The Price of Fairness

Dimitris Bertsimas, Vivek F. Farias, Nikolaos Trichakis

Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E40-139, Cambridge, Massachusetts 02139, {dbertsim@mit.edu, vivekf@mit.edu, nitric@mit.edu}

In this paper we study resource allocation problems that involve multiple self-interested parties or players, and a central decision maker. We introduce and study the price of fairness, which is the relative system efficiency loss under a "fair" allocation assuming that a fully efficient allocation is one that maximizes the sum of player utilities. We focus on two well accepted, axiomatically justified notions of fairness, viz. proportional fairness and max-min fairness. For these notions we provide a tight characterization of the price of fairness for a broad family of problems.

1. Introduction

In this paper we study the problem faced by a central decision maker of allocating a set of scarce resources among multiple self-interested parties or players. A solution that maximizes the sum of utilities of all the players might not be implementable, as some of the parties might consider it "unfair" in the sense that such a solution is achieved at the expense of some players. In many environments fairness might be more important than optimality. The overall objective of the paper is to study what we call "the price of fairness", that is the relative system efficiency loss under a "fair" allocation compared to the one that maximizes the sum of player utilities.

To concretely motivate the need for our study, let us consider the U.S. Federal Aviation Administration (FAA). The FAA is responsible for an important scheduling problem: it must generate precise schedules that determine not just when a particular flight might take off and land, but also what regions of U.S. airspace it might occupy over any given interval during its duration. The FAA must produce such a schedule for all flights and must dynamically adjust this schedule over the course of a given day to respond to unpredictable events, *e.g.*, inclement weather. Such a schedule allocates scarce resources, such as take-off and landing "slots" at airports, in a manner that respects flight plans. When a schedule must be recomputed due to an unforeseen event, this translates to ground and air-holding delays for flights. Since the estimated cost of such delays is very high (in the twelve-month period ending September 2008, 138 million system delay minutes drove an estimated USD 10 billion in direct operating costs for scheduled U.S. passenger airlines, see ATA (2008)), the importance of arriving at an effective schedule is apparent.

What do we mean by an effective schedule? Since delays (either on the ground or in the air) have well accepted dollar values, one natural notion of "effective" is a schedule that minimizes the total cost of delay to the airline industry. In fact, there is an extant body of research devoted to formulating and solving precisely this problem (see Odoni and Bianco (1987), Bertsimas and Stock-Patterson (1998), Bertsimas and Stock-Patterson (2000), Lulli and Odoni (2007)). While this work points at the possibility of dramatically reducing delay costs to the airline industry vis-a-vis current practice, the vast majority of these proposals remain un-implemented. The ostensible reason for this is *fairness*: the notion of equity is absent from consideration in the aforementioned proposals and while some of the stakeholders (namely, some airlines) clearly stand to gain from an

implementation of these proposals, other airlines might actually lose relative to the status quo. This apparent impasse, wherein a socially efficient solution, *i.e.*, one that maximizes the sum of utilities of individual players, is difficult to implement because it may be perceived as unfair to some of the stake-holders involved, is hardly unique to the air-traffic scheduling problem above. Indeed, issues of this sort arise in diverse scenarios ranging from the allocation of bandwidth in a communication network (see Bertsekas and Gallager (1987)) to the allocation of transaction costs among portfolios when a firm executes a large trade on behalf of multiple interested parties (see Fabozzi et al. (2007)). A great deal of thought has been invested in understanding, and axiomatically characterizing, what might constitute a "fair" allocation of resources. However, beyond qualitative economic analysis and with the exception of a handful of very special problems, there has been little work to quantitatively characterizing the tradeoffs inherent in employing these notions.

The present paper considers two axiomatically justified and well accepted notions of fairness in the context of general resource allocation problems whose solutions impact multiple players. We formulate the qualitative question on the *price of fairness* alluded to thus far, quantitatively: We take as our notion of socially optimal or efficient, an allocation that maximizes the natural *utilitarian* criterion (the sum of the utilities of individual players). We then define the price of fairness as the performance loss incurred relative to this criterion, in making allocations under either of the following fairness criteria: *max-min* fairness, and *proportional* fairness. We make the following contributions in regard to characterizing the price of fairness for general resource allocation problems:

1. We present bounds on the price of fairness for both max-min and proportional fairness that depend on a single parameter – the number of players. Our bounds are otherwise uniform over a broad class of resource allocation problems, namely, problems where the set of utilities individual players can simultaneously achieve is convex and compact.

2. Our bounds illustrate that a) The price of fairness as a function of the number of players is substantially smaller than a crude analysis might suggest, especially when the number of players is small and b) The price of proportional fairness is substantially smaller than the price of max-min fairness, especially when the number of players is large.

3. We show that our bounds on the price of fairness are *tight*; we do so by evaluating the price of fairness for examples of a well studied bandwidth allocation problem that arises in communication networks. These examples are by no means pathological. Further, we show that the class of resource allocation problems addressed by our bounds is, in a certain sense, the broadest class of problems we may hope to consider; the price of fairness for problems outside this class can be arbitrarily large.

To the best of our knowledge, the analysis undertaken here is the first of its kind. Our hope is that this analysis contributes to elucidating the precise tradeoffs one must make in allocating resources according to egalitarian criteria.

1.1. Relevant Literature

Applications: The importance of fairness issues in resource allocation problems has been recognized and well studied in a variety of settings. These range from engineering applications in communication networks (Bertsekas and Gallager (1987), Kelly et al. (1997), Luo et al. (2004), Luss (1999), Ogryczak et al. (2005), Radunovic and Boudec (2002), Radunovic and Boudec (2004)), the Air Traffic Flow Management problem (Bertsimas and Gupta (2010), Bertsimas et al. (2009a), Rios and Ross (2007), Soomer and Koole (2009), Vossen et al. (2003)), to financial applications and the multi-account optimization problem (Bertsimas et al. (2009b), Khodadadi et al. (2006), O'Cinneide et al. (2006)). In the communication network setting, where one must allocate bandwidth to flows in a network, a scheme that chooses to maximize throughput without regard to treating individual flows equitably is regarded as fully "efficient", and several studies address the efficiency loss due

to the incorporation of fairness considerations. These studies are typically numerical and focus on providing qualitative insights via studies of specific network topologies. Bonald and Massoulié (2001) introduce a number of network configurations where it is possible to derive performance results for proportional fairness, yet use simulation to assess max-min fairness. Radunovic and Boudec (2004) show that max-min fairness results in severe inefficiency for wireless networks in a limiting regime, and use numerical studies to validate that observation for practical situations. The impact of the fairness criteria utilized on the price of fairness has also received some attention: Mo and Walrand (2000) deal with this issue by studying a one parameter family of objectives that include both max-min fairness and proportional fairness as special cases. Our results imply a tight theoretical analysis of the loss in efficiency inherent in fair allocations of bandwidth in a communication network.

Worst Case Analysis & Approximation Algorithms: In recent work, Chakrabarty et al. (2009) seek to characterize what we refer to as the price of fairness for a specific class of resource allocation problems. In particular, that work shows that when the set of achievable "utilities" is a polymatroid, *all* Pareto resource allocations are efficient. This is, unfortunately, a somewhat restrictive condition and a general class of resource allocation problems that satisfy this condition is not known. In a similar vain, Butler and Williams (2002) show that the price of fairness is zero for a specialized facility location problem. Several pieces of work in the approximation algorithms literature have considered computing "approximately" fair solutions. Such work is motivated either by problems where fair solutions are difficult to compute, or else by the desire to simultaneously optimize several different objectives. For instance, Kleinberg et al. (1999) focus on the problem of approximating the max-min fair solution for routing and load balancing problems where the exact fair solution is hard to calculate. On the simultaneous optimization front, Kumar and Kleinberg (2000) discuss the existence of global c-approximation vectors (which are coordinate wise within a multiplicative factor of c of every other allocation) for bandwidth allocation, scheduling and facility location problems; the relevant value of c in each case is a function of problem primitives. The results of Goel et al. (2000) and Goel and Meyerson (2006) establish the existence of resource allocations that are simultaneously within a multiplicative factor of α for essentially all "fair" allocation criteria for general resource allocation problems of the type studied here; the authors show that α is logarithmic in the price of max-min fairness.

Commonly used notions of fairness, such as max-min fairness and proportional fairness arise from an appealing (and long standing) axiomatic characterization of what it means to be fair and an analogous characterization for "approximately" fair solutions is not available. It is thus difficult to judge what fairness properties (if any) such approximately fair solutions inherit. As an example, it is easily shown that by averaging the proportional fairness, max-min fairness, and utilitarian solutions to a resource allocation problem, one arrives at an allocation that is simultaneously within a multiplicative factor of 3 of the optimal solution for each those criteria for the class of resource allocation problems we consider; such a solution would be considered approximately fair in the aforementioned work but may not be Pareto. Seen in this light, our work studies the tradeoffs inherent in choosing a *fair* allocation as opposed to an approximation thereof.

Price of Anarchy: While in this work we assume that the utilities of the players are known and study the inefficiency that fairness constraints result in, another source of inefficiency could be the selfish behavior of players who do not truthfully reveal their utilities. The effect of selfish behavior has been studied as the price of anarchy in the literature. See Johari and Tsitsiklis (2004), Koutsoupias and Papadimitriou (1999), Papadimitriou (2001), Perakis (2007), Roughgarden and Tardos (2002) for more details.

Economic Theory: While we defer a thorough review of the literature in this area to Section 3, we mention for now that the question of what it means to be fair has been addressed extensively in the economics literature over the last century. In particular, see Young (1995) and Sen and Foster

(1997) for a thorough overview of this work. Fairness also plays a critical role in the selection of an appropriate social welfare function in welfare economics (see Mas-Colell et al. (1995)). The notions of fairness we focus on in this work are perhaps among the most prominently studied notions of fairness in the economics literature; for fundamental axiomatic characterizations of proportional and max-min fairness see Nash (1950) and Kalai and Smorodinsky (1975) respectively.

The structure of this paper is as follows. In Section 2, we introduce notation, focus on the socially optimal (taken as the sum of the utilities) and the fair allocations, and define the price of fairness. A general discussion on fairness schemes is included in Section 3. The main results of the paper are presented in Section 4, with illustrative examples given in Section 5. We conclude and point out interesting directions of future work in Section 6.

2. Problem Formulation

Consider a resource allocation problem involving n players and a central decision maker (CDM). There are some scarce resources that need to be allocated among the players by the CDM. According to her own preferences, each player derives a utility that depends on the allocation picked by the CDM. The preferences of each player are described by a *utility function*, which maps a feasible allocation into a utility level. We focus on problems where the CDM has complete knowledge of the preferences and possible constraints of the players, and has absolute control of the allocation decision.

