The price of imposing vertical equity through asymmetric outcome constraints

Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

The Price of Imposing Vertical Equity Through Asymmetric Outcome Constraints

Thomas Breugem, Luk N. Van Wassenhove

To cite this article:
Thomas Breugem, Luk N. Van Wassenhove (2022) The Price of Imposing Vertical Equity Through Asymmetric Outcome Constraints. Management Science

Published online in Articles in Advance 20 Jan 2022

. https://doi.org/10.1287/mnsc.2021.4287

http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–17
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

The Price of Imposing Vertical Equity Through Asymmetric Outcome Constraints

Thomas Breugem,[a] Luk N. Van Wassenhove[a]

[a] Technology and Operations Management, INSEAD, Fontainebleau 77300, France
Contact: thomas.breugem@insead.edu, https://orcid.org/0000-0002-1740-7697 (TB); luk.van-wassenhove@insead.edu, https://orcid.org/0000-0002-5573-1834 (LNVW)

Received: December 23, 2020
Revised: June 14, 2021
Accepted: September 15, 2021
Published Online in Articles in Advance: January 20, 2022

https://doi.org/10.1287/mnsc.2021.4287

Abstract. Vertical equity or fairness is an important aspect in many settings, yet has received relatively little attention in the literature. Recent developments underline the practical relevance (e.g., COVID-19 vaccination policies). It plays an important role in the performance evaluation of many (nongovernmental) organizations. For example, donors might require a family-planning organization to allocate a minimum fraction of the total utility (client volume) to a particular player (the "high-impact" subgroup of the population, e.g., young and poor clients). However, the price (decrease in client volume) of such requirements is not well-understood. Consequently, this price is not accounted for in decision making. We provide an analytical upper bound on the price (i.e., loss of overall utility) of vertical equity considerations in resource allocation. We assume that these concerns are expressed via outcome constraints, specifying a minimum percentage of the total utility for each player. Our set-up considers a decision maker maximizing total utility over a general convex set, subject to outcome constraints. The set-up is general and applicable to many practical problems. Our results depend only on high-level parameters and are therefore well-suited for strategic decision making. We conclude with two applications. First, we apply our results to practical instances in health delivery. We confirm that outcome constraints can entail a substantial price and analyze the factors driving this price close to the worst-case bound. Second, we analyze how our results can help bound the impact of prioritization in vaccine allocation.

History: Accepted by Jeannette Song, operations management.
Supplemental Material: Data files are available at https://doi.org/10.1287/mnsc.2021.4287.

Keywords: equity • fairness • health-delivery optimization • outcome constraints • resource allocation

1. Introduction

Resource allocation arises in many settings, including in healthcare and humanitarian operations (e.g., allocating scarce resources to patients or disaster-stricken populations). The trade-off between efficiency (maximizing overall utility) and fairness or equity (balancing the utility of the players) plays a central role in the allocation. Equity can be either horizontal (equal treatment of all players) or vertical (varying treatment among players). Vertical equity is an important aspect in many practical problems, yet has received comparatively little attention in literature. In this research, we aim to fill this gap and study the price (i.e., loss of utility) of vertical equity. We assume that vertical equity is imposed through outcome constraints, specifying a minimum percentage of the total utility for each player. For example, outcome constraints would require that a player receives at least 20% of the total utility, independent of whether the total utility is 100 or 1,000. Asymmetry in outcome constraints allows one to differentiate between players and ensure that resources go where most desired.

Efficiency and equity are often conflicting objectives. For instance, reaching poor rural areas with health services (vertical equity) may be time-consuming and substantially reduce the total volume of clients served (efficiency). Hence, imposing equity considerations has a very real impact on operations, and understanding this impact is fundamental to good decision making. Insights into the price of vertical equity help estimate the loss of efficiency due to prioritizing subgroups of recipients and provide a way to identify misalignments among performance measures (and the price associated with these misalignments).

Recent developments highlight the relevance of such insights. For example, prioritization of those most in need plays an important role in COVID-19 vaccination strategies.[1] Facing (severe) scarcity of vaccines, those most in need (e.g., the elderly) and/or those most pivotal in the response (e.g., health workers) are prioritized, meaning that they get the vaccine before the rest of the population. Other policies are also possible. For example, one might want to enlarge the population eligible for the first round of vaccination

(e.g., motivated by equity or social preferences). Prioritization policies can be seen as outcome constraints because they specify how utility (allocated vaccines or the benefit of these vaccines) should be distributed over the players (population segments). The price of vertical equity in such cases helps understand how costly (in terms of loss of efficiency/impact) different policies can be.

The motivation for this work is rooted in health-delivery optimization, in particular, access to family-planning services. Family planning is defined by the World Health Organization as "[allowing] individuals and couples to anticipate and attain their desired number of children and the spacing and timing of their births."[2] It plays an important role in increasing the well-being of women all around the globe. According to The United Nations' sexual and reproductive health agency, UNFPA, "family planning is central to gender equality and women's empowerment, and it is a key factor in reducing poverty."[3] Across family-planning organizations and initiatives, there is a clear consensus on providing services to those "most in need." The UNFPA, for example, states that "[p]oorer women and those in rural areas often have less access to family planning services. […] Particular attention must be paid to promoting their reproductive rights, access to family planning, and other sexual and reproductive health services."[4] Similarly, the Bill & Melinda Gates Foundation states that the foundation support "includes assessing family planning needs, particularly among the poorest and most vulnerable populations."[5]

The focus on marginalized groups is visible in the operations and reporting of family-planning organizations. Marie Stopes International, a nongovernmental organization operating in 37 countries, states in their 2019 impact report that "70% of all clients served [are] living in communities with the greatest need: adolescents for whom access is invariably heavily restricted; those using a method of contraception for the first time; and those living in extreme poverty, where often neither the public nor private sector provides affordable options."[6] Similarly, the International Planned Parenthood Federation (IPPF), operating in 164 countries, notes in their 2019 performance report that "IPPF delivered 252.3 million sexual and reproductive health services in 2019. […] Many of these services were delivered in rural and periurban areas where there are no other health facilities. An estimated 84 per cent of IPPF's clients were poor and vulnerable."[7]

Many (nongovernmental) organizations measure their success using performance indicators similar to those presented above. These targets are outcome constraints because they require a minimum fraction of the utility (client volume) to belong to a certain player (the "high-impact" subgroup of the population). The targets are often set by external parties, such as donors. However, the potential price (loss of client volume) they entail is not well understood. Consequently, this price is not accounted for in decision making. For example, organizations are expected to improve the percentage of clients from marginalized communities, while at the same time also increase overall client volume. It is important to understand this misalignment of objectives, so as to ensure that agreed-upon targets lead to desired outcomes.

The contributions of this paper are twofold. Firstly, we provide an analytical upper bound on the price of vertical equity for resource-allocation problems with asymmetric outcome constraints. Our set-up considers a decision maker maximizing total utility over a general convex set, specifying an outcome constraint (i.e, a percentage) for each player. We show how the worst-case price depends crucially on the percentages set by the decision maker. The set-up is general and is applicable to many practical problems. Secondly, our results provide a framework to assess the impact of vertical equity concerns for practical problems. Our results depend only on high-level parameters and are therefore well-suited to guide strategic decision making and better align objectives of organizations and external parties. We conclude with two applications. We apply our results to practical instances using data from a large family-planning organization and show how they can inform decision making.[8] We confirm that outcome constraints can lead to a substantial decrease in client volume and analyze which factors drive the price close to the worst-case bound. We also analyze how our results can help bound the impact of prioritization in vaccine allocation using publicly available risk factors for COVID-19.

The remainder of this paper is organized as follows. In Section 2, we discuss related work and embed our results in the literature. In Section 3, we formally define our problem setting, and we provide our main results in Section 4. In Section 5, we apply our results to health-delivery optimization based on data from the family-planning organization. We confirm that outcome constraints can entail a substantial price (loss of client volume) and analyze when our bound proves to be an appropriate tool for estimating this price. In Section 6, we consider publicly available risk factors for COVID-19 and analyze how our results can help bound the impact of prioritization in vaccination policies. Section 7 summarizes our conclusions.

2. Related Work

Equity in optimization problems has received considerable attention in the literature, especially in contexts where equity concerns "arise naturally." Important examples include public health (Hooker and Williams

2012), kidney exchange programs (Bertsimas et al. 2013, Dickerson et al. 2014, Dickerson et al. 2019), and humanitarian operations (McCoy and Lee 2014, Ares et al. 2016, and Eisenhandler and Tzur 2019, among others). Equity concerns are also common in scheduling and logistics. Recent work includes applications in air-traffic flow management (Barnhart et al. 2012 and Bertsimas and Gupta 2015), airline scheduling (e.g., Jacquillat and Vaze 2018, Androutsopoulos and Madas 2019, and Zografos and Jiang 2019), nurse rostering (see, for example, Martin et al. 2013), and crew rostering in public transportation (Nishi et al. 2014, Borndörfer et al. 2015, Breugem et al. 2021).

Yadav (2010) stresses the importance of incorporating equity in public health and discusses different measures. Marsh and Schilling (1994) provide a detailed overview of equity measures in the facility-location literature. Karsu and Morton (2015) give an extensive overview of research incorporating some notion of equity in the operations research literature. Authors distinguish between horizontal equity and vertical equity. Karsu and Morton (2015) refer to the latter type, where players are distinguishable, as balance concerns.

Theoretical analysis of the loss of utility from equity considerations was first considered in Bertsimas et al. (2011) and Caragiannis et al. (2012). This loss is commonly referred to as the price of fairness. The underlying framework of Bertsimas et al. (2011) is similar to the one considered in this paper. The authors analyze the price of fairness for a general convex utility set. Such sets arise, for example, from a resource-allocation problem over a convex set, under reasonable assumptions. Tight upper bounds are provided on the price of fairness for max-min and proportional fairness, two important fairness concepts tracing back to the work of Kalai and Smorodinsky (1975) and Nash (1950), respectively, and well-established in the literature (see, e.g., Mas-Colell et al. 1995 for an introduction). Bertsimas et al. (2012) provides a bound on the price of fairness for the efficiency-fairness trade-off curve resulting from the α-fairness framework (see Mo and Walrand 2000). Their analysis includes a bound on the price of fairness for the egalitarian fairness concept, first introduced in Kalai (1977), as a special case. Finally, Dickerson et al. (2014) provide a bound on the price of fairness in a kidney-exchange model when the fair solution focuses fully on matching highly sensitized (i.e., hard-to-match) patients.

Vertical equity concerns include heterogeneity of needs, preferences, or claims, and diversity requirements (Karsu and Morton 2015). Asymmetric outcome constraints are a way to incorporate vertical equity and are common in practice (as discussed in Section 1). In the literature, such constraints are used in, for example, Bertsimas et al. (2013) and Dickerson et al. (2014) to ensure that certain subgroups of the population are included in a sufficiently large percentage of kidney-exchange matches. The price of such constraints has not been adequately addressed in the literature, to the best of our knowledge. Work on the price of fairness (Bertsimas et al. 2011 and Bertsimas et al. 2012) considers fair allocations defined by picking appropriate welfare or objective functions, but does not consider fairness imposed via constraints. In this research, we bridge this gap and provide analytical bounds on the price of vertical equity concerns expressed as asymmetric outcome constraints.

