus on deterministic non-adaptive policies. This is without loss of generality, as the ex-post fairness of any randomized non-adaptive policy must be weakly dominated by the ex-post fairness of one of the deterministic policies that it randomizes over.

without loss of generality, it targets a FR of τ for agent 2 and obtains an expected minimum FR of $\frac{1}{2}\tau + \frac{1}{2}\min\{\tau, 1-\tau\}$ for small ϵ_1, ϵ_2, which attains its maximum equal to $\frac{1}{2}$ at any $\tau \in [\frac{1}{2}, 1]$.

In applications that allow for adaptivity, our PPA policy obtains the optimal ex-post fairness guarantee while also having the desirable properties of transparency and interpretability. However, adaptivity is not admissible in some practical scenarios—e.g., when the social planner should commit to an allocation plan in advance for even more transparency or due to legal restrictions.[19] Motivated by such scenarios, we study two simple and natural canonical classes of non-adaptive policies: those that fix the sequence of allocation decisions a priori, namely they specify one allocation vector $\vec{x}$, and "smarter" policies which fix the sequence of fill rates $\vec{\tau} \triangleq (\tau_1, \tau_2, \ldots, \tau_n)$ a priori. In Appendix A.4, we show that the ex-post fairness guarantee for the former subclass is vanishing as n gets large. Therefore, we focus on the latter subclass, which is formally defined as follows.

DEFINITION 2 (**Target-fill-rate Policies**). A *target-fill-rate (TFR)* policy is any policy π which pre-determines a target fill rate $\tau \in [0, 1]$. Then, for every arriving agent i, the policy π must either allocate sufficient supply to meet the target or allocate all remaining supply, i.e.,

$$\forall i \in [n]: \quad x_i(s_i, d_i) = \min\{\tau d_i, s_i\}.$$

In the following theorem, we provide a tight bound on the ex-post fairness guarantee (Definition 1) achievable by the *optimal* TFR policy—defined as the one that maximizes ex-post fairness for the given joint demand distribution. We remark that setting one threshold is without loss of generality because the ex-post fairness guarantee of a policy which pre-determines a *sequence* of target fill rates $\{\tau_i\}_{i \in [n]}$ is upper bounded by that of a TFR policy with the same target fill rate $\tau = \min_{i \in [n]}\{\tau_i\}$ for all agents. We also highlight that in addition to achieving a lower ex-post fairness guarantee compared to our adaptive policy, finding the best TFR policy requires full knowledge of the total demand distribution—in contrast to our PPA policy which only requires knowing the first conditional moments of the future total demand at each time.

THEOREM 3 (**Ex-post Fairness Guarantee of Optimal TFR Policy**). *Given any number of agents $n \in \mathbb{N} \setminus \{1\}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, the optimal TFR policy achieves an ex-post fairness guarantee (see Definition 1) of $\frac{\max\{1, \mu\}}{\mu + \sqrt{\mu^2 + 1}}$.*

In Figure 4, we compare the guarantee of the optimal TFR policy against our PPA policy for different model primitives, μ and n. First, we note that when n is not too large, our PPA policy achieves a considerably higher guarantee. Next, we highlight that the ex-post fairness guarantee

---

[19] As discussed in the introduction, the initial strategy for allocating medical supplies at the beginning of COVID-19 pandemic had the form of a target-fill-rate policy, which is a canonical non-adaptive strategy as we will discuss soon.
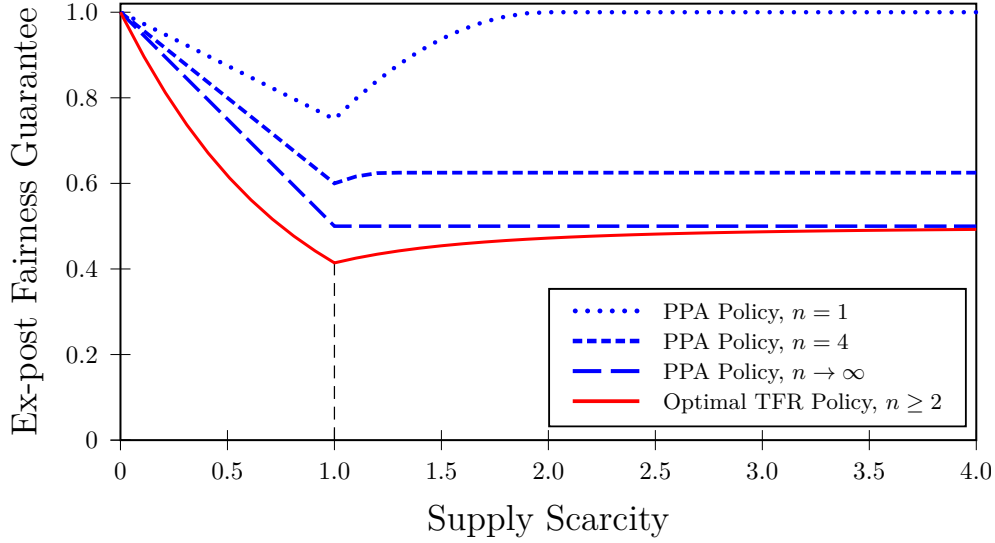
Figure 4 **Ex-post fairness guarantees of our PPA policy and the optimal TFR policy.**

for the optimal TFR policy does not depend on the number of agents n.[20] This is in contrast to the ex-post guarantee for the PPA policy $\kappa_{\mathrm{P}}(\mu, n)$, which worsens as the number of agents increases. Furthermore, the guarantee in Theorem 3 has a unique minimum of $\frac{1}{1+\sqrt{2}} \approx 0.41$ when $\mu = 1$. This once again suggests that the stochastic nature of demand is most harmful when expected demand exactly equals supply.

3.4.1. **Proof of Theorem 3** We begin by placing a lower bound on the performance of the optimal TFR policy, and we then demonstrate the existence of a matching upper bound.

Proof of lower bound: For any target fill rate τ, a TFR policy will achieve that fill rate if $\tau \left(\sum_{i \in [n]} d_i \right) \leq 1$. Let us define G as the cumulative distribution function (CDF) of the random variable $v \triangleq \frac{1}{\sum_{i \in [n]} d_i}$, where $\mathbf{E}_{v \sim G} \left[\frac{1}{v} \right] = \mu$. For ease of notation, we use $\overline{\Delta} (\mathbb{R}_{\geq 0}; \mu)$ to denote the domain of all such CDFs. Given a CDF G, a TFR policy with target fill rate τ achieves an expected minimum fill rate of at least $\tau(1 - G(\tau))$, which implies that the optimal TFR policy attains an expected minimum fill rate of at least $\max_{\tau \in [0,1]} \tau (1 - G(\tau))$. In the following lemma, we establish a lower bound on $\max_{\tau \in [0,1]} \tau (1 - G(\tau))$, which enables us to lower-bound the ex-post fairness guarantee that the optimal TFR policy achieves.

LEMMA 2 (**Tight Lower Bound for Optimal TFR Policy**). *Given a fixed number of agents $n \in \mathbb{N}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, the following holds:*

$$\inf_{G \in \overline{\Delta}\left(\mathbb{R}_{\geq 0}; \mu \right)} \left\{ \max_{\tau \in [0,1]} \tau \left(1 - G(\tau) \right) \right\} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}. \tag{13}$$

---

[20] We elaborate on the intuition behind this behavior when we present the hard instance for establishing the upper bound of Theorem 3.

This infimum is attained by the following CDF

$$\hat{G}(v) = \begin{cases} 0 & \text{if } v \in [0, \hat{q}) \\ 1 - \hat{q}/v & \text{if } v \in [\hat{q}, 1) \\ 1 - \hat{q} & \text{if } v \in [1, +\infty) \\ 1 & \text{if } v = +\infty \end{cases} \tag{14}$$

where $\hat{q} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}$.

Before presenting the proof of the above lemma in Section 3.4.2, which is the key step of the proof of Theorem 3, we establish a matching upper bound and complete the proof of the theorem.

Proof of upper bound: We show a matching upper bound by considering a two-agent instance. In this instance, only the first agent has stochastic demand. In particular, $d_1 = \frac{1-\epsilon}{v}$ for $v \sim \hat{G}$ (defined in eq. (14)) and $d_2 = \epsilon \mu$ deterministically. Note that $\mathbf{E}[d_1 + d_2] = \mu$. For any target fill rate $\tau' = (1 - \epsilon)\tau$ where $\tau \in [0, 1]$, supply will be exhausted before the arrival of agent 2 with probability $\hat{G}(\tau)$, in which case the minimum FR will be 0. Therefore, the expected minimum fill rate of the optimal TFR policy in this instance is at most

$$\max_{\tau \in [0,1]} (1 - \epsilon)^{-1} \tau \left(1 - \hat{G}(\tau)\right) = (1 - \epsilon)^{-1} \frac{1}{\mu + \sqrt{\mu^2 + 1}},$$

where the equality follows from Lemma 2.

By allowing $\epsilon \to 0$, we conclude that there exists an instance where the expected minimum fill rate of the optimal TFR policy is $\frac{1}{\mu + \sqrt{\mu^2 + 1}}$, which matches the lower bound from above. We remark that the construction of the above two-agent example clarifies why our upper bound does not depend on the number of agents: we can modify the example to an n-agent one where the total demand of the first $n - 1$ agents have correlated demand equal to $\frac{1-\epsilon}{v}$ for $v \sim \hat{G}$ and the last agent has a deterministic demand of $\epsilon \mu$.

With the above (matching) bounds, we complete the proof of Theorem 3 by scaling this tight bound by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$, to arrive at the guarantee stated in Theorem 3. □

3.4.2. Proof of Lemma 2 and Connections to Monopoly Pricing Having laid out the proof steps of Theorem 3, we now provide a constructive proof of the key lemma, i.e., Lemma 2. We do so by identifying properties of the worst-case distribution against the optimal TFR policy, which enables us to exactly characterize that distribution.

