Fairness and Utilization in Allocating Resources with Uncertain Demand

Kate Donahue

kdonahue@cs.cornell.edu

Department of Computer Science, Cornell University

Jon Kleinberg

Departments of Computer Science and Information

Science, Cornell University

kleinber@cs.cornell.edu

ABSTRACT

Resource allocation problems are a fundamental domain in which to evaluate the fairness properties of algorithms. The trade-offs between fairness and utilization have a long history in this domain. A recent line of work has considered fairness questions for resource allocation when the demands for the resource are distributed across multiple groups and drawn from probability distributions. In such cases, a natural fairness requirement is that individuals from different groups should have (approximately) equal probabilities of receiving the resource. A largely open question in this area has been to bound the gap between the maximum possible utilization of the resource and the maximum possible utilization subject to this fairness condition.

Here, we obtain some of the first provable upper bounds on this gap. We obtain an upper bound for arbitrary distributions, as well as much stronger upper bounds for specific families of distributions that are typically used to model levels of demand. In particular, we find — somewhat surprisingly — that there are natural families of distributions (including Exponential and Weibull) for which the gap is non-existent: it is possible to simultaneously achieve maximum utilization and the given notion of fairness. Finally, we show that for power-law distributions, there is a non-trivial gap between the solutions, but this gap can be bounded by a constant factor independent of the parameters of the distribution.

CCS CONCEPTS

• **Theory of computation** → *Design and analysis of algorithms*; • **Computing methodologies** → **Planning under uncertainty**.

KEYWORDS

resource allocation, algorithmic fairness, uncertainty, Weibull distribution, power law distribution

ACM Reference Format:

Kate Donahue and Jon Kleinberg. 2020. Fairness and Utilization in Allocating Resources with Uncertain Demand. In *Conference on Fairness, Accountability, and Transparency (FAT* '20)*, January 27–30, 2020, Barcelona, Spain. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3351095.3372847>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAT* '20, January 27–30, 2020, Barcelona, Spain

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6936-7/20/01...\$15.00

<https://doi.org/10.1145/3351095.3372847>

1 INTRODUCTION

The current interest in fairness properties of algorithms includes several distinct themes, one of which is the question of fairness in allocating scarce resources. Research on this question has a long history, with foundational work in the 1980s and 1990s on allocating resources in computer systems [2, 3, 6, 12, 13]. Allocation problems continue to form an important topic for fairness considerations, especially as automated systems make allocation decisions in a wide range of areas that reach far beyond the original computational settings of the problem [10, 14].

Recently, Elzayn et al. (2019) considered a novel allocation problem in this style. In their formulation, individuals are divided into groups, each of which has some probability distribution of *candidates* who desire the resource. There is a fixed amount of resource R and the question is how to allocate this resource across the groups. Such a set-up arises in many applications; as one concrete motivating example for purposes of discussion, suppose that we have a set number of doctors, each with a maximum number of people they can assist, and we would like to allocate them across a set of N geographically distributed regions. Each region contains a population with a potentially different probability distribution over the number of sick people in need of doctors. We know the distribution of demand for each region, but we will only know the actual demand — drawn from these distributions — once we have allocated the doctors by dividing them up in some fashion across the regions. If a community has too few doctors, some sick people will go unassisted; and if a community has more doctors than it needs for the number of sick people it has, some of the doctor's capacity to help will go unused. Formally, the expected number of candidates that receive a resource under a certain allocation is called the *utilization* of that allocation.

One allocation strategy could be to put the doctors where they are most likely to be used, with the aim of maximizing utilization. However, such an allocation could potentially “starve” certain regions, leaving them with very few doctors even though they have a non-trivial level of need. Such an allocation violates a natural definition of fairness: sick people in some regions would have a higher probability of receiving assistance than sick people in other regions, which conflicts with the premise that people are equally deserving of care from the doctors.

In this work, we thus say that an allocation is *fair* across multiple regions if a candidate who needs the resource has the same probability of receiving it, regardless of the region they belong to: the identity of the region doesn't impact the probability that those in need obtain assistance. At some points, we relax this definition slightly to say that an allocation is α -*fair* if the probability that a candidate receives the resource in different regions is within a

maximum difference of α , for $\alpha \in [0, 1]$ across all of the regions. Note that $\alpha = 0$ corresponds to our initial notion of fairness, and $\alpha = 1$ imposes no constraint at all on the allocation.

Elzayn et al. (2019) consider a version of this allocation problem in the scenario where the candidate probability distributions are unknown and must be learned. At each time step, a certain allocation is chosen and the feedback obtained reveals only the number of candidates who received the resource, not the true number. They adapt learning results from previous dark pool trading scenarios proposed in Ganchev et al. (2009). From these results, they construct algorithms that learn utilization-maximizing and α -fair utilization maximizing allocations through this type of censored feedback¹. They also define the *Price of Fairness* (PoF) as the ratio of the utilizations of the max-utilizing allocation and the max-utilizing α -fair allocation. From an empirical dataset, they calculate the PoF for various levels of α in practice. However, they leave largely open the following category of questions: can we obtain theoretical upper bounds on the price of fairness over different probability distributions? These questions appear to become quite complex even when the distributions are known.

The present work: The interaction of fairness and utilization. In this work we obtain bounds on the price of fairness for this family of resource allocation problems, both through general results that hold for all distributions and through stronger results that are specific to common families of distributions. We also distinguish between two versions of the problem — one in which we require that the allocation to each group be integer-valued, and one in which we allow resources to be allocated fractionally (or probabilistically, which yields equivalent results in expectation).

In the case of integer-valued resource allocation, we show through a constructive proof that the PoF can be unboundedly large. If we allow resources to be allocated fractionally, then we show that the PoF is bounded above by $\frac{1}{\alpha}$ for arbitrary distributions and $\alpha > 0$. In the case where $\alpha = 0$, the PoF can be unboundedly large for fractional allocations as well.

We then show that much stronger upper bounds can be obtained for large families of distributions that are often used to model levels of demand. First we show that certain families of natural distributions have PoF equal to 1, the smallest possible bound: for these distributions, there is no trade-off between fairness and maximum utilization. We show that distributions with this property include Exponential and Weibull distributions. We then consider the family of Power Law distributions for a constant number of groups; we show that these distributions have a PoF that can be strictly greater than 1, but is always bounded above by a universal constant. Table 3 in Appendix A contains a high-level summary of results explored in this paper.

Overall, our results reveal a rich picture in the trade-off between fairness and utilization for this type of allocation problem, through the different families of bounds for the price of fairness — ranging from distributions where perfect fairness can be achieved while maximizing utilization, to those in which the gap is bounded by a constant, to others for which the gap can be large. The techniques used for the analysis suggest further opportunities for reasoning

about how balancing the probability of service to different groups can both constrain and also be compatible with other objectives.

Important limitations. The assumptions and limitations of this model are discussed in greater detail in Section 3. However, it is worth emphasizing some of the specific choices inherent in creating this model. In particular, the notion of fairness that this paper uses is “equality in the chance of obtaining a resource, conditional on needing it”. Although we describe in the next section why this is a natural definition of fairness to use, it is certainly not the only possible definition of fairness, and alternative definitions could lead to different results. For example, a notion of fairness could take into account historical discrepancies in resources between locations and attempt to proactively rectify them.

Another fundamental assumption of this work is that there is a fixed amount of resources available to distribute. An alternative viewpoint could be to explore ways to increase the total amount of resources so that more people in need are able to obtain help. Increasing available resources is a critical goal, and one where advocacy and political action can play a key role in bringing this about for real-world allocation problems. However, in almost all cases it is only feasible to reduce, not fully eliminate, unmet need. Thus, even after the overall level of resources have been increased, fair resource allocation will still be a problem that needs to be addressed. In this way, we view both increasing the total amount of resources and fair resource allocation as important and complementary areas of study; in most cases, neither eliminates the need for the other.

2 MOTIVATING EXAMPLE

To motivate the problem and to make some of the types of calculations more clear, we start with the following example. Suppose there is a remote stretch of coastline with two small villages, *A* and *B*, each with a small number of houses. This stretch of coastline is prone to severe localized storms that lead to power outages. Each village has a slightly different probability distribution over storms - distributions that are known. In any particular week, the probability distribution over the number of houses *C* impacted by power outages in each village is as follows:

$$P_A(C = c) = \begin{cases} 0.6 & c = 0 \\ 0.4 & c = 2 \\ 0 & \text{otherwise} \end{cases} \quad P_B(C = c) = \begin{cases} 0.3 & c = 0 \\ 0.7 & c = 3 \\ 0 & \text{otherwise} \end{cases}$$

In our example, assume that each week’s power outage is independent of other weeks, and the villages are far enough apart that the power outages in one village aren’t correlated with the power outages in the other. A regional planning committee has at its disposal 2 generators, each of which can restore power to 1 house. They are trying to decide how to allocate these generators across the two villages. The generators cannot be transferred from one village to another after a storm strikes. If there are v generators allocated to a village and c houses in need of generators, the number of houses that receive generators in that village is $\min(v, c)$.