3. Problem Formulation and Definitions

Before stating our main results, we introduce the necessary definitions and terminology. In Section 3.1, we formalize the resource-allocation problem, and in Section 3.2, we define fairness schemes and the price of fairness. We conclude in Section 3.3 by presenting the class of fairness schemes based on outcome constraints.

3.1. Resource Allocation

We consider a resource-allocation problem among n players. Each player derives utility from the resources she is allocated. The resource-allocation problem is formalized by the utility set $U \subseteq \mathbb{R}^n_+$, defined as the set of all feasible utility allocations (i.e., the set of all possible utilities for the players given the resource set). Consider the following example of allocating a single resource between two players ($n = 2$). Let x_1 and x_2 denote the amount of resource allocated to the first and second player, respectively. Assume the resource set is given by all $[x_1, x_2] \in \mathbb{R}^2_+$, such that $3x_1 + 3x_2 \leq 7$. Both players derive utility from their allocated resources. Suppose the first player's utility function is given by $f_1(x_1) = \min\{1, x_1\}$ and the second player's utility function by $f_2(x_2) = \min\{1, 2x_2\}$. That is, the second player derives twice as much utility per allocated resource compared with the first player (and both functions are bounded by one). Figure 1 shows the utility set for this specific case. The set U is given by all $[u_1, u_2] \in [0,1]^2$, such that $3u_1 + 6u_2 \leq 7$. This is precisely the set of all $[f_1(x_1), f_2(x_2)] \in \mathbb{R}^2_+$ for which $[x_1, x_2]$ is a feasible resource allocation.

Many practical allocation problems fit the above framework. Furthermore, for many problems, the utility set is "well-behaved." To formalize this, we need to introduce some additional notation. We write $u \leq v$, for $u, v \in \mathbb{R}^n$, to indicate that $u_i \leq v_i$ for $i = 1, \ldots, n$. Furthermore, we say $u \in U$ is dominated if there exists a $v \in U$ such that $u_i \leq v_i$, for $i = 1, \ldots, n$, with at least one strict inequality. With slight abuse of notation, we write c_n for the vector in $\mathbb{R}^n_+$ with all elements equal to $c \in \mathbb{R}_+$; for example, 0_n and 1_n denote the n-dimensional vectors with all zeros and all ones, respectively.

Figure 1. Example of a Utility Set (Highlighted Area) for Two Players

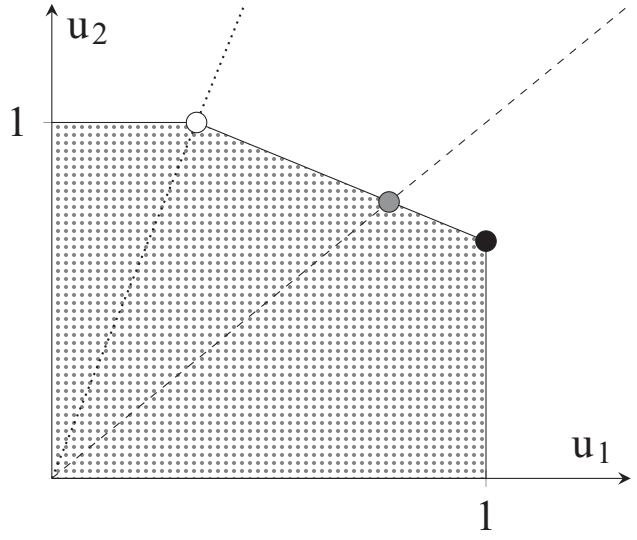

Notes. The efficient (black) and egalitarian (gray) allocation are shown. The white point indicates the allocation where the second player derives three times as much utility as the first player. That is, the first player derives 25%, and the second player derives 75% of the total utility.

When it is clear from the context, we drop the length indication.

We make the following assumption regarding U.

Assumption 1. *The utility set $U \subseteq \mathbb{R}^n_+$ satisfies the following properties*:

1. *U is compact and convex.*
2. *U is monotonic—that is, if $u \in U$, then $v \in U$ for all $0 \leq v \leq u$.*

Assumption 1 is similar to Bertsimas et al. (2011) and Bertsimas et al. (2012) and common in the literature. The above properties hold for resource-allocation problems for which the resource set is convex and utility functions satisfy some reasonable assumptions (such as concavity and boundedness; for details, see Bertsimas et al. 2011, proposition 1). For example, the utility set of Figure 1 satisfies Assumption 1. The monotonicity of U implies that the utility of a player can always be reduced. An informal interpretation of this property is that resources are freely disposable.

3.2. Fairness Schemes

We assume that the utility allocation is determined by a central decision maker aiming to balance fairness and efficiency. Fairness is formalized by using fairness schemes. Each fairness scheme maps a utility set to a vector within that set. This vector is the "fair" allocation with respect to the fairness scheme. Formally, a fairness scheme is a function $S : 2^{\mathbb{R}^n_+} \to \mathbb{R}^n_+$, and $S(U)$ is the fair allocation according to S. For example, the egalitarian fairness scheme maps a utility set to the allocation maximizing the utility for those worst off within that set (i.e., lexicographically largest, accounting for permutations). Figure 1 shows the egalitarian allocation for the depicted utility set in gray, given by $u = [7/9, 7/9]$.

The decision maker balances fairness and efficiency by selecting an appropriate fairness scheme. However, this generally comes at a cost (loss of efficiency). The price of fairness quantifies this loss. Formally, the price of fairness for a given fairness scheme is defined as the relative loss of utility between the maximum achievable utility (efficient allocation) and the sum of utilities from the fair solution obtained from the scheme. Let SYSTEM(U) denote the utility derived from the efficient allocation—that is:

$$\text{SYSTEM}(U) = \max\left\{\sum_{i=1}^{n} u_i \mid u \in U\right\}.$$

Furthermore, let FAIR(U, S) denote the utility derived under the fairness scheme S—that is:

$$\text{FAIR}(U, S) = \sum_{i=1}^{n} S(U)_i.$$

The price of fairness of a fairness scheme S, given a utility set U, denoted by POF(U, S), can now be defined as:

$$\text{POF}(U, S) = \frac{\text{SYSTEM}(U) - \text{FAIR}(U, S)}{\text{SYSTEM}(U)}.$$

For example, the egalitarian allocation of Figure 1 leads to a price of fairness of approximately 6%, given that the efficient allocation is $u = [1, 2/3]$.

3.3. Outcome Constraints

We now formalize the class of fairness schemes based on outcome constraints. Each outcome constraint specifies a minimum percentage of the overall utility ω_i for every player i. We also refer to ω as a priority vector (or simply priorities). The *outcome-based* fairness scheme S_ω maps U to the allocation maximizing overall utility whilst satisfying each outcome constraint. That is, the fairness scheme S_ω maps the utility set to the solution of the following problem:

$$\max \quad \sum_{i=1}^{n} u_i, \tag{1a}$$

$$\text{s.t.} \quad u_i \geq \omega_i \sum_{j=1}^{n} u_j \qquad i = 1, \ldots, n, \tag{1b}$$

$$u \in U. \tag{1c}$$

Figure 1 shows the allocation $S_\omega(U)$ for percentages $\omega = [1/4, 3/4]$. Outcome-based fairness schemes represent

a larger class of structurally similar fairness schemes, perhaps best referred to as *constraint-based* fairness schemes. This class "enforces" fairness by maximizing utility whilst explicitly constraining the utility set. Outcome constraints are one example, but a similar scheme could enforce that each player should derive at least a certain fraction of their maximum achievable utility.

Constraint-based fairness schemes are complementary to the "implicit" schemes based on social welfare functions well-studied in economic literature (e.g., proportional, max-min, and egalitarian fairness schemes). They are an important class of fairness schemes commonly used in practice, most likely due to their easy interpretability—that is, constraints represent tangible outcomes, and it is quickly checked whether they are satisfied or not. However, they do not necessarily satisfy some of the axioms that fairness schemes ideally satisfy (see Bertsimas et al. 2011). For example, the solution to (1a)–(1c), and, hence, $S_\omega(U)$, is not necessarily unique when percentages do not sum to one. Furthermore, the allocation $S_\omega(U)$ could be dominated (in fact, the solution corresponding to $\omega = [1,0]$ in Figure 1 would be an example). This can be resolved by slightly redefining S_ω. For example, one could first solve (1a)–(1c) and then determine a utility-maximizing allocation such that no player becomes worse off (in absolute terms) compared with this solution. Doing so would not change the results presented in this paper, as the presented bounds and worst-case instances do not exploit cases where S_ω maps to a dominated solution.

4. The Price of Asymmetric Outcome Constraints

We now turn to the main results of the paper and derive a bound on the price of fairness for the fairness scheme S_ω. All proofs can be found in the appendix. We consider a utility set U satisfying Assumption 1, where player i has a maximum achievable utility $\bar{u}_i > 0$. Define μ as the minimum of the percentages ω_i and define σ as the sum of the percentages ω_i. The price of fairness depends on the relative difference in maximum achievable utility among players. To capture this, we define $\bar{u}_{\min} = \min_i\{\bar{u}_i\}$ and $\bar{u}_{\max} = \max_i\{\bar{u}_i\}$ as the minimum and maximum among the maximum achievable utilities. The extent in which maximum achievable utility impacts the price of fairness depends on the percentages ω. Intuitively, prioritizing players with high maximum achievable utility should be less costly (in the worst-case) than prioritizing those with low maximum achievable utility. To account for this, we define (excluding the trivial case $\sigma = 0$):

$$\bar{u}_\omega = \left(\sum_{i=1}^{n} \frac{\omega_i}{\sigma \bar{u}_i}\right)^{-1}. \tag{2}$$

The value $\bar{u}_\omega \in [\bar{u}_{\min}, \bar{u}_{\max}]$ reflects how well the maximum achievable utilities and percentages are aligned. In particular, $\bar{u}_\omega = \bar{u}_{\min}$, respectively, $\bar{u}_\omega = \bar{u}_{\max}$, when $\omega_i > 0$ only for players with maximum achievable utility $\bar{u}_{\min}$, respectively, $\bar{u}_{\max}$. The value $\bar{u}_\omega$ represents the unique weighted sum of values $\bar{u}_i$, such that the corresponding convex combination satisfies the outcome constraints.

Besides maximum achievable utility, one can expect the price of fairness for S_ω to depend on two key characteristics. Firstly, it depends on the smallest percentage the decision maker specifies for a player. Intuitively, a small percentage (possibly zero) for a "profitable" player means that relatively few resources will be allocated to this player; hence, resources are not used efficiently. Secondly, it depends on the sum of the percentages. The percentages do not necessarily sum to one. For example, the decision maker could require one player to derive at least 30% of total utility, but be indifferent about the remaining 70%. The more the decision maker specifies, the higher the worst-case price will be. In this section, we formalize this intuition. We split the results in three parts. In the first two, we present the results for the cases $\sigma = 1$ and $\sigma \le 1$, respectively. In the third, we discuss the special case of $\mu = 0$, when some players are not assigned a minimum percentage, in more detail. We conclude with a managerial summary of our main results.

