Scarce Societal Resource Allocation and the Price of (Local) Justice

Quan Nguyen,[1] **Sanmay Das,** [2] **Roman Garnett** [1]

[1] Washington University in St. Louis
[2] George Mason University
quan@wustl.edu, sanmay@gmu.edu, garnett@wustl.edu

Abstract

We consider the allocation of scarce societal resources, where a central authority decides which individuals receive which resources under capacity or budget constraints. Several algorithmic fairness criteria have been proposed to guide these procedures, each quantifying a notion of *local justice* to ensure the allocation is aligned with the principles of the local institution making the allocation. For example, the efficient allocation maximizes overall social welfare, whereas the leximin assignment seeks to help the "neediest first." Although the "price of fairness" (PoF) of leximin has been studied in prior work, we expand on these results by exploiting the structure inherent in real-world scenarios to provide tighter bounds. We further propose a novel criterion – which we term LoINC (leximin over individually normalized costs) – that maximizes a different but commonly used notion of local justice: prioritizing those *benefiting the most* from receiving the resources. We derive analogous PoF bounds for LoINC, showing that the price of LoINC is typically much lower than that of leximin. We provide extensive experimental results using both synthetic data and in a real-world setting considering the efficacy of different homelessness interventions. These results show that the empirical PoF tends to be substantially lower than worst-case bounds would imply and allow us to characterize situations where the price of LoINC fairness can be high.

1 Introduction

Algorithmic fairness broadly refers to the impact of automated learning or decision-making methods on different subpopulations (*statistical fairness*) and/or on individuals of different types (*individual fairness*) (Dwork et al. 2012; Pleiss et al. 2017; Kearns et al. 2018). If we can measure this impact, we can audit existing systems for fairness violations and guarantee fairness by training learning algorithms subject to a fairness constraint. There is also a well-established literature on algorithmic *fair division* and, more generally, mechanisms for fair allocation of both divisible and indivisible goods, potentially to agents with different preferences (Bogomolnaia and Moulin 2004; Kurokawa, Procaccia, and Shah 2015). The allocation of scarce societal resources is related to these efforts but also has several unique features, making it important and interesting to study separately. First,

any notion of fairness in such domains is implicitly a notion of what the political philosopher Jon Elster terms *local justice:* the principles used by institutions to allocate scarce goods or necessary burdens (Elster 1992). Examples include the welfare-maximizing allocation or the allocation that prioritizes the most needy or vulnerable individuals. Operationalizing local justice involves explicitly spelling out rules for prioritization in the allocation of scarce resources, and those rules of course determine outcomes for different subgroups as well as individuals. Instead of fairness as an auditing mechanism or a constraint, justice becomes the primary concern of the allocation mechanism.

Second, specific structures of these allocation problems allow mechanisms simpler than in most problems of fair division or matching under preferences. All individuals need the resources, but there is limited supply, and the institution performing the allocation can be thought of as a benevolent planner who does not need to incentivize participation and can dictate allocations. Further, the types of resources available are typically limited in number and often well-ordered in terms of their added value independent of the individual (e.g., a public school with different class sizes: it is reasonable to assume that every student would do best in the smallest class, next best in the medium sized one, and worst in the largest).

These features allow one to usefully model many scenarios that arise in the allocation of scarce societal resources as assignment problems, where a central authority decides which individuals receive which resources under capacity or budget constraints. Different algorithmic fairness criteria in such problems operationalize different notions of local justice. The *efficient,* or lowest-cost allocation operationalizes social welfare maximization, whereas the *leximin* mechanism operationalizes the "neediest first" criterion. Although the general "price of fairness" (PoF) of leximin has been studied in prior work, here we exploit the structure inherent to many important instantiations of societal scarce resource allocation to provide substantially tighter bounds. Instead of $\mathcal{O}(n)$ where n is the number of agents, we show that the PoF can be upper bounded by functions of both n and the number of resource types, which are ultimately sub-linear in n.

We also propose a novel formalization of a fairness criterion, *LoINC* (for *Leximin over Individually Normalized Costs*), motivated by the local justice principle of prioritizing those who would *benefit the most* from receiving the scarce

resources (Elster 1992). This is distinct from either of the prior criteria and is commonly used in situations like medical triage. We derive analogous PoF bounds for LoINC, showing that the price of LoINC fairness is typically much lower than the price of leximin fairness. We then turn to analyzing the empirical price of fairness using both synthetic data and a dataset estimating real-world efficacy of different homelessness interventions. These results show that the empirical PoF tends to be substantially lower than the worst-case bounds would imply and also support our theoretical finding that the price of LoINC fairness is much lower than the price of leximin fairness. We also characterize the types of utility distributions for which the PoF can be high. For both leximin and LoINC, bimodal distributions, where agents can have either very low or very high costs, tend to lead to the highest prices of fairness, although leximin displays higher PoF for asymmetric distributions than LoINC.