We include this example only as a motivating case for our broader model: we do *not* intend that this model apply exclusively to disaster relief allocation. However, we do wish to make it clear that maximizing utilization and fairness in disaster relief resource allocation is an active area of study [15].

¹They consider a related, but different definition of fairness from the one used in this work. Later sections will expand on this difference.

$\{v_A, v_B\}$	$\mathbb{E}_{C \sim C_A} [\min(C, v_A)]$	$\mathbb{E}_{C \sim C_B} [\min(C, v_B)]$	$U(v_A, v_B)$
$\{0, 2\}$	$0.6 \min(0, 0) + 0.4 \min(2, 0) = 0$	$0.3 \min(0, 2) + 0.7 \min(3, 2) = 1.4$	1.4
$\{1, 1\}$	$0.6 \min(0, 1) + 0.4 \min(2, 1) = 0.4$	$0.3 \min(0, 1) + 0.7 \min(3, 1) = 0.7$	1.1
$\{2, 0\}$	$0.6 \min(0, 2) + 0.4 \min(2, 2) = 0.8$	$0.3 \min(0, 0) + 0.7 \min(3, 0) = 0$	0.8

Table 1: Utilization of resources across two groups in motivating example.

The regional planning committee first decides that it wishes to maximize *utilization*: the expected number of houses in need of generators that receive them. There are three options for allocations: $\{(v_A = 0, v_B = 2), (v_A = 1, v_B = 1), (v_A = 2, v_B = 0)\}$ where $v_i = j$ means that village i gets j generators. Formally, utilization can be written as:

$$U(v_A = n_A, v_B = n_B) = \mathbb{E}_{C \sim C_A} [\min(C, n_A)] + \mathbb{E}_{C \sim C_B} [\min(C, n_B)]$$

The utilizations across these three allocations are given in Table 1. The allocation that maximizes utilization is $(v_A = 0, v_B = 2)$.

However, something bothers the regional committee: it feels fundamentally unfair that none of the houses in village A ever receive generators. To formalize this, we could ask what fraction of houses in need of generators obtain them, on average, and aim to select an allocation that brings this proportion between the two villages as close to equality as possible. First, we derive the relevant probability for an arbitrary village:

$$\begin{aligned} & \frac{P(\text{house } k \text{ gets generator} \mid \text{house } k \text{ needs generator})}{P(\text{house } k \text{ gets generator} \cap \text{house } k \text{ needs generator})} \\ &= \frac{P(\text{house } k \text{ gets generator})}{P(\text{house } k \text{ needs generator})} \end{aligned}$$

We assume that there is a total number of houses $0 < K < \infty$ in the village and that all houses within the same village are interchangeable. The total number of houses in the village that need generators is given by $C \sim C$. Given that c houses need generators, there is a $\frac{c}{K}$ probability that a randomly selected house k will be in need. Given that a house is in need, the probability it obtains a generator equals $\min(\frac{c}{K}, 1)$. $P(\text{house } k \text{ needs generator})$ equals

$$\begin{aligned} &= \sum_{c=0}^K P(\text{house } k \text{ needs generator} \mid C = c) \cdot P(C = c) \\ &= \sum_{c=0}^K \frac{c}{K} P(C = c) = \sum_{c=1}^K \frac{c}{K} P(C = c) \end{aligned}$$

Similarly, we can use the K term to expand $P(\text{house } k \text{ gets generator})$:

$$\begin{aligned} &= P(k \text{ gets generator} \mid k \text{ needs generator}) P(k \text{ needs generator}) \\ &\quad + P(k \text{ gets generator} \mid k \text{ does NOT need generator}) \cdot P(k \text{ does NOT need generator}) \end{aligned}$$

$P(\text{house } k \text{ gets generator} \mid \text{house } k \text{ does NOT need generator}) = 0$. We expand the remaining term by conditioning on C . The sum

$\{v_A, v_B\}$	$q(v_A, C_A)$	$q(v_B, C_B)$	$ q(v_A, C_A) - q(v_B, C_B) $	$U(v_A, v_B)$
$\{0, 2\}$	$\frac{0}{0.8} = 0$	$\frac{1.4}{2.1} = 0.\bar{6}$	$0.\bar{6}$	1.4
$\{1, 1\}$	$\frac{0.4}{0.8} = 0.5$	$\frac{0.7}{2.1} = 0.\bar{3}$	$0.1\bar{6}$	1.1
$\{2, 0\}$	$\frac{0.8}{0.8} = 1$	$\frac{0}{2.1} = 0$	1	0.8

Table 2: Probability of receiving resources across two groups in motivating example. Utilization included for reference.

starts from $c = 1$ because if $c = 0$ there are 0 houses in need of generators.

$$\begin{aligned} & \sum_{c=1}^K P(k \text{ gets generator} \mid k \text{ needs generator}, C = c) \cdot P(k \text{ needs generator} \mid C = c) P(C = c) \\ &= \sum_{c=1}^K \min\left(\frac{v}{c}, 1\right) \cdot \frac{c}{K} \cdot P(C = c) \end{aligned}$$

Combining these results gives:

$$\begin{aligned} \frac{P(\text{house } k \text{ gets generator})}{P(\text{house } k \text{ needs generator})} &= \frac{\sum_{c=1}^K \min\left(\frac{v}{c}, 1\right) \cdot \frac{c}{K} \cdot P(C = c)}{\sum_{c=1}^K \frac{c}{K} \cdot P(C = c)} \\ &= \frac{\sum_{c=1}^K \min\left(\frac{v}{c}, 1\right) \cdot c \cdot P(C = c)}{\sum_{c=1}^K c \cdot P(C = c)} \\ &= \frac{\sum_{c=1}^K \min(v, c) P(C = c)}{\sum_{c=1}^K c \cdot P(C = c)} \\ &= \frac{\mathbb{E}_{C \sim C} \min(v, C)}{\mathbb{E}_{C \sim C} C} \\ &=: q(v_i, C_i) \end{aligned}$$

where in the last step we have defined the probability of receiving the resource, conditional on needing it, as $q(v_i, C_i)$, where i indexes the village. Note that this definition differs from Elzayn et al. (2019): that definition calculated the probability of candidates receiving the resource weighted by each time period. By contrast, our definition is based on the per-person probability of receiving the resource, which we believe is a more natural definition.

One goal might be to bring these probabilities across the villages to be as close as possible to each other. How does each allocation considered in Table 1 do under this definition? For reference, $\mathbb{E}_{C \sim C_A} C = 0.6 \cdot 0 + 0.4 \cdot 2 = 0.8$ and $\mathbb{E}_{C \sim C_B} C = 0.3 \cdot 0 + 0.7 \cdot 3 = 2.1$. The results are shown in Table 2.

Given these numbers, the regional committee could look at the utilization and fairness of each allocation and make their own decision about how resources should be allocated. In particular, they might decide that they want the $q(v_i, C_i)$ fractions between each village to be within a certain difference α and pick the allocation with maximal utilization subject to those constraints. It is this scenario that much of this paper will consider: comparing max-utilization with max-utilization subject to an α -fairness constraint.

3 APPLICATIONS AND FORMAL MODEL

Our model is best described as a two-stage stochastic problem in which resource allocation decisions are made in the first stage, and demand is drawn probabilistically in the second stage. We envision

our work being used in the following way. Suppose that the group identity is associated with some decomposition of the population into groups (whether by geographic location, or socioeconomic status, or attributes like race or ethnicity), such that we wish for resources to be fairly distributed across these groups. In this case, the PoF bound could be used to provide a general sense of the maximum trade-off between utilization and fairness that a decision-maker might need to face. The purpose of this paper is not to provide guidance for which allocation decision-makers should choose, but rather to provide general bounds on when this procedure will result in a large trade-off of utilization and fairness, and when the trade-off will be smaller.

Much of this paper will focus on cases with many candidates and many resources, which in the limit will be approximated by continuous probability distributions. This approximation is common: in many application areas in social sciences, it is known that certain phenomena are well-described by power law and exponential distributions, even when their exact parameters might be unclear [5] [4]. We have focused much of our analysis on these distributions with the goal that our results can be maximally useful.

3.1 Model detail

In this section, we make the motivating example more abstract. We assume a set of N groups. Each group i has a distribution C_i over the number of candidates C that are present during a particular unit of time. We assume that

$$\mathbb{E}_{C \sim C_i} C > 0 \quad \forall i \in [N] \text{ and } C_i \perp C_j \quad \forall i \neq j$$

and that each time period's number of candidates is independently and identically distributed. Much of this paper will focus on the continuous probability distribution case, but the results obtained largely do not depend on this factor.