To fix some notation, let $X \subset \mathbf{R}^m$ be the *resource set*, *i.e.*, the set of all feasible allocations of resources. An element $x \in X$ specifies a feasible allocation of resources among the players (*e.g.*, x might be the concatenation of n k-dimensional vectors that describe the quantities of k different resources allocated to each of n players; in this case $m = nk$). Note that the resource set also incorporates all constraints on allocations such as resource capacity constraints, individual limitations of the players or the central decision maker, etc. With the jth player, we associate a utility function $f_j : X \to \mathbf{R}_+$, for every $j = 1, \ldots, n$. If the CDM picks allocation x, the jth player derives a utility of $f_j(x)$. Finally, let U be the *utility set*, that is the set of all achievable utility allocations, or distributions:

$$U = \left\{ u \in \mathbf{R}_+^n \,\middle|\, \exists\, x \in X : f_j(x) = u_j, \,\forall j = 1, \ldots, n \right\}.$$

EXAMPLE 1. As a concrete example, consider two resources, denoted by A and B, being allocated among two players, denoted by 1 and 2 (*i.e.*, in this case $n = 2$). Let x_{A1} and x_{A2} be the fractions of the available resource A allocated to players 1 and 2 respectively; x_{B1} and x_{B2} are defined similarly. The resource set is then

$$X = \left\{ \begin{bmatrix} x_{A1} & x_{B1} & x_{A2} & x_{B2} \end{bmatrix}^T \in \mathbf{R}_+^4 \,\middle|\, x_{A1} + x_{A2} \leq 1, x_{B1} + x_{B2} \leq 1 \right\},$$

with $m = 4$. Assume that the utility derived by each player is equal to the square-root of the sum of the fractions of each resource allocated to him (*i.e.*, $f_j(x) = \sqrt{x_{Aj} + x_{Bj}}$ for $j = 1, 2$). The utility set in this case is thus

$$U = \left\{ \begin{bmatrix} \sqrt{x_{A1} + x_{B1}} & \sqrt{x_{A2} + x_{B2}} \end{bmatrix}^T \,\middle|\, x \in X \right\} = \left\{ \begin{bmatrix} u_1 & u_2 \end{bmatrix}^T \in \mathbf{R}_+^2 \,\middle|\, u_1^2 + u_2^2 \leq 2 \right\}.$$

Returning to our general formulation, the CDM's problem is to decide on a utility allocation among the players, $u \in U$. A good bit of research, notably in welfare economics, has dealt with the identification of the appropriate criteria that the CDM needs to take into account in order to make a decision (see Mas-Colell et al. (1995)). We next discuss the *utilitarian* criterion and fair criteria for such allocations.

2.1. Utilitarian solution

Under the classical utilitarian principle, the central decision maker picks an allocation that maximizes the sum of the utilities of the players. That is, the CDM decides on the allocation by solving the problem

$$\text{maximize } e^T u$$
$$\text{subject to } u \in U,$$

with variable $u \in \mathbf{R}^n$, where e is the vector of all ones. We denote the optimal value of this problem with SYSTEM(U), *i.e.*,

$$\text{SYSTEM}(U) = \sup \left\{ e^T u \,\middle|\, u \in U \right\}.$$

The resulting allocation is then the *utilitarian solution*. It is referred to also as the Bentham-Edgeworth solution in welfare economics, as the system optimum solution in engineering applications and as the best effort solution in telecommunications.

The utilitarian solution is a natural choice in applications where the sum of the utilities corresponds to some measure of system efficiency. For example, consider a communications network where a service provider controls the transmission rates allocated to clients, subject to capacity constraints. The service provider plays the role of the CDM, and the clients are the players in our setting. In case the utility that each player derives is equal to his transmission rate, the sum of the utilities corresponds to the total throughput of the network.

On the other hand, the sum of utilities is neutral towards potential inequalites in the utility distribution among the players. It is therefore possible that the utilitarian solution is achieved at the expense of some players.

2.2. Fair solutions

Alternatively to classical utilitarianism, the central decision maker might decide on the utility allocation incorporating fairness considerations. Depending on the nature of the problem and her own perception about fairness, the CDM picks a fairness scheme of her preference, that is, a set of rules or properties (*e.g.*, total equity, under which every player derives exactly the same utility). The selected allocation then needs to be compatible with the fairness scheme.

To make this more precise, we model a fairness scheme as a set of rules and a corresponding set function $\mathcal{S}: 2^{\mathbf{R}_+^n} \to \mathbf{R}_+^n$, that takes a utility set as an input, and maps it into an element of the utility set. Given a utility set U, $\mathcal{S}(U) \in U$ is then an allocation that abides to the set of rules of the fairness scheme in consideration.

By imposing a specific fairness scheme and deciding on a fair allocation, the sum of utilities in the system might, and in most cases will, decrease compared to the utilitarian solution. In case the sum of utilities corresponds to an efficiency measure of the system, fairness constraints might impose then a performance or efficiency loss. Let FAIR$(U; \mathcal{S})$ denote the sum of utilities under the fair allocation imposed by a fairness scheme $\mathcal{S}$, *i.e.*,

$$\text{FAIR}(U; \mathcal{S}) = e^T \mathcal{S}(U).$$

We define the *price of fairness*, denoted by POF$(U; \mathcal{S})$, for the problem involving the utility set U and the fairness scheme $\mathcal{S}$, to be the relative reduction in the sum of utilities under the fair solution $\mathcal{S}(U)$, compared to the utilitarian solution, *i.e.*,

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

Note that the price of fairness is a number between zero and one, since the sum of utilities under the utilitarian solution attains its maximum value. When the sum of utilities is an efficiency measure,

values closer to zero are preferable for the price of fairness, since the CDM can then combine high system efficiency and fairness.

The scope of this work is to quantify the price of fairness for a large family of problems. We first discuss fairness schemes in the next section, review the two most prominent schemes, proportional and max-min fairness, and then present our main results.

3. Fairness Schemes

Fairness in allocation problems has been extensively studied through the years in many areas, notably in social sciences, welfare economics and engineering. A plethora of fairness criteria have been proposed. Due to multiple (subjective) interpretations of the concepts of fairness, and the different characteristics of allocation problems, there is no single principle that is universally accepted. Nevertheless, there are general theories of justice and equity that figure prominently in the literature, on which most fairness schemes are based. Moreover, there has been a body of literature that deals with axiomatic foundations of the concepts of fairness. In this section, we briefly review the most important theories and axioms, and then focus on proportional and max-min fairness, the two criteria that emerge from the axiomatic foundations and are also widely used in practice. For more details, see Young (1995) and Sen and Foster (1997).

Among the most prominent, the oldest theory of justice is Aristotle's equity principle, according to which, resources should be allocated in proportion to some pre-existing claims, or rights to the resources that each player has. Another theory, widely considered in economics in the 19th century, is classical utilitarianism, which dictates an allocation of resources that maximizes the sum of utilities (see Section 2.1). A third approach is due to Rawls (1971). The key idea of Rawlsian justice is to give priority to the players that are the least well off, so as to guarantee the highest minimum utility level that every player derives. Finally, Nash introduced the Nash standard of comparison, which is the percentage change in a player's utility when he receives a small additional amount of the resources. A transfer of resources between two players is then justified, if the gainer's utility increases by a larger percentage than the loser's utility decreases.

Aristotle's equity principle is used in the majority of cases where players have specific pre-existing claims or rights to the resources (for example, split of profits among shareholders). In this work, we do not deal with such cases, hence the Aristotelian principle does not apply. The utilitarian principle has been criticized (see Young (1995)) since it is not clear that it is ethically sound: in maximizing the sum of utilities, the utility of some players might be greatly reduced in order to confer a benefit to the system. Finally, the two schemes that we will focus on are based on the Rawlsian justice and the Nash standard, which are in line with the common perception of equity and fairness.

In addition to using theories of justice and common perception, researchers have also established sets of axioms that a fairness scheme should ideally satisfy. The main work in this area is within the literature of fair bargains in economics (see Young (1995) and references therein). We now briefly present the most well studied set of axioms in the case of a two-player problem ($n = 2$). In the axioms that follow, we denote the utility set U and define the *maximum achievable utility* of the jth player, $u_j^\star$, according to

$$u_j^\star = \sup\{u_j \mid u \in U\}.$$

We utilize the notation $a \geq b$ for $a, b \in \mathbf{R}^n$ to denote $a_i \geq b_i$, $i = 1, \ldots, n$. Furthermore, if $g : \mathbf{R}^n \to \mathbf{R}^n$ is an operator and $A \subset \mathbf{R}^n$ is a set, then $g(A) = \{g(x) \mid x \in A\} \subset \mathbf{R}^n$.

Axiom 1 *(Pareto Optimality) The fair solution $\mathcal{S}(U)$ is Pareto optimal, that is, there does not exist an allocation $u \in U$, such that $u \geq \mathcal{S}(U)$ and $u \neq \mathcal{S}(U)$.*

Axiom 2 *(Symmetry)* If $\mathcal{I} : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ *is a permutation operator defined by* $\mathcal{I}((u_1, u_2)) = (u_2, u_1)$, *then the fair allocation under the permuted system,* $\mathcal{S}(\mathcal{I}(U))$, *is equal to the permutation of the fair allocation under the original system,* $\mathcal{I}(\mathcal{S}(U))$. *That is,* $\mathcal{S}(\mathcal{I}(U)) = \mathcal{I}(\mathcal{S}(U))$.

Axiom 3 *(Affine invariance)* If $A : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ *is an affine operator defined by* $A(u_1, u_2) = (A_1(u_1), A_2(u_2))$, *with* $A_i(u) = c_i u + d_i$ *and* $c_i > 0$, *then the fair allocation under the affinely transformed system,* $\mathcal{S}(A(U))$, *is equal to the affine transformation of the fair allocation under the original system,* $A(\mathcal{S}(U))$. *That is,* $\mathcal{S}(A(U)) = A(\mathcal{S}(U))$.

Axiom 4 *(Independence of irrelevant alternatives)* If U *and* W *are two utility sets such that* $U \subset W$, *and* $\mathcal{S}(W) \in U$, *then* $\mathcal{S}(U) = \mathcal{S}(W)$.