4.1. Bound for $\sigma = 1$

The first result considers $\sigma = 1$—that is, the elements of ω sum up to one. Here, the distribution of utility is "fully specified": The fraction player i receives is necessarily equal to ω_i. The proof of Theorem 1 builds upon the techniques used in Bertsimas et al. (2011) and Bertsimas et al. (2012), combined with several new ideas.

Theorem 1. *Consider an n-dimensional utility set U satisfying Assumption 1, where player i has a maximum achievable utility* $\bar{u}_i > 0$. *Let* $\bar{u}_{min} = \min_i\{\bar{u}_i\}$, $\bar{u}_{max} = \max_i\{\bar{u}_i\}$, *and define* $\bar{u}_\omega$ *according to* (2). *For a given priority vector* ω *with minimum element* μ *and* $\sum_{i=1}^{n}\omega_i = 1$, *it holds that*:

$$\mathrm{POF}(U, S_\omega) \le 1 - \min_{x \in [1, \bar{n}]} \frac{1}{\mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \tag{3}$$

with $\bar{n} = n - 1 + (\bar{u}_{min}/\bar{u}_{max})$ *and* $F(x) = \min\{1 - \mu x, (\bar{u}_{min}/\bar{u}_\omega) - (\bar{u}_{min}/\bar{u}_{max})\mu x\}$.

The bound of Theorem 1 is tight for all n and admissible μ—that is, all $\mu \in [0, 1/n]$—in the case of equal maximum achievable utilities. It is worth noting that the worst-case instances are relatively simple. The analysis of these instances, which are similar in nature to the worst-case instances considered in Bertsimas et al. (2011), is provided in the appendix.

The function $F(x)$ incorporates the information on $\bar{u}_\omega$ into the bound. Particularly insightful are the two most extreme cases—that is, the "worst-case," when $\bar{u}_\omega \approx \bar{u}_{\min}$, and the "best-case," when $\bar{u}_\omega \approx \bar{u}_{\max}$, respectively. In the former case, one has $F(x) = 1 - \mu x$ for all x, implying that the ratio $(\bar{u}_{\max}/\bar{u}_{\min})$ has a substantial impact on the price. However, in the latter case, one has $F(x) \approx (\bar{u}_{\min}/\bar{u}_{\max}) - (\bar{u}_{\min}/\bar{u}_{\max})\mu x$ for all x, implying that the ratio disappears from the bound—that is, the bound is computed as if all players have equal maximum achievable utility.

Theorem 1 shows how the minimum percentage impacts the worst-case ratio between the utility derived from the efficient and the fair allocation. We observe that the bound on the price of fairness is decreasing in μ, confirming the intuition that a smaller minimum percentage provides less guarantee that "profitable" players are included in the solution and, hence, a higher worst-case bound. Figure 2(a) shows the bound of Theorem 1 as a function of the number of players for various values of μ in case of equal maximum achievable utilities. Note that μ is necessarily restricted to the interval $[0, 1/n]$.

To compare the result of Theorem 1 with results in the literature, we note that in the special case that $\omega_i = 1/n$, for $i = 1, \ldots, n$, Theorem 1 provides a bound for the egalitarian fairness scheme and, hence, can be compared with results on the price of horizontal equity. In this case, (3) boils down to

$$1 - \min_{x \in [1, \bar{n}]} \frac{n}{x + (n - x)(\bar{u}_{\max}/\bar{u}_{\min})x},$$

which can be easily shown to be equivalent to the bound presented in Bertsimas et al. (2012).

4.2. Bound for $\sigma \leq 1$

Next, we consider the general case $\sigma \leq 1$—that is, the elements of ω do not necessarily sum up to one. This implies that Constraints (1b) do not completely determine the distribution of utility over the players and, hence, leave room to improve the overall derived utility.

Theorem 2. *Consider an n-dimensional utility set U satisfying Assumption 1, where player i has a maximum achievable utility $\bar{u}_i > 0$. Let $\bar{u}_{min} = \min_i\{\bar{u}_i\}$, $\bar{u}_{max} = \max_i\{\bar{u}_i\}$, and define $\bar{u}_\omega$ according to (2). Consider a given priority vector ω with minimum element μ and $\sum_{i=1}^{n} \omega_i = \sigma \leq 1$. It holds that*:

$$\mathrm{POF}(U, S_\omega) \leq 1 - \min_{x \in [1, \bar{n}]} \frac{1}{(1 - \sigma) + \mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \tag{4}$$

with $\bar{n} = n - 1 + (\bar{u}_{min}/\bar{u}_{max})$ and $F(x) = \min\{\sigma - \mu x, (\bar{u}_{min}/\bar{u}_\omega)\sigma - (\bar{u}_{min}/\bar{u}_{max})\mu x\}$. This bound can be strengthened in case we have as input a priority vector $\tilde{\omega}$, with $\sum_{i=1}^{n} \tilde{\omega}_i = \tilde{\sigma}$, such that $\tilde{\omega} \leq \omega$ and $\tilde{\omega} \leq (1 - \sigma + \tilde{\sigma})\omega^\star$, with $\omega^\star$ the priority vector for the efficient solution. In this case, one can instead compute the bound for the priority vector $\omega - \tilde{\omega}$.

We have not been able to find instances for which the bound of Theorem 2 is tight for $\sigma \neq 1$. However, our results in Section 5 empirically show there are instances for which the bound is close to the actual price of fairness.

Theorem 2 shows how the sum of the percentages σ influences the bound on the price of fairness. The bound on the price of fairness is increasing in σ, in line with the intuition that a smaller sum implies more freedom to maximize overall utility. It also shows how information on the utility-maximizing solution can be incorporated to strengthen the bound. This is particularly useful in practical applications, as we show in Section 5. It allows the central decision maker to calibrate the bound using additional information (or expert knowledge). In particular, if the

Figure 2. Bounds on the Price of Fairness as a Function of the Number of Players, for Different μ Values—That Is, the Minimum Percentage—and Different σ Values—That Is, the Total Sum of Percentages

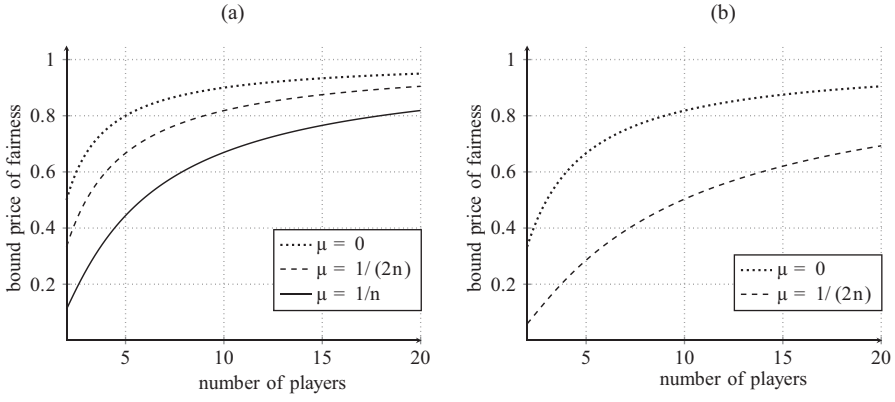

Notes. (a) $\sigma = 1$. (b) $\sigma = 1/2$.

percentages are close to the percentages for the efficient allocation, one can decrease the bound (via the priority vector $\tilde{\omega}$). Figure 2(b) shows the bound of Theorem 2 as a function of the number of players for different values of μ, for the case $\sigma = 1/2$. Note that in both Figure 2, (a) and (b), the depicted bound is lower for higher μ and that the curves for similar μ values are lower in Figure 2(b) compared with Figure 2(a).

4.3. Special Case $\mu = 0$

The special case of $\mu = 0$ is of particular interest. One can imagine that the percentage constraints are used to provide sufficient utility (e.g., service or access) to "vulnerable" players, which otherwise would not be considered in the efficient allocation. In such a context, there will generally be a subset of players (the "remaining" players), for which no minimum percentage is specified (i.e., the percentage is zero). For this specific case, Theorem 2 gives the following bound on the price of fairness, which is worth stating separately.

Corollary 1. *Consider an n-dimensional utility set U satisfying Assumption 1, where player i has a maximum achievable utility $\bar{u}_i > 0$. Let $\bar{u}_{min} = \min_i\{\bar{u}_i\}$, $\bar{u}_{max} = \max_i\{\bar{u}_i\}$, and define $\bar{u}_\omega$ according to (2). For a given priority vector ω with at least one element equal to zero and $\sum_{i=1}^n \omega_i = \sigma \leq 1$, it holds that*:

$$\mathrm{POF}(U, S_\omega) \leq 1 - \frac{\bar{u}_\omega}{(1-\sigma)\bar{u}_\omega + \sigma(\bar{u}_{\min} + (n-1)\bar{u}_{\max})}. \tag{5}$$

This bound can be strengthened in case we have as input a priority vector $\tilde{\omega}$, with $\sum_{i=1}^n \tilde{\omega}_i = \tilde{\sigma}$, such that $\tilde{\omega} \leq \omega$ and $\tilde{\omega} \leq (1 - \sigma + \tilde{\sigma})\omega^\star$, with $\omega^\star$ the priority vector for the efficient solution. In this case, one can instead compute the bound for the priority vector $\omega - \tilde{\omega}$.

Corollary 1 follows directly from Theorem 2 by noting that $F(x) = (\bar{u}_{\min}/\bar{u}_\omega)$ when $\mu = 0$ and setting $x = \bar{n}$. As Corollary 1 is a direct result of Theorem 2, we have not identified instances for which (5) is tight for $\sigma \neq 1$. We note, however, that it is relatively straightforward to construct instances for which the price of fairness equals:

$$1 - \frac{\bar{u}_\omega}{\sigma(n-1)\bar{u}_{\max} + \sigma\bar{u}_{\min}},$$

with $\bar{u}_\omega$ either $\bar{u}_{\min}$ or $\bar{u}_{\max}$, provided that σ is not too small.

Corollary 1 provides a simple expression for the bound in case $\mu = 0$. In particular, it gives an easy expression for the following question: What is the worst-case price due to enforcing σ percent of overall utility to be derived by a certain subset of players? Corollary 1 shows that this question can be answered by a relatively straightforward function of the sum of percentages σ, the number of players n, and the ratio between the maximum and minimum utility. Figure 3 shows the bound of Corollary 1 as a function of σ for different ratios of the maximum and minimum achievable utility in case the player with the smallest maximum achievable utility is prioritized ($\bar{u}_\omega = \bar{u}_{\min}$) for both two and four players. As is clear from Corollary 1, the bounds are increasing in both the ratio and the number of players.

4.4. Managerial Summary

The analytical results show that the worst-case price depends on (1) the number of players, (2) the difference in maximum achievable utility among players, (3) the smallest percentage the decision maker specifies for a player, and (4) the sum of the percentages specified by the decision maker. Figure 4 provides an overview of how these factors influence the price of fairness. We use values Low, Medium, and Large to indicate how the bound on the price of fairness increases/decreases when moving between quadrants. The exact values depend on the four factors outlined above and can be computed by using Theorem 2.

