SIAM J. COMPUT.
Vol. 38, No. 1, pp. 241–256

EQUITABLE COST ALLOCATIONS VIA PRIMAL–DUAL-TYPE ALGORITHMS*

KAMAL JAIN† AND VIJAY V. VAZIRANI‡

Abstract. Perhaps the strongest notion of truth-revealing in a cost sharing mechanism is group strategyproofness. However, matters are not so clear-cut on fairness, and many different, sometimes even conflicting, notions of fairness have been proposed which have relevance in different situations. We present a large class of group strategyproof cost sharing methods, for submodular cost functions, satisfying a wide range of fairness criteria, thereby allowing the service provider to choose a method that best satisfies the notion of fairness that is most relevant to its application. Our class includes the Dutta–Ray egalitarian method as a special case. It also includes a new cost sharing method, which we call the *opportunity egalitarian method*.

Key words. cost sharing methods, submodular cost functions, fairness in cost sharing, group strategyproof mechanism, opportunity egalitarian method

AMS subject classifications. 68Q25, 68W40, 91A12

DOI. 10.1137/060658448

1. Introduction.

Distributing the cost of a shared resource in a fair and truth-revealing manner is a central problem in cooperative game theory. Perhaps the strongest notion of truth-revealing is *group strategyproofness*, under which the dominant strategy of users is to reveal their true utility, even if they collude. However, matters are not so clear-cut on fairness—many different, sometimes conflicting, notions have been proposed which have relevance in different situations. Let us clarify that we are not necessarily postulating a service provider who is inherently "fair," but that in the long run, it is in the best interest of the service provider to subscribe to some form of fairness in choosing its cost allocations. We will assume that the cost function is submodular, a natural economies-of-scale condition. Equivalently, these results also apply to the situation of profit sharing under a convex transferable utility game; e.g., see [23]. In this paper, we will be concerned with fully budget-balanced methods; i.e., the total amount accrued from the users should be exactly equal to the cost of the shared resource.

As shown by Moulin [23], a *cross-monotone cost sharing method* for the given cost function gives rise to a group strategyproof mechanism, and for submodular cost functions, essentially the converse holds as well. Informally, a cost sharing method is cross-monotone, also called population monotone, if the cost share of any user can only decrease if a superset is being served.

Two well-known cross-monotone cost sharing methods for submodular cost functions are the Shapley value and the egalitarian method of Dutta and Ray [6] (the former requires that the cost function be *nondecreasing* as well; i.e., the cost of serving a set should not be larger than the cost of serving any of its supersets). Both these methods have been extensively studied; for the latter, see [12, 11, 13, 16, 21, 3, 2, 5, 7].

*Received by the editors April 28, 2006; accepted for publication (in revised form) November 21, 2007; published electronically March 28, 2008.
http://www.siam.org/journals/sicomp/38-1/65844.html
†Microsoft, One Microsoft Way, Redmond, WA 98052 (kamalj@microsoft.com).
‡Corresponding author. College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280 (vazirani@cc.gatech.edu). This author's research was supported by NSF grants CCR-0220343 and CCR-0515186.

These two methods satisfy different fairness criteria. The Shapley value charges higher amounts from users who are more expensive to serve. The egalitarian method attempts to charge equal amounts from all users subject to the coalition participation constraint that the method lies in the core of the game; i.e., no subset is charged more than its stand-alone cost (thereby precluding the possibility of its seceding).

Consider the following situation in which neither of these criteria appears to be fair. Suppose that two users, Fred Smith and Gill Bates, are proximally located so that they are equally expensive to serve. For concreteness, assume that the service provider is transmitting valuable financial data and needs to recover a cost of $2000, regardless of whether it serves one or both users. Fred Smith and Gill Bates derive widely different utilities on receiving this data—the former is a man of modest means, and the latter is a multimillionaire. Hence, they are also willing to pay different amounts for this data. In this situation, the Shapley value as well as the egalitarian method will assign cost shares of $1000 each for the service, an amount that is not acceptable to Fred Smith. However, Gill Bates considers this data useful for wisely managing his vast acquisitions and ends up paying the entire $2000 for the service. If, instead, the cost sharing method were to take into consideration the relative paying powers of the two users and charge differentially, it may be able to find an outcome that Pareto dominates the previous outcome. For instance, if it charges Fred Smith $100 and Gill Bates $1900, both accept the service and both are better off. In addition, the service provider is also better off since it has a larger and more satisfied pool of customers.