Axiom 5 *(Monotonicity)* Let U *and* W *be two utility sets, under which the maximum achievable utility of player* 1 *is equal, i.e.,* $u_1^\star = w_1^\star$. *If for every utility level that player* 1 *may demand, the maximum achievable utility that player* 2 *can derive simultaneously, is bigger or equal under* W, *then the utility level of player* 2 *under the fair allocation should also be bigger or equal under* W, *i.e.,* $\mathcal{S}(U)_2 \leq \mathcal{S}(W)_2$.

Pareto optimality ensures that there is no wastage. By symmetry, the central decision maker does not differentiate the players by their names. The affine invariance requirement means that the scheme is invariant to a choice of numeraire. According to the independence of irrelevant alternatives, preferring option A over option B is independent of other available options. Finally, by monotonicity, if for every utility level that player 1 may demand, the maximum utility level that player 2 can simultaneously derive is increased, then the utility level assigned to player 2 under the fair scheme should also be increased. For a more detailed discussion about monotonicity, see Kalai and Smorodinsky (1975).

The main result in this area is that, under mild assumptions on the utility set, there does not exist a scheme that satisfies all axioms; see Nash (1950) and Kalai and Smorodinsky (1975) for more details. Moreover, the unique scheme that satisfies Axioms 1-4 is the *Nash solution*; the unique scheme that satisfies Axioms 1-3, and 5 is the *Kalai-Smorodinsky solution*. Proportional and max-min fairness are direct generalizations of those schemes, and are studied next.

3.1. Proportional Fairness

Proportional fairness (PF) is the generalization of the Nash solution for a two-player problem. The Nash solution is the unique scheme that satisfies Axioms 1-4, and is based on the Nash standard of comparison. Under the Nash standard, a transfer of resources between two players is favorable and fair if the percentage increase in the utility of one player is larger than the percentage decrease in utility of the other player. Proportional fairness is the generalized Nash solution for multiple players. In that setting, the fair allocation should be such that, if compared to any other feasible allocation of utilities, the aggregate proportional change is less than or equal to zero. In mathematical terms,

$$\sum_{j=1}^{n} \frac{u_j - \mathcal{S}^{\mathrm{PF}}(U)_j}{\mathcal{S}^{\mathrm{PF}}(U)_j} \leq 0, \quad \forall u \in U.$$

In case U is convex, the fair allocation under proportional fairness $\mathcal{S}^{\mathrm{PF}}(U)$ can be obtained as the (unique) optimal solution of the problem

$$\begin{aligned} \text{maximize } & \sum_{j=1}^{n} \log u_j \\ \text{subject to } & u \in U, \end{aligned}$$

since the necessary and sufficient first order optimality condition for this problem is exactly the Nash standard of comparison principle for n players. Moreover, note that the proportionally fair allocation $\mathcal{S}^{\mathrm{PF}}(U)$, can also be obtained by (equivalently) maximizing the product of the utilities over U. This suggests that proportional fairness yields a Pareto optimal allocation and is also *scale-invariant*. In particular, we use the notation $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)$ to denote a diagonal matrix with entries $\sigma_1, \ldots, \sigma_n$ in the diagonal. We define the *scaled* utility set ΣU, with $\sigma_j > 0$ for all $j = 1, \ldots, n$, as

$$\Sigma U = \{\Sigma u \,|\, u \in U\}.$$

Then,

$$\mathcal{S}^{\mathrm{PF}}(\Sigma U) = \Sigma \mathcal{S}^{\mathrm{PF}}(U), \tag{1}$$

that is, the fair allocation under the scaled utility set, $\mathcal{S}^{\mathrm{PF}}(\Sigma U)$, is equal to the scaled fair allocation under the original utility set, $\Sigma \mathcal{S}^{\mathrm{PF}}(U)$.

Proportional fairness has been extensively studied and used in the areas of telecommunications and networks, especially after the paper of Kelly et al. (1997).

3.2. Max-Min Fairness

Max-min fairness (MMF) is a generalization of the Rawlsian justice and the Kalai-Smorodinsky (KS) solution in the two-player problem. The KS solution is the unique solution that satisfies Axioms 1-3, and 5. In settings where the maximum achievable utility levels of the two players are equal, the KS solution corresponds to maximizing the minimum utility the players derive simultaneously. Otherwise, the central decision maker decides on the allocation in the same way, but by considering a scaled, normalized system, under which the players have equal maximum achievable utility levels. In other words, under the KS solution the players simultaneously derive the largest possible equal fraction of their respective maximum achievable utilities. For simplicity, for the rest of this section, we deal with normalized problems where the players have equal maximum achievable utilities.

In a setting that involves more than two players, such an allocation may not be Pareto optimal, thus indicating a waste of resources. That can happen for instance in case there exist players that can derive higher utility levels without affecting the others, and their allocated resources are not optimized. Max-min fairness generalizes the above criteria to account for this potential loss of efficiency, and always yields Pareto optimal allocations.

Under max-min fairness, the central decision maker tries at the first step to maximize the lowest utility level among all the players. After ensuring that all players derive (at least) that level, the second lowest utility level among the players is maximized, and so on. The resulting allocation yields a distribution of utility levels among the players that has the following property: the distribution of the utility levels of any other allocation that achieves a strictly higher utility for a specific level, is such that there exists a lower level of utility that has been strictly decreased. In other words, any other allocation can only benefit the rich at the expense of the poor (in terms of utility).

Intuitively, max-min fairness maximizes the minimum utility that all players derive. In situations where an efficient allocation exists that results in equal utility for all players, MMF converges to this equitable allocation. In cases where some players can achieve higher utility levels, without depriving others of the minimum utility performance, MMF equitably and efficiently allows them to increase their utility, in a similar fashion, by maximizing a new minimum utility level that all improving players derive.

In mathematical terms, let $T : \mathbf{R}^n \to \mathbf{R}^n$ be the sorting operator, that is

$$T(y) = \left(y_{(1)}, \ldots, y_{(n)}\right), \quad y_{(1)} \leq \ldots \leq y_{(n)},$$

where $y_{(i)}$ is the ith smallest element of y. We say that $a \in \mathbf{R}^n$ is lexicographically larger than $b \in \mathbf{R}^n$ and write symbolically $a \succ_{\text{lex}} b$, if there exists an index $k \leq n$, such that $a_i = b_i$, $\forall i < k$, and $a_k > b_k$. Also, we write $a \succeq_{\text{lex}} b$ if $a \succ_{\text{lex}} b$ or $a = b$. The max-min fairness scheme then corresponds to lexicographically maximizing $T(u)$ over U, that is, finding an allocation $u^{\text{MMF}} \in U$ such that its resulting sorted utility distribution is lexicographically largest among all sorted feasible utility distributions. We then have

$$T(u^{\text{MMF}}) \succeq_{\text{lex}} T(u), \quad \forall u \in U.$$

The existence of a max-min fair allocation is guaranteed under mild conditions (*e.g.*, if U is compact), and the Pareto optimality of the allocation follows by its construction, see Radunovic and Boudec (2002) for more details. Efficient algorithms for computing an MMF allocation have also been developed and studied in the literature. The computations involve a sequential optimization procedure, that identifies the corresponding utility levels at each step. For more details, see Ogryczak et al. (2005).

Since the max-min fairness scheme deals with normalized utilities, it is also scale-invariant, hence if $\Sigma = \text{diag}(\sigma_1, \ldots, \sigma_n)$, with $\sigma_j > 0$ for all $j = 1, \ldots, n$, then

$$\mathcal{S}^{\text{MMF}}(\Sigma U) = \Sigma \mathcal{S}^{\text{MMF}}(U), \tag{2}$$

that is, the fair allocation under the scaled system, $\mathcal{S}^{\text{MMF}}(\Sigma U)$, is equal to the scaled fair allocation under the original system, $\Sigma \mathcal{S}^{\text{MMF}}(U)$.

Max-min fairness was first implemented in networking and telecommunications applications and has also initiated a lot of research in this area (see Bertsekas and Gallager (1987), Bonald and Massoulié (2001), Luss (1999)). It has many applications in bandwidth allocation, routing and load balancing problems, as well as in general resource allocation or multiobjective optimization problems.

4. The Price of Fairness

In this section, we present the main results of this paper, namely upper bounds for the price of fairness, under the proportional and the max-min fairness schemes, which depend only on the number of players, and their maximum achievable utilities (in case they are not equal).

Consider a resource allocation problem faced by a central decision maker as described in Section 2. We make the following assumption:

ASSUMPTION 1. *The utility set is compact and convex.*

Assumption 1 is very common in the literature of fair bargains. Compactness follows from bounded and continuous utility functions and a compact resource set. Convexity arises in cases of a randomization mechanism that may be employed, especially if the resources are indivisible, with the utility levels then corresponding to the expected utilities levels derived (see Young (1995)). Furthermore, convexity also follows in many cases where the utility functions are concave and non-decreasing. The following proposition suggests a rich family of problems that fits into this framework.

PROPOSITION 1. *Let the resource set $X \subset \mathbf{R}^m_+$ be compact, convex and monotone (a set $A \subset \mathbf{R}^m_+$ is called monotone if $\{b \in \mathbf{R}^m \mid 0 \leq b \leq a\} \subset A$, $\forall a \in A$). Suppose that the utility function of the jth player is such that $f_j(x) = \bar{f}_j(x_j)$, for all $x \in X$, with $\bar{f}_j : \mathbf{R}^{m_j} \to \mathbf{R}$, and $x^T = \begin{bmatrix} x_1^T & x_2^T & \ldots & x_n^T \end{bmatrix}$, where $m_1 + \ldots + m_n = m$. Moreover, $\bar{f}_j$ is non-decreasing in each argument, concave, bounded and continuous over X, and $\bar{f}_j(0) = 0$. Then, the resulting utility set U is compact, convex and monotone.*

Proof. Let $f: X \to \mathbf{R}_+^n$ be the vector of utility functions, *i.e.*,

$$f(x) = \left[\, \bar{f}_1(x_1) \; \bar{f}_2(x_2) \; \ldots \; \bar{f}_n(x_n) \right]^T, \quad \forall x \in X.$$

Since X is compact, f is continuous and bounded over X, it follows that U is compact.

To show monotonicity, let $u \in U$. Then, $\exists x \in X$, such that $f(x) = u$. Consider now an allocation u', such that $0 \le u' \le u$. For any j, let $g_j(\lambda) = \bar{f}_j(\lambda x_j)$, for $0 \le \lambda \le 1$. Since $\bar{f}_j$ is continuous and non-decreasing, so is g_j. Given also that $g_j(0) = 0$, $g_j(1) = u_j$, and that $0 \le u'_j \le u_j$, it follows that $\exists \lambda_j \in [0,1]$, such that $g_j(\lambda_j) = \bar{f}_j(\lambda_j x_j) = u'_j$. Note that by monotonicity of X, $z = \left[\lambda_1 x_1^T \; \lambda_2 x_2^T \; \ldots \; \lambda_n x_n^T \right]^T \in X$. But, $f(z) = u'$, which shows that $u' \in U$ and U is monotone.