Our results indicate that misalignment between efficiency and priorities is an important driver of loss. In particular, if prioritization reflects a second conflicting objective (i.e., when the relative utility of prioritized players is low), one ends up in the left side of Figure 4. This is the case in the family-planning example of Section 1, in which efficiency (client volume) and priorities (reaching marginalized communities) do not align (assuming marginalized communities are less easy to reach). Similar misalignment can occur when prioritizing highly sensitized patients in kidney-exchange models, because highly sensitized patients have fewer matches, and, hence, enforcing a minimum percentage of matches to include these patients negatively influences efficiency (maximizing the amount of kidney exchanges). If priorities align with the efficiency metric, one ends up in the right side of Figure 4, where the price is expected to be (substantially) less. This is the case in, for example, prioritization policies in vaccine allocation, where efficiency and priorities are both aimed at maximizing the impact per vaccine. Hence, decision makers should carefully consider how they define efficiency and why they set certain priorities. For many practical problems, we believe one can appropriately assess this misalignment based on high-level parameters (e.g., average client volume or benefit per vaccine), without explicitly computing the maximum achievable utilities.

Our results also provide intuition on how problem-specific factors, such as diminishing returns, influence the price of fairness. One can expect the efficient allocation to be skewed (i.e., a small fraction of players receives all resources) in the case of low diminishing returns and balanced (i.e., all players receive some

Figure 3. Bounds on the Price of Fairness as a Function of the Sum of the Percentages, for Different Ratios Between the Maximum and Minimum Derivable Utility, and Different n—That Is, the Number of Players

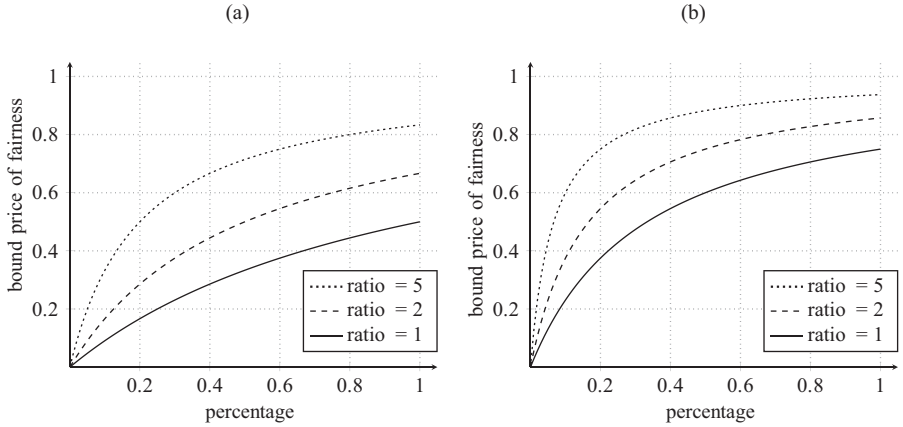

Notes. (a) $n = 2$. (b) $n = 4$.

resources) in the case of high diminishing returns. If priorities match with the (expected) efficient allocation, our results indicate that one can adjust the estimates in Figure 4 downward.

In short, the results help identify drivers of the price of fairness, thereby informing decision makers when (not) to worry about efficiency loss when setting prioritization requirements. The bound provides a (conservative) estimate for the price of fairness. In case one has suitable approximations for maximum achievable utilities, one can use the bound as an easily computed "ballpark" estimate, possibly adjusted to incorporate problem-specific information (e.g., knowledge on diminishing returns, as discussed above). However, one should be careful when making decisions based on the bound alone, especially when the indicated price is high (and thus potentially overestimated). Ideally, the bound is complemented by a confirmatory (more computationally extensive) exact study when the estimated price is high. In Sections 5 and 6, we analyze empirically the difference between the bound and the true price of fairness for two practical applications.

Figure 4. Overview of How the Relative Utility of Prioritized Players ($\bar{u}_\omega/\bar{u}_{\max}$) and Minimum Percentage (μ) Affect the Bound on the Price of Fairness

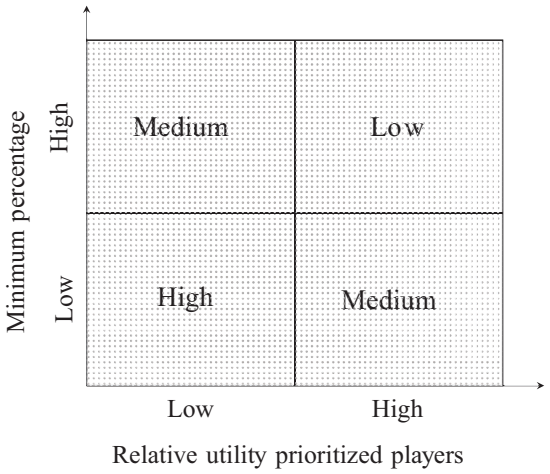

Note. These values increase in the number of players (n) and the sum of specified percentages (σ).

5. An Application in Health-Delivery Optimization

First, we apply our methodology to optimization in health-delivery programs. Our experiments are based on field operations of a large nonprofit Family-Planning Organization, to which we simply refer as "FPO" to preserve confidentiality. Outreach is an important part of FPO's operations. It consists of teams of healthcare workers traveling the country and providing services to women living in remote and rural locations and urban low-income areas, where access to family-planning services is difficult.

As discussed in Section 1, organizations in health-delivery measure their success using a collection of performance indicators. For example, the proportion of served clients that are young women is an important performance indicator for almost all family-planning organizations. The importance of reaching these women is clear. UNFPA states that "[a]dolescents have higher rates of unintended pregnancies than any other age group, and the unmet need for contraception is much higher among adolescents compared with all women ages 15-49 (57 per cent versus 24 per cent). […] Roughly half of the pregnancies among adolescents aged 15-19 in developing regions are unintended, and about half of these end in abortions, most

of which are unsafe."[9] However, reaching young women is difficult due to stigmatization and other access barriers.

We assume that an external party (e.g., a donor) specifies a minimum proportion of young clients as the target for FPO—that is, as a condition for the funds to be allocated. Note that these output constraints are quite common in practice. The organization has a number of ways to reach this target: marketing, capacity building, and tailored site-visit policies, among others. Our focus is on the latter. Optimizing site-visit policies can, in essence, be modeled as a resource-allocation problem, with the resource representing the amount of visits the organization can plan. The minimum proportion of young clients can be taken into account as an outcome constraint. Focusing on the proportion of young clients forces the outreach teams toward sites that comply with this metric—that is, toward sites with a high proportion of young women. However, overprioritizing sites can lead to demand saturation, and client volume might be negatively correlated with the proportion of young clients. Furthermore, it might be conflicting with other objectives, such as providing access for all. Donors do not often have a clear view about the price in terms of total number of clients reached (i.e., the usual overall performance objective) of the targets they impose. Our results provide a framework to estimate this price.

5.1. Mathematical Model

We model the problem as follows. We consider k sites and a fixed capacity for visits during the planning period (without loss of generality standardized to one). The decision variable x_j, for $j = 1, \ldots, k$, denotes the number of visits to site j. We refer to the collection of x_j variables as a visiting schedule. The cost of a visit to site j is denoted by c_j. This cost could represent the time necessary for a visit. The number of visits per site is modeled as a continuous decision variable. The assumption of continuity can be justified, as the planning period is often a year, and sites are visited several times per year. Furthermore, fractional visit schedules can be interpreted as visit schedules that vary over planning periods.

We assume that site visits exhibit diminishing returns. The diminishing returns could capture demand saturation at a site. They could also capture the organization's attitude toward access. For example, the organization might strongly prefer visiting two different sites one time over visiting one site twice, even if the latter site represents more clients. The level of diminishing returns is captured by the parameter $\theta \in [0, 1)$, which can be set by the organization. The number of effective site visits to site j is given by $x_j^{1-\theta}$. When diminishing returns are limited (θ close to zero), the number of effective site visits will be close to the actual number of site visits. However, when returns decrease rapidly (θ close to one), the number of effective site visits will be close to one, even if the number of site visits is large.

We assume the clients are categorized into n priority groups. For example, the population is grouped into two priority groups, representing high- and low-impact clients, with high-impact clients being young women. Let u_i, for $i = 1, \ldots, n$, denote the number of clients of priority group i. Furthermore, define v_{ij} as the amount of clients belonging to priority group i per site visit at site j. Finally, let ω_i be the desired percentage for each priority group. The problem now reads as follows:

$$\max \quad \sum_{i=1}^{n} u_i, \tag{6a}$$

$$\text{s.t.} \quad u_i \geq \omega_i \sum_{j=1}^{n} u_j \qquad i = 1, \ldots, n, \tag{6b}$$

$$u_i \leq \sum_{j=1}^{k} v_{ij} x_j^{1-\theta} \qquad i = 1, \ldots, n, \tag{6c}$$

$$\sum_{j=1}^{k} c_j x_j \leq 1, \tag{6d}$$

$$u_i \geq 0 \qquad i = 1, \ldots, n, \tag{6e}$$

$$x_j \geq 0 \qquad j = 1, \ldots, k. \tag{6f}$$

The Objective (6a) maximizes client volume. Constraint (6b) ensures the minimum percentage for each priority group, and Constraints (6c) specify the relation between client volume and site visits. Constraint (6d) ensures the visiting schedule respects the planning capacity, and Constraints (6e) and (6f) specify the domains of the decision variables.

It is readily seen that the set of achievable client volume is compact, convex, and monotonic; hence, Assumption 1 is satisfied. Furthermore, Assumption 1 would still hold if one would impose additional constraints on the visiting schedule (e.g., lower and upper bounds, site-dependent restrictions, or site allocations to teams). This means our methodology is applicable to the outreach problem, as expressed by (6a)–(6f) and also variations of this problem.

5.2. Computational Results

We calibrate our model using data from FPO's operations in one of the larger African countries. Operations are split into regions. We consider two regions, referred to as regions A and B, corresponding to a more populated semiurban and fairly remote poor rural area, respectively. We consider two priority groups (high- and low-impact clients). We assume that the distinction between high- and low-impact clients is made at the site level. In particular, high-impact clients are from sites with a high proportion of

young women, and low-impact clients are from sites with a low proportion. The distinction high/low proportion is based on the median—that is, we have an equal number of sites of high and low proportion. Albeit different (more complex) definitions can also be considered, we believe this definition appropriately captures the potential misalignment between maximizing client volume and percentage-based performance indicators.

Table 1 gives an overview of key characteristics of both regions. The client volume per visit and the proportion of young clients for each site are estimated by using historical data of FPO. The data in Table 1 are representative. Exact data are not shown for confidentiality reasons. We observe that for both regions, there is a substantial difference, approximately a factor of two, between the high- and low-impact sites in average proportion of young clients. For region A, we see a substantial decrease in average client volume between the high- and low-impact sites (approximately (approx.) 13%). For region B, however, we observe an increase in average client volume (approximately 22%). Furthermore, region A has substantial higher average client volume compared with region B.