To aid in this proof, we introduce a one-to-one mapping of each target fill rate τ into the quantile space, such that quantile q corresponds to TFR τ if and only if there is sufficient supply to meet a fraction τ of demand with probability exactly q. We start by describing notation for this transformation, along with some basic properties, in the following definition. For simplicity of exposition, we assume all the distributions playing the role of G are non-atomic.[21]

---

[21] This assumption is without loss of generality, as one can always add an infinitesimal continuous perturbation to each distribution, which does not change any of the arguments in this proof.
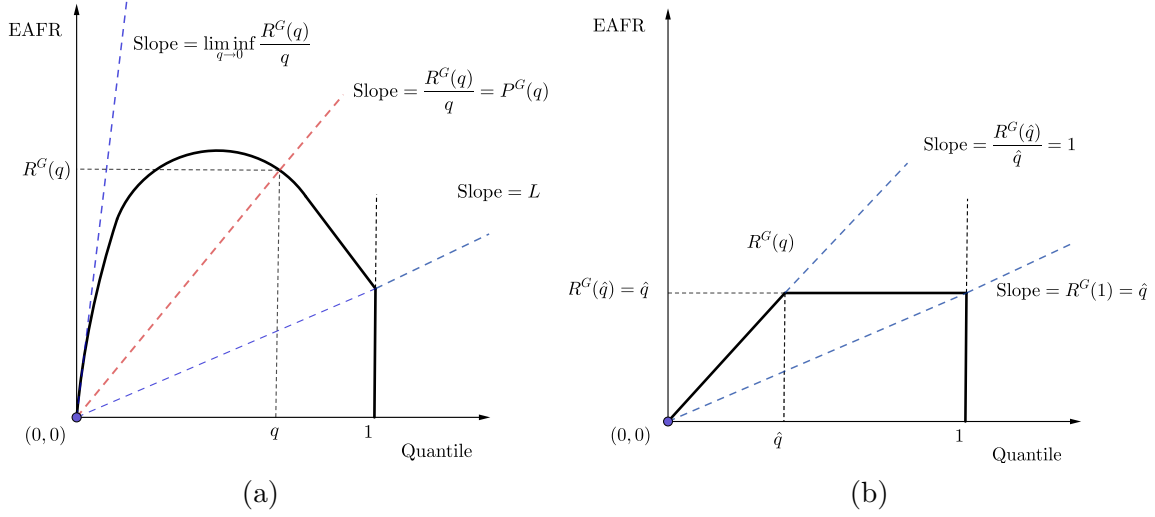
Figure 5 **(a) The expected achievable fill rate (EAFR) curve, fill rates, and support in quantile space. (b) The EAFR curve of the worst-case distribution, as defined in eq. (14), in quantile space.**

DEFINITION 3 (**TFR in Quantile Space**). Given a (non-atomic) CDF $G : \mathbb{R}_{\geq 0} \to [0,1]$ and inverse total demand $v \sim G$, we define the following mappings.

- *TFR-to-quantile map Q^G:* The quantile corresponding to TFR $\tau \in [0,1]$ is $Q^G(\tau) \triangleq 1 - G(\tau)$. In words, the probability of being able to meet a fraction τ of total demand is $Q^G(\tau)$. This map is monotone non-increasing.

- *Quantile-to-TFR map T^G:* The TFR corresponding to quantile $q \in [0,1]$ is $T^G(q) \triangleq G^{-1}(1-q)$. In words, $T^G(q)$ is the TFR for which the probability of being able to meet a fraction $T^G(q)$ of total demand is q. This map is monotone non-increasing and is the inverse of the TFR-to-quantile map, i.e., $T^G = (Q^G)^{-1}$.

- *The expected achievable fill rate (EAFR) curve R^G:* For $q \in [0,1]$, $R^G(q) \triangleq q \cdot G^{-1}(1-q)$ is the EAFR when the probability of meeting demand (given the TFR) is exactly equal to $q \in [0,1]$, i.e., the EAFR obtained by targeting a fill rate $T^G(q)$.

REMARK 2. In light of the above transformation, we remark that there is a reduction from our setup to a single-parameter Bayesian mechanism design problem in which a monopolistic seller has an item to sell to a single buyer with private valuation $v \sim G$, where G is the common prior valuation distribution. See Alaei et al. (2019) for an example of such a setting; also refer to Hartline (2013) for more details on monopoly pricing. In this reduction, target fill rates correspond to prices and the EAFR corresponds to the expected revenue in monopoly pricing (accordingly, the EAFR curve also corresponds to the revenue curve). The problem in this parallel monopoly pricing setting is identifying the worst-case distribution G satisfying $\mathbf{E}_{v \sim G}[1/v] = \mu$, so that we minimize the maximum revenue obtained from selling the item at prices constrained to be in the interval $[0,1]$.

According to Definition 3, $\tau(1 - G(\tau))$ is equivalent to $R^G(Q^G(\tau))$ for any TFR $\tau \in [0,1]$. Based on this insight, to prove Lemma 2, it is sufficient to show that

$$\inf_{G \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)} \left\{ \max_{q \in [0,1]: T^G(q) \in [0,1]} R^G(q) \right\} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}. \tag{15}$$

Consider all cumulative distribution functions $G \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$. We first identify two additional constraints on G that do not change the infimum in eq. (15). These constraints enable us to find the worst-case distribution that achieves the infimum value which establishes the desired result.

Before proceeding, we develop intuition using an illustrative example of the EAFR curve shown in Figure 5(a). In general, if one draws $R^G(q)$ as a function of $q \in [0,1]$ (i.e., in the quantile space), then the slope of the line connecting the point $(0,0)$ to $(q, R^G(q))$ is equal to $T^G(q) = R^G(q)/q$. This slope is monotone non-increasing in q for any CDF G according to Definition 3. Hence, given the EAFR curve $R^G(q)$, the support of the feasible fill rates is equal to $[L, H]$, where $L \triangleq R^G(1)$ and $H \triangleq \min\left\{1, \liminf_{q \to 0} R^G(q)/q\right\}$. The two constraints that we will add below, as stated in Claims 1 and 2, imply that the outer optimization problem in eq. (15) will remain unchanged if we require the EAFR curve to be (i) flat over quantiles corresponding to target fill rates in $[L, 1]$, i.e., quantiles in the interval $[Q^G(1), 1]$, and (ii) a straight line with slope 1 for quantiles in the interval $[0, Q^G(1)]$. With these two additional constraints, in Claim 3 we find the worst-case CDF, which has an EAFR curve as shown in Figure 5(b).

CLAIM 1 (**Equal EAFR**). *Adding the constraint $R^G(q) = R^G(q'), \forall q, q' \in [Q^G(1), 1]$ to the outer optimization in eq. (15) does not change its infimum value.*

We prove Claim 1 by contradiction: we show that for any CDF $G \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$, if the above condition does not hold, we can slightly modify G to design a new distribution $\tilde{G} \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$ which has an EAFR curve with a lower maximum value. The details are presented in Appendix A.5.1. The above claim readily implies that we can focus on distributions for which the EAFR curve is flat in the interval $[Q^G(1), 1]$.

Next, we claim that we can restrict our attention to distributions where there is no probability mass for $v \in (1, +\infty)$. Said differently, the support of inverse demand is $(0, 1] \cup \{+\infty\}$.

CLAIM 2 (**Restricted Support for Inverse Demand**). *Adding the constraint $G(v) = G(1)$ for all $v \in [1, +\infty)$ and $\lim_{v \to +\infty} G(v) = 1$ to the outer optimization in eq. (15) does not change its infimum value.*

We also prove Claim 2 by contradiction: we show that for any CDF $G \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$, if there is probability mass on $v \in (1, +\infty)$, we can construct a CDF $\tilde{G} \in \overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$ which has an EAFR curve with a lower maximum value by shifting that mass to $+\infty$. The details are presented in

Appendix A.5.2. Again, note that this claim implies that we can focus on distributions for which the EAFR curve starts with a straight line up to quantile $Q^G(1)$.

Given the two claims above, the distribution that attains the infimum in eq. (15) must satisfy the two constraints introduced. Figure 5(b) summarizes the effect of these two restrictions on $R^G(q)$.

CLAIM 3 (**Worst-case CDF**). *For any $\mu \in \mathbb{R}_{\geq 0}$, the distribution $\hat{G}$ given in eq. (14) is the unique distribution in $\overline{\Delta}(\mathbb{R}_{\geq 0}; \mu)$ satisfying the constraints introduced in Claims 1 and 2. Therefore, this distribution attains the infimum in eq. (15).*

We prove Claim 3 in Appendix A.5.3. Since the EAFR curve $R^{\hat{G}}(q)$ has a maximum value of $\hat{q} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}$, we have shown that the optimal TFR policy always achieves an EAFR of at least $\hat{q}$. This completes the proof of Lemma 2. □

3.5. Ex-ante Fairness

In this section, we study our second notion of fairness, namely, ex-ante fairness. As we did for ex-post fairness, we first establish an upper bound on the ex-ante fairness guarantee achievable by any policy. More importantly, we then show that our PPA policy achieves this worst-case ex-ante fairness bound. The following theorem establishes our matching upper and lower bounds on the ex-ante fairness guarantee.

THEOREM 4 (**Ex-ante Fairness Guarantee of PPA Achieves Upper Bound**). *Given a fixed number of agents $n \in \mathbb{N}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, no sequential allocation policy obtains an ex-ante fairness guarantee (see Definition 1) greater than $\kappa_{\mathrm{A}}(\mu, n)$, defined as*

$$\kappa_{\mathrm{A}}(\mu, n) = \begin{cases} 1 - \frac{\mu}{4}, & \mu \in [0, 1) \\ \mu\left(1 - \frac{\mu}{4}\right), & \mu \in [1, 2) \\ 1, & \mu \in [2, +\infty) \end{cases} . \tag{16}$$

Further, the PPA policy achieves an ex-ante fairness guarantee of at least $\kappa_{\mathrm{A}}(\mu, n)$.

Like its counterpart for ex-post fairness, $\kappa_{\mathrm{A}}(\mu, n)$ depends on the supply scarcity, μ, and is at its lowest when expected demand equals supply, which highlights the loss due to stochasticity when trying to achieve efficiency and equity ex ante. However, unlike the bound for ex-post fairness, the ex-ante fairness bound is independent of the number of agents. In fact, this bound is identical to the ex-post fairness bound in the single-agent case, i.e. $\kappa_{\mathrm{A}}(\mu, n) = \kappa_{\mathrm{P}}(\mu, 1)$ (which is shown by the dotted line in Figure 4).

For intuition about this relationship, note that one feasible policy is to allocate supply to each agent proportional to their expected demand. Since the ex-ante problem only depends on marginal FRs, this reduces ex-ante fairness to the minimum ex-ante fairness across n single-agent instances (where in each instance, the supply scarcity is μ). In a single-agent instance, ex-ante fairness is

equal to ex-post fairness, which implies that any lower bound on single-agent ex-post fairness also serves as a lower bound on ex-ante fairness with n agents. Furthermore, since demands can be perfectly correlated, any single-agent instance can be expressed as an instance with n agents for any $n \in \mathbb{N}$. This implies that any upper bound on single-agent ex-post fairness also serves as a upper bound on ex-ante fairness with n agents.

To prove the upper bound in Theorem 4, we build on the hard instances from the proof of Theorem 1. To prove the lower bound, we use ideas similar to the proof of Theorem 2. We show that when following the PPA policy, the expected FR for each agent is a decreasing and convex function of the ratio of expected remaining demand to remaining supply upon their arrival. We inductively place an upper bound on the ex-ante expected value of that ratio for each agent, which enables us to provide a lower bound on ex-ante fairness. See Appendix A.6 for a detailed proof.