Related work. The price of fairness measure was first introduced by Bertsimas, Farias, and Trichakis (2011) and Caragiannis et al. (2012) independently to analyze various fair assignment objectives such as proportionality, envy-freeness, or equitability. PoF was then studied under a wide range of settings, including budget allocation (Naldi et al. 2016; Nicosia, Pacifici, and Pferschy 2017), machine scheduling (Bilò et al. 2014), and kidney exchange (Dickerson, Procaccia, and Sandholm 2014; McElfresh and Dickerson 2018). Some recent work focused on PoF of allocation algorithms when demand is uncertain (Elzayn et al. 2019; Donahue and Kleinberg 2020). Most closely related to our result is the work of Bei et al. (2019), who established the upper bound for the price of leximin under indivisible goods, assuming each agent's maximum utility, achieved when all goods are assigned to them, is 1 (normalized utilities). Our work presents the analogous bound for *unit-demand* leximin from the cost-centric perspective without this assumption of normalized utilities/costs. Typically considered a function of the number of agents (Bertsimas, Farias, and Trichakis 2011; Caragiannis et al. 2012; Bei et al. 2019; Suksompong 2019; Bouveret, Chevaleyre, and Maudet 2016), PoF may also be analyzed with respect to the number of resources k. For example, Kurz (2016) showed that the price of envy-freeness is much lower if k is small. Our analyses exhibit similar implications about both leximin and LoINC.

Other lines of research focus on the asymptotic cost of various assignment objectives when costs are drawn from a probability distribution and the number of agents $n \to \infty$. Perhaps the most famous result is the limit of $\pi^2/6$ of the expected minimum total cost under uniformly distributed costs, proven by Linusson and Wästlund (2004) and Nair, Prabhakar, and Sharma (2005) independently. More generally, under mild assumptions about the cost-generating distribution, Olin (1992) showed that this minimum total cost is bounded by a constant dependent on the distribution. Results relevant to the asymptotic cost of leximin are limited to the expectation of the bottleneck (the minimized largest cost any agent receives) under the assignment. Most notably, Pferschy (1995) showed that this bottleneck asymptotically tends to zero, and Spivey (2011) computed the asymptotic moments

of the bottleneck under various distribution families. To our knowledge, no result on the asymptotic cost of leximin has been established.

Fair resource allocation within the context of homelessness services has been studied from various perspectives (Fitzpatrick and Pleace 2012; Azizi et al. 2018; Kube, Das, and Fowler 2019). Most related to our work, Kube, Das, and Fowler (2019) formulated the assignment problem as an integer program with a fairness constraint on the maximum increase in cost an agent incurs when compared to a baseline assignment. This maximum-increase constraint was an adjustable parameter of the integer program and efficiency loss was studied as a function of this parameter.

Others have additionally studied allocation objectives that balance efficiency and fairness. Specifically, Hooker and Williams (2012) proposed a hybrid objective that optimizes for leximin at the beginning of the allocation computation and switches to efficiency when fairness becomes costly. McElfresh and Dickerson (2018) applied this approach to the kidney exchange setting, whereas Chen and Hooker (2020) extended it to incorporate a stronger notion of fairness. This objective requires a parameter to specify the efficiency-fairness trade-off, but a principled choice of this parameter has not been explored.

2 Problem statement and preliminaries

An assignment problem $\langle [n], [k], C \rangle$ is defined by a set of $n \in \mathbb{N}$ agents $[n] = \{1, 2, ..., n\}$, a set of $k \in \mathbb{N}$ resources $[k] = \{1, 2, ..., k\}$, and a matrix $C = (c_{i,j}) \in \mathbb{R}^{n \times k}$ denoting the cost of assigning agent i to resource j. We will also refer to the k resources as *interventions*. We further assume that $c_{i,j} \in [0, 1]$, which can be made true via various normalization techniques. In many contexts, we assume that $c_{i,j}$ are i.i.d. samples from a probability distribution to synthetically simulate cost matrices in our experiments and study asymptotic behavior.

Assignment. We consider assignments of unit demands (i.e., each agent is assigned to exactly one intervention). A member $A = (a_1, a_2, ..., a_n)$ of the set of valid assignments M is a vector of length n, where each element $a_i \in [k]$ denotes the intervention agent i is assigned to under A. We assume that each intervention $j \in [k]$ has a capacity limit $u_j \in \mathbb{N}$, specifying the maximum number of agents that may be assigned to it. To ensure each agent is assigned to an intervention, we assume $\sum_{j \in [k]} u_j = n$.

One of the most well-studied assignment objectives is social welfare or efficiency, which produces the assignment minimizing the sum of agents' assigned costs. We call this the *efficient assignment*, denoted as E. Defining $\mathrm{C}(A)$ as the total cost of an assignment A, we have $\mathrm{E} = \arg\min_{A \in M} \mathrm{C}(A)$, which can be efficiently computed using the Kuhn-Munkres algorithm (Kuhn 2005; Munkres 1957).

In contrast, the *leximin* objective minimizes individual agents' costs in lexicographic order: it first minimizes the largest cost any agent incurs, then, subject to that constraint, it minimizes the second largest cost, and so on; among others, Bogomolnaia and Moulin (2004) and Kurokawa, Procaccia, and Shah (2015) present the mathematical formulation of the

leximin objective. Arising from a Rawlsian notion of fairness that prioritizes the most disadvantaged (agents with largest costs) (Rawls 1971, 1993), leximin is a popular fair assignment objective with various desiderata such as envy-freeness up to any good and Pareto-optimality (Plaut and Roughgarden 2020). The leximin-optimal assignment is denoted as L and can be computed in polynomial time under our unit-demand assignment model (Sokkalingam and Aneja 1998). Efficient computations of leximin have also been studied within other specialized contexts such as constraint networks (Bouveret and Lemaître 2009) or under dichotomous preferences (Bogomolnaia and Moulin 2004; Kurokawa, Procaccia, and Shah 2015).