We consider two levels of diminishing returns: low ($\theta = 0.25$) and high ($\theta = 0.75$). These two levels lead to a fundamentally different visit schedule. For example, doubling the number of visits to a site results in an effective increase of about 70% when the level is low, but only 20% when the level is high. As noted in Section 5.1, the levels can represent different types of services (with different demand saturation), but also different attitudes toward access and horizontal equity. In the latter case, the organization and external party would agree upon a parameter value that reflects their views. At FPO, the cost of visiting sites is relatively uniform. We therefore set the cost for each site equal to one.

The result of our analysis is shown in Figure 5. For each region and each level of diminishing returns, we compute the price of fairness and the bound as a function of the minimum percentage of high-impact clients. We also compute the adjusted bound, which is calibrated by using the volume-maximizing solution (i.e., without outcome constraints), according to Corollary 1. We observe a loss of client volume for both regions and both levels of diminishing returns. When fully prioritizing high-impact clients, the price ranges from 10% to as much as 50%. Not surprisingly, the price is highest for region A, where average client volume is lower for the high-impact sites compared with low-impact sites. The results show that the level of diminishing returns is a key driver of price. Prioritizing high-impact clients leads to a substantial price when diminishing returns are high. However, we observe only a limited price for low diminishing returns. The results also show the benefit of the adjusted bound over the nonadjusted bound. The adjusted bound provides a good approximation when the level of diminishing returns is high. The nonadjusted bound, however, significantly overestimates the price for all but percentages close to one.

In case of low diminishing returns, the loss in client volume is mostly caused by large relative differences between high- and low-impact sites. Such relative differences are small in our instances (see Table 1), which explains the limited price. To further analyze the case of low diminishing returns, we vary the cost per site visit between high- and low-impact sites. The most natural interpretation for this is that visits to high-impact sites require substantially more travel time compared with low-impact sites (e.g., rural versus urban). As noted, this effect is limited the in case of FPO. However, one can think of applications where travel time has a major influence on visiting schedules (consider, for example, the case study at Riders for Health by McCoy and Lee 2014). Alternatively, one can imagine high-impact sites as sites where no system for outreach is yet in place. We consider region B, where most people live in rural areas, and analyze the price of fairness in case site visits to high-impact sites are a factor of two and a factor of five more costly. The results are shown in Figure 6. We indeed observe that the price approaches the bound when the cost increases.

Table 1. Key Characteristics for Each Region

				Low-impact		High-impact	
Region	Sites	Clients	Young (%)	Clients	Young (%)	Clients	Young (%)
A	300	30	10	32	5	28	15
B	400	20	10	18	5	22	15

Notes. For each region, we show the number of sites, the average number of clients per site, and the average percentage of young clients per site. The average number of clients and average percentage of young clients per site are also shown separately for the high- and low-impact sites.

5.3. Insights and Recommendations

Our experiments confirmed a substantial loss of client volume when prioritizing high-impact clients. The price of fairness can be almost 50%, depending on the region and the level of diminishing returns. Even for region B, where the average client volume for high-impact sites is higher, focusing too much on high-impact clients can entail a high price. Performance targets should therefore be set carefully. Our analysis offers a way to discuss these performance targets with donors, who may not always fully realize the potential impacts of putting earmarks on how their funds are being used.

Figure 5. The Price of Fairness for Each Region and Level of Diminishing Returns as a Function of the Minimum Percentage of High-Impact Clients

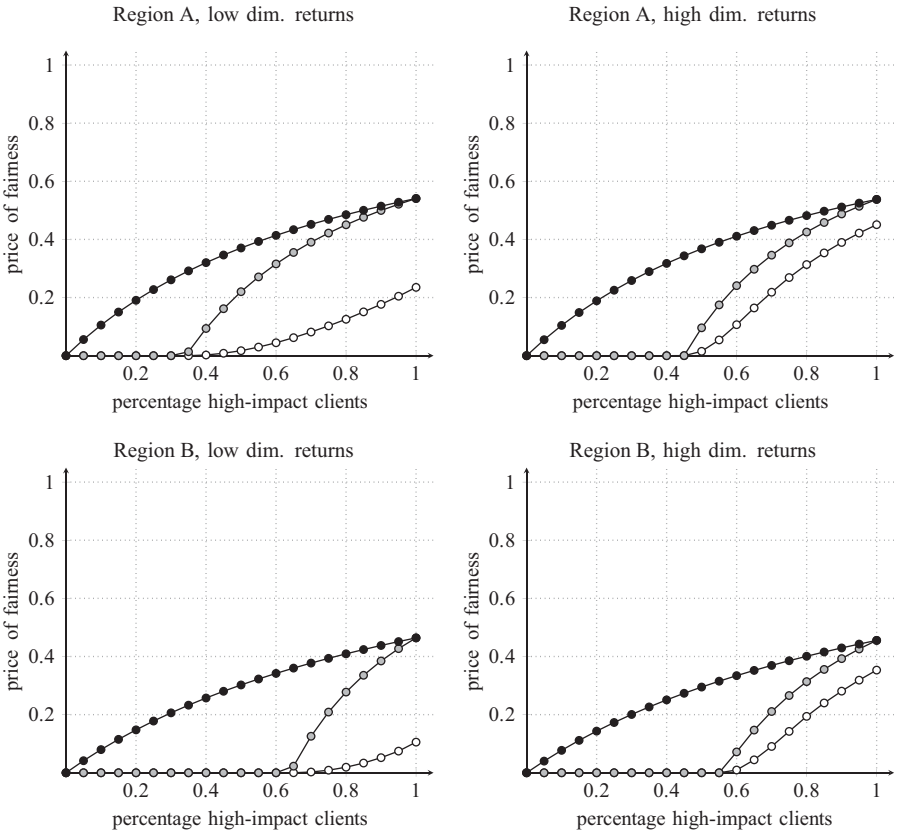

Notes. For each region and each level of diminishing (dim.) returns, we show the price of fairness (white), the bound (black), and the adjusted bound (gray), which is calibrated by using the volume-maximizing solution.

We observed two key drivers of price. Firstly, services that exhibit a high level of diminishing returns are bound to have a substantial price in case of prioritization. Secondly, large relative differences between high- and low-impact sites lead to a high price, even in case of low diminishing returns. The price is smallest when the relative difference between high- and low-impact sites is small and the level of diminishing

Figure 6. The Price of Fairness for Different Efficiency and Level of Diminishing Returns as a Function of the Minimum Percentage of High-Impact Clients

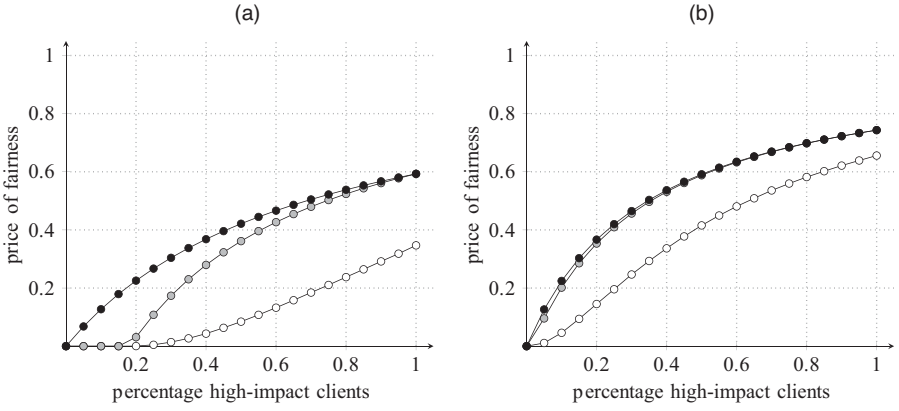

Notes. The cost-increase factor indicates how more costly high-impact site visits are compared with low-impact site visits. For each factor, we show the price of fairness (white), the bound (black), and the adjusted bound (gray), which is calibrated by using the volume-maximizing solution. (a) Region B, low diminishing (dim.) returns, cost increase factor two; (b) Region B, low dim. returns, cost increase factor five.

Table 2. Different Outcomes for Age Groups

	Age groups								
Outcome	0–4	5–17	18–29	30–39	40–49	50–64	65–74	75–84	85+
Infection	<1	1	2	2	2	2	1	1	2
Hospitalization	2	1	6	10	15	25	40	65	95
Death	1	1	10	45	130	440	1,300	3,200	8,700

Notes. The values are relative to age group 5–17. For each age group, we show the infection, hospitalization, and death rates.

returns is low. Having a good understanding of these problem-specific factors is important to assess the potential price of fairness.

Our results show that the bound provides a useful estimate of client loss. The conservatism of this estimate depends on the level of diminishing returns and the difference between high- and low-impact sites. Information regarding the volume-maximizing solution can be important for calibrating the bound. The organization should therefore carefully incorporate expert knowledge and information from current operations. Because the bound relies only on a small number of high-level parameters, it is well suited for ballpark estimates and high-level discussions—for example, with donors.

6. An Application in Vaccine Allocation

Next, we consider prioritization policies in vaccine allocation. Many countries are using prioritization policies in their COVID-19 vaccination strategies. We consider a decision maker deciding upon the vaccination strategy for a given population. We assume that the population is segmented based on age, and needs/risks differ among segments. For example, the elderly are much more likely to die from a COVID-19 infection than young adults. Considering only a limited amount of available vaccines, the decision maker has to decide which people to vaccinate (and which ones later)—for example, does one allow all people to get their vaccine, or is it better to first prioritize those 85+? Or is it perhaps best to allow all people age 75+? Clearly, this is a difficult decision depending on many factors related to disease and social dynamics, but also access/equity concerns and societal preferences. For such complex problems, our results help provide a "first estimate" for the price of fairness for prioritization policies, using only key problem parameters on the risk per age group.

6.1. Computational Results

We model the problem as a resource-allocation problem with each player representing an age group. We take as basis the publicly available risk factors provided by the Centers for Disease Control and Prevention,[10] shown in Table 2. We take as utility per allocated vaccine the death rate for each age group (i.e., expected number of deaths averted, not taking other factors into account). We assume that vaccines are sufficiently scarce so that the population of each group exceeds the amount available. We approximate the (ratios of) maximum achievable utility using the utility per vaccine. For example, the ratio between maximum achievable utilities of age groups 85+ and 75–84 is approximately 2.7 (calculated as 8,700/3,200). We remark that other factors can also be incorporated. For example, one can account for economic loss (due to lost future wages) in the utility coefficients (see, e.g., Camelo et al. 2021 and references therein). We model the prioritization of age groups via outcome constraints. We assume that, among prioritized groups, the outcome constraints sum to one (i.e., all vaccines go to the prioritized age groups) and are proportional to the utility coefficients. For example, if age groups 75–84 and 85+ are prioritized, then we set the percentage for age group 75–84 equal to 27% (calculated as 3,200 divided by 3,200 + 8,700) and 85+ equal to 73%. Intuitively, this means that both groups receive an equal amount of vaccines (they have equal access).