In the model, we assume we have a number $R \in \mathbb{N}$ of *resources* that we can allocate across these groups. We assume that allocations must be selected before the number of candidates is realized each time period. If v resources and c candidates are present a given group, then $\min(v, c)$ candidates will receive the resource. If $v > c$, some resources will go unused, and if $v < c$, some candidates will go without resources. We assume that the only distinguishing characteristic candidates have is which group they come from: otherwise, they are interchangeable.

We have two different objectives: *utilization* and *fairness*.

$$\begin{aligned} \text{Utilization: } U(\mathbf{v}, \{C_i\}) &:= \sum_{i=1}^N \mathbb{E}_{C \sim C_i} \min(C, v_i) \\ \text{Fairness: } Q(\mathbf{v}, \{C_i\}) &:= \max_{i,j} |q(v_i, C_i) - q(v_j, C_j)| \\ &\quad \text{with } q(v, C) := \frac{\mathbb{E}_{C \sim C} \min(C, v)}{\mathbb{E}_{C \sim C} C} \end{aligned}$$

The fairness term bounds the difference between groups in the probability of a candidate receiving a resource. An allocation $\mathbf{v}$ is α -fair if the fairness objective has value $\leq \alpha$, so all $q(v_i, C_i)$ fractions

are within α of each other. Next, we define the Price of Fairness:

$$\begin{aligned} \text{PoF}(\alpha) &= \frac{\text{Unconstrained utilization}}{\text{Maximum utilization that is } \alpha\text{-fair}} \\ &= \frac{\max_{\mathbf{v}} U(\mathbf{v}, \{C_i\}) \text{ s.t. } \sum_i v_i = R}{\max_{\mathbf{v}} U(\mathbf{v}, \{C_i\}) \text{ s.t. } Q(\mathbf{v}, \{C_i\}) \leq \alpha \text{ and } \sum_i v_i = R} \end{aligned}$$

If there exists an allocation that simultaneously maximizes utilization and is $\alpha = 0$ fair, then the PoF is 1 and there is no tension between our two objectives. Otherwise, the PoF is > 1 and we will need to trade off between these objectives in order to make a decision about which allocation to choose. We will use the standard notation of having $f(x)$ represent the probability density function (or probability mass function) of a distribution and $F(x)$ represent the cumulative distribution function.

Assumptions. There are a few assumptions in this model that are important to recognize. To start, we require that any allocation always use its entire budget R . It might be possible that an α -fair allocation might prefer to use less than its entire budget in order to achieve fairness: this is a case that our model specifically disallows. Additionally, it would often be the case that increasing the budget R would make it easier to achieve any given fairness metric. We prove in Lemma 14 that increasing the budget only improves fairness and utilization metrics, but overall the area of increasing or decreasing the available resources is not the focus of this paper.

Secondly, our definition of fairness revolves around the maximum difference in probability of receiving the resource. A different metric could have used the average difference in probability, for example. A more ambitious multi-time period extension could take into account historical differences in resource allocation and have a fairness metric that attempts to rectify past inequalities. While these (and other) different models might also be valid and interesting future avenues to explore, they would all be built on the necessary foundation this paper provides. The rest of this paper will use the assumptions and model listed above.

3.2 Application areas

Despite the motivating example in Section 2, the focus of this paper is not to model any one specific application such as disaster relief aid. Instead, we view this model as representing the core features of allocation problems in a variety of contexts. A few of these are listed below.

- **Doctors:** As described in the introduction, one example could be allocating doctors across regions with different probabilities of ill patients. It is worth noting that the presence of doctors would likely influence the level of illnesses.
- **Blood banks:** Donated blood needs to be allocated across different hospitals with potentially different needs for blood. This area has the added complication that blood donations are not interchangeable: candidates can only receive a certain subset of the available types of blood.
- **Schools:** A public school district might try to allocate certain resources (teachers, computers, specialized classes) across schools of different size, and therefore different distributions of need.

We have noted that in each case, the potential application area would have key features that are not already captured in the above

model. This is intentional because the goal of this paper is to provide broad results for a broad class of problems. Both Elzayn et al. (2019) and Ensign et al. (2018) use related models in the contexts of allocating police officers to districts with differing crime rates. While this is also an area of application for this model, it raises additional issues that we do not model here, such as the likelihood that the presence of police officers will have an impact on the rate of crime.

Fairness in resource allocation work has been studied in various application areas. For example, [2, 6, 12, 13] study the issue of fair allocation of bandwidth, speed of service, and buffer space in networking. These works differ from ours in the selected model: it is assumed that actors can modify the amount of resources requested at will, which differs from our probabilistic resource generation method. Wang et al. ([n. d.]) studies the specific case of disaster relief aid. In this case, it is assumed that the amount of need is determined before the allocation decision needs to be made, which again differs from our model.

4 GENERAL BOUNDS ON PRICE OF FAIRNESS

One natural question we might ask is, “What is the maximum amount of utilization we give up by requiring fairness?” This amount depends on the set of demand distributions $\{C_i\}$ and our budget R . If we have these quantities, we can directly calculate the PoF by first calculating the unconstrained max-utilization, then the α -fair max-utilization. Elzayn et al. (2019) provides algorithms to calculate both of these quantities (so long as the allocations are integers). However, in many cases we might not know the specific $\{C_i\}$ in advance, but we might still want to have a bound on the PoF. In this section, we explore general bounds on the PoF when we do not know the candidate distributions in advance.

4.1 Discrete resource allocation has unbounded PoF

Suppose we allocate resources in integer units: for example, one generator at a time. In this formulation, we find that *the PoF is unbounded: achieving α fairness could require giving up arbitrarily large amounts of utilization*. This proof is presented in detail in Appendix B. A rough sketch is below.

Suppose we are given a desired $\alpha < 1$ and a maximum PoF ρ we are willing to accept. (We require $\alpha < 1$ because if $\alpha = 1$ the fairness constraint has no effect, so the PoF will be 1 always.) Then, it is possible to construct a set of candidate distributions $\{C_i\}$ and a set of resources R over this such that the PoF is $> \rho$. The proof involves creating two different candidate distributions with parameters depending on α and ρ and then showing that the PoF of allocating resources across those groups is always $> \rho$.

4.2 PoF bounded by $1/\alpha$ under fractional allocation

A critical reader might wonder about the PoF if we were allowed to allocate resources fractionally. Such allocations would offer more flexibility and might allow for lower PoF. In the motivating example of allocating generators across villages, for example, this might correspond to allocating 0.7 generators to Village A and 1.3 generators to Village B. For divisible resources (such as potable water

stores), this type of allocation makes intuitive sense. For indivisible resources, like generators, we can view the allocation as *probabilistic*: An allocation of $v = 1.3$, for example, could be viewed as a 70% chance of $v = 1$ and a 30% chance of $v = 2$.

The expected utilization is equal to $\mathbb{E}_{V \sim \mathcal{V}} \mathbb{E}_{C \sim \mathcal{C}} \min(V, C)$. Both the “deterministic, fractional” and “probabilistic, integer-valued” interpretations give in the same expected utilization (proven in Appendix B). We will switch between these two interpretations (probabilistic and deterministic) depending on which makes more sense for the problem at hand.

Given the ability to allocate resources fractionally, we will find that the PoF is no longer unbounded. In particular, the PoF is upper bounded by $\frac{1}{\alpha}$. Note that this proof relies on an algorithm to calculate a max-utilizing allocation in the fractional case, as well as other related proofs, all present in Appendix C.

THEOREM 1. *Given a set of candidate distributions $\{C_i\}$, level of resources R , and ability to allocate resources fractionally, it is possible to find an α -fair allocation with PoF at most $\frac{1}{\alpha}$, though this allocation may not use all of the resources.*

PROOF. First, we find a max-utilization allocation and divide the groups into two categories:

$$A = \{i \mid q(v_i, C_i) > \alpha\} \quad B = \{i \mid q(v_i, C_i) \leq \alpha\}$$

If $A = \emptyset$, then the max-utilization allocation is already α -fair and $\text{PoF} = 1 < \frac{1}{\alpha}$. Otherwise, modify the allocation $v_i \rightarrow v'_i$ for each $i \in A$ so $q(v'_i, C_i) = \alpha$ exactly. We are able to do this because we know that $q(v_i, C_i)$ is continuous (proven in Appendix B) and increases monotonically from 0 to 1.