To show convexity, let $\tilde{u} \in U$; then, $\exists \tilde{x} \in X$, such that $f(\tilde{x}) = \tilde{u}$. Let $\theta \in [0,1]$. By convexity of X, $\theta x + (1 - \theta)\tilde{x} \in X$. By concavity of $\bar{f}_j$,

$$f(\theta x + (1 - \theta)\tilde{x}) \ge \theta f(x) + (1 - \theta)f(\tilde{x}) = \theta u + (1 - \theta)\tilde{u} \ge 0.$$

Since U is monotone, it follows that $\theta u + (1 - \theta)\tilde{u} \in U$, and U is convex. $\square$

Note that, although not exhaustive, Proposition 1 shows that a general class of problems satisfy Assumption 1. The compactness and convexity requirements on the resource set X are common and support a broad family of problems. The monotonicity requirement is satisfied in case of freely disposable physical resources, *i.e.*, when an allotted resource can be reduced or nullified, without necessarily affecting the rest of the allocation. Finally, the requirements on the utility function depending only on the allocation of each player, being concave and non-decreasing are also well studied in the literature (see Mas-Colell et al. (1995), Young (1995)).

Examples in the next section indicate that in the absense of the convexity assumption, the worst case price of fairness can get arbitrarily close to 1, even for a two-player problem.

We now provide the main results of this paper, for the case where the maximum achievable utilities of the players are equal and for the case they are not.

4.1. Equal maximum achievable utilities

The following Theorem provides upper bounds for the price of fairness, in case of equal maximum achievable utilities.

THEOREM 1. *Consider a resource allocation problem with n players, with $n \ge 2$. Let the utility set, denoted by $U \subset \mathbf{R}_+^n$, satisfy Assumption 1. If all players have equal maximum achievable utilities, which are greater than zero,*

(a) the price of proportional fairness is bounded by

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{PF}}) \le 1 - \frac{2\sqrt{n} - 1}{n},$$

(b) the price of max-min fairness is bounded by

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{MMF}}) \le 1 - \frac{4n}{(n+1)^2}.$$

Moreover, the bound under proportional fairness is tight if $\sqrt{n} \in \mathbf{N}$, and the bound under max-min fairness is tight for all n.

Proof. By assumption, the players have equal maximum achievable utilities. We assume further that they are equal to 1, *i.e.*,

$$u_j^\star = \max \left\{ u_j \mid u \in U \right\} = 1, \quad \forall j = 1, \ldots, n. \tag{3}$$

Figure 1   An example of a two-dimensional utility set, with the points of interest and the associated supporting hyperplanes used in the proof of Theorem 1.
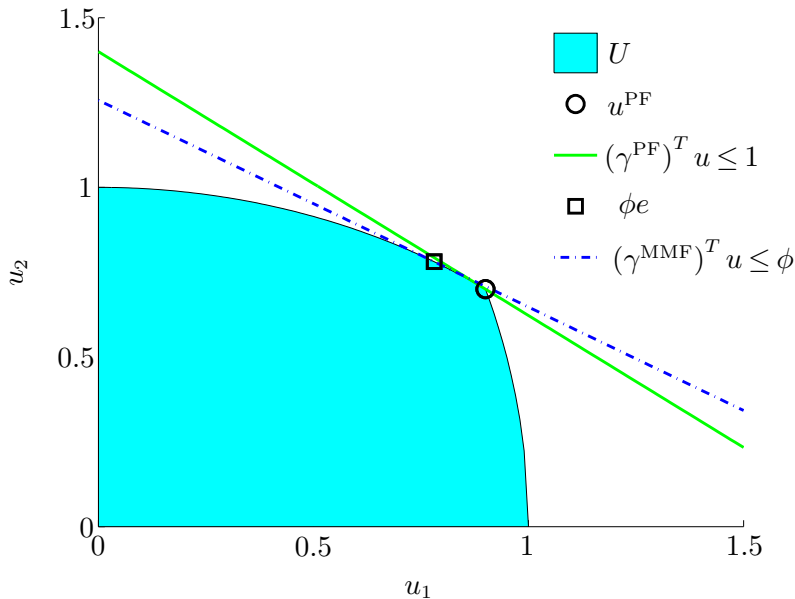
This is without loss of generality, and can be achieved simply by scaling. As a result,

$$0 \leq u \leq e, \quad \forall u \in U. \tag{4}$$

Without loss of generality, we assume that U is monotone. This is because all schemes we consider, namely utilitarian, proportional and max-min fairness yield Pareto optimal allocations. In particular, suppose there exist allocations $a \in U$ and $b \notin U$, with allocation a dominating allocation b, *i.e.*, $0 \leq b \leq a$. Note that allocation b can thus not be Pareto optimal. Then, we can equivalently assume that $b \in U$, since b cannot be selected by any of the schemes.

Note that the monotonicity assumption and (3) also imply that $0 \in U$ and $e_j \in U$ for all $j = 1, \ldots, n$. By Assumption 1, we also have $\frac{1}{n} e \in U$ (by convexity).

(a) *Proportional fairness.* Let $u^{\mathrm{PF}} \in U$ be the utility distribution under the proportionally fair solution. By definition, we have

$$\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{PF}}) = e^T \mathcal{S}^{\mathrm{PF}}(U) = e^T u^{\mathrm{PF}}. \tag{5}$$

By the first order optimality condition (see Section 3.1), we have

$$\sum_{j=1}^{n} \frac{u_j - u_j^{\mathrm{PF}}}{u_j^{\mathrm{PF}}} \leq 0, \quad \forall u \in U.$$

Equivalently,

$$\left(\gamma^{\mathrm{PF}}\right)^T u \leq 1, \quad \forall u \in U, \tag{6}$$

where

$$\gamma_j^{\mathrm{PF}} = \frac{1}{n u_j^{\mathrm{PF}}}. \tag{7}$$

This defines a hyperplane that supports U at u^{PF}. Figure 1 illustrates u^{PF} and the hyperplane in the case of a two-dimensional example.

Since $u^{\mathrm{PF}} \in U$, using (4) we have that $u_j^{\mathrm{PF}} \leq 1 \Rightarrow \gamma_j^{\mathrm{PF}} \geq \frac{1}{n}$, for all j. Moreover, since $e_j \in U$ for all j, using (6) we have $(\gamma^{\mathrm{PF}})^T e_j \leq 1 \Rightarrow \gamma_j^{\mathrm{PF}} \leq 1$. Without loss of generality, we also assume that the elements of γ^{PF} are ordered. To summarize, we have

$$\frac{1}{n} \leq \gamma_1^{\mathrm{PF}} \leq \ldots \leq \gamma_n^{\mathrm{PF}} \leq 1. \tag{8}$$

The supporting hyperplane we identified can now be used to bound the sum of utilities under the utilitarian solution. In particular, using (4) and (6) we get that

$$\begin{aligned}
\mathrm{SYSTEM}(U) &= \max\left\{ e^T u \,\middle|\, u \in U \right\} \\
&\leq \max\left\{ e^T u \,\middle|\, 0 \leq u \leq e, \left(\gamma^{\mathrm{PF}}\right)^T u \leq 1 \right\},
\end{aligned} \tag{9}$$

where the right hand side is the optimal value of the linear relaxation of the well-studied knapsack problem, a version of which we review next.

Let $w \in \mathbf{R}^n$ and $B \in \mathbf{R}$ be such that $0 \leq w_1 \leq \ldots \leq w_n \leq B$, $e^T w \geq 1$, $\frac{1}{n} \leq B \leq 1$. Then, one can show (see Bertsimas and Tsitsiklis (1997)) that the linear program

$$\begin{aligned}
\text{maximize } & e^T y \\
\text{subject to } & w^T y \leq B \\
& 0 \leq y \leq e,
\end{aligned} \tag{10}$$

has an optimal value equal to $\ell(w, B) + \delta(w, B)$, where

$$\ell(w, B) = \max\left\{ i \,\middle|\, \sum_{j=1}^{i} w_j \leq B, \, i \leq n - 1 \right\} \in \{1, \ldots, n-1\} \tag{11}$$

$$\delta(w, B) = \frac{B - \sum_{j=1}^{\ell(w,B)} w_j}{w_{\ell(w,B)+1}} \in [0, 1]. \tag{12}$$

Using this observation, we can rewrite (9) as

$$\mathrm{SYSTEM}(U) \leq \ell(\gamma^{\mathrm{PF}}, 1) + \delta(\gamma^{\mathrm{PF}}, 1). \tag{13}$$

We can now provide an upper bound to the price of fairness:

$$\begin{aligned}
\mathrm{POF}(U; \mathcal{S}^{\mathrm{PF}}) &= \frac{\mathrm{SYSTEM}(U) - \mathrm{FAIR}(U; \mathcal{S}^{\mathrm{PF}})}{\mathrm{SYSTEM}(U)} \\
&= 1 - \frac{\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{PF}})}{\mathrm{SYSTEM}(U)} \\
&= 1 - \frac{\sum_{j=1}^{n} z_j^{\mathrm{PF}}}{\mathrm{SYSTEM}(U)} && \text{(from (5))} \\
&= 1 - \frac{\sum_{j=1}^{n} \frac{1}{n \gamma_j^{\mathrm{PF}}}}{\mathrm{SYSTEM}(U)} && \text{(from (7))} \\
&\leq 1 - \frac{\sum_{j=1}^{n} \frac{1}{n \gamma_j^{\mathrm{PF}}}}{\ell(\gamma^{\mathrm{PF}}, 1) + \delta(\gamma^{\mathrm{PF}}, 1)}. && \text{(from (13))}
\end{aligned}$$

Let $g : \mathbf{R}^n \to \mathbf{R}$ be defined as

$$g(\gamma) = \frac{\sum_{j=1}^{n} \frac{1}{n \gamma_j}}{\ell(\gamma, 1) + \delta(\gamma, 1)}.$$

Using this definition and (8), the bound can now be rewritten as

$$\text{POF}(U; \mathcal{S}^{\text{PF}}) \le 1 - g\left(\gamma^{\text{PF}}\right) \le 1 - \inf_{\frac{1}{n} \le \gamma_1 \le \ldots \le \gamma_n \le 1} g(\gamma),$$

and it suffices to show that

$$F_1 = \inf_{\frac{1}{n} \le \gamma_1 \le \ldots \le \gamma_n \le 1} g(\gamma) \ge \frac{2\sqrt{n} - 1}{n}.$$