Given the above, the bound on the price of fairness for different prioritization policies can be calculated without additional assumptions on the utility set (besides Assumption 1). Table 3 shows the bound for different age groups and different prioritization policies. We consider two instances, focusing on the two—respectively, three—highest-impact groups (age groups ranging 75–85+ and 65–85+, respectively). We consider different prioritization policies for each instance, based on subsets of age groups. For example, we prioritize age groups 75–84 and 85+ (indicated by "Priority 75–85+" in Table 3). For the instance considering three age groups ("Age 65–85+"), this would mean no vaccines are allocated to age group 65–74.

To help assess the quality of the bound, we solve a stylized allocation problem based on the model presented in Section 5.1. We model each age group as a priority group (consisting of a single "site"). Coefficients' "client volumes" are given by the utility coefficients. The result is a stylized allocation problem that captures the impact of potential diminishing returns and large differences in utility coefficients among age groups, without taking additional nonlinearities into account (e.g., it omits disease dynamics and other

Table 3. The Bound on the Price of Fairness Together with the Realized Price Fairness for a Stylized Allocation Problem

	Priority 65–85+ (%)		Priority 75–85+ (%)		Priority 85+ (%)	
Instance	No Dim.	Low Dim.	No Dim.	Low Dim.	No Dim.	Low Dim.
Age 75–85+						
Price			31.6	19.0	0.0	0.5
Bound			46.4	46.4	26.9	26.9
Age 65–85+						
Price	49.4	33.7	31.6	19.0	0.0	0.5
Bound	75.0	75.0	68.2	68.2	53.5	53.5

Notes. We consider instances focusing on the two (75–85+) and three highest age groups (65–85+), each for different prioritization policies. Dim., diminishing.

interaction effects). We consider both low ($\theta = 0.25$) and no ($\theta = 0$) diminishing returns. Low diminishing returns could account for, for example, vaccine refusal or heterogeneity in needs among recipients within the same age group (modelled using a declining benefit per administered vaccine). Evidence for (potential) refusal can be found in, for example, Schwarzinger et al. (2021), whose experiments show refusal percentages ranging between 19.8% and 36.5% among age groups in the working population in France. The price of fairness for this model is also shown in Table 3.

The results in Table 3 show that the performance of the bound varies. The bound substantially overestimates the price of fairness when only age group 85+ is prioritized. In this case, the allocation is close to optimal, yet the estimated price is substantial. Note that one can incorporate knowledge on the efficient allocation in the bound (cf. the discussion in Section 4.4 and results in Section 5) to reduce this effect. The gap decreases when the number of prioritized groups increases. In this case, the relative utility of prioritized groups decreases, and the prioritization policy moves away from the efficient allocation. This effect is strongest in case of no diminishing returns (when the efficient allocation is most skewed). When all groups are prioritized, the bound overestimates the price by approximately a factor of 1.5. Although this is still substantially conservative, it is the same order of magnitude. We observe that for a given prioritization policy, the bound increases in the number of age groups (players), whereas the price remains the same (up to one decimal precision). This is due to low diminishing returns—that is, including age group 65–74 in the allocation has little additional benefit. However, the bound increases in the number of players.

6.2. Insights and Recommendations

The results show that the bound on the price of fairness can provide estimates for the efficiency loss due to prioritization policies. The conservatism of the bound depends heavily on the prioritization policy. The bound performs best when the number of players is small. This highlights the importance of correctly calibrating the problem before computing the bound and incorporating additional information (e.g., exclude age groups if one expects no efficient solution to allocate resources to these groups). It helps to gain intuition in drivers of loss without explicitly computing the utility set. At the same time, it is not a substitute for a detailed analysis. The bound should be seen as complementary to more computationally intensive methods.

7. Conclusion

In this paper, we analyzed the loss of utility from vertical equity concerns in resource-allocation problems. We studied the price of fairness for resource-allocation problems where vertical equity concerns are modeled using asymmetric outcome constraints. These constraints specify a minimum percentage of the total utility that should be derived by each player. The asymmetry resides in the fact that the percentages are not necessarily the same for all players. Asymmetry allows one to differentiate between players and ensure that resources go where most needed or most desired. We motivated the study of this type of resource-allocation problem in the context of health delivery and vaccine allocation.

We derived an analytical upper bound on the price of fairness. We considered a decision maker maximizing total utility over a general convex set, specifying an outcome constraint (i.e., a percentage) for each player. This set-up is general and applies to many practical problems. We showed that the worst-case price depends crucially on the percentages set by the decision maker. In particular, we showed that it depends on (1) the number of players, (2) the difference in maximum achievable utility among players, (3) the smallest percentage the decision maker specifies for a player, and (4) the sum of the percentages specified by the decision maker. Each of these parameters is high-level, making the bound well-suited for ballpark estimates.

We applied our results to two practical cases. Firstly, we considered health-delivery optimization using data from a large nonprofit family-planning organization. We confirmed that outcome constraints

can lead to a substantial decrease in client volume and analyzed which factors drive the price close to the worst-case bound. The experiments showed that our results are a helpful tool to guide strategic decision making. We also pinpointed when one can expect the estimated price to be (too) conservative and when not. Secondly, we considered vaccine-allocation policy making based on publicly available risks factors on COVID-19. We showed that our results provide a first estimate for the loss due to prioritization policies, the usefulness of which depends on different problem parameters. Our results help to gain intuition in drivers of loss and are complementary to more accurate (computationally intensive) methods.

We propose different avenues for further research. Firstly, one could analyze the price of fairness for different practical problems (such as the health-delivery and vaccine-allocation problems). This would help further improve the understanding of which type of problems exhibit a high price and provide insight into when (not) to focus on vertical equity. For example, in humanitarian and development operations, donors very often try to impose earmarking constraints specifying how the money can be used. Recipient organizations know these earmarked funds often limit their performance in the field, but have a hard time proving this. They also fear to upset donors, especially in these budget-constrained times. Secondly, one could incorporate a level of uncertainty or approximation in the analysis to acknowledge that data are often incomplete and difficult to obtain. This would help convey the importance of good data-collection systems, among other things. Finally, we believe that important work remains to be done on simple decision rules, as optimizing over the utility set can be difficult in resource-constrained environments. Our results provide insight into the price of fairness, but do not give any guidance on how to achieve vertical equity. Well-developed decision rules can bridge this gap.

Appendix. Mathematical Proofs

A.1. Proof of Theorem 1

Theorem 1. *Consider an n-dimensional utility set U satisfying Assumption 1, where player i has a maximum achievable utility $\bar{u}_i > 0$. Let $\bar{u}_{min} = \min_i\{\bar{u}_i\}$, $\bar{u}_{max} = \max_i\{\bar{u}_i\}$, and define $\bar{u}_\omega$ according to (2). For a given priority vector ω with minimum element μ and $\sum_{i=1}^n \omega_i = 1$, it holds that:*

$$\mathrm{POF}(U, S_\omega) \le 1 - \min_{x\in[1,\bar{n}]} \frac{1}{\mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \tag{3}$$

with $\bar{n} = n - 1 + (\bar{u}_{min}/\bar{u}_{max})$ and $F(x) = \min\{1 - \mu x, (\bar{u}_{min}/\bar{u}_\omega) - (\bar{u}_{min}/\bar{u}_{max})\mu x\}$.

Proof. Let ϕ be the largest value such that $\phi\omega \in U$. Note that $\phi\omega$ is the unique optimal solution to (1), because $\sum_{i=1}^n \omega_i = 1$. The existence of ϕ follows from the compactness of U. Define $e_i \in \mathbb{R}^n_+$ as the vector with i-th element equal to one and all other elements equal to zero.

We have:

$$\mathrm{FAIR}(U, S_\omega) = \sum_{i=1}^n (\phi\omega)_i = \phi, \tag{A.1}$$

because $\sum_{i=1}^n \omega_i = 1$.

By the supporting hyperplane theorem, there exists a $\gamma \in \mathbb{R}^n$, with $\gamma \neq 0$ such that:

$$\sum_{i=1}^n \gamma_i u_i \le \sum_{i=1}^n \gamma_i (\phi\omega)_i \quad \forall u \in U. \tag{A.2}$$

Note that $\sum_{i=1}^n \gamma_i \omega_i \ge 0$, because $0 \in U$. Furthermore, $\sum_{i=1}^n \gamma_i \omega_i \neq 0$, because, observing that $\bar{u}_i e_i \in U$, this would imply that $\gamma_i \le 0$, for $i = 1, \ldots, n$, contradicting that $\sum_{i=1}^n \gamma_i \omega_i \ge 0$. This implies that, without loss of generality, we can assume $\sum_{i=1}^n \gamma_i \omega_i = 1$, reducing (A.2) to:

$$\sum_{i=1}^n \gamma_i u_i \le \phi \quad \forall u \in U. \tag{A.3}$$

Furthermore, because $\bar{u}_i e_i \in U$ and $\omega - (1/2)\omega_i e_i \in U$ for every $i = 1, \ldots, n$, it is readily seen that:

$$0 \le \gamma_i \le \phi/\bar{u}_i. \tag{A.4}$$

Without loss of generality, we assume the γ_i are in ascending order. To bound SYSTEM(U), we first note that by definition of $\bar{u}_{\min}$ and $\bar{u}_{\max}$, we have SYSTEM(U) $\le (n-1)\bar{u}_{\max} + \bar{u}_{\min}$. Next, we consider the auxiliary utility set V defined as:

$$V = \left\{ v \in \mathbb{R}^n \mid 0 \le v \le \hat{u}, \sum_{i=1}^n \gamma_i v_i \le \phi \right\}, \tag{A.5}$$

with:

$$\hat{u}_i = \min\{\bar{u}_{\max}, \phi/\gamma_i\} \ge \min\{\bar{u}_i, \phi/\gamma_i\} = \bar{u}_i. \tag{A.6}$$

Here, the last inequality follows from (A.4). Because every $u \in U$ satisfies $0 \le u \le \bar{u}$ and $\sum_{i=1}^n \gamma_i u_i \le \phi$, it follows that $U \subseteq V$ and, hence:

$$\mathrm{SYSTEM}(U) \le \min\{(n-1)\bar{u}_{\max} + \bar{u}_{\min}, \mathrm{SYSTEM}(V)\}. \tag{A.7}$$

As such, we have:

$$\begin{aligned} &\frac{\mathrm{FAIR}(U, S_\omega)}{\mathrm{SYSTEM}(U)} \\ &\ge \frac{\phi}{\mathrm{SYSTEM}(U)} \\ &\ge \frac{\phi}{\min\{(n-1)\bar{u}_{\max} + \bar{u}_{\min}, \mathrm{SYSTEM}(V)\}}. \end{aligned} \tag{A.8}$$