This new allocation is α -fair: $q(v_i, C_i) \leq \alpha$. It uses at most the same amount of resources that the optimal allocation used so it is achievable. The only case where it uses exactly the same amount of resources is when A is the empty set. For $i \in A$, we know that:

$$q(v'_i, C_i) = \alpha = \frac{\mathbb{E}_C [\min(C, v'_i)]}{\mathbb{E}_C [C]}$$

$$\Rightarrow \mathbb{E}_C [\min(C, v'_i)] = \alpha \cdot \mathbb{E}_C [C] \geq \alpha \cdot \mathbb{E}_C [\min(C, v_i)]$$

This shows that the utilization of the groups in A is greater than or equal to α multiplied by the utilization of the unconstrained maximum utilization. The utilization of groups in B is identical to their unconstrained utilization because their allocation has not been changed. Then:

$$\begin{aligned} \text{PoF} &= \frac{\sum_{i \in A} \mathbb{E}_C \min(C, v_i) + \sum_{i \in B} \mathbb{E}_C \min(C, v_i)}{\sum_{i \in A} \mathbb{E}_C \min(C, v'_i) + \sum_{i \in B} \mathbb{E}_C \min(C, v_i)} \\ &\leq \frac{\sum_{i \in A} \mathbb{E}_C \min(C, v_i) + \sum_{i \in B} \mathbb{E}_C \min(C, v_i)}{\alpha \cdot \sum_{i \in A} \mathbb{E}_C \min(C, v_i) + \sum_{i \in B} \mathbb{E}_C \min(C, v_i)} \\ &\leq \frac{\sum_{i \in A} \mathbb{E}_C \min(C, v_i)}{\alpha \cdot \sum_{i \in A} \mathbb{E}_C \min(C, v_i)} = \frac{1}{\alpha} \end{aligned}$$

where the inequality in the second to last step comes from the fact that B could be the empty set. $\square$

Note that we have required that all allocations use all of the resources, so we will need to convert the allocation calculated above into one with $\sum_i v_i = R$. The supporting lemma below (proved in the Appendix) provides this final step.

LEMMA 1. *If there exists a fractional allocation $\mathbf{v}$ over candidate distributions $\{C\}$ such that $\sum_i v_i = V' < R$ that is α -fair, there also exists an allocation $\mathbf{v}'$ that is α -fair and has utilization at least as large as $\mathbf{v}$, but additionally has the property that $\sum_i v_i = R$.*

Taken together, these two lemmas show that, for every optimal allocation, we can find an allocation that is α -fair, uses all of the resources R , and has PoF less than $\frac{1}{\alpha}$. This tells us that the PoF for the max-utilization α -fair allocation could be at most $\frac{1}{\alpha}$ for $\alpha > 0$.

4.3 PoF unbounded for $\alpha = 0$ under fractional allocation

However, this $\frac{1}{\alpha}$ bound is undefined for $\alpha = 0$. The proof below shows that we can achieve arbitrarily high PoF with only $n = 2$ groups, even in the case where we increase the budget R by a constant factor in the $\alpha = 0$ case.

LEMMA 2. *Suppose that in the case requiring $\alpha = 0$ allocation, the budget R is multiplied by $k \geq 1$. For any $\rho > 1$, we can create a set of $n = 2$ groups such that requiring $\alpha = 0$ fairness involves a $\text{PoF} > \rho$, assuming fractional allocation of resources.*

PROOF. We create 2 groups, one with distribution C_1 and one with distribution C_2 , described below:

$$\begin{aligned} C \sim C_1 \quad P_1(C = c) &= \begin{cases} 1 - p_1 & c = 0 \\ p_1 & c = n_1 \\ 0 & \text{otherwise} \end{cases} \\ C \sim C_2 \quad P_2(C = c) &= \begin{cases} 1 - p_2 & c = 0 \\ p_2 & c = n_2 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Set $0 < p_1 < 1$, $p_2 = \frac{p_1}{n_1}$, $n_2 = n_1^2$, and $n_1 > 2k \cdot \rho - 1$. Finally, set $R = n_1$.

Because $p_1 > p_2$, candidates are more likely to be present in Group 1, so the optimal allocation is $v_1 = n_2$, $v_2 = 0$.

The $\alpha = 0$ fair allocation has $q(v_1, C_1) = q(v_2, C_2)$. In the case that $v_i \leq n_i$, the q probabilities can be calculated as follows:

$$q(v_i, C_i) = \frac{\mathbb{E}_{C_i} \min(v_i, C)}{\mathbb{E}_{C_i} C} = \frac{v_i \cdot p_i}{n_i \cdot p_i} = \frac{v_i}{n_i}$$

As stated above, this equation only holds for $v_i \leq n_i$. If $v_i > n_i$, $q(v_i, C_i) = 1$ always. However, we must have $v_i \leq n_i$ in the $\alpha = 0$ fair allocation: we have insufficient resources to have $q(v_1, C_1) = q(v_2, C_2) = 1$. This implies that:

$$\frac{v_1}{n_2} = \frac{v_2}{n_2} \quad v_1 + v_2 = k \cdot R = k \cdot n_1$$

where we have incorporated the fact that the budget is k times larger. The solution to this system of equations is:

$$v_1 = \frac{k \cdot n_1}{n_1 + n_2} \cdot n_1 \quad v_2 = \frac{k \cdot n_1}{n_1 + n_2} \cdot n_2$$

This gives us a PoF as follows:

$$\begin{aligned} \text{PoF} &= \frac{p_1 \cdot n_1}{\frac{k \cdot n_1}{n_1 + n_2} \cdot n_1 \cdot p_1 + \frac{k \cdot n_1}{n_1 + n_2} \cdot n_2 \cdot p_2} = \frac{p_1}{\frac{k \cdot n_1}{n_1 + n_2} \cdot p_1 + \frac{k \cdot n_2}{n_1 + n_2} \cdot p_2} \\ &= \frac{p_1}{\frac{k \cdot n_1}{n_1 + n_1^2} \cdot p_1 + \frac{k \cdot n_1^2}{n_1 + n_1^2} \cdot p_2} \end{aligned}$$

where we have used the fact that $n_2 = n_1^2$. Simplifying further gives:

$$\text{PoF} = \frac{p_1}{\frac{k \cdot n_1}{n_1 + n_1^2} \cdot (p_1 + n_1 \cdot p_2)} = \frac{p_1}{\frac{k}{1 + n_1} \cdot (p_1 + n_1 \cdot p_2)} = \frac{p_1 \cdot (1 + n_1)}{k(p_1 + n_1 \cdot p_2)}$$

Using the facts that $p_2 = \frac{p_1}{n_1}$ and $n_1 > 2k \cdot \rho - 1$, the bound becomes:

$$\text{PoF} = \frac{p_1 \cdot (1 + n_1)}{2k \cdot p_1} = \frac{1 + n_1}{2k} > \frac{1 + 2k \cdot \rho - 1}{2k} = \rho$$

as desired. $\square$

Note that p_1 is left as a free parameter. In particular, it could be set to something decreasing in ρ , such as $\frac{1}{\rho}$. This demonstrates that even in the case of tail probabilities decreasing in ρ , it is still possible to create an example with PoF greater than ρ .

5 POF EQUAL TO 1 FOR CERTAIN FAMILIES OF DISTRIBUTIONS

The previous section showed that with fractional allocation, the upper bound on the Price of Fairness is $\frac{1}{\alpha}$. This could still be a large price to pay: if the desired level of fairness is $\alpha = 0.1$, this could mean we might be obliged to reduce utilization by 90% in order to achieve the desired level of equity. Depending on the context, this could be a very high tradeoff. However, the upper bound obtained in the previous section was derived *without any reliance on the characteristics of the candidate distributions $\{C\}$* . In this section, we will show that for certain reasonable distributions, the Price of Fairness is much lower. In some cases, the max-utilization allocation is already $\alpha = 0$ fair!

5.1 Illustrative example for exponential distribution

First, we will work out an example illustrating this phenomenon for $N = 2$ groups, both with exponential distributions. As proved in Appendix C, for this case there exists only one max-utilizing allocation, which has $F_1(v_1) = F_2(v_2) = \tau$, with $v_1 + v_2 = R$. This can be written as $1 - e^{-\lambda_1 \cdot v_1} = 1 - e^{-\lambda_1 \cdot (R - v_1)}$, which implies:

$$v_1 = \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot R \quad \text{and} \quad v_2 = \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot R$$

This allocation was calculated to be max-utilizing. We will also show that it is $\alpha = 0$ fair, which means that $\text{PoF} = 1$. To do this, we will calculate $q(v_i, C_i)$ for both groups. In general,

$$\begin{aligned} q(v, C) &= \frac{\int_0^v c \cdot f(c) dc + v \cdot (1 - F(v))}{\mathbb{E}_{C \sim \exp(\lambda)} C} = \frac{\int_0^v c \cdot \lambda e^{-\lambda c} dx + v \cdot e^{-\lambda v}}{\frac{1}{\lambda}} \\ &= \frac{\frac{1}{\lambda} \cdot (1 - e^{-\lambda \cdot v})}{\frac{1}{\lambda}} = 1 - e^{-\lambda \cdot v} \end{aligned}$$

Then, we know that $q(v_1, C_1) = q(v_2, C_2)$ because

$$1 - e^{-\lambda_1 \cdot \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot R} = 1 - e^{-\lambda_2 \cdot \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot R}$$

This demonstrates that, for the given scenario, the max-utilizing allocation already satisfies $\alpha = 0$ fairness. However, this property is true more broadly than just in this illustrative example. In the next section, we will describe a property of probability distributions such that the max-utilizing allocation is also $\alpha = 0$ fair.