We also introduce LoINC, a modification of leximin where the *increases in cost* each agent incurs, relative to their individual lowest costs, are considered. In particular, this objective computes the leximin assignment on matrix $C' = (c'_{i,j})$, the *normalized* version of C where $c'_{i,j} = c_{i,j} - \min_j c_{i,j}$ (each element subtracted by its row minimum). Intuitively, LoINC prioritizes those who benefit most from the interventions by minimizing the largest degradation from agents' individual best allocations. These normalized costs can be further interpreted as individual agents' *regrets* with respect to their own best allocations, in which case LoINC lexicographically minimizes these regrets. LoINC also preserves the mentioned fairness-related characteristics of leximin but with respect to the normalized costs, not the actual costs themselves.

Price of fairness. A natural question to ask regarding a fair assignment objective that might be orthogonal to efficiency is how much additional cost is incurred. We adopt the PoF measure from Bertsimas, Farias, and Trichakis (2011) and Caragiannis et al. (2012) for a given fair assignment $A \in M$ as the ratio between the cost of A and the minimized cost of the efficient assignment E under a common cost matrix:

$$\text{Price of fairness of } A = \text{PoF}(A) = \frac{\text{C}(A)}{\text{C}(E)}.$$

If $\text{C}(E) > 0$, $\text{PoF}(A) \geq 1$. If $\text{C}(E) = 0$ under a given cost matrix, E assigns zero cost to all agents, and both leximin and LoINC coincide with this assignment; as neither leximin incurs any additional cost, we define PoF to be 1 in this case. So, $\text{PoF}(A) \geq 1, \forall A \in M$ under any cost matrix P. In this work, we study the behavior of $\text{PoF}(L)$ and $\text{PoF}(\text{LoINC})$ within different contexts that define the cost matrix. Intuitively, LoINC de-prioritizes agents who will incur a high cost regardless of which intervention they are assigned to; as we will show, this behavior typically keeps PoF relatively lower than under leximin, while enforcing a commonly used notion of fairness.

3 Price of fairness upper bounds

We first establish worst-case upper bounds for these two quantities as follows:

Proposition 1. $\text{PoF}(L) \leq n; \quad \text{PoF}(\text{LoINC}) \leq n - 1$.

We refer readers to the supplementary material for the proof. We note that the first result is analogous to the $\mathcal{O}(n)$

bound established by Bei et al. (2019) for leximin allocations of indivisible goods with *normalized utilities* (where the maximum utility of each agent when they receive *all* goods equals 1). In our model where each agent is to be assigned to exactly one intervention and subsequently incurs a cost, we omit this normalization assumption but nonetheless obtain the same PoF bound. Moreover, these bounds can be shown to be essentially tight (see the supplementary material); for example, Equation $(*)$ denotes an $n \times n$ cost matrix C^* where $\text{PoF}(L) \to n$ as $\varepsilon \to 0^+$.

$$C^* = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & 1-\varepsilon \\ 0 & 0 & 0 & \cdots & 1-\varepsilon & 1 \\ 0 & 0 & 0 & \cdots & 1 & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 1-\varepsilon & 1 & \cdots & 1 & 1 \\ 1-\varepsilon & 1 & 1 & \cdots & 1 & 1 \end{bmatrix} \quad (*)$$

Under C^*, $\text{E} = (n-1, n-2, ..., 2, 1, n)$ assigns agent n a cost of 1 and the other agents costs of 0. On the other hand, to minimize the largest cost of any agent (which is at least $1-\varepsilon$), $\text{L} = (n, n-1, ..., 1)$ assigns all agents costs of $1-\varepsilon$, raising its total cost to $n(1-\varepsilon)$. In this case, leximin sacrifices agents who would otherwise enjoy zero cost to lower agent n's cost by ε. This is often described as the *bottomless pit problem* (considerable amounts of a limited resource spent on marginal improvements) (Veatch 1991; Brock and Wikler 2006). However, our later analyses indicate that this situation rarely arises under mild assumptions.

4 Row-sorted cost matrices

We now turn to considering cost models that reflect reality in the space of scarce societal resource allocation. Specifically, we assume that within the cost matrix that defines an assignment problem, elements in each row are always in the same order. This reflects the idea that interventions can be sorted by intensity, and everyone benefits more from more intensive interventions. The characteristic of being roughly ordered by effectiveness, albeit a strong assumption, is present in resources of many applications: in the homelessness intervention data, studied by Kube, Das, and Fowler (2019) and examined in our later experiments, the four services (homelessness prevention, emergency shelter, rapid rehousing, and transitional housing) provide different support levels, increasing in intensity and cost, with their benefit presumably increasing in the same order (although in practice there is some heterogeneity – see discussion in Section 5); in Covid patient care, resources such as home care, hospital beds, and ICU beds have monotonically increasing value in helping patients survive.

Under the assumption, we rearrange the matrix columns in the order of increasing cost such that for every agent, the first intervention is better (has lower cost) than the second, which is better than the third, and so on. This arrangement, formalized below, is assumed for the rest of this section.

$$c_{i,1} \leq c_{i,2} \leq \cdots \leq c_{i,k}, \forall i \in [n].$$

We refer to such matrices as *row-sorted,* and note that if a matrix C is row-sorted, then so is its normalized form C',

as normalization preserves the order of the elements in each row. The matrices used to show the tightness of the PoF bounds in Proposition 1 are row-sorted themselves (see the supplementary material), so this assumption alone does not diminish the worst-case PoF. However, it allows us to derive more expressive bounds that will in fact have lower values in most cases. To do this, we consider the following proposition.