Let $\bar{v}$ be an efficient allocation in V. It is readily seen that $\bar{v}$ is of the form $\bar{v}_i = \hat{u}_i$ for $i \le \ell$, $\bar{v}_{\ell+1} = d\hat{u}_{\ell+1}$ and $\bar{v}_i = 0$ for $i > \ell + 1$ for suitable picked $\ell \in \{1, \ldots, n-1\}$ and $d \in [0, 1]$. Without loss of generality, we can assume that $d > 0$ whenever $\ell + d > 1$. We claim that $\sum_{i=1}^n \bar{v}_i = \ell\bar{u}_{\max} + d\hat{u}_{\ell+1}$. If $\ell + d = 1$, this is trivial, as $\hat{u}_1 = \bar{u}_{\max}$ by definition. Suppose $\ell + d > 1$, and let $j \in \{1, \ldots, \ell\}$ be such that $\hat{u}_j < \bar{u}_{\max}$. This

implies that $\hat{u}_j = \phi/\gamma_j$. We have:

$$\sum_{i=1}^{\ell} \gamma_i \hat{u}_i + \gamma_{\ell+1}(d\hat{u}_{\ell+1}) \geq \gamma_j \hat{u}_j + \gamma_{\ell+1}(d\hat{u}_{\ell+1}) = \phi + \gamma_{\ell+1}(d\hat{u}_{\ell+1}) > \phi, \tag{A.9}$$

which implies $\bar{v} \in V$, a contradiction. To bound ϕ/SYSTEM (V), we define the set Q as all pairs (λ, δ), with $\lambda \in \{1, \ldots, n-1\}$ and $\delta \in [0,1]$, such that:

$$\gamma_1 + \ldots + \gamma_\lambda + \gamma_{\lambda+1}\delta \leq \phi/\bar{u}_{\max}. \tag{A.10}$$

Note that $(\ell, d\hat{u}_{\ell+1}/\bar{u}_{\max}) \in Q$; hence:

$$\max_{(\lambda,\delta)\in Q} \{\lambda \bar{u}_{\max} + \delta \bar{u}_{\max}\} \geq \text{SYSTEM}(V). \tag{A.11}$$

Consider $\bar{Q} \subseteq Q$ defined as:

$$\bar{Q} = \{(\lambda, \delta) \in Q \mid \lambda + \delta \leq \bar{n}\}. \tag{A.12}$$

Note that if $(\lambda, \delta) \in Q \setminus \bar{Q}$, then $\lambda \bar{u}_{\max} + \delta \bar{u}_{\max} > (n-1)\bar{u}_{\max} + \bar{u}_{\min}$, Hence, it follows that

$$\max_{(\lambda,\delta)\in\bar{Q}} \{\lambda \bar{u}_{\max} + \delta \bar{u}_{\max}\} \geq \min\{(n-1)\bar{u}_{\max} + \bar{u}_{\min}, \text{SYSTEM}(V)\}. \tag{A.13}$$

Next, we derive a lower bound on ϕ for $(\lambda, \delta) \in \bar{Q}$ as follows:

$$1 = \sum_{i=1}^{n} \gamma_i \omega_i, \tag{A.14}$$

$$= \mu\left(\sum_{i=1}^{\lambda} \gamma_i + \gamma_{\lambda+1}\delta\right) + \sum_{i=1}^{\lambda} \gamma_i(\omega_i - \mu) + \gamma_{\lambda+1}(\omega_{\lambda+1} - \mu\delta) + \sum_{i=\lambda+2}^{n} \gamma_i \omega_i, \tag{A.15}$$

$$\leq \mu(\phi/\bar{u}_{\max}) + \sum_{i=1}^{\lambda} \gamma_i(\omega_i - \mu) + \gamma_{\lambda+1}(\omega_{\lambda+1} - \mu\delta) + \sum_{i=\lambda+2}^{n} \gamma_i \omega_i, \tag{A.16}$$

$$\leq \mu(\phi/\bar{u}_{\max}) + \sum_{i=1}^{\lambda} (\phi/\bar{u}_i)(\omega_i - \mu) + (\phi/\bar{u}_{\lambda+1})(\omega_{\lambda+1} - \mu\delta) + \sum_{i=\lambda+2}^{n} (\phi/\bar{u}_i)\omega_i, \tag{A.17}$$

$$= \mu(\phi/\bar{u}_{\max}) + \sum_{i=1}^{n} (\phi/\bar{u}_i)\hat{\omega}_i, \tag{A.18}$$

where we define $\hat{\omega}$ as $\hat{\omega}_i = \omega_i - \mu$ for $i \leq \lambda$, $\hat{\omega}_{\lambda+1} = \omega_{\lambda+1} - \mu\delta$, and $\hat{\omega}_i = \omega_i$ otherwise, for notational convenience. Inequalities (A.16) and (A.17) follow directly from (A.10) and (A.4), respectively. Note that:

$$\sum_{i=1}^{n} (\phi/\bar{u}_i)\hat{\omega}_i \leq \sum_{i=1}^{n} (\phi/\bar{u}_{\min})\hat{\omega}_i, \tag{A.19}$$

$$= (\phi/\bar{u}_{\min})(1 - \mu(\lambda + \delta)), \tag{A.20}$$

where the inequality follows from the definition of $\bar{u}_{\min}$, and we use that $\sum_{i=1}^{n} \hat{\omega}_i = 1 - \mu(\lambda + \delta)$. Furthermore, it holds that:

$$\sum_{i=1}^{n} (\phi/\bar{u}_i)\hat{\omega}_i = \sum_{i=1}^{n} (\phi/\bar{u}_i)\omega_i - \sum_{i=1}^{\lambda} (\phi/\bar{u}_i)\mu - (\phi/\bar{u}_{\lambda+1})\mu\delta, \tag{A.21}$$

$$= (\phi/\bar{u}_\omega) - \sum_{i=1}^{\lambda} (\phi/\bar{u}_i)\mu - (\phi/\bar{u}_{\lambda+1})\mu\delta, \tag{A.22}$$

$$\leq (\phi/\bar{u}_\omega) - \sum_{i=1}^{\lambda} (\phi/\bar{u}_{\max})\mu - (\phi/\bar{u}_{\max})\mu\delta, \tag{A.23}$$

$$= (\phi/\bar{u}_\omega) - (\phi/\bar{u}_{\max})\mu(\lambda + \delta), \tag{A.24}$$

where the second equality follows from the definition of $\bar{u}_\omega$ and the inequality from the definition of $\bar{u}_{\max}$. Combining the above two bounds, it follows that:

$$1 \leq \mu(\phi/\bar{u}_{\max}) + (\phi/\bar{u}_{\min})F(\lambda + \delta), \tag{A.25}$$

where $F(x)$ is defined as:

$$F(x) = \min\{1 - \mu x, (\bar{u}_{\min}/\bar{u}_\omega) - (\bar{u}_{\min}/\bar{u}_{\max})\mu x\}. \tag{A.26}$$

Concluding, we obtain the following lower bound on ϕ:

$$\phi \geq \frac{1}{(\mu/\bar{u}_{\max}) + (1/\bar{u}_{\min})F(\lambda + \delta)}. \tag{A.27}$$

Combining all the above, we obtain:

$$\frac{\text{FAIR}(U, S_\omega)}{\text{SYSTEM}(U)} \geq \frac{\phi}{\max_{(\lambda,\delta)\in\bar{Q}}\{\lambda \bar{u}_{\max} + \delta \bar{u}_{\max}\}}, \tag{A.28}$$

$$\geq \frac{1}{\max_{(\lambda,\delta)\in\bar{Q}}\{(\lambda + \delta)(\mu + (\bar{u}_{\max}/\bar{u}_{\min})F(\lambda + \delta))\}}. \tag{A.29}$$

To finalize the proof, we note that:

$$\max_{(\lambda,\delta)\in\bar{Q}} \{(\lambda + \delta)(\mu + (\bar{u}_{\max}/\bar{u}_{\min})F(\lambda + \delta)\} \leq \max_{x\in[1,\bar{n}]} \{\mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x\},$$

where the final inequality follows from the fact that:

$$\{\lambda + \delta \mid (\lambda, \delta) \in \bar{Q}\} \subseteq [1, \bar{n}]. \tag{A.30}$$

□

To show that the bound of Theorem 1 is tight, we consider the following instance. The resource set is given by all $x \in \mathbb{R}^n_+$ satisfying

$$\sum_{i=1}^{n} w_i x_i \leq 1,$$
$$x_i \leq 1 \quad \forall i = 1, \ldots, n,$$
$$x_i \geq 0 \quad \forall i = 1, \ldots, n,$$

where w is defined as

$$w_i = \begin{cases} 1/\theta & \text{if } i \leq k \\ 1 & \text{otherwise,} \end{cases} \tag{A.31}$$

and $k \in \{1, \ldots, n-1\}$ and $\theta \geq 1$ will be picked accordingly. The utility of the i-th player, denoted by u_i, is given by $u_i = c_i x_i$, with $c_i = \bar{u}_{\max}$ for $i \leq k$ and $c_i = \bar{u}_{\min}$ for $i \geq k+1$. Let U denote the resulting utility set. It is readily seen that

$$\text{SYSTEM}(U) = k\bar{u}_{\max} + (1 - k/\theta)\bar{u}_{\min}.$$

Let ω be with minimum element μ and such that $\omega_i = \mu$ for $i \leq k$. It is readily checked that

$$\text{FAIR}(U, S_\omega) = \left(\sum_{i=1}^{k} \frac{\omega_i}{\theta \bar{u}_{\max}} + \sum_{i=k+1}^{n} \frac{\omega_i}{\bar{u}_{\min}}\right)^{-1} = \left(\frac{k\mu}{\theta \bar{u}_{\max}} + \frac{1 - k\mu}{\bar{u}_{\min}}\right)^{-1}.$$

Next, define $\phi = k + (\bar{u}_{\min}/\bar{u}_{\max})(1-k/\theta)$. Note that

$$\text{SYSTEM}(U) = \bar{u}_{\max}\phi,$$

and

$$\text{FAIR}(U, S_\omega) = \left(\frac{1}{\bar{u}_{\min}} - \frac{\mu}{\bar{u}_{\min}}\left(\phi - \frac{\bar{u}_{\min}}{\bar{u}_{\max}}\right)\right)^{-1}.$$

It follows that

$$\begin{aligned}\frac{\text{FAIR}(U, S_\omega)}{\text{SYSTEM}(U)} &= \left(\frac{\bar{u}_{\max}}{\bar{u}_{\min}}\phi - \frac{\mu\bar{u}_{\max}}{\bar{u}_{\min}}\phi\left(\phi - \frac{\bar{u}_{\min}}{\bar{u}_{\max}}\right)\right)^{-1} \\ &= \left(\phi\mu + (1-\mu\phi)\left(\frac{\bar{u}_{\max}}{\bar{u}_{\min}}\right)\phi\right)^{-1},\end{aligned}$$

which is exactly the expression of Theorem 1 when $\bar{u}_\omega \approx \bar{u}_{\min}$. Noting that all $\phi \in [1, \bar{n})$ can be achieved for suitable k and θ in the case of equal maximum achievable utilities, this shows that the worst-case performance can be achieved up to arbitrary precision for this case.