5.2 Proof of general property

THEOREM 2. Consider a set of continuous candidate distributions with $F_i(0) = 0$ and $f_i(v) > 0 \forall v \geq 0$. Suppose the set of candidate distributions $\{C_i\}$ has the following property:

$$F_i(v) = F_j \left(v \cdot \frac{\mathbb{E}_{C \sim C_j}[C]}{\mathbb{E}_{C \sim C_i}[C]} \right) \quad \forall v \in [0, \infty) \forall i, j$$

Then, under fractional allocation of resources, the max-utilization allocation is already $\alpha = 0$ fair. In other words, $\forall i, j$

$$F_i(v_i) = F_j(v_j) \Rightarrow \frac{\mathbb{E}_{C \sim C_i} \min(C, v_i)}{\mathbb{E}_{C \sim C_i}[C]} = \frac{\mathbb{E}_{C \sim C_j} \min(C, v_j)}{\mathbb{E}_{C \sim C_j}[C]}$$

Intuitively, this property means that the CDFs of all of the candidate distributions $\{C_i\}$ are versions of the same function, scaled by the ratio of their expected values. If this property holds, it will turn out mathematically that the max-utilizing allocation is already $\alpha = 0$ fair.

First, we can verify that the exponential distribution satisfies these properties. The exponential distribution has the given expectation and CDF:

$$\mathbb{E}_{C \sim C_j}[C] = \frac{1}{\lambda_j} \quad \mathbb{E}_{C \sim C_i}[C] = \frac{1}{\lambda_i} \quad F_j(v) = 1 - e^{-\lambda_j \cdot v}$$

Substitution shows that the premise of the lemmas above holds:

$$F_j \left(v \cdot \frac{\mathbb{E}_{C \sim C_j}[C]}{\mathbb{E}_{C \sim C_i}[C]} \right) = 1 - e^{-\lambda_j \cdot v \cdot \frac{\lambda_i}{\lambda_j}} = 1 - e^{-\lambda_i \cdot v} = F_i(v)$$

Next, we will explore why this property of probability distributions leads to $PoF = 1$. As shown in Appendix C, the max-utilization allocation has the property that

- $F_i(v_i) > \tau$ only if $F_i(0) \geq \tau$. In this case, $v_i = 0$.
- $F_i(v_i) = \tau$ otherwise.

We require $F_i(0) = 0$, so $v_i > 0 \forall i$. We additionally require $f_i(v) > 0 \forall v \geq 0$, which tells us that there exists only one max-utilizing allocation. Next, we examine the fairness constraint and rewrite it as:

$$\begin{aligned} & \frac{\mathbb{E}_{C \sim C_i} [C \mid C \leq v_i] \cdot P_i(C \leq v_i)}{\mathbb{E}_{C \sim C_i}[C]} + \frac{v_i \cdot P_i(C > v_i)}{\mathbb{E}_{C \sim C_i}[C]} = \\ & \frac{\mathbb{E}_{C \sim C_j} [C \mid C \leq v_j] \cdot P_j(C \leq v_j)}{\mathbb{E}_{C \sim C_j}[C]} + \frac{v_j \cdot P_j(C > v_j)}{\mathbb{E}_{C \sim C_j}[C]} \quad \forall i, j \end{aligned}$$

For notational convenience, we denote $m_{ij} = \frac{\mathbb{E}_{C \sim C_i}[C]}{\mathbb{E}_{C \sim C_j}[C]}$. We will prove this theorem by considering subterms of each side independently and proving that they are equal.

LEMMA 3.

$$\begin{aligned} F_i(v) &= F_j \left(\frac{v}{m_{ij}} \right) \quad \forall v \in [0, \infty) \forall i, j \\ \Rightarrow \frac{v \cdot P_i(C > v)}{\mathbb{E}_{C \sim C_i}[C]} &= \frac{\frac{v}{m_{ij}} \cdot P_j \left(C > \frac{v}{m_{ij}} \right)}{\mathbb{E}_{C \sim C_j}[C]} \quad \forall v \in [0, \infty) \forall i, j \end{aligned}$$

PROOF.

$$\begin{aligned} \frac{\frac{v}{m_{ij}} \cdot P_j \left(C > \frac{v}{m_{ij}} \right)}{\mathbb{E}_{C \sim C_j}[C]} &= \frac{\frac{v}{m_{ij}} \cdot P_i(C > v)}{\mathbb{E}_{C \sim C_j}[C]} = \frac{v \cdot P_i(C > v)}{m_{ij} \cdot \mathbb{E}_{C \sim C_j}[C]} \\ &= \frac{v \cdot P_i(C > v)}{\mathbb{E}_{C \sim C_i}[C]} \end{aligned}$$

where the first step comes from the fact that $P_j(C > \frac{v}{m_{ij}}) = 1 - F_j(\frac{v}{m_{ij}}) = 1 - F_i(v) = P_i(C > v)$ $\square$

The max-utilizing allocation has the property that $F_i(v_i) = F_j(v_j)$. Having the additional property that $F_i(v) = F_j \left(\frac{v}{m_{ij}} \right)$ implies that $v_j = \frac{v_i}{m_{ij}}$. This allows us to rewrite the above lemma in terms of the max-utilizing allocations:

$$F_i(v_i) = F_j(v_j) \Rightarrow \frac{v_i \cdot P_i(C > v_i)}{\mathbb{E}_{C \sim C_i}[C]} = \frac{v_j \cdot P_j(C > v_j)}{\mathbb{E}_{C \sim C_j}[C]}$$

Next, we consider the second subterm in the equality:

LEMMA 4.

$$\begin{aligned} F_i(v) &= F_j \left(\frac{v}{m_{ij}} \right) \quad \forall v \in [0, \infty) \forall i, j \Rightarrow \\ \frac{\mathbb{E}_{C \sim C_i} [C \mid C \leq v_i] \cdot P_i(C \leq v_i)}{\mathbb{E}_{C \sim C_i}[C]} &= \frac{\mathbb{E}_{C \sim C_j} [C \mid C \leq v_j] \cdot P_j(C \leq v_j)}{\mathbb{E}_{C \sim C_j}[C]} \end{aligned}$$

PROOF. What is $\mathbb{E}[C \mid C < v]$ written out in terms of the CDF of a distribution? We will consider $C \mid C < v$ as a different probability distribution. It has PDF and CDF as shown below:

$$f^*(t) = \begin{cases} \frac{f(t)}{F(v)} & 0 \leq t \leq v \\ 0 & \text{otherwise} \end{cases} \quad F^*(t) = \begin{cases} \frac{F(t)}{F(v)} & 0 \leq t \leq v \\ 1 & \text{otherwise} \end{cases}$$

We can rewrite one of the terms above as:

$$\begin{aligned} \mathbb{E}[C \mid C < v] &= \int_0^v (1 - F^*(t)) dt = \int_0^v (1 - \frac{F(t)}{F(v)}) dt \\ &= \int_0^v \left(1 - \frac{F(t)}{F(v)} \right) dt = v - \frac{1}{F(v)} \int_0^v F(t) dt \end{aligned}$$

We will use this fact to evaluate the following term, substituting $\frac{v}{m_{ij}}$ as the input v_j :

$$\begin{aligned} \mathbb{E}_{C \sim C_j} [C \mid C \leq v] \cdot \frac{F_j \left(\frac{v}{m_{ij}} \right)}{\mathbb{E}_{C \sim C_j}[C]} &= \left(\frac{v}{m_{ij}} - \frac{\int_0^{v/m_{ij}} F_j(t) dt}{F_j \left(\frac{v}{m_{ij}} \right)} \right) \frac{F_j \left(\frac{v}{m_{ij}} \right)}{\mathbb{E}_{C \sim C_j}[C]} \\ &= \left(\frac{v \cdot F_j \left(\frac{v}{m_{ij}} \right) - m_{ij} \cdot \int_0^{v/m_{ij}} F_j(t) dt}{m_{ij} \cdot F_j \left(\frac{v}{m_{ij}} \right)} \right) \frac{F_j \left(\frac{v}{m_{ij}} \right)}{\mathbb{E}_{C \sim C_j}[C]} \\ &= \frac{v \cdot F_j \left(\frac{v}{m_{ij}} \right) - m_{ij} \int_0^{v/m_{ij}} F_j(t) dt}{m_{ij} \cdot \mathbb{E}_{C \sim C_j}[C]} \\ &= \frac{v \cdot F_i(v) - m_{ij} \int_0^{v/m_{ij}} F_j(t) dt}{\mathbb{E}_{C \sim C_i}[C]} \end{aligned}$$

where we have used the assumption of equality in the CDFs in the last step. Next, we note that $F_j(t) = F_i(g(t))$ where $g(t) = m_{ij} \cdot t$.