Let $p : \mathbf{R}^2 \to \mathbf{R}$ be defined as

$$p(y) = \frac{\frac{y_1}{y_2} + n - y_1}{n y_1},$$

and

$$F_2 = \inf_{\substack{y_1 y_2 \le 1 \\ 1 \le y_1 \le n \\ \frac{1}{n} \le y_2 \le 1}} p(y).$$

We will first show that $F_1 \ge F_2$. To do that, it is sufficient to show that for any γ such that $\frac{1}{n} \le \gamma_1 \le \ldots \le \gamma_n \le 1$, there exists a $y \in \mathbf{R}^2$, such that $y_1 y_2 \le 1$, $1 \le y_1 \le n$, $\frac{1}{n} \le y_2 \le 1$, and $g(\gamma) \ge p(y)$. Let $y_1 = \ell(\gamma, 1) + \delta(\gamma, 1)$. By the ranges of $\ell(\gamma, 1)$ and $\delta(\gamma, 1)$, it follows that $1 \le y_1 \le n$. Moreover, let

$$y_2 = \frac{y_1}{\frac{1}{\gamma_1} + \ldots + \frac{1}{\gamma_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}}}.$$

Since $\gamma_j \ge \frac{1}{n}$, we get

$$y_2 = \frac{y_1}{\frac{1}{\gamma_1} + \ldots + \frac{1}{\gamma_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}}} \ge \frac{y_1}{n(\ell(\gamma,1) + \delta(\gamma,1))} = \frac{1}{n}.$$

A similar argument utilizing that $\gamma_j \le 1$ shows that $y_2 \le 1$. To show that $y_1 y_2 \le 1$, consider the following convex optimization problem:

$$\begin{aligned} \text{minimize} \quad & \frac{1}{v_1} + \ldots + \frac{1}{v_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{v_{\ell(\gamma,1)+1}} \\ \text{subject to} \quad & v_1 + \ldots + v_{\ell(\gamma,1)} + \delta(\gamma,1) v_{\ell(\gamma,1)+1} = 1 \\ & v \ge 0, \end{aligned}$$

with variable $v \in \mathbf{R}^{\ell(\gamma,1)+1}$. Note that γ is feasible for this problem, since by (12) we have

$$\gamma_1 + \ldots + \gamma_{\ell(\gamma,1)} + \delta(\gamma,1) \gamma_{\ell(\gamma,1)+1} = 1.$$

We will show that

$$\bar{v} = \frac{1}{\ell(\gamma,1) + \delta(\gamma,1)} e$$

is an optimal solution. Feasibility is immediate, and the necessary and sufficient first order optimaltiy conditions are also satisfied: Noting that $\bar{v}_1 = \bar{v}_j$ for all $j = 1, \ldots, \ell(\gamma,1) + 1$, we have that for any $v \ge 0$, with $v_1 + \ldots + v_{\ell(\gamma,1)} + \delta(\gamma,1) v_{\ell(\gamma,1)+1} = 1$,

$$\sum_{j=1}^{\ell(\gamma,1)} \frac{(\bar{v}_j - v_j)}{\bar{v}_j^2} + \frac{\delta(\gamma,1)\left(\bar{v}_{\ell(\gamma,1)+1} - v_{\ell(\gamma,1)+1}\right)}{\bar{v}_{\ell(\gamma,1)+1}^2} =$$

$$\frac{1}{\bar{v}_1^2}\left(\left(\bar{v}_1 + \ldots + \bar{v}_{\ell(\gamma,1)} + \delta(\gamma,1)\bar{v}_{\ell(\gamma,1)+1}\right) - \left(v_1 + \ldots + v_{\ell(\gamma,1)} + \delta(\gamma,1)v_{\ell(\gamma,1)+1}\right)\right) = 0.$$

Since γ is feasible and $\bar{v}$ optimal, it follows that

$$
\begin{aligned}
\frac{y_1}{y_2} &= \frac{1}{\gamma_1} + \ldots + \frac{1}{\gamma_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}} \\
&\geq \frac{1}{\bar{v}_1} + \ldots + \frac{1}{\bar{v}_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{\bar{v}_{\ell(\gamma,1)+1}} \\
&= \frac{\ell(\gamma,1) + \delta(\gamma,1)}{\bar{v}_1} = (\ell(\gamma,1) + \delta(\gamma,1))^2 = y_1^2.
\end{aligned}
$$

Finally,

$$
\begin{aligned}
g(\gamma) &= \frac{\sum_{j=1}^n \frac{1}{n\gamma_j}}{\ell(\gamma,1) + \delta(\gamma,1)} \\
&= \frac{\frac{1}{\gamma_1} + \ldots + \frac{1}{\gamma_{\ell(\gamma,1)}} + \frac{\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}} + \frac{1-\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}} + \frac{1}{\gamma_{\ell(\gamma,1)+2}} + \ldots + \frac{1}{\gamma_n}}{n\left(\ell(\gamma,1) + \delta(\gamma,1)\right)} \\
&= \frac{\frac{y_1}{y_2} + \frac{1-\delta(\gamma,1)}{\gamma_{\ell(\gamma,1)+1}} + \frac{1}{\gamma_{\ell(\gamma,1)+2}} + \ldots + \frac{1}{\gamma_n}}{ny_1} \\
&\geq \frac{\frac{y_1}{y_2} + n - \ell(\gamma,1) - \delta(\gamma,1)}{ny_1} \qquad\qquad\qquad \text{(from (8))} \\
&\geq \frac{\frac{y_1}{y_2} + n - y_1}{ny_1} = p(y).
\end{aligned}
$$

We now evaluate F_2:

$$
F_2 = \inf_{\substack{y_1 y_2 \leq 1 \\ 1 \leq y_1 \leq n \\ \frac{1}{n} \leq y_2 \leq 1}} \frac{\frac{y_1}{y_2} + n - y_1}{ny_1} = \inf_{\substack{y_1 y_2 \leq 1 \\ 1 \leq y_1 \leq n \\ \frac{1}{n} \leq y_2 \leq 1}} \left(\frac{1}{ny_2} + \frac{1}{y_1} - \frac{1}{n} \right).
$$

Clearly, the infimum is attained, and at the optimum $y_1 y_2 = 1$, *i.e.*, $\frac{1}{y_2} = y_1$, and

$$
F_2 = \inf_{1 \leq y_1 \leq n} \left(\frac{y_1}{n} + \frac{1}{y_1} - \frac{1}{n} \right) = \frac{2\sqrt{n} - 1}{n}.
$$

The proof is complete by noting that $F_1 \geq F_2$. Section 5 includes examples that show that the bound is tight in case $\sqrt{n} \in \mathbf{N}$.

(b) *Max-min fairness.* Consider the ray re, $r \geq 0$. Since $0 \in U$ and $\frac{1}{n}e \in U$, by convexity of U we have that $re \in U$, for $0 \leq r \leq \frac{1}{n}$. Since $U \subset [0,1]^n$ is compact, there exists a $\phi \in \left[\frac{1}{n}, 1\right]$ such that $\phi e \in \mathbf{bd}(U)$, the boundary of the set U. Note that ϕ corresponds to the maximum minimum achievable utility level that all players can derive simultaneously. Under max-min fairness, the utility derived by all players is at least ϕ, as discussed in Section 3.2, that is,

$$
\mathcal{S}^{\mathrm{MMF}}(U) \geq \phi e. \tag{14}
$$

We can thus use ϕ to bound the sum of utilities under the max-min fair allocation,

$$
\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{MMF}}) = e^T \mathcal{S}^{\mathrm{MMF}}(U) \geq e^T (\phi e) = n\phi. \tag{15}
$$

Similarly to the derivation for proportional fairness, we will identify a hyperplane that supports U at ϕe. In particular, since U is convex and $\phi e \in \mathbf{bd}(U)$, by the supporting hyperplane theorem, $\exists \gamma^{\mathrm{MMF}} \in \mathbf{R}^n \setminus \{0\}$ such that

$$
\left(\gamma^{\mathrm{MMF}}\right)^T u \leq \left(\gamma^{\mathrm{MMF}}\right)^T (\phi e), \quad \forall u \in U. \tag{16}
$$

Applying the above equation to $0 \in U$,

$$0 \in U \Rightarrow \left(\gamma^{\mathrm{MMF}}\right)^T 0 \le \left(\gamma^{\mathrm{MMF}}\right)^T (\phi e) \Rightarrow e^T \gamma^{\mathrm{MMF}} \ge 0.$$

Suppose that $e^T \gamma^{\mathrm{MMF}} = 0$. Combining this fact with (16) for every $e_j \in U$, we get

$$e_j \in U \Rightarrow \left(\gamma^{\mathrm{MMF}}\right)^T e_j \le \left(\gamma^{\mathrm{MMF}}\right)^T (\phi e) \Rightarrow \gamma_j^{\mathrm{MMF}} \le 0.$$

Together with the assumption $e^T \gamma^{\mathrm{MMF}} = 0$, that leads to $\gamma^{\mathrm{MMF}} = 0$, a contradiction. Hence, $e^T \gamma^{\mathrm{MMF}} > 0$, and we can assume without loss that

$$e^T \gamma^{\mathrm{MMF}} = 1.$$

The equation that defines the supporting hyperplane to U, (16), can now be rewritten as

$$\left(\gamma^{\mathrm{MMF}}\right)^T u \le \phi, \quad \forall u \in U. \tag{17}$$

Figure 1 again illustrates the point ϕe and the supporting hyperplane in the case of a two-dimensional example.

We will now show that $\gamma^{\mathrm{MMF}} \ge 0$. Suppose that $\gamma_j^{\mathrm{MMF}} < 0$, and let $y = \phi e - \frac{\phi}{2} e_j$. Since $0 \le y \le \phi e$, we have $y \in U$, by monotonicity of U. But,

$$\left(\gamma^{\mathrm{MMF}}\right)^T y = \left(\gamma^{\mathrm{MMF}}\right)^T \left(\phi e - \frac{\phi}{2} e_j\right) = \phi - \frac{\phi}{2} \gamma_j^{\mathrm{MMF}} > \phi,$$

a contradiction to (17), since $y \in U$. Hence, $\gamma^{\mathrm{MMF}} \ge 0$.