4. Numerical Results

We complement our theoretical developments with an illustrative case study motivated by the allocative challenges that FEMA faced when rationing COVID-19 medical supplies in the early days of the pandemic in the US. First, we provide background on the sequential and heterogeneous nature of the demand from different states. Then, we study this dynamic rationing problem within our framework and illustrate the effectiveness of our PPA policy by comparing it to its theoretical guarantee, to the optimal TFR policy, and to a DP approach.

4.1. Background and Model Primitives

The U.S. federal government is equipped with a stockpile of medical resources which can be used to alleviate excessive demands on states' local resources. However, the stockpile is intended to address isolated emergencies in a small number of states, not a nationwide epidemic. Thus, due to COVID-19, the total demand for certain resources was expected to far exceed the federal government's stockpile. Given the uneven spread of COVID-19 toward the outset of the pandemic, different states' needs were expected to realize at different times. Furthermore, the size of the demands were likely to be correlated across states. When states with early outbreaks began requesting supplies, FEMA had to determine how much of the federal stockpile they should allocate and how much they should save to meet the projected future needs of other states.[22]

To capture the setting faced by FEMA, we draw upon the projections made by the Institute for Health Metrics and Evaluation (IHME) at the University of Washington as of April 1, 2020.[23] These projections were cited by Dr. Deborah Birx (the White House Coronavirus Response Coordinator)

---

[22] In mid-March, FEMA took control of the stockpile and was tasked with distributing medical supplies.

[23] The data can be found at http://www.healthdata.org/covid/data-downloads.

as influencing decision-making at that time (Washington Post 2020b). We use each state's projected peak excess demand for ICU beds as a proxy for their demand for medical supplies. The scale and timing of the projected peak demand varied significantly from state to state. To succinctly present the temporal aspect of states' demands, we partition them into four groups based on when their peak demand was projected to occur: April 1st-7th, April 8th-14th, April 15th-21st, and after April 21st. We visualize this partitioning in the map in the left panel of Figure 6(a).

Next, we describe how we construct the demand distributions for the aforementioned groups. We remark that the IHME's projections only consist of the estimated mean as well as estimated 5th and 95th percentiles for each individual state. Said differently, the model does not specify a joint (or even marginal) distribution for the demand. Consequently, we augment the model by assuming that the aggregate demand in each group $i \in [4]$ is Normally distributed with mean equal to the IHME's total estimated mean for the states belonging to group i. The standard deviation for group i's demand distribution is set such that the 5th percentile of demand is equal to the 5th percentile of projected peak excess demands for ICU beds from all states in group i. Summary statistics about the demand distributions are presented in the right panel of Figure 6(b).[24] We highlight the significant heterogeneity in both the mean and the standard deviation across these four groups. Finally, we construct the joint distribution for these four groups by letting all of the pairwise correlation coefficients equal ρ, where $\rho \in \{0, 0.25, 0.5, 0.75\}$.[25] The limited number of groups and simple model of correlated demand allow us to compare the PPA policy against a DP solution with manageable computing time, as we discuss in the following subsection.

4.2. Policy Evaluation

For the setup described above, we numerically evaluate the performance of various policies. In the simulations, we consider three values for supply scarcity (1, 2, and 4) by setting the initial supply relative to the sum of expected demand across the four groups according to the IHME model. Then, fixing the correlation coefficient $\rho \in \{0, 0.25, 0.5, 0.75\}$, we generate 10,000 sample paths for the purpose of Monte-Carlo simulation.

For each sample path, we compute the minimum FR of the following policies: (i) our PPA policy, (ii) the optimal TFR policy (defined in Section 3.4), and (iii) the solution to a DP computed by discretizing the state space, which consists of the remaining supply, the current demand, the minimum FR, and the distribution of future (i.e., remaining) demand given the observed demand thus far. We choose the discretization level such that each DP can be solved in under three hours

---

[24] In rare cases, the demand drawn from the described Normal distributions can be negative. In such cases, we assume the demand is 0.

[25] In this context, we focus on different levels of positive correlation to capture network effects during the pandemic, e.g., virus transmission across state borders.

Group	Mean Demand	Std. Dev.
1	15,366	3287
2	2193	710
3	9719	3600
4	3831	1763

(a)   (b)
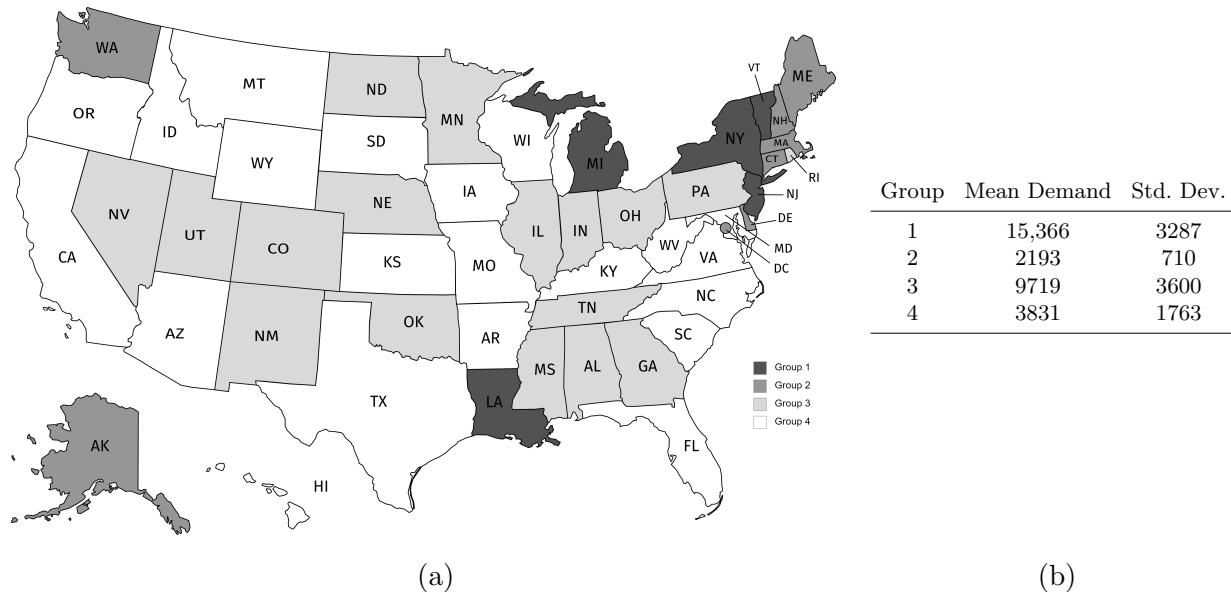
Figure 6   **(a) The allocation group for each state, which is based on the projected timing of excess demand for ICU beds. (b) Summary of the demand distribution for each group of states.**

on Google's Compute Engine. In Table 1, we present the average ex-post fairness for the aforementioned policies along with the PPA policy's theoretical guarantee for ex-post fairness (as given in Theorem 2).[26] We make three key observations:

- **PPA vs. Guarantee:** In all cases, our policy performs more than 33% better than its theoretical guarantee. In addition, in over-demanded settings our policy performs better when the correlation is higher because it leverages the information provided by realized demands. However, when supply is equal to expected demand, the extra information can be outweighed by the strain that correlated demands place on the available supply.

- **PPA vs. Optimal TFR:** Our PPA policy outperforms the optimal TFR policy by between 7% and 28%. As we would expect, the gap is largest when there is correlation between demands since TFR policies do not take advantage of the information that can be extracted from realized demands. Correlation between demands is likely strong during a nationwide epidemic, implying that our PPA policy is better suited than the optimal TFR policy in such situations.

- **PPA vs. DP:** The solution to the DP, despite being challenging to compute even in this simple setting, exhibits nearly identical performance to the PPA policy.

Moving beyond the performance comparison, we highlight that the PPA policy can be implemented simply by knowing the expected remaining demand, which was frequently updated by the

---

[26] We remark that ex-post fairness is well-concentrated around its average value for all of the considered policies.

IHME. In contrast, determining the optimal TFR policy or solving the DP requires *full* distributional information, which was not provided by the IHME. Further, regarding the DP solution, we remark that our distributional assumptions (four agents and a simple correlation structure across demands) make it possible to solve such a DP with a granular state space in a reasonable amount of time. However, we emphasize that in other settings, particularly those with a more complicated correlation structure, such an approach may not be practical. The solution to the DP also suffers from two additional shortcomings: it achieves consistently worse ex-ante fairness than the PPA policy (by an average of 6%), and its allocation is not transparent. Transparency is particularly important in this setting, as numerous states questioned the allocation procedures implemented by the federal government (see Footnote 1 for one such example). Similar to TFR policies, the PPA policy follows a strategy that can be easily explained to stakeholders.

	Supply Scarcity (μ)											
	1				2				4			
Demand Correlation (ρ)	0.00	0.25	0.50	0.75	0.00	0.25	0.50	0.75	0.00	0.25	0.50	0.75
PPA Policy	0.85	0.85	0.85	0.86	0.84	0.85	0.87	0.92	0.84	0.85	0.87	0.91
PPA Guarantee		0.600				0.625				0.625		
Opt. TFR Policy	0.79	0.76	0.74	0.72	0.79	0.76	0.74	0.72	0.79	0.76	0.74	0.72
DP Solution	0.85	0.85	0.85	0.86	0.85	0.87	0.89	0.92	0.86	0.86	0.89	0.92

Table 1 Ex-post fairness of three policies across 10,000 simulations of the setting described in Figure 6(b), as well as the ex-post fairness guarantee of the PPA policy.

5. Concluding Remarks, Extensions, and Future Directions

We conclude the paper by first summarizing our main findings, and then elaborating on a few extensions of our base framework. We finish by listing a few future directions.

5.1. Summary

In this paper, we initiate the theoretical study of fair dynamic rationing by introducing a simple yet fundamental and well-motivated framework. In a nutshell, we design sequential policies for allocating limited supply to a sequence of arbitrarily correlated demands given an objective which encompasses the dual goals of efficiency and equity. Based on our formalized notions of ex-post and ex-ante fairness, we establish upper bounds on the fairness guarantees achievable by any sequential allocation policy which depend on the supply's scarcity level and the number of demanding agents. More importantly, we show that our simple PPA policy achieves the "best of both worlds" by attaining the upper bound on both the ex-post and ex-ante fairness guarantees. In addition to enjoying optimal fairness guarantees, our PPA policy is practically appealing: it is interpretable

as well as computationally efficient since it does not rely on distributional knowledge beyond the conditional first moments.

Our framework lends itself to extensions such as considering generalized objectives and rationing multiple types of resources. More broadly, it serves as a base model for theoretically studying sequential allocation problems with an objective beyond utility maximization, which in turn opens several new research directions. In the rest of this section, we first discuss the aforementioned extensions of our base model and then finish the paper by discussing future directions.