Proposition 2. Under leximin (resp. LoINC), the c_k agents with the lowest costs (resp. normalized costs) in intervention k, the last and least effective intervention, are assigned to that intervention.

We again refer readers to the supplementary material for the proof. The key idea is that, if one of the agents with the lowest costs in intervention k was not assigned to the intervention, there would exist a pair of agents whose assignments violate the leximin objective.

Efficient leximin assignment. A direct corollary of Proposition 2 is a more efficient method of computing the leximin assignment on a row-sorted cost matrix. In particular, we first assign the agents with the lowest costs in intervention k to that intervention. Proposition 2 then applies for the remaining agents with the lowest costs in intervention $k - 1$, and subsequently for the remaining agents in intervention $k - 2$, and so on. In short, leximin is equivalent to the assignment that greedily fills each intervention to capacity with agents having the corresponding lowest costs, iterating from the last to the first intervention. This procedure is described in Algorithm 1, where the returned value is the leximin assignment on a row-sorted matrix C. To compute LoINC, the normalized matrix C' is used as input.

Algorithm 1 Efficient leximin

Input: matrix C with sorted rows
Output: leximin assignment on P
1: Initialize $A = (a_1, a_2, ..., a_n)$ as $(0, 0, ..., 0)$.
2: **for** $j = k, ..., 1$ **do**
3:     **while** $u_j > 0$ **do**
4:         $i^* = \arg\min_{i, a_i = 0} c_{i,j}$
5:         $a_{i^*} \leftarrow j$
6:         $u_j \leftarrow u_j - 1$
7:     **end while**
8: **end for**
9: **return** A

At each iteration $j \in [k]$, the u_j remaining agents with the lowest costs in intervention j can be simultaneously found in $\mathcal{O}(n)$ time in the worst case by the selection algorithm. The entire procedure thus has an $\mathcal{O}(nk)$, or linear, running time complexity. Compared with the optimized $\mathcal{O}(n^2 k^2)$ running time of a generic unit-demand leximin algorithm, established by Sokkalingam and Aneja (1998), this algorithm enjoys a significantly lower running time.

Remark. We assume that there are no elements of equal value in the input cost matrix to exclude the need for tie-breaking during the identification of agents having the lowest costs in individual interventions. With the exception of

number instability, the event that values i.i.d. drawn from a probability distribution without atoms are exactly equal *almost surely* does not occur.

Improved price of fairness upper bounds. By selecting the agents with the lowest costs, L (resp. LoINC) minimizes the total cost (resp. normalized cost) incurred by agents assigned to intervention k. This allows us to refine the upper bounds for the PoF of the two assignments as follows.

Proposition 3.

$$\mathrm{PoF(L)} \le 1 + \frac{b_u}{\underline{b}} \cdot \frac{n - u_k}{u_k};$$

$$\mathrm{PoF(LoINC)} \le 1 + \frac{b_n}{\underline{b}} \cdot \frac{n - u_k - u_1}{u_k}.$$

Here u_1 and u_k are the capacities of the first and last interventions, b_u (resp. b_n) is the largest cost (resp. normalized cost) of any agent, minimized under L (resp. LoINC). $\underline{b}$ is the lowest cost in intervention k. The proof for these bounds can be found in the supplementary material. Further, if all k interventions have equal capacity ($u_1 = \cdots = u_k = n/k$), the bounds become:

$$\mathrm{PoF(L)} \le 1 + \frac{b_u}{\underline{b}}(k-1); \quad \mathrm{PoF(LoINC)} \le 1 + \frac{b_n}{\underline{b}}(k-2).$$

When costs are drawn from a probability distribution, $(b_u/\underline{b})$ and $(b_n/\underline{b})$ are random variables dependent on that distribution and n. Analyzing the induced distributions of these variables is quite challenging, but the dependence on the number of interventions k suggests that these new bounds are significantly lower than the worst-case bounds when k is relatively small, which is a natural assumption to make in real-life scenarios. In the extreme case where $k = n$, the two bounds approach the respective worst-case bounds. Moreover, the bound for $\mathrm{PoF(L)}$ is at least one unit larger than that for $\mathrm{PoF(LoINC)}$, as $b_u \ge b_n$; in practice, we expect this difference to be substantially larger. In later experiments, we will observe empirical values of these bounds under randomly generated costs.

Remark. Substituting $k = 2$, we have $\mathrm{PoF(LoINC)} \le 1$. Since $\mathrm{PoF(LoINC)}$ is defined to be at least 1, this shows that when there are only two interventions in the assignment problem, LoINC incurs no additional cost and is equivalent to the efficient assignment E. That LoINC coincides with E when $k = 2$ can in fact be proven even without the assumption of sorted rows.

Proposition 4. $\mathrm{PoF(LoINC)} = 1$ if $k = 2$.

The proof involves constructing a specific assignment under any two-column cost matrix and showing that it optimizes both efficiency and the LoINC objective, and is included in the supplementary material. It is also worth noting that this lack of additional cost when $k = 2$ does not hold for traditional leximin.