Furthermore, since $e_j \in U$ for all j, using (17) we have

$$\left(\gamma^{\mathrm{MMF}}\right)^T e_j \le \phi \Rightarrow \gamma_j^{\mathrm{MMF}} \le \phi.$$

Without loss, we can assume similarly to the proportional fairness case, that the elements of γ^{MMF} are ordered. To summarize, if we let

$$C = \left\{ (y, B) \in \mathbf{R}^n \times \mathbf{R} \,\middle|\, 0 \le y_1 \le \ldots \le y_n \le B, \, e^T y = 1, \, \frac{1}{n} \le B \le 1 \right\},$$

then $(\gamma^{\mathrm{MMF}}, \phi) \in C$.

Similar to the analysis for the case of proportional fairness, using (4), (17) and the analysis of (10) we get

$$\mathrm{SYSTEM}(U) \le \max \left\{ e^T u \,\middle|\, 0 \le u \le e, \left(\gamma^{\mathrm{MMF}}\right)^T u \le \phi \right\}$$
$$= \ell(\gamma^{\mathrm{MMF}}, \phi) + \delta(\gamma^{\mathrm{MMF}}, \phi). \tag{18}$$

It follows that

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{MMF}}) = 1 - \frac{\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{MMF}})}{\mathrm{SYSTEM}(U)}$$

$$\le 1 - \frac{n\phi}{\mathrm{SYSTEM}(U)} \qquad \text{(from (15))}$$

$$\le 1 - \frac{n\phi}{\ell(\gamma^{\mathrm{MMF}}, \phi) + \delta(\gamma^{\mathrm{MMF}}, \phi)} \qquad \text{(from (18))}$$

$$\le 1 - \inf_{(\gamma, \phi) \in C} \frac{n\phi}{\ell(\gamma, \phi) + \delta(\gamma, \phi)}.$$

We will show that

$$\ell(\gamma,\phi) + \delta(\gamma,\phi) \le n+1-\frac{1}{\phi}, \quad \forall\,(\gamma,\phi) \in C.$$

That will imply that for any such γ and ϕ,

$$\frac{n\phi}{\ell(\gamma,\phi)+\delta(\gamma,\phi)} \ge \frac{n\phi}{n+1-\frac{1}{\phi}} \ge \frac{4n}{(n+1)^2},$$

and the proof will be complete. Note that the last inequality follows by simply minimizing over $\phi \in \left[\frac{1}{n}, 1\right]$. Moreover, that will also demonstrate that

$$\frac{\text{FAIR}(U; \mathcal{S}^{\text{MMF}})}{\text{SYSTEM}(U)} \ge \frac{n\phi}{\text{SYSTEM}(U)} \ge \frac{4n}{(n+1)^2}. \tag{19}$$

Fix any $(\gamma,\phi) \in C$. If $\ell(\gamma,\phi) + \delta(\gamma,\phi) < n$, let

$$y = \frac{(1-\delta(\gamma,\phi))\gamma_{\ell(\gamma,\phi)+1} + \gamma_{\ell(\gamma,\phi)+2} + \ldots + \gamma_n}{n - \ell(\gamma,\phi) - \delta(\gamma,\phi)}.$$

Note that since $\gamma_j \le \phi$, we get $y \le \phi$. Then,

$$
\begin{aligned}
1 &= e^T\gamma \\
&= \gamma_1 + \ldots + \gamma_{\ell(\gamma,\phi)} + \delta(\gamma,\phi)\gamma_{\ell(\gamma,\phi)+1} + (1-\delta(\gamma,\phi))\gamma_{\ell(\gamma,\phi)+1} + \gamma_{\ell(\gamma,\phi)+2} + \ldots + \gamma_n \\
&= \phi + (1-\delta(\gamma,\phi))\gamma_{\ell(\gamma,\phi)+1} + \gamma_{\ell(\gamma,\phi)+2} + \ldots + \gamma_n \\
&= \phi + (n - \ell(\gamma,\phi) - \delta(\gamma,\phi))y \\
&\le \phi + (n - \ell(\gamma,\phi) - \delta(\gamma,\phi))\phi,
\end{aligned}
$$

which demonstrates that $\ell(\gamma,\phi) + \delta(\gamma,\phi) \le n+1-\frac{1}{\phi}$. If $\ell(\gamma,\phi) + \delta(\gamma,\phi) = n$, we get $1 = e^T\gamma = \phi$, and hence $\ell(\gamma,\phi) + \delta(\gamma,\phi) = n = n+1-\frac{1}{\phi}$, and the proof is complete.

Section 5 includes examples that show that the bound is tight for all $n \ge 2$. $\quad\square$

Table 1   Bounds on the price of fairness, under
Assumption 1 and equal maximum achievable
utilities for all players, for the proportional and
max-min fairness schemes, for a small number of
players n.

	Proportional Fairness	Max-min Fairness
$n=2$	0.086	0.111
$n=3$	0.179	0.25
$n=4$	0.25	0.36
$n=5$	0.306	0.444

At this point, it serves us to pause and remark on the result we have established:

• The bounds we have established depend only on the number of players involved in the resource allocation; they are independent of the shape of the utility set, as long as it is compact and convex, and the players have equal maximum achievable utilities. Note that the assumption of equal maximum achievable utilities is not overly restrictive: the utility levels of the players are commonly normalized in a variety of settings, so that the comparison between them is meaningful. Under normalization, the maximum achievable utility of each player is typically equal to 1.

Figure 2   Bounds on the price of fairness, under Assumption 1 and equal maximum achievable utilities for all players, for the proportional (PF) and max-min fairness (MMF) schemes, against the number of players.
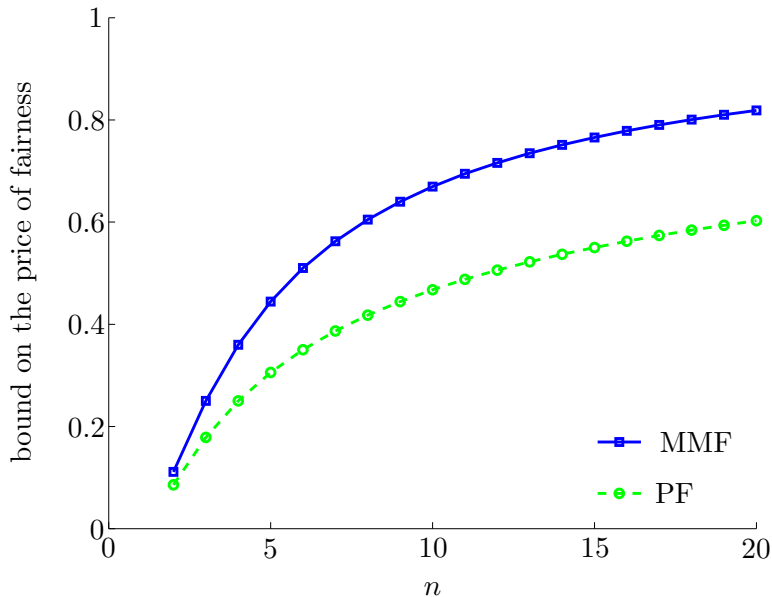
• Our results show that for a small number of players, the price of fairness stays relatively low. In particular, these results establish that for Nash's original two player bargaining game (*i.e.*, $n = 2$), the price of fairness is at most 8.6% for proportional fairness and 11.1% for max-min fairness! For $n = 5$, these numbers are 30.6% and 44.4% respectively. This suggests that in cases with a relatively small number of players, the central decision maker can achieve fair allocations, without incurring a high reduction in the sum of utilities. For illustration purposes, Table 1 lists the values for the worst case bounds under the two schemes, for a small number of players.

• Figure 2 depicts the bounds as a function of the number of players. Note that the worst case price of fairness strictly increases with the number of players, under both schemes, and approaches 1 asymptotically. However, proportional fairness bears a significantly lower price compared to max-min fairness in the worst case; this is especially so for large numbers of players. Those observations are in line with intuition and provide a sound theoretical basis to prior empirical work in the literature (see Radunovic and Boudec (2004) and Tang et al. (2004)).

4.2. Unequal maximum achievable utilities

We now generalize the result of the previous section for the case where the players potentially have unequal maximum achievable utilities. The following Theorem provides upper bounds for the price of fairness. Recall that the maximum achievable utility of the jth player is defined as

$$u_j^\star = \sup \left\{ u_j \,|\, u \in U \right\}.$$

THEOREM 2. *Consider a resource allocation problem with n players; $n \geq 2$. Let the utility set, denoted by $U \subset \mathbf{R}_+^n$, satisfy Assumption 1. If all players have maximum achievable utilities greater than zero,*

   (a) the price of proportional fairness is bounded by

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{PF}}) \leq 1 - \frac{2\sqrt{n} - 1}{n} \frac{\min_{j \in \{1,\dots,n\}} u_j^\star}{\max_{j \in \{1,\dots,n\}} u_j^\star} - \frac{1}{n} + \frac{\min_{j \in \{1,\dots,n\}} u_j^\star}{\sum_{j=1}^n u_j^\star},$$

(b) the price of max-min fairness is bounded by

$$\text{POF}(U; \mathcal{S}^{\text{MMF}}) \leq 1 - \frac{4n}{(n+1)^2} \frac{\frac{1}{n} \sum_{j=1}^{n} u_j^\star}{\max_{j \in \{1,\ldots,n\}} u_j^\star}.$$

Proof. To ease notation, define

$$u_{\max}^\star = \max_{j \in \{1,\ldots,n\}} u_j^\star, \quad u_{\min}^\star = \min_{j \in \{1,\ldots,n\}} u_j^\star.$$

Let

$$\Sigma = \text{diag}(u_1^\star, \ldots, u_n^\star)$$

be a diagonal scaling matrix. Consider the normalized problem, with utility set

$$\overline{U} = \Sigma^{-1} U.$$

Note that $\overline{U}$ satisfies Assumption 1, and has also the property that the maximum achievable utilities for all players are equal to one.