5.2. Extensions

5.2.1. Generalized Social Welfare Objective Functions

Throughout the paper, we have focused on the minimum FR as a social welfare objective that combines elements of equity and efficiency (i.e., $U(\vec{x}) = \min_{i \in [n]} \{\frac{x_i}{d_i}\}$). However, this social welfare function—which is also known as the Rawlsian social welfare function thanks to the philosophical work of John Rawls (1973)—is only a special case of a more general class of social welfare functions that we call the *weighted power mean (WPM) social welfare* family of functions. More precisely, this family is parameterized by $\alpha \in [0, +\infty)$ and defined as

$$U_\alpha(\vec{x}) \triangleq \begin{cases} \left(\frac{1}{\sum_{i \in [n]} d_i} \sum_{i \in [n]} d_i \left(\frac{x_i}{d_i} \right)^{1-\alpha} \right)^{1/(1-\alpha)}, & \alpha \neq 1 \\ \prod_{i \in [n]} \left(\frac{x_i}{d_i} \right)^{d_i / \sum_{i \in [n]} d_i}, & \alpha = 1 \end{cases} \tag{17}$$

Note that the above is a weighted version of the celebrated power mean functions, introduced in Atkinson et al. (1970), that provides a broad class of social welfare functions which balance equity and efficiency to varying degrees. Having weights proportional to the demands in eq. (17) ensures that equity is measured *relative* to demand and not simply based on the absolute allocation.[27] By varying the parameter α from 0 to $+\infty$, the focus of the planner is shifted from extreme efficiency towards more equitable allocations. When $\alpha = 0$, a utilitarian allocation (i.e., any allocation without waste) maximizes social welfare. In the limit as $\alpha \to 1$, proportional fairness (i.e., a generalization of the Nash bargaining solution) maximizes social welfare. Finally, in the limit as $\alpha \to +\infty$, maximizing the minimum FR maximizes social welfare. In fact, we highlight that the value of this social welfare function exactly approaches our objective in the base model (i.e., the minimum FR, or equivalently, the Rawlsian social welfare function).

For any parameter α (including when $\alpha \to +\infty$, which corresponds to our base model), the optimal policy when demands are deterministic is to allocate the supply proportionally. Such a

---

[27] Further, this family of functions has a one-to-one relationship with the α-fairness social welfare functions introduced in Mo and Walrand (2000). In fact, the two families of functions are the same up to a transformation via a one-to-one, increasing function, which means that the maximizing vectors are identical for a given α.

policy achieves the optimal social welfare of $\overline{W}$ (defined in eq. (1)).[28] However, the value of the parameter α impacts the optimal policy when demands are stochastic. To study this impact, we naturally generalize our notion of ex-post fairness to WPM social welfare functions, i.e., ex-post fairness is given by $\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}[U_\alpha] / \overline{W}$. We remark that in the limit as $\alpha \to +\infty$, this is equivalent to the notion of ex-post fairness introduced in Section 2. In the following corollary of Theorem 2, we establish that our PPA policy achieves an ex-post fairness guarantee of at least $\kappa_{\mathrm{P}}(\mu, n)$ for any $\alpha \in [0, +\infty)$.

COROLLARY 1 (**PPA's Guarantee for WPM Objectives**). *Given a fixed number of agents $n \in \mathbb{N}$, supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, and any $\alpha \in [0, +\infty)$, the PPA policy achieves an ex-post fairness guarantee of at least $\kappa_{\mathrm{P}}(\mu, n)$ (defined in eq. (2)) when social welfare is measured by a WPM function (defined in eq. (17)) with parameter α.*

We prove Corollary 1 in Appendix B.1.

5.2.2. Rationing Multiple Types of Resources In our base model, we assume that agents have demand for only a single type of resource. However, in many of the motivating applications that we consider, agents may have concurrent demand for multiple types of resources. For example, states may need many different types of medical supplies during the peak of a pandemic.

Our setup readily extends to the sequential allocation of m different resource types where an arriving agent *simultaneously* demands m types of supply. We allow the demands to be correlated across agents as well as resource types. For the sake of brevity, we refrain from repeating the setup and we simply augment our notation for various quantities (e.g., supply, demand, and allocation) by adding a superscript $j \in [m]$. In this generalized model, we define agent i's utility to be their *weighted* FR, defined as:

$$\sum_{j \in [m]} \lambda^j \frac{x_i^j}{d_i^j}, \tag{18}$$

where we normalize the weights λ^j to satisfy $\sum_{j \in [m]} \lambda^j = 1$. A simple corollary of Theorem 2, as stated below, ensures that independently following our PPA policy for each resource achieves a lower bound on the expected minimum weighted FR which is a weighted sum of the expected minimum FR guaranteed by the PPA for one resource, i.e. $\kappa_{\mathrm{P}}(\mu^j, n) \max\{1, \mu^j\}$ for resource j.

---

[28] Note that in a deterministic setting, maximizing any social welfare objective function as in eq. (17) is a concave maximization. By writing KKT conditions it is not hard to see that any such function attains its maximum at a feasible proportional allocation, i.e., $x_i = \min\left\{d_i, \frac{d_i}{\sum_{j \in [n]} d_j}\right\}$, under deterministic demands. The maximum is then equal to $\overline{W} = \min\{1, 1/\mu\}$.

COROLLARY 2 (**PPA's Guarantee on Expected Minimum Weighted FR**). *Consider any instance with $n \in \mathbb{N}$ agents and $m \in \mathbb{N}$ resources, where the initial supply for resource j is $s^j \in \mathbb{R}_{\geq 0}$. For any joint demand distribution over all agents and resources $\vec{\mathcal{F}} \in \Delta(\mathbb{R}_{\geq 0}^{n \times m})$, independently following the PPA policy for each resource achieves an expected minimum weighted FR (as defined in eq. (18)) of at least $\sum_{j \in [m]} \lambda^j \kappa_{\mathrm{P}} \left(\frac{\mu^j}{s^j}, n \right) \max \left\{ 1, \frac{\mu^j}{s^j} \right\}$, where $\mu^j \in \mathbb{R}_{\geq 0}$ is the expected total demand for resource j.*

In addition, since demand can be correlated across agents, we can re-use the hard instances of Theorem 1 to construct a joint distribution which establishes an upper-bound on the performance of any policy matching the lower bound given in Corollary 2. We state this upper bound as a corollary of Theorem 1 below.

COROLLARY 3 (**Upper Bound on Expected Minimum Weighted FR**). *For any $n \in \mathbb{N}$ agents, $m \in \mathbb{N}$ resources, and any initial supply for resource j of $s^j \in \mathbb{R}_{\geq 0}$, there exists a joint demand distribution over all agents and resources $\vec{\mathcal{F}} \in \Delta(\mathbb{R}_{\geq 0}^{n \times m})$ for which no policy can achieve an expected minimum weighted FR greater than $\sum_{j \in [m]} \lambda^j \kappa_{\mathrm{P}} \left(\frac{\mu^j}{s^j}, n \right) \max \left\{ 1, \frac{\mu^j}{s^j} \right\}$, where $\mu^j \in \mathbb{R}_{\geq 0}$ is the expected total demand for resource j.*

Together, these two corollaries (which we prove in Appendix B) establish that independently following the PPA policy for each resource j provides the best possible guarantee on the expected minimum weighted FR. Consequently, we can use our PPA policy to shed light on how the social planner can prepare for demand across multiple types of resources. If the initial endowment of different resource types is not exogenously set, then the social planner can solve an outer endowment optimization problem to maximize the guarantee on the expected minimum weighted FR subject to a budget constraint.

We remark that such an endowment optimization problem is a max-max-min problem where the social planner first optimizes the initial endowment across resource types subject to a budget constraint (the outer problem). Then, given the initial endowment, the social planner maximizes over policies the minimum over demand distributions of our objective, i.e., the expected minimum weighted FR among agents. We solve the outer maximization of the multiple resource-type problem by determining the optimal initial endowment across resource types when the social planner independently follows the PPA policy for each resource. To be concrete, suppose the social planner has a fixed budget B that can be used to procure an initial endowment $\vec{s} = (s^1, s^2, \ldots, s^m)$. Further, suppose the per unit cost of resource $j \in [m]$ is c^j. Then, the outer endowment optimization problem can be formulated as follows:

$$\max_{\vec{s} \in \mathbb{R}_{\geq 0}^m} \quad \sum_{j \in [m]} \lambda^j \kappa_{\mathrm{P}} \left(\frac{\mu^j}{s^j}, n \right) \max \left\{ 1, \frac{\mu^j}{s^j} \right\} \qquad \text{s.t.}$$

(ENDOWMENT OPTIMIZATION)

$$B \geq \sum_{j \in [m]} c^j s^j$$

We highlight that to formulate this optimization problem, we crucially use the parameterized characterization of the ex-post fairness guarantee in Theorem 2, as opposed to the worst-case guarantee for any set of parameters. Further, we remark that the above maximization problem is linearly separable and concave in the decision variables $\vec{s}$, meaning that it can be solved efficiently.[29] Based on Corollaries 2 and 3, the optimal solution $\vec{s}^*$, combined with independently implementing the PPA policy for each resource, is indeed the optimal solution of the max-max-min multiple resource-type problem.

5.3. Future Directions

Our paper can be viewed as an analog of the classic prophet inequality problem (Krengel and Sucheston 1978, Samuel-Cahn et al. 1984) for equitably allocating divisible goods. As such, similar to prophet inequalities, many interesting variants of our setting arise. We discussed two such variants above, and for both, we established an achievable lower bound by employing our PPA policy. However, in the former variant, we do not establish a matching upper bound. Establishing tight bounds on the achievable performance in such a setting—which may require the use of a different policy—is an interesting direction for future research. Further, understanding the inefficiency (unused supply) which may occur in sequential allocation due to our focus on an egalitarian objective is a fruitful research direction that we plan to pursue. Finally, here we made no assumption about the correlation structure underlying the demand sequence. It would be compelling to investigate whether including a (well-motivated) correlation structure can result in improved fairness guarantees.

References

Nikhil Agarwal, Itai Ashlagi, Michael A Rees, Paulo J Somaini, and Daniel C Waldinger. An empirical framework for sequential assignment: The allocation of deceased donor kidneys. Technical report, National Bureau of Economic Research, 2019.

Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1405–1424. SIAM, 2014.

Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510, 2019.

Jerry Anunrojwong, Krishnamurthy Iyer, and Vahideh Manshadi. Information design for congested social services: Optimal need-based persuasion. *arXiv preprint arXiv:2005.07253*, 2020.

---

[29] It is not difficult to check that the function $\kappa_{\mathrm{P}}\left(\frac{\mu^j}{s^j}, n\right) \max\left\{1, \frac{\mu^j}{s^j}\right\}$ is concave in s^j for any choice of μ^j and n; check eq. (2) for a definition of $\kappa_{\mathrm{P}}(\cdot, \cdot)$. We omit this purely algebraic proof for brevity.

Nick Arnosti and Peng Shi. Design of lotteries and wait-lists for affordable housing allocation. *Management Science*, 66(6):2291–2307, 2020. doi: 10.1287/mnsc.2019.3311.