By the reverse chain rule, $\int_a^b f'(g(t)) \cdot g'(t) dt = [f(g(t))]_a^b$. Here, $g'(t) = m_{ij}$, so if we denote the anti-derivative of F_i by H_i , then

$$\int_0^{v/m_{ij}} F_j(t) \cdot m_{ij} dt = \int_0^{v/m_{ij}} F_i(g(t)) \cdot m_{ij} dt = [H_i(g(t))]_0^{v/m_{ij}} = [H_i(m_{ij} \cdot t)]_0^{v/m_{ij}} = H_i(v) - H_i(0) = \int_0^v F_i(t) dt$$

This allows us to rewrite the item above as

$$\frac{v \cdot F_i(v) - \int_0^v F_i(t) dt}{\mathbb{E}_{C \sim C_i}[C]}$$

This gives a value for the righthand side of the equality we are trying to show. For the lefthand side, we use results of our previous analysis of $\mathbb{E}[C \mid C \leq v]$ to find that the two terms are equivalent:

$$\begin{aligned} \frac{\mathbb{E}_{C \sim C_i}[C \mid C \leq v] \cdot F_i(v)}{\mathbb{E}_{C \sim C_i}[C]} &= \left(v - \frac{1}{F_i(v)} \int_0^v F_i(t) dt \right) \cdot \frac{F_i(v)}{\mathbb{E}_{C \sim C_i}[C]} \\ &= \frac{v \cdot F_i(v) - \int_0^v F_i(t) dt}{\mathbb{E}_{C \sim C_i}[C]} \end{aligned}$$

□

Taken together, these two lemmas tell us that

$$\begin{aligned} F_i(v) &= F_j\left(\frac{v}{m_{ij}}\right) \quad \forall v \in [0, \infty) \\ \Rightarrow \frac{\mathbb{E}_{C \sim C_i} \min(C, v_i)}{\mathbb{E}_{C \sim C_i}[C]} &= \frac{\mathbb{E}_{C \sim C_j} \min(C, v_j)}{\mathbb{E}_{C \sim C_j}[C]} \quad \forall i, j \end{aligned}$$

so any max-utilization allocation is already $\alpha = 0$ fair.

5.3 Distributions that satisfy the given premise

We have previously shown that exponential distributions satisfy the given premise, which is perhaps not very surprising. Exponential distributions have certain nice properties: they are memoryless, for example, and are not heavy-tailed. At first guess, it could be reasonable to guess that this phenomenon of PoF equal to 1 is only present in exponential distributions.

However, there exist further families of distributions that also satisfy the premise above and therefore have PoF equal to 1; these include Weibull distributions with the same k parameter. They have expectation and CDF as given below:

$$\begin{aligned} \mathbb{E}_{C \sim C_j}[C] &= \lambda_j \cdot \Gamma\left(1 + \frac{1}{k}\right) \quad \mathbb{E}_{C \sim C_i}[C] = \lambda_i \cdot \Gamma\left(1 + \frac{1}{k}\right) \\ F_j(v) &= 1 - e^{-\left(\frac{v}{\lambda_j}\right)^k} \end{aligned}$$

Again, substitution shows that the premise of the above lemmas holds:

$$F_j\left(v \cdot \frac{\mathbb{E}_{C \sim C_j}[C]}{\mathbb{E}_{C \sim C_i}[C]}\right) = 1 - e^{-\left(\frac{v}{\lambda_j} \cdot \frac{\lambda_j}{\lambda_i}\right)^k} = 1 - e^{-\left(\frac{v}{\lambda_i}\right)^k} = F_i(v)$$

Thus, the theorem proved in this section shows that Weibull distributions with the same k parameter have PoF = 1. This result is somewhat surprising, given that Weibull distributions in general have fewer nice properties than exponential distributions. However, Weibull distributions, along with power law distributions, which will be discussed in the next section, tend to be better models of real-life phenomena. For example, natural disasters and terrorist

attacks are commonly believed to follow heavy-tailed distributions in severity [4, 5]. This result shows us that there are plausibly realistic models of natural phenomena that have no trade-off between fair and max-utilizing allocations.

6 POF BOUNDED BY SMALL CONSTANT FOR POWER LAW DISTRIBUTION

So far, Section 4 has shown that PoF can be arbitrarily large, and Section 5 has shown that multiple distributions have $PoF = 1$. A next logical step is to explore probability distributions with small gaps between the max-utilizing and α -fair max utilizing allocations. Specifically, we consider power law distributions. There are multiple distributions that are commonly described as “power law” distributions. The one we will use, also called the “Lomax distribution”, has PDF and CDF as below. Note that it has support for all $x \geq 0$ and a parameter $b > 1$ controlling its shape.

$$f_C(v) = \frac{b}{(v+1)^{b+1}} \quad F_C(v) = 1 - \frac{1}{(v+1)^b} \quad \mathbb{E}_C[C] = \frac{1}{b-1}$$

THEOREM 3. *Set the number of groups N to some constant h . Then, the Price of Fairness from allocating R resources across these groups is upper bounded by*

$$h \cdot H_h$$

where H_h is the h th harmonic number.

Note that this bound is independent of the fairness parameter α as well as the budget R and any distribution parameters b_i !

We will prove this theorem by breaking it up into several smaller lemmas. For convenience, we order the groups so that $b_1 \geq b_2 \dots \geq b_h$. First, we write out the full PoF term:

$$\begin{aligned} PoF &= \frac{\text{Unconstrained max-utilization}}{\text{Utilization constrained to be } \alpha\text{-fair}} \\ &= \frac{\sum_{i=1}^h \left[\int_0^{v_i} x \cdot f_i(x) dx + v_i \cdot (1 - F_i(v_i)) \right]}{\sum_{i=1}^h \left[\int_0^{v'_i} x \cdot f_i(x) dx + v'_i \cdot (1 - F_i(v'_i)) \right]} \end{aligned}$$

As proved in Appendix C, because the power law distribution has positive support everywhere in the space we are considering ($v \geq 0$), it has a unique max-utilizing allocation and α -fair max-utilizing allocation. The PoF is largest when $\alpha = 0$, so that is the case we will upper-bound.

We will denote the unconstrained max-utilization allocations by $\{v_i\}$ and the α -fair allocations by $\{v'_i\}$, for $i \in [h]$.

LEMMA 5. $v_i \leq \frac{R}{h-i+1} \quad \forall i$.

PROOF. As proved in Appendix C, the max-utilizing allocation $\{v_i\}$ has the property that $F_i(v_i) = F_j(v_j) \quad \forall i, j$. For the power law distributions, this property can be written as:

$$\frac{1}{(1+v_i)^{b_i}} = \frac{1}{(1+v_j)^{b_j}} \quad \Rightarrow \quad (1+v_i)^{b_i} = (1+v_j)^{b_j}$$

For $b_i \leq b_j$, this implies $v_i \geq v_j$, so $v_1 \leq v_2 \dots \leq v_h$. Any max-utilizing allocation over power law distributions of this form must use the entire resource R : because the support of the distribution continues to ∞ , any allocation using $R - \epsilon$ resources can be strictly improved by allocating the remaining ϵ resources. This tells us that $v_1 + v_2 + \dots + v_h = R$.

We can prove the upper bound by working backwards from h . $v_h \leq R$ because the allocation to any one group cannot exceed the total amount of resources available. $v_{h-1} \leq v_h$ and together both of them cannot sum to more than R , so $v_{h-1} \leq \frac{R}{2}$. Similarly, for any v_i , $v_i \leq v_{i+1} \leq \dots v_h$, with $v_i + v_{i+1} + \dots v_h \leq R$, so $v_i \leq \frac{R}{h-i+1}$. $\square$

LEMMA 6. $v'_h \geq \frac{R}{h}$.