5 Numerical experiments

Homelessness reentry probabilities. We consider the homelessness reentry probability dataset, introduced by Kube,
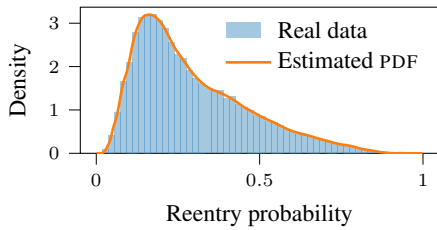
Figure 1: Homelessness reentry probability distribution

Table 1: Assignment results on the homelessness reentry probability dataset

Criterion	E	L	LoINC
Total cost	3682	3774	3692
PoF	1	1.025	1.003
Largest cost	0.930	0.897	0.930
Largest cost increase	0.080	0.183	0.075

Das, and Fowler (2019), which includes information about $n = 13940$ households and $k = 4$ homelessness interventions. Each value $c_{i,j} \in [0,1]$ in the dataset denotes the estimated probability household i will reenter the homelessness services system if assigned to intervention j. Figure 1 shows the histogram distribution of these probabilities and the probability density function (PDF) of a kernel density estimate (KDE) of this distribution (the supplementary material includes more detail on the KDE). The dataset also specifies the capacities of the four interventions. As reentry is a negative outcome, this data is directly used as the cost matrix that defines the assignment problem, where the cost of an assignment is the expected number of reentering households given that assignment. This matrix is not row-sorted, but not all possible orderings of intervention effectiveness are equally likely: some orderings are not present; the most common ordering applies to 38.0 percent of the households, while the second most common applies to another 31.1 percent (Kube, Das, and Fowler 2019).

Table 1 summarizes the result of the efficient and leximin assignments with respect to various cost- and fairness-related criteria. While having a low relative PoF, L incurs an additional cost of 92 households estimated to reenter the homelessness services system compared with E. In the context of scarce resources where costs have serious implications such as reentry into homelessness, this additional cost may be considered inefficient. Conversely, LoINC enjoys a lower PoF with only 10 additional households reentering the system. In terms of fairness, L minimizes b_u, the largest reentry probability for all households, to be 89.7 percent. However, this is arguably not a significant improvement from the same statistic in E and LoINC, which are coincidentally both 93 percent. On the other hand, LoINC minimizes the largest probability increase b_n to be 7.5 percent, whereas it is 18.3 percent under L.

The reason b_u under L is not significantly lower than that under E or LoINC is the existence of agents whose lowest probabilities are already high, who will inevitably incur large costs regardless of the intervention they are assigned to. As a result, even their minimized costs under L remain relatively large. These agents correspond to rows in the cost matrix with large minimum values. The large values of these lowest costs could also cause the normalization of the matrix to drastically change the ordering among its elements, partially contributing to the difference between L and LoINC via their inputs. The effect of these high-cost agents on the distinction between L and LoINC is further explored in later

experiments.

Empirical price of fairness with synthetic data. We draw i.i.d. samples from Beta distributions to generate random cost matrices as instances of the assignment problem. By adjusting the parameters α and β, we can simulate various distribution shapes from which costs $c_{i,j} \in [0,1]$ are drawn. Samples are also drawn from the KDE of the distribution of the homelessness reentry probabilities. This allows us to observe empirical PoF values when the costs follow the same distribution as those in the homelessness reentry dataset, and compare them with the results from the experiments with the Beta distributions.

We observe the growth of PoF with respect to the number of agents n. The top portion of Figure 2 shows the average, the 25th and 75th percentiles of empirical PoF(L) and PoF(LoINC) as functions of $n \in \{30, 40, ..., 90, 100\}$. Each subplot corresponds to the distribution from which individual costs were drawn: $U(0,1)$ (or $\text{Beta}(1,1)$), $\text{Beta}(0.5, 0.5)$, and the KDE of homelessness reentry probabilities. For each combination of the cost distribution and n, we ran 500 experiments with $k = 5$ interventions of equal capacity, and recorded the resulting PoF values. As expected, these values fall significantly below the worst-case bounds. Moreover, PoF(L) tends to be higher than PoF(LoINC), similar to the result from the homelessness reentry dataset. It is also interesting to note that all lines have a downward trend as n increases.

We are also interested in the PoF bounds that depend on the number of interventions k, established in Proposition 3 under the assumption of sorted rows. We repeated the experiments above; this time, copies of the cost matrices were made, sorted row-wise, and used as input to the assignment problem. Empirical values of the two upper bounds were then recorded. Our initial intuition was that these k-dependent bounds are lower than the linear worst-case bounds, especially when k is small. This is confirmed by the bottom portion of Figure 2, which shows the growth of the average, the 25th and 75th percentiles of these values as functions of n, under the same three cost-generating distributions.