Kenneth Joseph Arrow. *Social Choice and Individual Values*. Yale University Press, 1963.

Itai Ashlagi, Patrick Jaillet, and Vahideh H Manshadi. Kidney exchange in dynamic sparse heterogenous pools. *arXiv preprint arXiv:1301.3509*, 2013.

Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019.

Anthony B Atkinson et al. On the measurement of inequality. *Journal of economic theory*, 2(3):244–263, 1970.

Yossi Azar, Niv Buchbinder, and Kamal Jain. How to allocate goods in an online market? In *European Symposium on Algorithms*, pages 51–62. Springer, 2010.

Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. The best of many worlds: Dual mirror descent for online allocation problems. *arXiv preprint arXiv:2011.10124*, 2020.

Santiago R Balseiro and David B Brown. Approximations to stochastic dynamic programs via information relaxation duality. *Operations Research*, 67(2):577–597, 2019.

Mohammad Hossein Bateni, Yiwei Chen, Dragos Florin Ciocan, and Vahab Mirrokni. Fair resource allocation in a volatile marketplace. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 819–819, 2016.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59 (1):17–31, 2011.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

William Cai, Johann Gaebler, Nikhil Garg, and Sharad Goel. Fair allocation through selective information acquisition. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 22–28, 2020.

Andre P Calmon, Florin D Ciocan, and Gonzalo Romero. Revenue management with repeated customer interactions. *Management Science*, 2020.

Semih Cayci, Swati Gupta, and Atilla Eryilmaz. Group-fair online allocation in continuous time. *Advances in Neural Information Processing Systems*, 33, 2020.

Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. Efx exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 1–19, 2020.

Dragos Florin Ciocan and Vivek Farias. Model predictive control for dynamic resource allocation. *Mathematics of Operations Research*, 37(3):501–525, 2012.

Maxime Cohen, Adam N Elmachtoub, and Xiao Lei. Price discrimination with fairness constraints. *Available at SSRN 3459289*, 2019.

Kate Donahue and Jon Kleinberg. Fairness and utilization in allocating resources with uncertain demand. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 658–668, 2020.

Ohad Eisenhandler and Michal Tzur. The humanitarian pickup and distribution problem. *Operations Research*, 67(1):10–32, 2019.

Ezekiel J. Emanuel, Govind Persad, Ross Upshur, Beatriz Thome, Michael Parker, Aaron Glickman, Cathy Zhang, Connor Boyle, Maxwell Smith, and James P. Phillips. Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, 382(21):2049–2055, 2020. doi: 10.1056/NEJMsb2005114. URL https://doi.org/10.1056/NEJMsb2005114.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.

Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. *arXiv preprint arXiv:2005.14122*, 2020.

Aram Grigoryan. Effective, fair and equitable pandemic rationing. *Available at SSRN 3646539*, 2020.

Swati Gupta and Vijay Kamble. Individual fairness in hindsight. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 805–806, 2019.

Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

Nicole Immorlica, Sahil Singla, and Bo Waggoner. Prophet inequalities with linear correlations and augmentations. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 159–185, 2020.

Stefanus Jasin and Sunil Kumar. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345, 2012.

Jiashuo Jiang, Shixin Wang, and Jiawei Zhang. Achieving high individual service-levels without safety stock? optimal rationing policy of pooled resources. *Optimal Rationing Policy of Pooled Resources (May 2, 2019). NYU Stern School of Business*, 2019.

Edward H Kaplan. *Managing the Demand for Public Housing*. PhD thesis, Massachusetts Institute of Technology, 1984.

Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

Jacob D. Leshno. Dynamic matching in overloaded waiting lists. Technical report, SSRN:2967011, 2017. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2967011.

Retsef Levi, Elisabeth Paulson, and Georgia Perakis. Optimal interventions for increasing healthy food consumption among low income households. *Available at SSRN 3486292*, 2019.

Robert W Lien, Seyed MR Iravani, and Karen R Smilowitz. Sequential resource allocation for nonprofit operations. *Operations Research*, 62(2):301–317, 2014.

Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.

Will Ma and Pan Xu. Group-level fairness maximization in online bipartite matching. *arXiv preprint arXiv:2011.13908*, 2020.

Aranyak Mehta. Online matching and ad allocation. *Theoretical Computer Science*, 8(4):265–368, 2012.

Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on networking*, 8(5):556–567, 2000.

Vedant Nanda, Pan Xu, Karthik Abhinav Sankararaman, John Dickerson, and Aravind Srinivasan. 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, pages 2210–2217, 2020.

Irem Sengul Orgut, Julie S Ivy, Reha Uzsoy, and Charlie Hale. Robust optimization approaches for the equitable and effective distribution of donated food. *European Journal of Operational Research*, 269 (2):516–531, 2018.

Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.

Parag A Pathak, Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez. Fair allocation of vaccines, ventilators and antiviral treatments: leaving no ethical value behind in health care rationing. *arXiv preprint arXiv:2008.00374*, 2020.

John Rawls. A theory of justice, 1973.

Yosef Rinott and Ester Samuel-Cahn. Optimal stopping values and prophet inequalities for some dependent random variables. *Lecture Notes-Monograph Series*, 22:343–358, 1992. ISSN 07492170. URL `http://www.jstor.org/stable/4355751`.

Ester Samuel-Cahn et al. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, 12(4):1213–1216, 1984.

Sean R Sinclair, Gauri Jain, Siddhartha Banerjee, and Christina Lee Yu. Sequential fair allocation of limited resources under stochastic demands. *arXiv preprint arXiv:2011.14382*, 2020.

Senay Solak, Christina Scherrer, and Ahmed Ghoniem. The stop-and-drop problem in nonprofit food distribution networks. *Annals of Operations Research*, 221(1):407–426, 2014.

Van-Anh Truong and Xinshang Wang. Prophet inequality with correlated arrival probabilities, with application to two sided matchings. *arXiv preprint arXiv:1901.02552*, 2019.

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

Washington Post. Desperate for medical equipment, states encounter a beleaguered national stockpile, 2020a. URL `https://www.washingtonpost.com/national/health-science/desperate-for-medical-equipment-states-encounter-a-beleaguered-national-stockpile/2020/03/28/1f4f9a0a-6f82-11ea-aa80-c2470c6b2034_story.html`.

Washington Post. Trump projects up to 240,000 coronavirus deaths in u.s., even with mitigation efforts, 2020b. URL `https://www.washingtonpost.com/politics/trump-white-house-projects-up-to-240000-coronavirus-deaths-in-us-even-with-mitigation-efforts/2020/03/31/62df5344-7367-11ea-87da-77a8136c1a6d_story.html`.

Appendix

A. Missing Proofs and Discussions of Section 3

A.1. Proof of Theorem 1 (Section 3.1)

We prove the theorem by considering two separate cases corresponding to the over-demanded regime ($\mu \geq 1 + \frac{1}{n}$) and the under-demanded regime ($\mu < 1 + \frac{1}{n}$). For each regime, we provide an instance of the problem under which no sequential allocation policy obtains ex-post fairness larger than $\kappa_{\mathrm{P}}(\mu, n)$ restricted to that regime.

Over-demanded regime ($\mu \geq 1 + \frac{1}{n}$): Consider an instance with n equally likely scenarios, where in scenario σ all agents $i \in [\sigma]$ have demand $d_i = \frac{2\mu}{n+1}$. This instance is depicted in Figure 2(a). We remark that the total expected demand is equal to μ, simply because

$$\sum_{\sigma \in [n]} \frac{1}{n} \cdot \frac{2\mu}{n+1} \cdot \sigma = \mu.$$

In such a setting, whenever agent i has non-zero demand, every agent j where $j < i$ must also have had non-zero demand. Since the policy cannot distinguish among scenarios $i, i+1, \ldots, n$, its allocation decision must be independent of the scenario. Therefore, any policy can be described by a set of allocations $\vec{y} = (y_1, y_2, \ldots, y_n)$, such that if agent i has non-zero demand, then they receive an expected allocation y_i. Furthermore, when making the allocation decision for agent n, there is only one possible history: every other agent also had non-zero demand. Thus, any feasible sequential allocation policy must respect the constraint $\sum_{i \in [n]} y_i \leq 1$.

Let us define r_σ as the expected minimum FR of the given policy in scenario σ (i.e., if only the first σ agents have non-zero demand). By convention, we set $r_0 = 1$ and we must have $r_\sigma \leq r_{\sigma-1}$ by definition. In addition, r_σ must be less than the expected FR of agent σ, so $r_\sigma \leq \frac{(n+1)}{2\mu} y_\sigma$. Given $\vec{r}$, the expected minimum FR in this instance is equal to $\frac{1}{n} \sum_{\sigma \in [n]} r_\sigma$. Based on these constraints and the objective, we can formulate a linear program whose optimal solution is an upper bound on the expected minimum FR achievable by any feasible sequential allocation policy. This linear program was originally presented in Section 3.1, but we replicate it here (PRIMAL-LP1) along with its dual program (DUAL-LP1).

(PRIMAL-LP1)

$$\max_{\vec{y}, \vec{r} \in \mathbb{R}^n_{\geq 0}} \quad \frac{1}{n} \sum_{\sigma \in [n]} r_\sigma \qquad \text{s.t.}$$

$$r_\sigma \leq \frac{(n+1)y_\sigma}{2\mu} \qquad \sigma \in [n]$$

$$r_\sigma \leq r_{\sigma-1}, \quad r_0 = 1 \qquad \sigma \in [n]$$

$$\sum_{i \in [n]} y_i \leq 1$$

(DUAL-LP1)

$$\min_{\substack{\vec{\gamma}, \vec{\delta} \in \mathbb{R}^n_{\geq 0} \\ \omega \in \mathbb{R}_{\geq 0}}} \quad \omega + \gamma_1 \qquad \text{s.t.}$$

$$\omega \geq \frac{(n+1)}{2\mu} \delta_i \qquad i \in [n]$$

$$\delta_i \geq \frac{1}{n} + \gamma_{i+1} - \gamma_i \qquad i \in [n-1]$$

$$\delta_n \geq \frac{1}{n} - \gamma_n$$

To upper-bound the value of the program PRIMAL-LP1, we find a feasible assignment for the dual program DUAL-LP1. Consider the assignment where $\delta_i = \frac{1}{n}$ and $\gamma_i = 0$ for all $i \in [n]$, and where $\omega = \frac{n+1}{2n\mu}$. Under this
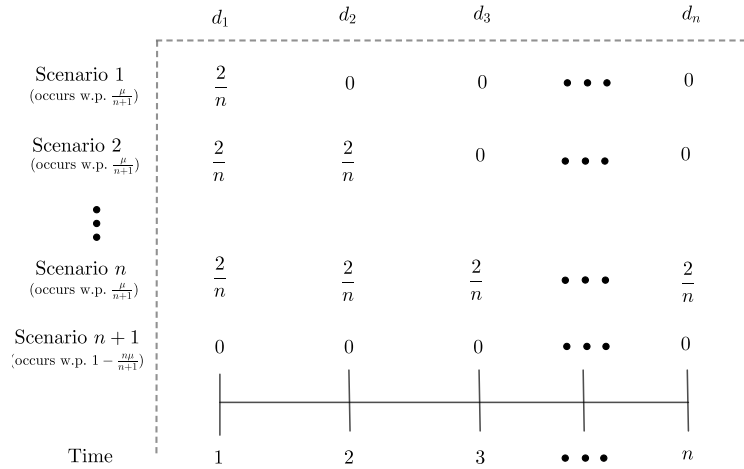
Figure 7   The instance which establishes an upper bound of $\kappa_{\mathsf{p}}(\mu, n)$ when $\mu < 1 + \frac{1}{n}$ (which we refer to as the under-demanded regime).

assignment, all dual variables are non-negative and all constraints are satisfied (in fact, all are tight). Thus, this assignment is feasible in DUAL-LP1. It also attains an objective value of $\frac{n+1}{2n\mu}$. By weak duality, this represents an upper bound on the optimal value of PRIMAL-LP1, and hence an upper bound on the expected minimum FR of any policy in the over-demanded regime.