For all $u \in U$, and the corresponding $\bar{u} = \Sigma^{-1} u \in \overline{U}$, we have

$$e^T u = e^T \Sigma \bar{u} \leq u_{\max}^\star e^T \bar{u} \leq u_{\max}^\star \text{SYSTEM}(\overline{U}).$$

As a result,

$$\text{SYSTEM}(U) \leq u_{\max}^\star \text{SYSTEM}(\overline{U}). \tag{20}$$

Moreover,

$$\text{SYSTEM}(U) \leq \sum_{j=1}^{n} u_j^\star = e^T \Sigma e. \tag{21}$$

(a) *Proportional fairness.* Using Theorem 1,

$$\frac{\text{FAIR}(\overline{U}; \mathcal{S}^{\text{PF}})}{\text{SYSTEM}(\overline{U})} = \frac{e^T \mathcal{S}^{\text{PF}}(\overline{U})}{\text{SYSTEM}(\overline{U})} \geq \frac{2\sqrt{n}-1}{n}. \tag{22}$$

Moreover, by (7) and (8), we have that $\mathcal{S}^{\text{PF}}(\overline{U}) \geq \frac{1}{n} e$. Hence,

$$\mathcal{S}^{\text{PF}}(\overline{U}) = \frac{1}{n} e + q, \tag{23}$$

for some $q \geq 0$. By utilizing this expression and (22) we get

$$\frac{e^T \mathcal{S}^{\text{PF}}(\overline{U})}{\text{SYSTEM}(\overline{U})} = \frac{1 + e^T q}{\text{SYSTEM}(\overline{U})} \geq \frac{2\sqrt{n}-1}{n}. \tag{24}$$

We can now bound the sum of utilities under the proportionally fair allocation for the problem involving U:

$$\begin{aligned}
\text{FAIR}(U; \mathcal{S}^{\text{PF}}) &= e^T \mathcal{S}^{\text{PF}}(U) \\
&= e^T \mathcal{S}^{\text{PF}}(\Sigma \overline{U}) \\
&= e^T \Sigma \mathcal{S}^{\text{PF}}(\overline{U}) && \text{(from (1))} \\
&= e^T \Sigma \left(\frac{1}{n} e + q \right) && \text{(from (23))} \\
&= \frac{1}{n} e^T \Sigma e + e^T \Sigma q \\
&\geq \frac{1}{n} e^T \Sigma e + u_{\min}^\star e^T q. && \text{(since } q \geq 0) \tag{25}
\end{aligned}$$

We then have

$$\frac{\text{FAIR}(U;\mathcal{S}^{\text{PF}})}{\text{SYSTEM}(U)} \geq \frac{\frac{1}{n}e^T\Sigma e + u^\star_{\min}e^T q}{\text{SYSTEM}(U)} \qquad \text{(from (25))}$$

$$= \frac{\frac{1}{n}e^T\Sigma e - u^\star_{\min}}{\text{SYSTEM}(U)} + \frac{u^\star_{\min}(1+e^T q)}{\text{SYSTEM}(U)}$$

$$\geq \frac{\frac{1}{n}e^T\Sigma e - u^\star_{\min}}{\text{SYSTEM}(U)} + \frac{u^\star_{\min}(1+e^T q)}{u^\star_{\max}\text{SYSTEM}(\overline{U})} \qquad \text{(from (20))}$$

$$\geq \frac{\frac{1}{n}e^T\Sigma e - u^\star_{\min}}{e^T\Sigma e} + \frac{u^\star_{\min}(1+e^T q)}{u^\star_{\max}\text{SYSTEM}(\overline{U})} \qquad \text{(from (21))}$$

$$\geq \frac{1}{n} - \frac{u^\star_{\min}}{\sum_{j=1}^n u^\star_j} + \frac{2\sqrt{n}-1}{n}\frac{u^\star_{\min}}{u^\star_{\max}}. \qquad \text{(from (24))}$$

(b) *Max-min fairness.* We apply Theorem 1 for the normalized problem that involves $\overline{U}$. Let ϕ be the maximum minimum utility for $\overline{U}$. Then,

$$\text{FAIR}(U;\mathcal{S}^{\text{MMF}}) = e^T\mathcal{S}^{\text{MMF}}(U)$$

$$= e^T\mathcal{S}^{\text{MMF}}(\Sigma\overline{U})$$

$$= e^T\Sigma\mathcal{S}^{\text{MMF}}(\overline{U}) \qquad \text{(from (2))}$$

$$\geq e^T\Sigma(\phi e) \qquad \text{(from (14))}$$

$$= \left(\frac{1}{n}\sum_{j=1}^n u^\star_j\right)n\phi. \qquad (26)$$

We therefore have,

$$\frac{\text{FAIR}(U;\mathcal{S}^{\text{MMF}})}{\text{SYSTEM}(U)} \geq \frac{\left(\frac{1}{n}\sum_{j=1}^n u^\star_j\right)n\phi}{\text{SYSTEM}(U)} \qquad \text{(from (26))}$$

$$\geq \frac{\frac{1}{n}\sum_{j=1}^n u^\star_j}{u^\star_{\max}}\frac{n\phi}{\text{SYSTEM}(\overline{U})} \qquad \text{(from (20))}$$

$$\geq \frac{\frac{1}{n}\sum_{j=1}^n u^\star_j}{u^\star_{\max}}\frac{4n}{(n+1)^2}. \qquad \text{(from (19))} \quad \square$$

Theorem 2 extends the results of Theorem 1 in case of a problem where players have unequal maximum achievable utilities. In general, assymetric maximum achievable utilities may result (although not necessarily) in higher price of fairness. Theorem 2 characterizes the way in which the worst case bounds are affected. In the next section we address a natural question that arises in response to the results of our Theorems, namely, how loose are our bounds? The surprising answer, is that our bounds are in fact tight; they are achieved by several realistic examples.

5. Examples

This section addresses two natural questions that arise in the context of our analysis of the price of fairness. The first concerns the tightness of our bounds. To that end we will study a problem

Figure 3   The network flow topology in case of $k = 3$, for the example in Section 5.1.
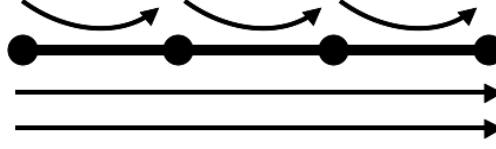
of bandwidth allocation for a communication network wherein our bounds on the price of fairness are, in fact, achieved. The next question one may ask regards our assumptions on the structure of the utility set, namely Assumption 1. Here we show via an example, that if Assumption 1 is violated, the price of fairness can be arbitrarily large, even for a small number of players.

5.1. A communication network

We illustrate the tightness of our bounds for a problem of bandwidth allocation on a communication network. The network consists of hubs (nodes) that are connected via capacitated links (edges). Clients, or flows, wish to establish transmission from one hub to another over the network, via a pre-specified and fixed route. The network administrator needs to decide on the transmission rate assigned to each flow, subject to capacity constraints. The resources to be allocated in this case are the available capacities of the links, the players are the flows, and the central decision maker is the network administrator. We now fix some notation, and specify the problem data more precisely.

We have a network with k links of unit capacity. There are in total $n = 2k - 1$ flows in the network, each of which is associated with a fixed route, *i.e.*, some subset of the k links. The network is assumed to be a line-graph with k links. The routes of the first k flows are disjoint and they all occupy a single (distinct) link. The remaining $k - 1$ flows have routes that utilize all k links. The described network topology is shown in Figure 3, for $k = 3$. Each flow has a nonnegative rate, which we denote $x_1, \dots, x_n$. The first k flows derive M units of utility for every unit rate they are assigned (*i.e.*, $f_j(x) = Mx_j$, for $j = 1, \dots, k$), with $M \geq 1$. The remaining $k - 1$ flows derive utility equal to their rates (*i.e.*, $f_j(x) = x_j$, for $j = k + 1, \dots, n$).

The routing matrix $R \in \mathbf{R}^{k \times n}$, defined as

$$R_{ij} = \begin{cases} 1, & \text{flow } j\text{'s route passes over link } i, \\ 0, & \text{otherwise}, \end{cases}$$

is then such that, its ith row is of the form $\begin{bmatrix} e_i^T & e^T \end{bmatrix}$, where e_i is the unit vector in $\mathbf{R}^k$, with the ith component equal to 1. The resource set can be expressed as

$$X = \{x \in \mathbf{R}^n \,|\, Rx \leq e, \, x \geq 0\}.$$

For the case of $k = 3$ (depicted in Figure 3), we have

$$X = \left\{ x \in \mathbf{R}^5 \,\middle|\, \begin{bmatrix} 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_5 \end{bmatrix} \leq e, \, x \geq 0 \right\}.$$

Accordingly, the utility set is

$$U = \left\{ \begin{bmatrix} Mx_1 & \dots & Mx_k & x_{k+1} & \dots & x_n \end{bmatrix}^T \in \mathbf{R}^n \,\middle|\, Rx \leq e, \, x \geq 0 \right\}.$$

Note that the utility set is convex and compact. In particular, Assumption 1 is satisfied.

Furthermore, since all links have unit capacity, all flows can be assigned a maximum rate of 1. As a result, the maximum achievable utility for each of the first k flows is $u_j^\star = M$, $j = 1, \ldots, k$, and for each of the remaining flows is $u_j^\star = 1$, $j = k+1, \ldots, n$. Theorem 1 then applies, only in case of $M = 1$. If we apply Theorem 2, we get

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{MMF}}) \leq 1 - \frac{4n}{(n+1)^2} \frac{\frac{1}{n} \sum_{j=1}^n u_j^\star}{\max_{j \in \{1, \ldots, n\}} u_j^\star} = 1 - \frac{4}{(n+1)^2} \frac{kM + k - 1}{M}.$$

For the utilitarian solution, the central decision maker assigns unit rate to the first k flows, and achieves a throughput of kM, *i.e.*,

$$\mathrm{SYSTEM}(U) = kM.$$

Under the max-min fairness allocation, a rate of $\frac{1}{k}$ is assigned to each flow, hence

$$\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{MMF}}) = \frac{kM + k - 1}{k}.$$

Thus, by substituting for the above expressions and for $k = \frac{n+1}{2}$,

$$\begin{aligned}
\mathrm{POF}(U; \mathcal{S}^{\mathrm{MMF}}) &= 1 - \frac{\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{MMF}})}{\mathrm{SYSTEM}(U)} \\
&= 1 - \frac{kM + k - 1}{Mk^2} \\
&= 1 - \frac{1}{\left(\frac{n+1}{2}\right)^2} \frac{kM + k - 1}{M} \\
&= 1 - \frac{4}{(n+1)^2} \frac{kM + k - 1}{M},
\end{aligned}$$

which is exactly the upper bound we derived from Theorem 2. In case $M = 1$, we get

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{MMF}}) = 1 - \frac{4n}{(n+1)^2},$$

which is the upper bound from Theorem 1.

This example illustrates the tightness of our bounds for the max-min fairness scheme for an odd number of players. Similar tight bounds can be derived for an even number of players, by studying the utility set

$$W = \left\{ u \in \mathbf{R}_+^n \,\middle|\, \frac{1}{n} u_1 + \ldots + \frac{1}{n} u_{n/2} + u_{n/2+1} + \ldots + u_n \leq 1, \, u \leq e \right\}.$$

In order to obtain a tight upper bound for the case of proportional fairness, we study a similar setup, but with additional long flows. In particular, let the number of long flows be equal to $k^2 - k$ (instead of $k - 1$). Thus, there are now $n = k^2$ flows. Let also $M = 1$.