We finally consider the effect of row-sorted cost matrices on the behavior of PoF by observing PoF(L) under a random matrix and its row-sorted version in each experiment above. Figure 3 includes KDE plots of the distributions of PoF(L) across all experiments: the left shows the distribution of PoF(L) when matrices were row-sorted and that when they were not; the right shows the distribution of the one-to-one difference between the former and the latter in
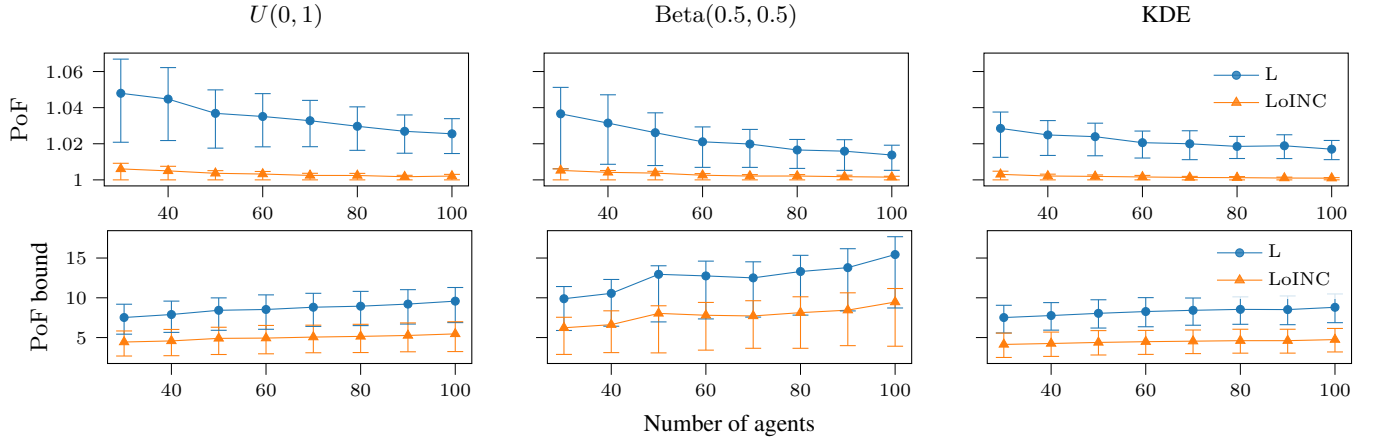
Figure 2: Empirical PoF (top) and improved upper bounds under row-sorted cost matrices (bottom) as functions of n.

each experiment. Intuitively, under a random matrix, there is no intervention that universally leads to bad outcomes, and L has more freedom to minimize the costs of its agents; under a row-sorted matrix, c_k agents are assigned to the least effective intervention and incur high costs. A random matrix thus presents less "competition" among its agents than a row-sorted one (equivalently, more agents can be satisfied when their preferences are more diverse, as has been established in, for example, two-sided matching (Anshelevich, Das, and Naamad 2013)), and likely yields a lower PoF(L). (However, this is not always the case: the one-to-one difference in PoF(L) is negative when C(E) increases by a larger portion than C(L) under the row-sorted matrix; see the supplementary material for more discussion.) In these experiments, PoF(LoINC) also exhibits the same trend, albeit with lower values.

Overall, the conducted experiments offer empirical evidence indicating that the prices of leximin and of LoINC tend to be much lower than their worst-case bounds. Moreover, if not row-sorted, a cost matrix is more likely to yield a low PoF, further suggesting that the linear worst-case bounds are rarely achieved when costs are i.i.d. samples of a probability distribution.

Price of fairness under Beta distributed costs. We observe empirical PoF(L) and PoF(LoINC) under an extensive set of cost-generating Beta distributions. Each cell of each heat map in Figure 4 denotes the average value of PoF(L) (left) and PoF(LoINC) (right) across 1000 experiments where costs are drawn from a Beta distribution, whose parameters are denoted in the row and column labels. Individual rows in each matrix were sorted before the assignments were computed. We fixed $n = 30$, $k = 5$, and $u_j = 6, \forall j \in [5]$. In both cases, PoF tends to increase as α and β grow smaller. Moreover, under any distribution, the average price of leximin is greater than that of LoINC, and the difference is more pronounced when $\beta < \alpha < 1$. This distinction relates back to high-cost agents. Specifically, Beta distributions with $\beta < \alpha < 1$ have more probability mass around areas close to 1, and thus tend to generate agents with high-valued lowest costs: an agent's lowest cost is the

first order statistic of k samples from the generating distribution; when the distribution has more probability mass around 1, this first order statistic tends to be large. The rows whose smallest elements are large correspond to high-cost agents, an indication of a large difference between PoF(L) and PoF(LoINC).

It is interesting to note in the first heat map that for each value of β, the average PoF(L) peaks at a certain value of α and subsequently decreases as α increases. This situation is quite similar to what we have observed from the effect of row-sorted matrices: as α increases, so do both C(E) and C(L); at first, C(L) increases faster than C(E), leading to higher PoF, but at a certain point, the latter starts to increase faster, thus lowering the resulting PoF. More detailed discussions are included in the supplementary material.

Comparing the two leximin assignments. The effect of high-cost agents is made most tangible by matrix C^* in (∗). Again, under C^*, L minimizes the bottleneck cost by assigning agent n to the most effective intervention, incurring the cost of $n(1 - \varepsilon)$. However, the normalized cost array of agent n is $(0, \varepsilon, \varepsilon, ..., \varepsilon)$, whose low values de-prioritize the agent under LoINC, which in this case coincides with E and $C(\text{LoINC}) = 1$. By prioritizing agent n, a high-cost agent whose lowest cost is as high as $1 - \varepsilon$, L sacrifices its most effective intervention and consequently raises its PoF to essentially n. On the other hand, by focusing on the normalized costs, LoINC does not exhibit this behavior. This is the core difference between the two notions of leximin, highlighted by the presence of these high-cost agents.