Under-demanded regime $(\mu < 1 + \frac{1}{n})$: Consider an instance with $n + 1$ scenarios, where the first n scenarios each occur with equal probability of $\frac{\mu}{n+1}$ and scenario $n + 1$ occurs with probability $1 - \frac{n\mu}{n+1}$. In scenario $n + 1$, there is no demand. In scenario σ for $\sigma \in [n]$, all agents $i \in [\sigma]$ have demand $d_i = \frac{2}{n}$. This instance is depicted in Figure 7. We remark the total expected demand is equal to μ, simply because

$$\sum_{\sigma \in [n]} \frac{\mu}{n+1} \cdot \frac{2}{n} \cdot \sigma = \mu.$$

As was the case in the over-demanded regime, any sequential allocation policy can be described by a set of allocation decisions such that if agent i has non-zero demand, then they receive an expected allocation y_i. We again define r_σ as the expected minimum FR in scenario σ, and we note that $r_{n+1} = 1$. Thus, the expected minimum FR in this instance is equal to $\frac{\mu}{n+1} \sum_{\sigma \in [n]} r_\sigma + 1 - \frac{n\mu}{n+1}$. By imposing the constraints described for PRIMAL-LP1, we can formulate a slightly different linear program whose optimal solution is an upper bound on the expected minimum FR of any feasible sequential allocation policy in the under-demanded regime. This linear program (PRIMAL-LP2), along with its dual program (DUAL-LP2), is presented below.

(Primal-LP2)

$$\max_{\vec{y}, \vec{r} \in \mathbb{R}^n_{\geq 0}} \quad \frac{\mu}{n+1} \sum_{\sigma \in [n]} r_\sigma + 1 - \frac{n\mu}{n+1} \qquad \text{s.t.}$$

$$r_\sigma \leq \frac{ny_\sigma}{2} \qquad\qquad \sigma \in [n]$$

$$r_\sigma \leq r_{\sigma-1}, \quad r_0 = 1 \qquad \sigma \in [n]$$

$$\sum_{i \in [n]} y_i \leq 1$$

(Dual-LP2)

$$\min_{\substack{\vec{\gamma}, \vec{\delta} \in \mathbb{R}^n_{\geq 0} \\ \omega \in \mathbb{R}_{\geq 0}}} \quad \omega + \gamma_1 + 1 - \frac{n\mu}{n+1} \qquad \text{s.t.}$$

$$\omega \geq \frac{n}{2}\delta_i \qquad\qquad i \in [n]$$

$$\delta_i \geq \frac{\mu}{n+1} + \gamma_{i+1} - \gamma_i \qquad i \in [n-1]$$

$$\delta_n \geq \frac{\mu}{n+1} - \gamma_n$$

To upper-bound the value of the program Primal-LP2, we find a feasible assignment for its dual. Consider $\delta_i = \frac{\mu}{n+1}$ and $\gamma_i = 0$ for all $\sigma \in [n]$, and $\omega = \frac{n\mu}{2(n+1)}$. Under this dual assignment, all dual variables are non-negative and all constraints are satisfied (and again tight). Thus, this assignment is feasible in Dual-LP2. It also attains an objective value of $\frac{n\mu}{2(n+1)}1 - \frac{n\mu}{n+1} = 1 - \frac{n\mu}{2(n+1)}$. By weak duality, this represents an upper bound on the optimal value of Primal-LP2, and hence an upper bound on the expected minimum FR attainable by any sequential allocation policy when $\mu < 1 + \frac{1}{n}$.

We conclude the proof by scaling the obtained upper bounds on the expected minimum FR by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$. This establishes an upper bound of $\kappa_{\text{P}}(\mu, n)$ on the ex-post fairness guarantee (see Definition 1) achievable by any sequential allocation policy. $\qquad\square$

A.2. Example Illustrating Limitations of Dynamic Programming (Section 3.1)

To highlight some of the shortcomings of a DP approach, we will use the following example, which is in fact a perturbed version of a particular instance of the class of hard examples illustrated in Figure 2(a).

EXAMPLE 1. Consider an instance with two agents ($n = 2$) where expected demand is almost twice the amount of supply ($\mu = 2 + \epsilon$ for small $\epsilon > 0$). The demand sequence is either $(\frac{4}{3} + \epsilon, \frac{4}{3})$ or $(\frac{4}{3} + \epsilon, 0)$ with equal probabilities, that is, the first agent has deterministic demand $d_1 = \frac{4}{3} + \epsilon$ and the second agent either has no demand or has demand $d_2 = \frac{4}{3}$.

Given the objective of maximizing the expected minimum FR in such a setting, it is not hard to see that the optimal online policy—which is the solution of the ex-post DP—simply allocates $\frac{4+3\epsilon}{8+3\epsilon}$ of the supply to the first agent and the remaining supply to the second agent, assuming that agent's demand is realized. This achieves an expected minimum FR of $\frac{3}{8+3\epsilon}$. Note that the minimum expected FR is no higher than the expected minimum FR, as the FR of first agent is always the minimum FR when following this DP solution. We now make the following two observations:

(i) *The DP solution achieves a sub-optimal minimum expected FR.* As mentioned earlier, the minimum expected FR of the DP solution is $\frac{3}{8+3\epsilon}$. Since $\mu = 2 + \epsilon$, this corresponds to ex-ante (and ex-post) fairness of $\frac{6+3\epsilon}{8+3\epsilon}$. We highlight that the ex-ante fairness of the DP solution is less than the ex-ante fairness (and the ex-ante fairness guarantee) of our PPA policy on this instance (for small enough ϵ) by a constant multiplicative factor of $\approx \frac{3}{4}$. Note that the PPA policy targets an FR of $\approx \frac{1}{4/3+2/3} = \frac{1}{2}$ for the first agent, resulting in the second agent having a higher expected FR of $\approx \frac{1/3}{4/3} \times \frac{1}{2} + 1 \times \frac{1}{2} = \frac{5}{8}$. So,

the PPA policy achieves a minimum expected FR of $\approx \frac{1}{2}$ and an ex-ante fairness of ≈ 1, which for small ϵ matches the ex-ante fairness guarantee of the PPA policy established in Theorem 4 when $\mu = 2 + \epsilon$.

(ii) *The DP solution is sensitive to small perturbations in agents' demand distributions.* To see this, we show that the above DP solution can vary significantly if the demand distribution is slightly different. Suppose that we perturb this instance by decreasing the first agent's demand by 2ϵ. In this case, the DP solution is to allocate all of the supply to the first agent. The expected minimum FR is essentially unchanged, but the first agent receives almost twice as much supply as before. Not only does this demonstrate that the DP solution is highly sensitive, but it also highlights that the DP solution can suffer from a lack of interpretability: the first agent receives more supply in this case, even though the only change was a deterministic and negligible decrease in that agent's demand.

A.3. Simple Backward Dynamic Programming for Optimum Offline (Section 3.2)

Suppose the planner has access to all the demand realizations $\vec{d}$, and let f_i be the minimum FR of the policy at the end of period $i - 1$, i.e., $f_1 = 1$, $f_i = \min\{f_{i-1}, \frac{x_{i-1}}{d_{i-1}}\}$ for $i \in [2 : n+1]$. We will show via backward induction that for any remaining supply s_i and minimum FR f_i, the maximum-achievable minimum FR is $\min\{f_i, \frac{s_i}{\sum_{j \in [i:n]} d_j}\}$, which is achieved by a policy $x_i^* = \min\left\{d_i, s_i \frac{d_i}{\sum_{j \in [i:n]} d_j}\right\}$.

Clearly, this is true when $i = n$, as the optimal policy is to allocate as much supply as possible, i.e. $x_n^* = \min\{d_n, s_n\}$. This policy achieves a minimum FR of

$$f_{n+1} = \min\left\{f_n, \frac{x_n}{d_n}\right\} = \min\left\{f_n, \frac{\min\{d_n, s_n\}}{d_n}\right\} = \min\left\{f_n, \frac{s_n}{d_n}\right\}.$$

We now assume this is true for $i > k$. In that case, given an allocation to agent k of x_k, the minimum FR at the end of period k is given by $f_{k+1} = \min\left\{f_k, \frac{xk}{d_k}\right\}$ and the remaining supply is $s_{k+1} = s_k - x_k$. Based on our inductive hypothesis, the maximum-achievable minimum FR is thus

$$\min\left\{f_{k+1}, \frac{s_{k+1}}{\sum_{j \in [k+1:n]} d_j}\right\} = \min\left\{f_k, \frac{xk}{d_k}, \frac{s_k - x_k}{\sum_{j \in [k+1:n]} d_j}\right\}.$$

This is maximized when the second and third terms are equal, which occurs when $x_k = s_k \frac{d_k}{\sum_{j \in [k:n]} d_j}$. If this allocation is infeasible $\left(\text{i.e., if } s_k \frac{d_k}{\sum_{j \in [k:n]} d_j} \geq d_k\right)$, then an allocation of $x_k = d_k$ is optimal because this allocation ensures that f_k must be the minimum of the three terms. This completes the proof by backward induction that the maximum-achievable minimum FR is $\min\left\{f_i, \frac{s_i}{\sum_{j \in [i:n]} d_j}\right\}$, which is achieved by a policy $x_i^* = \min\left\{d_i, s_i \frac{d_i}{\sum_{j \in [i:n]} d_j}\right\}$.

A.4. Analysis of Non-adaptive Fixed-allocation Policies (Section 3.4)

In this appendix section, we consider another class of non-adaptive policies which we call fixed-allocation policies. A fixed-allocation policy is one which pre-determines an allocation x_i for each agent $i \in [n]$.[30] The optimal fixed-allocation policy is the policy which, given a joint demand distribution $\vec{\mathcal{F}}$, pre-determines a vector of allocations $\vec{x} = (x_1, x_2, \ldots, x_n)$ which maximizes $\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\min_{i \in [n]}\left\{\frac{x_i}{d_i}\right\}\right]$.