PROOF. As proved in Appendix B, any α -fair allocation $\mathbf{v}_i$ with $\sum_i v_i < R$ can be increased to an α -fair allocation $\mathbf{v}'_i$ with $\sum_i v'_i = R$ and no decrease in utilization. This tells us that the α -fair max-utilizing allocation must use all of the resources, so $v'_1 + v'_2 + \dots v'_h = R$. Under the definition of $q(v_i, C_i)$,

$$\begin{aligned} q(v_i, C_i) &= \frac{\int_0^v x \cdot f(x) dx + v \cdot (1 - F(v))}{\mathbb{E}_{C \sim C} [C]} \\ &= \frac{\frac{1-b \cdot v \cdot (v+1)^{-b} - (v+1)^{-b}}{b-1} + v \cdot (v+1)^{-b}}{\frac{1}{b-1}} = 1 - \frac{1}{(v+1)^{b-1}} \end{aligned}$$

For $\alpha = 0$, we require $q(v'_i, C_i) = q(v'_j, C_j) \forall i, j$, which can be written as:

$$1 - \frac{1}{(v'_i + 1)^{b_i-1}} = 1 - \frac{1}{(v'_j + 1)^{b_j-1}} \Rightarrow (v'_i + 1)^{b_i-1} = (v'_j + 1)^{b_j-1}$$

For $b_i \geq b_j$, this implies $v'_i \leq v'_j$. Given that we have h terms that sum up to R and v'_h is the largest of them, we must have $v'_h \geq \frac{R}{h}$. $\square$

While Lemma 5 produces bounds for $v_i \forall i \in [h]$, Lemma 6 only bounds the v_h case. The following lemma will demonstrate that other bounds aren't necessary, since we can rewrite all of the α -fair utilization in terms of v'_h .

LEMMA 7. The α -fair max-utilizing allocation can be rewritten entirely in terms of v'_h :

$$\begin{aligned} &\sum_{i=1}^h \left[\int_0^{v'_i} x \cdot f_i(x) dx + v'_i \cdot (1 - F_i(v'_i)) \right] \\ &= (b_h - 1) \cdot \left(\sum_{i=1}^h \frac{1}{b_i - 1} \right) \cdot \left(\int_0^{v'_h} x \cdot f_h(x) dx + v'_h \cdot (1 - F_h(v'_h)) \right) \end{aligned}$$

PROOF. $\alpha = 0$ fairness requires $q(v'_i, C_i) = q(v'_h, C_h)$, which can be written as:

$$\frac{\int_0^{v'_i} x f_i(x) dx + v'_i (1 - F_i(v'_i))}{\frac{1}{b_i-1}} = \frac{\int_0^{v'_h} x f_h(x) dx + v'_h (1 - F_h(v'_h))}{\frac{1}{b_h-1}}$$

Simple rearrangement shows that $\int_0^{v'_i} x \cdot f_i(x) dx + v'_i (1 - F_i(v'_i))$ is equal to $\frac{b_h-1}{b_i-1} \left(\int_0^{v'_h} x \cdot f_h(x) dx + v'_h (1 - F_h(v'_h)) \right)$. This procedure could be repeated to obtain rewritten terms for every v'_i term with $i \neq h$. $\square$

The above lemma allows us to rewrite the entire PoF term as:

$$\frac{\sum_{i=1}^h \left[\int_0^{v_i} x \cdot f_i(x) dx + v_i \cdot (1 - F_i(v_i)) \right]}{(b_h - 1) \cdot \left(\sum_{i=1}^h \frac{1}{b_i-1} \right) \left(\int_0^{v'_h} x \cdot f_h(x) dx + v'_h \cdot (1 - F_h(v'_h)) \right)}$$

Next, our approach for bounding the PoF will be to bound each term within the sum in the numerator independently.

LEMMA 8. $PoF_h \leq h$ where

$$PoF_h := \frac{\int_0^{v_h} x \cdot f_h(x) dx + v_h \cdot (1 - F_h(v_h))}{(b_h - 1) \cdot \left(\sum_{i=1}^h \frac{1}{b_i-1} \right) \left(\int_0^{v'_h} x \cdot f_h(x) dx + v'_h \cdot (1 - F_h(v'_h)) \right)}$$

PROOF. We can apply Lemmas 5 and 6 to obtain $v_h \leq R$ and $v'_h \geq \frac{R}{h}$. We additionally use the fact that

$$\int_0^v x \cdot f_i(x) dx + v \cdot (1 - F_i(v)) = \frac{1}{b_i - 1} \left(1 - \frac{1}{(1+v)^{b_i-1}} \right) = \frac{1}{b_i - 1} q(v, C_i)$$

These facts allow us to bound the equation as:

$$\begin{aligned} PoF_h &\leq \frac{\frac{1}{b_h-1} \left(1 - \frac{1}{(1+R)^{b_h-1}} \right)}{(b_h - 1) \cdot \left(\sum_{j=1}^h \frac{1}{b_j-1} \right) \cdot \frac{1}{b_h-1} \left(1 - \frac{1}{(1+\frac{R}{h})^{b_h-1}} \right)} \\ &= \frac{1 - \frac{1}{(1+R)^{b_h-1}}}{(b_h - 1) \cdot \left(\sum_{j=1}^h \frac{1}{b_j-1} \right) \cdot \left(1 - \frac{1}{(1+\frac{R}{h})^{b_h-1}} \right)} \\ &\leq \frac{1 - \frac{1}{(1+R)^{b_h-1}}}{1 - \frac{1}{(1+\frac{R}{h})^{b_h-1}}} \end{aligned}$$

In the last step, to over-bound PoF_h , we under-bounded a term in the denominator, replacing $(b_h - 1) \cdot \left(\sum_{i=1}^h \frac{1}{b_i-1} \right)$ with 1.

At $R = 0$, this fraction is undefined. We denote the numerator as s , the denominator as t , and use L'Hôpital's rule to investigate the limit as $R \rightarrow 0$. We note that $s' = (b_h - 1) \frac{1}{(1+R)^{b_h}}$ and $t' = \frac{b_h-1}{h} \frac{1}{(1+\frac{R}{h})^{b_h}}$. As $R \rightarrow 0$, $s' \rightarrow b_h - 1$ and $t' \rightarrow \frac{b_h-1}{h}$, so $\frac{s'}{t'} \rightarrow h$, which tells us that $PoF_h \rightarrow h$ as $R \rightarrow 0$.

Next, we will show that PoF_h is decreasing in R , proving that h is an upper bound on PoF_h everywhere. To do this, we will use L'Hôpital's Monotone rule [9], which states the following:

LEMMA 9. Given $-\infty < a < c < \infty$, and s, t continuous functions differentiable on (a, c) with $s(a) = t(a) = 0$ or $s(c) = t(c) = 0$ and $g(x) \neq 0 \forall x \in (a, c)$, then if $\frac{s'}{t'}$ is increasing (decreasing) on (a, c) , then $\frac{s}{t}$ is also increasing (decreasing).

The functions in PoF_h satisfy this rule: s, t are both continuous and differentiable, and for $a = 0$, $s(a) = t(a) = 0$. Using this rule:

$$\frac{s'}{t'} = \frac{(b_h - 1) \frac{1}{(1+R)^{b_h}}}{\frac{b_h-1}{h} \frac{1}{(1+\frac{R}{h})^{b_h}}} = \frac{h \cdot \left(1 + \frac{R}{h} \right)^{b_h}}{(1+R)^{b_h}} = h \cdot \left(\frac{1 + \frac{R}{h}}{1+R} \right)^{b_h}$$

This term is decreasing in R because the denominator $1 + R$ is increasing faster than the numerator $1 + \frac{R}{h}$ is increasing. This tells us that $\frac{s'}{t'}$ is decreasing in R , so PoF_h is decreasing as well. The upper bound of $R = 0$ is an upper bound for all of PoF_h . $\square$

Next, we consider a bound for any $i \neq h$.