The utilitarian solution remains unchanged in this case, with the central decision maker allocating unit rate to the first k flows.

Under proportional fairness, we have $u_j^{\mathrm{PF}} = x_j^{\mathrm{PF}} = \frac{1}{k}$ for $j = 1, \ldots, k$, and $u_j^{\mathrm{PF}} = x_j^{\mathrm{PF}} = \frac{1}{k^2}$ for the remaining long flows $j = k+1, \ldots, n$, since this point satisfies the first order optimality condition (see Section 3.1). In particular, for any $u \in U$,

$$\sum_{j=1}^n \frac{u_j - u_j^{\mathrm{PF}}}{u_j^{\mathrm{PF}}} = \sum_{j=1}^k \frac{u_j - \frac{1}{k}}{\frac{1}{k}} + \sum_{j=k+1}^n \frac{u_j - \frac{1}{k^2}}{\frac{1}{k^2}}$$

$$= k \sum_{j=1}^{k} u_j + k^2 \sum_{j=k+1}^{n} u_j - k^2$$
$$= k \left(e^T R u - k \right)$$
$$\leq k \left(e^T e - k \right) = 0.$$

Thus,

$$\mathrm{FAIR}(U; \mathcal{S}^{\mathrm{PF}}) = k\frac{1}{k} + \left(k^2 - k \right) \frac{1}{k^2} = 2 - \frac{1}{k},$$

and

$$\mathrm{POF}(U; \mathcal{S}^{\mathrm{PF}}) = 1 - \frac{2 - \frac{1}{k}}{k} = 1 - \frac{2\sqrt{n} - 1}{n},$$

which is again exactly the upper bound from Theorem 1. We are unable to establish that our bound on the price of proportional fairness is tight in the event that maximum achievable utilities are unequal (*i.e.*, $M > 1$ for the communication network example).

5.2. Non convex utility set

Here we consider what happens if one were to relax the requirements of Assumption 1. Consider a setup with two players (*i.e.*, $n = 2$), in which the central decision maker has the option of allocating all resources to one of the players, or splitting them equally among them. In the case one player receives all resources, she derives a utility of 1, while the other player derives a utility of 0. If the resources are split, both players derive a utility of ϵ. The utility set is thus

$$U = \{e_1, e_2, \epsilon e\}.$$

Note that U is discrete, in particular non convex. As a result, Assumption 1 is violated.

It is easy to check that for $\epsilon \ll 1$, the utilitarian solution corresponds to one player receiving all resources, and the corresponding sum of utilities is equal to 1. Under the max-min fairness scheme, the CDM splits the resources among the players, thus resulting in aggregate utility of 2ϵ and price of fairness of $1 - 2\epsilon$. We can thus see that for non convex utility sets, the price of max-min fairness can get arbitrarily close to 1, even for two players.

Note that in this case there does not exist a feasible allocation that satisfies the Nash standard (see Section 3.1). If we allow the PF allocation to be the one that maximizes the sum of logarithms of the utilities (see Section 3.1), then the CDM again splits the resources among the players, and similar obervations to MMF apply for PF.

A practical situation under which we might obtain non convex utility sets, is the power control problem in a wireless cellular system under severe interference effects (see Goldsmith (2005)).

6. Conclusions

This paper has attempted to quantify the "price" one has to pay in demanding that an allocation of resources is fair. In particular, we presented results on the relative efficiency loss incurred in using either of two widely accepted and axiomatically justified notions of fairness – max-min fairness and proportional fairness. It is our belief that the "price" of fairness is effectively inescapable if the allocations prescribed by a given scheme are to be ethically acceptable and implementable. Our analysis has yielded two primary insights. First, it has given us an understanding of when this "price" is likely to be small; this will be the case when the number of players is small. Second, we have presented a quantitive distinction between max-min fairness and proportional fairness, showing that the latter is a substantially cheaper notion than the former, especially when the number of players is large. Our analysis is tight and addresses a vast swath of resource allocation problems.

Moving forward, we believe that one fruitful direction for future research is identifying specialized families of utility sets (within the family considered here) that admit smaller prices of fairness. Good work in this direction will yield a succinct characterization of resource allocation problems for which fair allocations are close to efficient. It is of course important that the classes of problems so identified be relevant; for instance, a condition that guaranteed that all Pareto solutions are equally efficient (which is true if the utility set is a poly-matroid) while interesting is perhaps too narrow to be relevant to practice.

On the practical front, there are a number of important (and real) resource allocation problems wherein it is highly desirable that allocations are fair, for example the air traffic flow management problem, alluded to in the Introduction. We are currently evaluating the performance of "fair" allocation schemes for real world instances of such problems (see Bertsimas et al. (2009a)).

Acknowledgments

References
ATA. 2008. URL http://www.airlines.org/economics/cost+of+delays/.

Bertsekas, D., R. Gallager. 1987. *Data Networks*. Prentice-Hall.

Bertsimas, D., V. F. Farias, N. Trichakis. 2009a. Max-min fairness in the air traffic flow management problem. *Working Paper* .

Bertsimas, D., S. Gupta. 2010. A two-stage model for network air traffic flow management incorporating fairness and slot trading. *Operations Research, submitted for publication* .

Bertsimas, D., D. A. Iancu, N. Trichakis. 2009b. Fairness in multi-account optimization. *Working Paper* .

Bertsimas, D., S. Stock-Patterson. 1998. The air traffic flow management problem with enroute capacities. *Operations Research* **46**(3) 406–422.

Bertsimas, D., S. Stock-Patterson. 2000. The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach. *Transportation Science* **34**(3) 239–255.

Bertsimas, D., J. N. Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific.

Bonald, T., L. Massoulié. 2001. Impact of fairness on internet performance. *SIGMETRICS Perform. Eval. Rev.* **29**(1) 82–91. doi:http://doi.acm.org/10.1145/384268.378438.

Butler, M., H. P. Williams. 2002. Fairness versus efficiency in charging for the use of common facilities. *The Journal of the Operational Research Society* **53**(12) 1324–1329. URL http://www.jstor.org/stable/822721.

Chakrabarty, D., G. Goel, V. V. Vazirani, L. Wang, C. Yu. 2009. Some computational and game-theoretic issues in Nash and nonsymmetric bargaining games. *Working Paper* .

Fabozzi, F., P. Kolm, D. Pachamanova, S. Focardi. 2007. *Robust portfolio optimization and management*. Wiley.

Goel, A., A. Meyerson. 2006. Simultaneous optimization via approximate majorization for concave profits or convex costs. *Algorithmica* **44**(4) 301–323.

Goel, A., A. Meyerson, S. Plotkin. 2000. Combining fairness with throughput: online routing with multiple objectives. *Proceedings of the 32nd annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 670–679.

Goldsmith, A. 2005. *Wireless Communications*. Cambridge University Press, New York, NY, USA.

Johari, R., J. N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* **29**(3) 407–435.

Kalai, E., M. Smorodinsky. 1975. Other solutions to Nash's bargaining problem. *Econometrica* **43** 510–18.

Kelly, F. P., A. Maulloo, D. Tan. 1997. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* **49** 237–252.

Khodadadi, A., R. Tütüncü, P. Zangari. 2006. Optimisation and quantitative investment management. *Journal of Asset Managemen* **7** 83–92.

Kleinberg, J., Y. Rabani, E. Tardos. 1999. Fairness in routing and load balancing. *Proceedings of the 40th annual symposium on Foundations of Computer Science*. 568–578.

Koutsoupias, E., C. Papadimitriou. 1999. Worst-case equilibria. *In 16th Annual Symposium on Theoretical Aspects of Computer Science*. Trier, Germany, 404–413.

Kumar, A., J. Kleinberg. 2000. Fairness measures for resource allocation. *Proceedings of the 41st annual symposium on Foundations of Computer Science*. 75–85.

Lulli, G., A. R. Odoni. 2007. The European Air Traffic Flow Management Problem. *Transportation Science* **41**(4) 431–443.

Luo, H., S. Lu, V. Bharghavan, J. Cheng, G. Zhong. 2004. A packet scheduling approach to qos support in multihop wireless networks. *Mobile Networks and Applications* **9**(3) 193–206.

Luss, H. 1999. On equitable resource allocation problems: a lexicographic minimax approach. *Operations Research* **47**(3) 361–378.

Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press.

Mo, J., J. Walrand. 2000. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**(5) 556–567. doi:http://dx.doi.org/10.1109/90.879343.

Nash, J. 1950. The bargaining problem. *Econometrica* **18** 155–62.

O'Cinneide, C., B. Scherer, X. Xu. 2006. Pooling trades in a quantitative investment process. *Journal of Portfolio Management* **32**(4) 33–43.

Odoni, A. R., L. Bianco. 1987. *Flow Control of Congested Networks*, chap. The Flow Management Problem in Air Traffic Control. Springer-Verlag, Berlin.

Ogryczak, W., M. Pióro, A. Tomaszewski. 2005. Telecommunications network design and max-min optimization problem. *Journal of Telecommunications and Information Technology* 43–56.

Papadimitriou, C. 2001. Algorithms, games, and the internet. *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 749–753.

Perakis, G. 2007. The price of anarchy when costs are non-separable and asymmetric. *Mathematics of Operations Research* **32** 614–628.

Radunovic, B., J.-Y. Le Boudec. 2002. A unified framework for max-min and min-max fairness with applications. *Proceedings of the Annual Allerton Conference on Communication Control and Computing* **40**(2) 1061–1070.

Radunovic, B., J.-Y. Le Boudec. 2004. Rate performace objectives of multihop wireless networks. *IEEE Transactions on Mobile Computing* **3**(4) 334–349.

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

Rios, J., K. Ross. 2007. Delay optimization for airspace capacity management with runtime and equity considerations. *AIAA Guidance, Navigation and Control Conference and Exhibit*. Hilton Head, South Carolina.

Roughgarden, T., E. Tardos. 2002. How bad is selfish routing? *J. ACM* **49**(2) 236–259.

Sen, A., J. E. Foster. 1997. *On Economic Inequality*. Oxford University Press.

Soomer, M. J., G. M. Koole. 2009. Fairness in the aircraft landing problem. *Working Paper* .

Tang, A., J. Wang, S. H. Low. 2004. Is fair allocation always inefficient. *Proceedings of IEEE Infocom* .

Vossen, T., M. Ball, R. Hoffman. 2003. A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs. *Air Traffic Control Quarterly* **11** 277–292.

Young, H. Peyton. 1995. *Equity: In Theory and Practice*. Princeton University Press.