---

[30] As explained in Footnote 18, we focus on deterministic policies without loss of generality.

PROPOSITION 1 **(Ex-post Fairness Guarantee of the Optimal Fixed-allocation Policy)**. *Given a fixed number of agents $n \in \mathbb{N}$ and supply scarcity $\mu \in \mathbb{R}_{\geq 0}$, the optimal fixed-allocation policy achieves an ex-post fairness guarantee of*

$$\kappa_{\text{FA}}(\mu, n) = \max\{1, \mu\} \begin{cases} 1 - \frac{n\mu}{4}, & \mu n \in [0, 2) \\ \frac{1}{n\mu}, & \mu n \in [2, +\infty) \end{cases}. \tag{19}$$

We remark that the ex-post fairness guarantee $\kappa_{\text{FA}}(\mu, n)$ tends to 0 as the number of agents n gets large. This is in stark contrast to the guarantees provided by the PPA policy and the optimal TFR policy, which are lower-bounded by a constant regardless of the number of agents.

A.4.1. Proof of Proposition 1 We prove this proposition by first showing that there exists a distribution where no fixed-allocation policy achieves ex-post fairness greater than $\kappa_{\text{FA}}(\mu, n)$, which thus serves as an upper bound on the ex-post fairness guarantee of the optimal fixed-allocation policy. We then show that there exists a fixed-allocation policy that achieves ex-post fairness of at least $\kappa_{\text{FA}}(\mu, n)$ for any demand distribution, which means the bound $\kappa_{\text{FA}}(\mu, n)$ is tight.

Upper bound: We prove the hardness result by considering two separate cases corresponding to $n\mu < 2$ and $n\mu \geq 2$. For each case, we provide an instance of the problem under which any fixed-allocation policy obtains ex-post fairness no larger than $\kappa_{\text{FA}}(\mu, n)$.

(i) If $n\mu < 2$, consider a joint demand distribution such that with probability $1 - \frac{n\mu}{2}$ there is no demand, and with probability $\frac{n\mu}{2}$ one agent chosen uniformly at random has demand $\frac{2}{n}$. In this case, with probability $1 - \frac{n\mu}{2}$, the minimum FR is 1, and with probability $\frac{n\mu}{2}$, the minimum FR is equal to the allocation of a randomly selected agent (which is at most $1/n$) divided by the total demand $2/n$. Therefore, the minimum expected FR for this instance is upper-bounded by $1 - \frac{n\mu}{4}$.

(ii) If $n\mu \geq 2$, consider a joint demand distribution where one agent chosen uniformly at random has demand equal to the expected total demand μ. In this instance, the minimum expected FR is upper-bounded by the allocation of a randomly selected agent (which is at most $1/n$) divided by the total demand μ.

Taken together, these instances provide an upper bound on the expected minimum FR one can hope to achieve with a fixed-allocation policy. We then scale each instance by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$, which provides an upper bound of $\kappa_{\text{FA}}(\mu, n)$ on the ex-post fairness guarantee (see Definition 1) of any fixed-allocation policy.

Lower bound: Consider a policy which allocates an equal amount of supply to each agent, i.e., $x_i = \frac{1}{n}$ for all $i \in [n]$. In that case, the minimum FR is lower-bounded by $\min\left\{1, \frac{1}{n\sum_{i \in [n]} d_i}\right\}$. Further,

$$\min\left\{1, \frac{1}{n\sum_{i \in [n]} d_i}\right\} \geq \begin{cases} 1 - \frac{n\sum_{i \in [n]} d_i}{4}, & n\sum_{i \in [n]} d_i \in [0, 2) \\ \frac{1}{n\sum_{i \in [n]} d_i}, & n\sum_{i \in [n]} d_i \in [2, +\infty) \end{cases}$$

We note that the right hand side of the above inequality is convex in $\sum_{i \in [n]} d_i$. Therefore, using Jensen's inequality, the expected minimum FR must be at least

$$\begin{cases} 1 - \frac{n\mu}{4}, & n\mu \in [0, 2) \\ \frac{1}{n\mu}, & n\mu \in [2, +\infty) \end{cases}$$

We then scale each this lower bound on the expected minimum FR by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$. This provides a lower bound on the ex-post fairness guarantee (see Definition 1) that is equal to $\kappa_{\text{FA}}(\mu, n)$. Thus, we have shown that $\kappa_{\text{FA}}(\mu, n)$ is a tight bound on the ex-post fairness guarantee of the optimal fixed-allocation policy. $\square$

A.5.   Proof of Claims in Lemma 2 (Section 3.4)

In this section, we present the proofs of the three claims which appear in the proof of Lemma 2.

A.5.1.   Proof of Claim 1   Suppose we have a distribution G such that $R^G(q)$ is not flat in $(Q^G(1), 1]$, that is, $\exists\; q_1, q_2 \in (Q^G(1), 1]$ such that $R^G(q_1) \neq R^G(q_2)$. Let $R^G(q)$ attain its maximum in $[Q^G(1), 1]$ at the quantile $\bar{q}$. Now consider a quantile $q' \in (Q^G(1), 1]$ so that $R^G(q') < R^G(\bar{q})$. Let $\delta \triangleq q' - Q^G\left(T^G(q') + \epsilon\right)$ be the total probability mass in the TFR interval $(T^G(q'), T^G(q') + \epsilon]$. Then pick a small enough $\epsilon > 0$ such that:

   (i) $\epsilon + \delta + \epsilon \cdot \delta < R^G(\bar{q}) - R^G(q')$,

   (ii) $\epsilon \leq 1 - T^G(q')$,

   (iii) $\bar{q} \notin [Q^G\left(T^G(q') + \epsilon\right), q']$.

Now consider a distribution $\bar{G}$ that is generated from G by moving all the δ probability mass in $(T^G(q'), T^G(q') + \epsilon]$ to the point $T^G(q') + \epsilon$. With this modification in the distribution, the EAFR of every TFR $\tau \in [0, T^G(q')] \cup (T^G(q') + \epsilon, 1]$ remains the same. Moreover, the maximum EAFR in the interval $(T^G(q'), T^G(q') + \epsilon]$ is also achieved at $T^G(q') + \epsilon$. Given this target fill rate, the EAFR of the distribution $\bar{G}$ is equal to

$$R^{\bar{G}}(T^G(q') + \epsilon) = (T^G(q') + \epsilon)\left(Q^G\left(T^G(q') + \epsilon\right) + \delta\right)$$
$$< T^G(q') \cdot Q^G\left(T^G(q') + \epsilon\right) + \epsilon + \delta + \epsilon \cdot \delta$$
$$< R^G(q') + (R^G(\bar{q}) - R^G(q')) = R^G(\bar{q}) \ .$$

Therefore, the maximum EAFR over all possible TFRs in $[0, 1]$ is the same for G and $\bar{G}$, i.e.,

$$\max_{q \in [0,1]: T^G(q) \in [0,1]} R^G(q) = \max_{q \in [0,1]: T^{\bar{G}}(q) \in [0,1]} R^{\bar{G}}(q) = R^G(\bar{q}).$$

However, $\bar{\mu} = \mathbf{E}_{v \sim \bar{G}}\left[\frac{1}{v}\right] < \mathbf{E}_{v \sim G}\left[\frac{1}{v}\right] = \mu$. Now let $\tilde{G}$ be the distribution of the random variable $(\bar{\mu}/\mu) \cdot v$ where $v \sim \bar{G}$. We have:

$$\max_{\tau \in [0,1]} \tau(1 - \tilde{G}(\tau)) = \max_{\tau \in [0,1]} \tau\left(1 - \bar{G}(\tau \cdot \mu / \bar{\mu})\right) \leq \max_{\tau \in [0,1]} \tau\left(1 - \bar{G}(\tau)\right) = \max_{\tau \in [0,1]} \tau\left(1 - G(\tau)\right) \ .$$

Also, $\mathbf{E}_{v \sim \tilde{G}}\left[\frac{1}{v}\right] = \mathbf{E}_{v \sim G}\left[\frac{1}{v}\right] = \mu$. Therefore, dropping such a distribution G from the feasible set in the outer optimization of eq. (15) does not change the infimum value, which proves Claim 1.   □

A.5.2.   Proof of Claim 2   Suppose we have a distribution G that has non-zero total mass in the interval $(1, +\infty)$. Now shift all the probability mass in $(1, +\infty)$ to $+\infty$. Let $\bar{G}$ be the resulting distribution. Note that the maximum EAFR among targets in $[0, 1]$ is the same for G and $\bar{G}$, as the EAFR for any target in $[0, 1]$ remains the same. However, $\bar{\mu} = \mathbf{E}_{v \sim \bar{G}}\left[\frac{1}{v}\right] < \mathbf{E}_{v \sim G}\left[\frac{1}{v}\right] = \mu$, as we have moved the probability mass of v towards larger values (equivalently, the probability mass of demand to lower values). By using the same trick as above in Section A.5.1, we conclude that dropping such a distribution G from the feasible set in the outer optimization of eq. (15) does not change the infimum value, which proves Claim 2.   □

A.5.3.   Proof of Claim 3   Consider an inverse demand distribution $\bar{G}$ that satisfies the two constraints given by Claim 1 and Claim 2, namely (i) $R^G(q) = R^G(q'), \forall q, q' \in [Q^G(1), 1]$, and (ii) $G(v) = G(1)$ for all $v \in [1, +\infty)$. Let us define $\bar{q}$ such that $Q^{\bar{G}}(1) = \bar{q}$, or equivalently, $T^{\bar{G}}(\bar{q}) = 1$.

The EAFR curve for $\bar{G}$ by definition attains a value of $R^{\bar{G}}(\bar{q}) = \bar{q}$. Since $\bar{G}$ has a constant EAFR curve in the interval $[\bar{q}, 1]$, its EAFR curve must be constantly equal to $\bar{q}$ over that interval. Futher, since the EAFR curve can also be expressed as $q \cdot \bar{G}^{-1}(1-q)$, we must have $\bar{G}(\bar{q}/q) = 1 - q$ for all $q \in [\bar{q}, 1]$. Equivalently, using a change of variable $v = \bar{q}/q$, we must have $\bar{G}(v) = 1 - \bar{q}/v$ for all $v \in [\bar{q}, 1]$ (which implies $\bar{G}(\bar{q}) = 0$).