Overall, our analyses suggest that PoF(L) tends to be higher than PoF(LoINC): the former has larger theoretical upper bounds as well as empirical values; LoINC even maximizes efficiency under any $n \times 2$ cost matrix (Proposition 4). However, LoINC still enforces a meaningful fairness notion by minimizing an agent's cost increase with respect to their own costs. Agents with low-valued lowest costs are prioritized similarly by L and LoINC. A high-cost agent, on the other hand, is less prioritized by LoINC than by L; LoINC recognizes the agent's inevitably bad outcome and thus tends
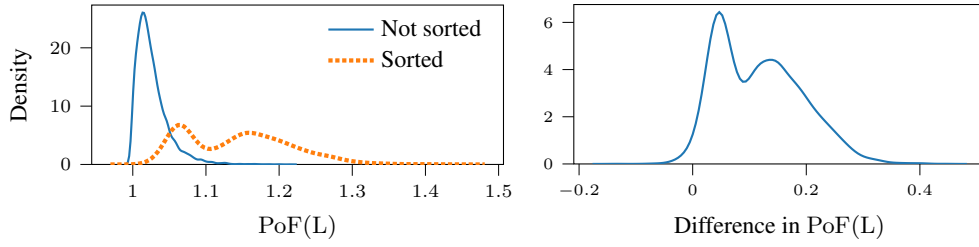
Figure 3: Distributions of empirical PoF(L) under random and row-sorted cost matrices (left) and their one-to-one differences (right).
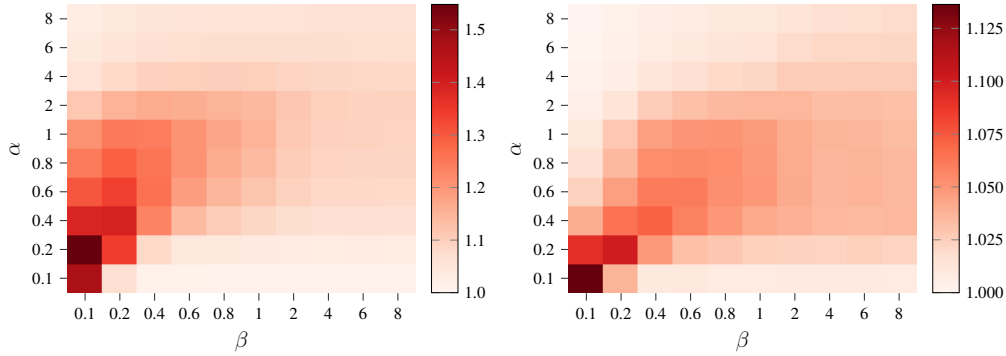

Figure 4: Average PoF(L) (left) and PoF(LoINC) (right) under different Beta distributions. Note that the color range differs between the plots; the PoF values have a smaller range in the second case.

to reserve effective interventions for those that benefit more from them. As such, LoINC is less likely to encounter the bottomless pit problem. We argue that LoINC naturally balances between a Rawlsian fairness notion and efficiency by not blindly prioritizing agents with inevitably bad outcomes.

Leximin under bimodal cost distributions. We noted from Figure 4 that as α and β grow closer to 0 (i.e., when the generating Beta distribution becomes increasingly bimodal), empirical PoF tends to increase. We hypothesize that this is because a matrix generated from such a distribution contains costs very close to 1 (matrix C^* in (∗) is an example where costs are close to 1, taking on values of $1 - \varepsilon$), and leximin chooses among them the lowest-valued to assign to its agents. This process might assign effective interventions to agents with inevitably high costs, thus lowering the overall efficiency. In contrast, the efficient assignment can save effective interventions for its other agents and enjoy a lower cost. So, the bimodality of the cost-generating distribution may lead to a high expected PoF.

Interestingly, if we take the Beta distribution to the limit where $\alpha \to 0^+$ and $\beta \to 0^+$ (at potentially different rates) and consider a Bernoulli distribution for the costs, the behavior of PoF sharply changes. Each Bernoulli cost matrix contains only zeros and ones; under such a matrix, E and both leximin assignments coincide, as all three aim to minimize the number of agents assigned to cost 1. The price of either leximin equals 1 in this case, even without the assumption of sorted rows. This change of behavior of PoF is due to the fact that under a Bernoulli matrix, the ordering among

the costs (which are now simply zeros and ones) becomes binary. The deliberation of choosing the minimum out of the costs that are very close to 1 by leximin described above is no longer necessary, since all large costs equal 1 and are equivalent from the leximin perspective. A high-cost agent in this case has a cost vector of $(1, 1, ..., 1)$, and both leximin assignments are indifferent among these costs.

6 Conclusion

We show that while having a linear worst-case bound, the price of leximin fairness under unit-demand assignment problems takes on significantly lower empirical values. By assuming homogeneity among the agents' preferences, we are able to prove a tighter bound when the number of interventions is small. Moreover, we propose a novel objective (LoINC) inspired by the local justice principle of prioritizing those who would be helped most by receiving a resource, and show that PoF(LoINC) has lower-valued upper bounds as well as lower empirical values compared with traditional leximin. We finally characterize settings where the two fair objectives are costly.

Ethics statement

This work studies the cost of applying Rawlsian notions of fairness to the allocation of scarce societal resources, and presents an optimistic message regarding the efficiency of the allocation. Further, we introduce a novel formal formulation of an allocation objective (LoINC) that corresponds to a different common notion of local justice and show that

under many circumstances allocations corresponding to this objective are closer to the efficient allocation than the leximin allocation, one operationalization of the Rawlsian notion. It is worth noting that the LoINC allocation can suffer from the problem of abandonment, "giving up" on individuals who are unlikely to benefit from the allocation of resources in a way the leximin allocation does not. The characterization of settings where both of these fair allocations suffer from significant efficiency loss that we establish can also prove useful in practical situations. Perhaps the most major caution to be aware of as research in this area begins to be taken to practice is that, regardless of which notion of local justice is used (which is usually determined by a social process), any valuation-based resource allocation algorithm is subject to the efficacy of the methods used in determining those valuations. For example, measuring health-related benefits/costs in the health care setting is generally difficult, and may involve balancing multiple factors such as accuracy, intrusiveness, and monetary cost (Brock and Wikler 2006).