A.2. Proof: Theorem 2

Theorem 2. *Consider an n-dimensional utility set U satisfying Assumption 1, where player i has a maximum achievable utility $\bar{u}_i > 0$. Let $\bar{u}_{min} = \min_i\{\bar{u}_i\}$, $\bar{u}_{max} = \max_i\{\bar{u}_i\}$, and define $\bar{u}_\omega$ according to (2). Consider a given priority vector ω with minimum element μ and $\sum_{i=1}^n \omega_i = \sigma \le 1$. It holds that*

$$\text{POF}(U, S_\omega) \le 1 - \min_{x\in[1,\bar{n}]} \frac{1}{(1-\sigma) + \mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \quad (4)$$

with $\bar{n} = n - 1 + (\bar{u}_{min}/\bar{u}_{max})$ and $F(x) = \min\{\sigma - \mu x, (\bar{u}_{min}/\bar{u}_\omega)\sigma - (\bar{u}_{min}/\bar{u}_{max})\mu x\}$. This bound can be strengthened in case we have as input a priority vector $\tilde{\omega}$, with $\sum_{i=1}^n \tilde{\omega}_i = \tilde{\sigma}$, such that $\tilde{\omega} \le \omega$ and $\tilde{\omega} \le (1-\sigma+\tilde{\sigma})\omega^\star$, with $\omega^\star$ the priority vector for the efficient solution. In this case, one can instead compute the bound for the priority vector $\omega - \tilde{\omega}$.

Proof. First, we note that in the case that $\sigma = 0$ the bound is trivially zero. Hence, it remains to prove the result for $\sigma > 0$. We consider the priority vector $\hat{\omega} = \omega + (1-\sigma)\omega^\star$, with $\omega^\star$ the priority vector obtained from the utilitarian solution. Let ϕ be the largest value such that $\phi\omega \in U$, and, analogously, let $\phi^\star$ be the largest value such that $\phi^\star\omega^\star \in U$. By construction we have $\phi^\star = \text{SYSTEM}(U)$.

Note that for every $\alpha \in [0,1]$, we have $\alpha\phi\omega + (1-\alpha)\phi^\star\omega^\star \in U$, because U is convex. We pick α such that

$$\alpha\phi\omega + (1-\alpha)\phi^\star\omega^\star = \psi(\omega + (1-\sigma)\omega^\star) = \psi\hat{\omega}, \quad (A.32)$$

for some $\psi \ge 0$. We solve

$$\alpha\phi = (1-\alpha)\phi^\star/(1-\sigma), \quad (A.33)$$

which gives

$$\alpha = \frac{\phi^\star}{(1-\sigma)\phi + \phi^\star}, \quad (A.34)$$

and hence

$$\psi = \alpha\phi = \frac{\phi^\star\phi}{(1-\sigma)\phi + \phi^\star} = \frac{\phi^\star}{(1-\sigma) + \phi^\star/\phi}. \quad (A.35)$$

Using that $\phi^* = \text{SYSTEM}(U)$, it follows that

$$\begin{aligned}\frac{\text{FAIR}(U, S_\omega)}{\text{SYSTEM}(U)} &\ge \frac{\text{FAIR}(U, S_{\hat{\omega}})}{\text{SYSTEM}(U)} \ge \frac{\psi}{\text{SYSTEM}(U)} \\ &= \frac{1}{(1-\sigma) + (\text{SYSTEM}(U)/\phi)}.\end{aligned} \quad (A.36)$$

Let $\hat{\mu} = \mu/\sigma$. Note that $\bar{u}_{(1/\sigma)\omega} = \bar{u}_\omega$. Because $(1/\sigma)\omega$ sums to one, we know from the proof of Theorem 1 that

$$\frac{\sigma\phi}{\text{SYSTEM}(U)} \ge \min_{x\in[1,\bar{n}]} \frac{1}{\hat{\mu}x + \hat{F}(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \quad (A.37)$$

with

$$\hat{F}(x) = \min\{1 - \hat{\mu}x, (\bar{u}_{\min}/\bar{u}_\omega) - (\bar{u}_{\min}/\bar{u}_{\max})\hat{\mu}x\}. \quad (A.38)$$

Hence, we have:

$$\frac{\text{SYSTEM}(U)}{\sigma\phi} \le \max_{x\in[1,\bar{n}]} \hat{\mu}x + \hat{F}(x)(\bar{u}_{\max}/\bar{u}_{\min})x. \quad (A.39)$$

It follows that

$$\begin{aligned}&\frac{1}{(1-\sigma) + (\text{SYSTEM}(U)/\phi)} \\ &\ge \frac{1}{(1-\sigma) + \sigma\left(\max_{x\in[1,\bar{n}]} \hat{\mu}x + \hat{F}(x)(\bar{u}_{\max}/\bar{u}_{\min})x\right)},\end{aligned} \quad (A.40)$$

$$= \min_{x\in[1,\tilde{n}]} \frac{1}{(1-\sigma) + \mu x + F(x)(\bar{u}_{\max}/\bar{u}_{\min})x}, \quad (A.41)$$

with

$$F(x) = \min\{\sigma - \mu x, (\bar{u}_{\min}/\bar{u}_\omega)\sigma - (\bar{u}_{\min}/\bar{u}_{\max})\mu x\}. \quad (A.42)$$

To see how additional information can be incorporated, we reason as follows. Instead of $\hat{\omega} = \omega + (1-\sigma)\omega^\star$, we take a linear combination with the priority vector $\frac{1-\sigma+\tilde{\sigma}}{1-\sigma}\omega^\star - \frac{1}{1-\sigma}\tilde{\omega}$. Note that this is a well-defined priority vector because $\tilde{\omega} \le (1-\sigma+\tilde{\sigma})\omega^\star$. It follows that $\hat{\omega} = \omega + (1-\sigma)\left(\frac{1-\sigma+\tilde{\sigma}}{1-\sigma}\omega^\star - \frac{1}{1-\sigma}\tilde{\omega}\right) = \omega + (1-\sigma+\tilde{\sigma})\omega^\star - \tilde{\omega} = (\omega - \tilde{\omega}) + (1-\sigma+\tilde{\sigma})\omega^\star$. Hence, we can consider the analysis for the priority vector $\omega - \tilde{\omega}$. The result then follows directly from the reasoning above. □

Endnotes

[1] For example, the European Commission on priority groups in the COVID-19 vaccination strategy: https://ec.europa.eu/health/sites/health/files/vaccination/docs/2020_strategies_deployment_en.pdf.

[2] Source: https://www.who.int/pmnch/media/news/2010/20100322_d_shaw_oped/en/.

[3] Source: https://www.unfpa.org/family-planning.

[4] Source: https://www.unfpa.org/family-planning.

[5] Source: https://www.gatesfoundation.org/what-we-do/global-development/family-planning.

[6] Source: https://www.mariestopes.org/resources/global-impact-report-2019/.

[7] Source: https://www.ippf.org/resource/2019-annual-performance-report-apr.

[8] In order to preserve confidentiality, all references to the organization have been removed.

[9] Source: https://www.unfpa.org/resources/contraception-adolescents-and-youth.

[10] https://www.cdc.gov/coronavirus/2019-ncov/downloads/covid-data/hospitalization-death-by-age.pdf.

References

Androutsopoulos KN, Madas MA (2019) Being fair or efficient? A fairness-driven modeling extension to the strategic airport slot scheduling problem. *Transportation Res. Part E Logist. Transportation Rev.* 130:37–60.

Ares JN, De Vries H, Huisman D (2016) A column generation approach for locating roadside clinics in Africa based on effectiveness and equity. *Eur. J. Oper. Res.* 254(3):1002–1016.

Barnhart C, Bertsimas D, Caramanis C, Fearing D (2012) Equitable and efficient coordination in traffic flow management. *Transportation Sci.* 46(2):262–280.

Bertsimas D, Gupta S (2015) Fairness and collaboration in network air traffic flow management: An optimization approach. *Transportation Sci.* 50(1):57–76.

Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.

Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.

Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Oper. Res.* 61(1):73–87.

Borndörfer R, Reuther M, Schlechte T, Schulz C, Swarat E, Weider S (2015) Duty rostering in public transport-facing preferences, fairness, and fatigue. *Conf. Adv. Systems Public Transport*. http://www.rotterdam2015.caspt.org/proceedings/paper118.pdf.

Breugem T, Dollevoet T, Huisman D (2021) Is equality always desirable? Analyzing the trade-off between fairness and attractiveness in crew rostering, *Management Sci.*, ePub ahead of print September 8, https://doi.org/10.1287/mnsc.2017.4005.

Camelo S, Ciocan D, Iancu D, Warnes X, Zoumpoulis S (2021) Quantifying the benefits of targeting for pandemic response. Working paper, Stanford University, Stanford, CA.

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

Dickerson JP, Procaccia AD, Sandholm T (2014) Price of fairness in kidney exchange. *AAMAS'14 Proc. 2014 Internat. Conf. Autonomous Agents Multi-Agent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 1013–1020.

Dickerson JP, Procaccia AD, Sandholm T (2019) Failure-aware kidney exchange. *Management Sci.* 65(4):1768–1791.

Eisenhandler O, Tzur M (2019) The humanitarian pickup and distribution problem. *Oper. Res.* 67(1):10–32.

Hooker JN, Williams HP (2012) Combining equity and utilitarianism in a mathematical programming model. *Management Sci.* 58(9):1682–1693.

Jacquillat A, Vaze V (2018) Interairline equity in airport scheduling interventions. *Transportation Sci.* 52(4):941–964.

Kalai E (1977) Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica* 45(7):1623–1630.

Kalai E, Smorodinsky M (1975) Other solutions to Nash's bargaining problem. *Econometrica* 43(3):513–518.

Karsu Ö, Morton A (2015) Inequity averse optimization in operational research. *Eur. J. Oper. Res.* 245(2):343–359.

Marsh MT, Schilling DA (1994) Equity measurement in facility location analysis: A review and framework. *Eur. J. Oper. Res.* 74(1):1–17.

Martin S, Ouelhadj D, Smet P, Vanden Berghe G, ÖZcan E (2013) Cooperative search for fair nurse rosters. *Expert Systems Appl.* 40(16):6674–6683.

Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic Theory*, vol. 1 (Oxford University Press, New York).

McCoy J, Lee H (2014) Using fairness models to improve equity in health delivery fleet management. *Production Oper. Management* 23(6):965–977.

Mo J, Walrand J (2000) Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Networking* 8(5):556–567.

Nash J (1950) The bargaining problem. *Econometrica* 18(2):155–162.

Nishi T, Sugiyama T, Inuiguchi M (2014) Two-level decomposition algorithm for crew rostering problems with fair working condition. *Eur. J. Oper. Res.* 237(2):465–473.

Schwarzinger M, Watson V, Arwidson P, Alla F, Luchini S (2021) COVID-19 vaccine hesitancy in a representative working-age population in France: A survey experiment based on vaccine characteristics. *Lancet Public Health* 6(4):e210–e221.

Yadav P (2010) Improving public health in developing countries through operations research. Cochran JJ, Cox LA, Keskinocak P, Kharoufeh JP, Smith JC, eds. *Wiley Encyclopedia of Operations Research and Management Science* (Wiley, Hoboken, NJ), https://doi.org/10.1002/9780470400531.eorms0401.

Zografos KG, Jiang Y (2019) A bi-objective efficiency-fairness model for scheduling slots at congested airports. *Transportation Res. Part C Emerging Tech.* 102:336–350.