Further, since the CDF $\bar{G}$ pushes all the probability mass in the interval $(1, +\infty)$ to $+\infty$, $\bar{G}$ must be constant in the interval $[1, +\infty)$. Thus, we have uniquely described $\bar{G}$, up to a constant $\bar{q}$:

$$\bar{G}(v) = \begin{cases} 0 & \text{if } v \in [0, \bar{q}) \\ 1 - \bar{q}/v & \text{if } v \in [\bar{q}, 1) \\ 1 - \bar{q} & \text{if } v \in [1, +\infty) \\ 1 & \text{if } v = +\infty \end{cases}$$

For this distribution to have an expected demand of μ, consider the corresponding CDF for demand $\bar{F} : \mathbb{R}_{\geq 0} \to [0, 1]$ for the random variable $x \triangleq \frac{1}{v}$ where $v \sim \bar{G}$. We have:

$$\bar{F}(x) = 1 - \bar{G}(1/x) = \begin{cases} 0 & \text{if } x = 0 \\ \bar{q} & \text{if } x \in (0, 1] \\ \bar{q}x & \text{if } x \in (1, \frac{1}{\bar{q}}] \\ 1 & \text{if } x \in (\frac{1}{\bar{q}}, +\infty) \end{cases}$$

As a result,

$$\mathbf{E}_{v \sim G}\left[\frac{1}{v}\right] = \int_0^\infty \left(1 - \bar{F}(x)\right) dx = 1 - \bar{q} + \int_1^{\frac{1}{\bar{q}}} (1 - \bar{q}x) dx = \frac{1}{2}\left(\frac{1}{\bar{q}} - \bar{q}\right).$$

The unique solution to $\frac{1}{2}\left(\frac{1}{\bar{q}} - \bar{q}\right) = \mu$ satisfying $\bar{q} \geq 0$ is $\bar{q} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}$.

We highlight that $\bar{G}$ with $\bar{q} = \frac{1}{\mu + \sqrt{\mu^2 + 1}}$ is identical to the distribution $\hat{G}$ defined in eq. (14). This shows that $\hat{G}$ is the unique worst-case distribution. $\qquad\square$

A.6.   Proof of Theorem 4 (Section 3.5)

We prove this theorem by first providing two instances which together show that no policy achieves ex-ante fairness greater than $\kappa_\text{A}(\mu, n)$. We then show that our PPA policy achieves ex-ante fairness of at least $\kappa_\text{A}(\mu, n)$ for any demand distribution, which means the bound is tight.

Upper bound: We prove the hardness result by considering two separate cases corresponding to $\mu < 2$ and $\mu \geq 2$. For each case, we provide an instance of the problem under which no policy can obtain ex-ante fairness larger than $\kappa_\text{A}(\mu, n)$.

(i) If $\mu < 2$, consider a joint demand distribution for an arbitrary number of agents such that with probability $1 - \frac{\mu}{2}$ there is no demand, and with probability $\frac{\mu}{2}$ the total demand is 2 (arbitrarily and deterministically split among agents). In this case, with probability $1 - \frac{\mu}{2}$, each agent achieves an FR of 1, and with probability $\frac{\mu}{2}$, the expected FR cannot exceed $\frac{1}{2}$ for at least one agent. Therefore, in this instance the minimum expected FR is upper-bounded by $1 - \frac{\mu}{4}$.

(ii) If $\mu \geq 2$, consider a deterministic demand distribution where total demand is equal to its expectation μ (arbitrarily split among n agents). In this case, the minimum expected FR is clearly upper-bounded by $\frac{1}{\mu}$.

Taken together, these instances provide an upper bound on the minimum expected FR one can hope to achieve with any policy. We then scale each instance by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$, which provides an upper bound of $\kappa_{\mathrm{A}}(\mu, n)$ on the ex-ante fairness guarantee (see Definition 1) of any policy. $\qquad\square$

Lower bound: We now show that the PPA policy achieves an ex-ante fairness guarantee of $\kappa_{\mathrm{A}}(\mu, n)$. First, we show via induction that after the arrival of any agent $i \in [n]$, the ratio $\frac{d_i + \mu_{i+1}}{s_i}$, i.e., expected total remaining demand to remaining supply, is in expectation at most μ. We then place a lower bound on the expected FR of agent i when following the PPA policy which is a decreasing and convex function of the ratio $\frac{d_i + \mu_{i+1}}{s_i}$. We conclude by applying Jensen's inequality to lower bound the expected FR of agent i. Since this lower bound is the same for each agent, it constitutes a lower bound on the minimum expected FR.

CLAIM 4 (**Upper Bound on Demand-to-Supply Ratio**). *When following the PPA policy, for all $i \in [n]$, $\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{d_i + \mu_{i+1}}{s_i}\right] \leq \mu$.*

Proof: We proceed by induction. Clearly, when $i = 1$, $\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{d_1 + \mu_2}{s_1}\right] = \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{\mu_1}{s_1}\right] = \mu$. We now assume that this holds for $i = k$ and attempt to prove the claim for $i = k + 1$. According to the PPA policy, $x_k \leq s_k \frac{d_k}{d_k + \mu_{k+1}}$. Thus, s_{k+1} is at least $s_k \frac{\mu_{k+1}}{d_k + \mu_{k+1}}$. Consequently,

$$\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{d_{k+1} + \mu_{k+2}}{s_{k+1}}\right] = \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{\mu_{k+1}}{s_{k+1}}\right] \geq \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{\mu_{k+1}}{s_k \frac{\mu_{k+1}}{d_k + \mu k + 1}}\right] = \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{d_k + \mu_{k+1}}{s_k}\right] \geq \mu.$$

The final inequality comes from our inductive hypothesis, which completes the proof by induction. $\qquad\square$

Given a current demand d_i and expected future demand μ_{i+1}, the FR of agent i is $\min\left\{1, \frac{s_i}{d_i + \mu_{i+1}}\right\}$. It is straightforward to show that this FR is lower-bounded by the following function of the ratio $\frac{d_i + \mu_{i+1}}{s_i}$:

$$h\left(\frac{d_i + \mu_{i+1}}{s_i}\right) = \begin{cases} 1 - \frac{1}{4}\frac{d_i + \mu_{i+1}}{s_i}, & \frac{d_i + \mu_{i+1}}{s_i} < 2 \\ \frac{s_i}{d_i + \mu_{i+1}}, & \frac{d_i + \mu_{i+1}}{s_i} \geq 2 \end{cases}$$

We remark that h is a decreasing and convex function of its argument. Hence, by Jensen's inequality and Claim 4,

$$\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[h\left(\frac{d_i + \mu_{i+1}}{s_i}\right)\right] \quad \geq \quad h\left(\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\frac{d_i + \mu_{i+1}}{s_i}\right]\right) \quad \geq \quad h(\mu).$$

Since this lower bound on the expected FR holds for each agent i, we have shown a lower bound on the minimum expected FR when following the PPA policy. When scaled by our benchmark for deterministic demand, namely $\overline{W} = \min\{1, 1/\mu\}$, this lower bound exactly matches the upper bound of $\kappa_{\mathrm{A}}(\mu, n)$ established above, and thus completes the proof of Theorem 4. $\qquad\square$

B. Missing Proofs of Section 5.2

B.1. Proof of Corollary 1

Suppose the allocation given by the PPA policy is $\vec{x}$. By Theorem 2, the PPA policy achieves ex-post fairness of $\kappa_{\mathrm{P}}(\mu, n)$ when the social welfare function is the minimum FR, which implies $U_{+\infty}(\vec{x}) = \min_{i \in [n]}\{\frac{x_i}{d_i}\} \geq \kappa_{\mathrm{P}}(\mu, n)\overline{W}$.

Now consider a new allocation vector $\vec{x}'$ such that for all $i \in [n]$, $\frac{x_i'}{d_i} = \min_{j \in [n]}\{\frac{x_j}{d_j}\}$. We remark that $x_i \geq x_i'$ for all $i \in [n]$. Further, it is easy to verify that $U_\alpha(\vec{x}') = U_{+\infty}(\vec{x})$ for any $\alpha \in [0, +\infty)$. Since $U_\alpha(\vec{x})$ is non-decreasing in each x_i, $U_\alpha(\vec{x}) \geq U_{+\infty}(\vec{x})$.

Therefore, the expected WPM social welfare when following the PPA policy is weakly greater than the minimum FR (regardless of the fairness parameter α). Scaling by the achievable social welfare when demand is deterministic, we have shown that ex-post fairness, i.e., $\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}[U_\alpha(\vec{x})]/\overline{W}$, must be at least $\kappa_{\mathrm{P}}(\mu, n)$. $\square$

B.2. Proof of Corollary 2

Theorem 2 establishes that the PPA policy guarantees an expected minimum FR for resource j of at least $\kappa_{\mathrm{P}}\left(\frac{\mu^j}{s^j}, n\right) \max\left\{1, \frac{\mu^j}{s^j}\right\}$. Consequently, we can place a lower bound on the expected minimum weighted FR:

$$\mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\min_{i \in [n]} \sum_{j \in [m]} \lambda^j \frac{x_i^j}{d_i^j}\right] \geq \sum_{j \in [m]} \lambda^j \mathbf{E}_{\vec{d} \sim \vec{\mathcal{F}}}\left[\min_{i \in [n]} \frac{x_i^j}{d_i^j}\right] \geq \sum_{j \in [m]} \lambda^j \kappa_{\mathrm{P}}\left(\frac{\mu^j}{s^j}, n\right) \max\left\{1, \frac{\mu^j}{s^j}\right\}. \qquad \square$$

B.3. Proof of Corollary 3

Consider a correlated distribution for demands—correlated across agents and resource types—where the marginal distribution of demands for each resource type matches the worst-case joint distribution of the single-type problem described in the proof of Theorem 1. These marginal distributions are then coupled such that the last-arriving agent with non-zero demand for resource j is the same as the last-arriving agent with non-zero demand for resource j', for any resources $j, j' \in [m]$ for which at least one agent has non-zero demand. Given the marginal distributions, each agent is equally likely to be this last-arriving agent. For any sample path drawn from this distribution, it is without loss of generality to only consider policies where the allocation is decreasing (i.e., where this last-arriving agent has the worst FR) for every resource.[31]

If supply of resource j is s^j, Theorem 1 establishes an upper-bound of $\kappa_{\mathrm{P}}\left(\frac{\mu^j}{s^j}, n\right) \max\left\{1, \frac{\mu^j}{s^j}\right\}$ on the expected minimum FR of any policy in the single-type problem corresponding to resource j. Since for every sample path the agent with the worst FR (i.e., the last-arriving agent) is the same across resources, aggregating these bounds establishes an upper-bound of $\sum_{j \in [m]} \lambda^j \kappa_{\mathrm{P}}\left(\frac{\mu^j}{s^j}, n\right) \max\left\{1, \frac{\mu^j}{s^j}\right\}$ on the expected minimum weighted FR. $\square$

---

[31] For intuition as to why this is without loss of generality, note (i) each agent with non-zero demand for resource $j \in [m]$ has identical demand for that resource, (ii) each agent linearly aggregates their FRs by the same weights $\{\lambda^j\}_{j \in [m]}$, and (iii) each agent is equally likely to be the last-arriving agent. Consequently, if agent i has a strictly larger allocation than agent i' (where $i' < i$) for any resource j, a policy which switches their allocations for that resource will have a weakly greater expected minimum weighted FR.