Acknowledgments

References

Anshelevich, E.; Das, S.; and Naamad, Y. 2013. Anarchy, Stability, and Utopia: Creating Better Matchings. *Autonomous Agents and Multi-Agent Systems* 26(1): 120–140.

Azizi, M. J.; Vayanos, P.; Wilder, B.; Rice, E.; and Tambe, M. 2018. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 35–51. Springer.

Bei, X.; Lu, X.; Manurangsi, P.; and Suksompong, W. 2019. The Price of Fairness for Indivisible Goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 81–87.

Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2011. The price of fairness. *Operations Research* 59(1): 17–31.

Bilò, V.; Fanelli, A.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2014. The price of envy-freeness in machine scheduling. In *International Symposium on Mathematical Foundations of Computer Science*, 106–117. Springer.

Bogomolnaia, A.; and Moulin, H. 2004. Random matching under dichotomous preferences. *Econometrica* 72(1): 257–279.

Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair Allocation of Indivisible Goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press.

Bouveret, S.; and Lemaître, M. 2009. Computing leximin-optimal solutions in constraint networks. *Artificial Intelligence* 173(2): 343–364.

Brock, D.; and Wikler, D. 2006. Ethical issues in resource allocation, research, and new product development. In Jamison, D. T.; et al., eds., *Disease Control Priorities in Developing Countries*, volume 2, 259–60. Oxford University Press and the World Bank New York.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2012. The efficiency of fair division. *Theory of Computing Systems* 50(4): 589–610.

Chen, V.; and Hooker, J. 2020. A Just Approach Balancing Rawlsian Leximax Fairness and Utilitarianism. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 221–227.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2014. Price of fairness in kidney exchange. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, 1013–1020.

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

Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 214–226.

Elster, J. 1992. *Local Justice: How Institutions Allocate Scarce Goods and Necessary Burdens*. Russell Sage Foundation.

Elzayn, H.; Jabbari, S.; Jung, C.; Kearns, M.; Neel, S.; Roth, A.; and Schutzman, Z. 2019. Fair algorithms for learning in allocation problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 170–179.

Fitzpatrick, S.; and Pleace, N. 2012. The statutory homelessness system in England: A fair and effective rights-based model? *Housing Studies* 27(2): 232–251.

Hooker, J. N.; and Williams, H. P. 2012. Combining equity and utilitarianism in a mathematical programming model. *Management Science* 58(9): 1682–1693.

Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *International Conference on Machine Learning*, 2564–2572.

Kube, A.; Das, S.; and Fowler, P. J. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 622–629.

Kuhn, H. W. 2005. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)* 52(1): 7–21.

Kurokawa, D.; Procaccia, A. D.; and Shah, N. 2015. Leximin allocations in the real world. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 345–362.

Kurz, S. 2016. The price of fairness for a small number of indivisible items. In *Operations Research Proceedings 2014*, 335–340. Springer.

Linusson, S.; and Wästlund, J. 2004. A proof of Parisi's conjecture on the random assignment problem. *Probability Theory and Related Fields* 128(3): 419–440.

McElfresh, D. C.; and Dickerson, J. P. 2018. Balancing lexicographic fairness and a utilitarian objective with application to kidney exchange. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 1161–1168.

Munkres, J. 1957. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics* 5(1): 32–38.

Nair, C.; Prabhakar, B.; and Sharma, M. 2005. Proofs of the Parisi and Coppersmith-Sorkin random assignment conjectures. *Random Structures & Algorithms* 27(4): 413–444.

Naldi, M.; Nicosia, G.; Pacifici, A.; and Pferschy, U. 2016. Maximin fairness in project budget allocation. *Electronic Notes in Discrete Mathematics* 55: 65–68.

Nicosia, G.; Pacifici, A.; and Pferschy, U. 2017. Price of fairness for allocating a bounded resource. *European Journal of Operational Research* 257(3): 933–943.

Olin, B. 1992. *Asymptotic properties of random assignment problems*. Ph.D. thesis, Kungliga Tekniska Högskolan, Stockholm, Sweden.

Pferschy, U. 1995. The random linear bottleneck assignment problem. In *International Conference on Integer Programming and Combinatorial Optimization*, 145–156. Springer.

Plaut, B.; and Roughgarden, T. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics* 34(2): 1039–1068.

Pleiss, G.; Raghavan, M.; Wu, F.; Kleinberg, J.; and Weinberger, K. Q. 2017. On fairness and calibration. In *Advances in Neural Information Processing Systems*, 5680–5689.

Rawls, J. 1971. *A Theory of Justice*. Harvard University Press.

Rawls, J. 1993. *Political Liberalism*. Columbia University Press.

Sokkalingam, P.; and Aneja, Y. 1998. Lexicographic bottleneck combinatorial problems. *Operations Research Letters* 23(1-2): 27–33.

Spivey, M. Z. 2011. Asymptotic moments of the bottleneck assignment problem. *Mathematics of Operations Research* 36(2): 205–226.

Suksompong, W. 2019. Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics* 260: 227–236.

Veatch, R. M. 1991. Equality, justice, and rightness in allocating health care: A response to James Childress. In Kogan, B. S., ed., *A Time to Be Born and a Time to Die: The Ethics of Choice*, 205. de Gruyter Hawthorne, NY.