LEMMA 10. For $i \in [h-1]$, $PoFi \leq \frac{h}{h-i+1}$, where

$$PoFi := \frac{\int_0^{v_i} x \cdot f_i(x) dx + v_i(1 - F_i(v_i))}{(b_h - 1) \cdot \left(\sum_{j=1}^h \frac{1}{b_j - 1} \right) \left(\int_0^{v'_h} x \cdot f_h(x) dx + v'_h(1 - F_h(v'_h)) \right)}$$

PROOF. Using bounds $v_i \leq \frac{R}{h-i+1}$ and $v'_h \geq \frac{R}{h}$ from Lemmas 5 and 6, we substitute and rewrite:

$$\begin{aligned} PoFi &\leq \frac{\frac{1}{b_i - 1} \left(1 - \frac{1}{\left(1 + \frac{R}{h-i+1}\right)^{b_i - 1}} \right)}{(b_h - 1) \cdot \left(\sum_{j=1}^h \frac{1}{b_j - 1} \right) \cdot \frac{1}{b_h - 1} \left(1 - \frac{1}{\left(1 + \frac{R}{h}\right)^{b_h - 1}} \right)} \\ &= \frac{1 - \frac{1}{\left(1 + \frac{R}{h-i+1}\right)^{b_i - 1}}}{(b_i - 1) \cdot \left(\sum_{j=1}^h \frac{1}{b_j - 1} \right) \cdot \left(1 - \frac{1}{\left(1 + \frac{R}{h}\right)^{b_h - 1}} \right)} \\ &< \frac{1 - \frac{1}{\left(1 + \frac{R}{h-i+1}\right)^{b_i - 1}}}{\left(\frac{b_i - 1}{b_h - 1} + 1 \right) \cdot \left(1 - \frac{1}{\left(1 + \frac{R}{h}\right)^{b_h - 1}} \right)} \end{aligned}$$

At $R = 0$, this is again undefined. Using the same process of rewriting the fraction as $\frac{s}{t}$, we get

$$s' = \frac{b_i - 1}{h - i + 1} \frac{1}{\left(1 + \frac{R}{h-i+1}\right)^{b_i}} \quad t' = \frac{b_h - 1 + b_i - 1}{h} \frac{1}{\left(1 + \frac{R}{h}\right)^{b_h}}$$

Using L'Hôpital's rule again, as $R \rightarrow 0$, $\frac{s'}{t'}$ goes to a ratio of:

$$\frac{\frac{b_i - 1}{h - i + 1}}{\frac{b_h - 1 + b_i - 1}{h}} = \frac{b_i - 1}{b_h - 1 + b_i - 1} \cdot \frac{h}{h - i + 1} = \frac{1}{\frac{b_h - 1}{b_i - 1} + 1} \cdot \frac{h}{h - i + 1}$$

$\frac{b_h - 1}{b_i - 1}$ is decreasing in b_i , so this overall term is increasing in b_i . $b_i \rightarrow \infty$ upper-bounds this fraction by $\frac{1}{0+1} \cdot \frac{h}{h-i+1} = \frac{h}{h-i+1}$ which provides the limit as $R \rightarrow 0$. Again, we use L'Hôpital's Monotone rule to show that $\frac{s}{t}$ is decreasing in R and that the $R = 0$ upper bound is sufficient $\forall R$. This tells us that $\frac{s'}{t'}$ equals

$$\frac{\frac{b_i - 1}{h - i + 1} \frac{1}{\left(1 + \frac{R}{h-i+1}\right)^{b_i}}}{\frac{b_h - 1 + b_i - 1}{h} \frac{1}{\left(1 + \frac{R}{h}\right)^{b_h}}} = \frac{b_i - 1}{h - i + 1} \cdot \frac{h}{b_h - 1 + b_i - 1} \cdot \frac{\left(1 + \frac{R}{h}\right)^{b_h}}{\left(1 + \frac{R}{h-i+1}\right)^{b_i}}$$

The term in the denominator has a base that grows more quickly in R (because $1 + \frac{R}{h-i+1}$ increases in R faster than $1 + \frac{R}{h}$) and a larger exponent (because $b_i \geq b_h$). Because of this, the denominator grows more quickly in R , so $\frac{s'}{t'}$ is decreasing in R . L'Hôpital's Monotone rule tells us that $\frac{s}{t}$ is decreasing in R as well, and so the over bound at $R = 0$ is an over bound for the entire function. $\square$

These lemmas, taken together, give us a bound for the overall PoF in terms of H_h , the h th harmonic number:

$$h + \sum_{i=1}^{h-1} \left\lceil \frac{h}{h-i+1} \right\rceil = h \cdot H_h$$

We also investigated this bound empirically by directly calculating the PoF for groups with different power law distributions. We found that the PoF is typically much lower than this theoretical bound, frequently around 1.1. This result tells us that, in general, the most max-utilizing allocations are already very close in utilization to $\alpha = 0$ fair solutions. Overall, this provides an optimistic outlook for resource allocation in the numerous real-world scenarios modeled by power law distributions.

7 CONCLUSION

In this work we have analyze a problem in fair resource allocation from the recent literature [7]. Starting from a set of open questions around the relationship between fair allocations and maximum-utilization allocations, we obtained upper bounds on the gap between the utilization of optimal and fair solutions – formalized as the *Price of Fairness* (PoF). In addition to bounds for general distributions, we showed that a number of natural families of distributions exhibit no gap at all: there exist optimal allocations that also achieve perfect fairness guarantees. Finally, we established constant upper bounds for power-law distributions.

There are a number of further directions suggested by this work. To begin with, it would be interesting to try characterizing the set of distributions for which the Price of Fairness is equal to 1; this is a fundamental distributional property that applies more broadly than it initially appears. It would also be valuable to extend the set of distributions for which we have the techniques to prove constant upper bounds on the Price of Fairness.

Additionally, future work could explore different definitions of fairness. If we consider models where candidates are not interchangeable and persist from one time state to another, it might be considered more “fair” to allocate resources preferentially to those who have been waiting longest for resources. Finally, the model so far largely assumes a fixed amount of resources. In some cases, though, the amount of resources could be increased or decreased and it could be enlightening to consider how modifying the availability of resources might impact the tradeoff between utilization and fairness. By developing a richer understanding of the behavior of different classes of distributions in this model, we can further our understanding of the interplay between fairness and utilization when we allocate resources with this type of uncertainty.

ACKNOWLEDGMENTS

This work was in part supported by a Simons Investigator Award, a grant from the MacArthur Foundation, and NSF grants DGE-1650441, CCF-1740822, and SES-1741441. We are grateful to Manish Raghavan, the AI in Policy and Practice working group at Cornell, and Raunak Kumar for invaluable discussions. We would also like to thank four anonymous reviewers for their kind and insightful feedback.

REFERENCES

- [1] Stephen Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.
- [2] Dah-Ming Chiu and Raj Jain. 1989. Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. *Comput. Netw. ISDN Syst.* 17, 1 (June 1989), 1–14. [https://doi.org/10.1016/0169-7552\(89\)90019-6](https://doi.org/10.1016/0169-7552(89)90019-6)
- [3] Dah Ming Chiu and Raj Jain. 1989. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems* 17 (06 1989), 1–14. [https://doi.org/10.1016/0169-7552\(89\)90019-6](https://doi.org/10.1016/0169-7552(89)90019-6)
- [4] Aaron Clauset. 2018. Trends and fluctuations in the severity of interstate wars. *Science Advances* 4, 2 (2018). <https://doi.org/10.1126/sciadv.aao3580> arXiv:<https://advances.sciencemag.org/content/4/2/eaao3580.full.pdf>
- [5] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. 2007. Power-law distributions in empirical data. *arXiv e-prints*, Article arXiv:0706.1062 (Jun 2007), arXiv:0706.1062 pages. arXiv:physics.data-an/0706.1062
- [6] A. Demers, S. Keshav, and S. Shenker. 1989. Analysis and Simulation of a Fair Queueing Algorithm. *SIGCOMM Comput. Commun. Rev.* 19, 4 (Aug. 1989), 1–12. <https://doi.org/10.1145/75247.75248>
- [7] Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. 2019. Fair Algorithms for Learning in Allocation Problems. *Proceedings of the Conference on Fairness, Accountability, and Transparency- FAT*19* (2019). <https://doi.org/10.1145/3287560.3287571>
- [8] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway Feedback Loops in Predictive Policing. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research)*, Sorelle A. Friedler and Christo Wilson (Eds.), Vol. 81. PMLR, New York, NY, USA, 160–171. <http://proceedings.mlr.press/v81/ensign18a.html>
- [9] Ricardo Estrada and Miroslav Pavlovic. 2016. L'Hôpital's Monotone rule, Gromov's theorem, and operations that preserve the monotonicity of quotients. *Publications de l'Institut Mathématique* 101(115), 2017, 11–24 101 (12 2016). <https://doi.org/10.2298/PIM1715011E>
- [10] Virginia Eubanks. 2018. *Automating Inequality*. St. Martin's Press.
- [11] Kuzman Ganchev, Michael Kearns, Yuriy Nevmyvaka, and Jennifer Wortman Vaughan. 2009. Censored exploration and the dark pool problem. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 185–194.
- [12] Jeffrey M. Jaffe. 1981. Bottleneck flow control. *IEEE Transactions on Communications* 29 (1981), 954–962.
- [13] Frank P Kelly, Aman K Maulloo, and David KH Tan. 1998. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society* 49, 3 (1998), 237–252.
- [14] Ariel D. Procaccia. 2013. Cake cutting: not just child's play. *Commun. ACM* 56, 7 (2013), 78–87. <https://doi.org/10.1145/2483852.2483870>
- [15] Yanyan Wang, Vicki M. Bier, and Baiqing Sun. [n. d.]. Measuring and Achieving Equity in Multiperiod Emergency Material Allocation. *Risk Analysis* 0, 0 ([n. d.]). <https://doi.org/10.1111/risa.13342> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.13342>