This form of differential pricing, sometimes also called price discrimination, is widely resorted to and is in fact crucial to the survival of many industries [25, 27, 28, 29]. For instance, it provides mechanisms to the airline industry to charge higher fares from business travelers, who can afford to pay higher fares, than from casual travelers. Clearly, the fate of the airline industry, which has been on the brink of bankruptcy numerous times, would be dire without such a mechanism. Another common example is differential subscription rates for journals charged from students, professionals, and institutions.

Can the service provider resort to differential pricing and still ensure that the mechanism is strategyproof or, better, group strategyproof? In this paper, we provide a formal setting to accomplish this. We present a large class of group strategyproof mechanisms for submodular cost functions, satisfying a wide range of fairness criteria—hence the name "equitable." Our class includes the Dutta–Ray egalitarian method as a special case. It also includes a new cost sharing method, which we call the *opportunity egalitarian method*. Assuming that individual utilities are drawn from probability distributions which are known to the service provider, this method finds cost shares that attempt to equalize the users' probabilities of accepting the service, subject to core constraints. The above-stated examples of differential cost shares can be viewed as approximations of the opportunity egalitarian method.

Each equitable cost sharing method is parameterized by *n equalizing functions* which encode the fairness criterion chosen. The method ensures that w.r.t. the chosen criterion, the cost shares satisfy min-max as well as max-min fairness; i.e., no one underpays, and no one overpays. Thus, in the case of the egalitarian method (opportunity egalitarian method), among all cost allocations in the core, the chosen allocation minimizes the maximum cost shares (probability of accepting service) as well as maximizes the minimum cost shares (for fairness). Precise definitions appear in sections 3, 5.1, and 6. Max-min fairness has been used in the networking community for tackling issues of bandwidth allocation [4, 15] and has also been algorithmically studied in the context of routing in networks [22, 1, 20]. Approximate versions of this notion have also been studied [9].

Our algorithms are inspired by the primal-dual schema from the field of approximation algorithms (see [30]). In the latter setting, it is natural to view the dual program as "paying" for the primal, and the algorithm as progressive bidding to get access to a shared resource (this viewpoint is particularly clear in the primal-dual algorithm presented in [18]). The equalizing functions determine the rates at which individual users increase their "bids." Each iteration of our algorithm runs in polynomial time. We utilize recently discovered polynomial time algorithms for the minimization of a submodular function [14, 26]. The precise number of iterations depends upon the accuracy needed.

We have recently found some rather unexpected applications of the results of this paper. The cross-monotonic cost sharing method developed in this paper for submodular cost functions has been used for proving competition monotonicity for submodular utility allocation markets in [19]. Our max-min and min-max fairness results for these cost sharing methods have also been used in [19] for establishing that the equilibrium allocations for submodular utility allocation markets are max-min fair.

It is interesting to note that independently, though somewhat preceding our work, Hokari [11] generalized Dutta–Ray solutions to give a class of cost sharing methods that turns out to be identical to ours. He calls his methods *monotone path cost allocations*. Hokari's formalization and point of view are quite different from ours—the definitions of the cost sharing methods are strikingly different,[1] and so are the algorithms for computing them (Hokari does not address issues of algorithmic efficiency). We believe that this class deserves further study—in the past, notions derived from diverse considerations have turned out to be particularly robust and canonical.

Mutuswami [24] proved the following interesting fact about the Dutta–Ray egalitarian method. If the utilities of individual users are independently and identically distributed (i.i.d.) (and the distribution satisfies the monotone hazard rate condition; see section 7 for a formal definition), then for every set S of users, the egalitarian method maximizes the probability that all members of S accept the service, among all cost sharing methods in the core. We generalize this result by removing the restriction that all utilities come from the *same* distribution. We show that for each choice of distributions from which the utilities are picked (provided they satisfy the strict monotone hazard rate condition), there is an equitable cost sharing method that maximizes, for every set $S \subseteq U$, the probability that all members of S accept the service, among all cost sharing methods for S in the core.

2. Basic definitions.

Let $U = \{1, \ldots, n\}$ denote the set of users and cost : $2^U \to \mathbf{R}^+$ denote the function that gives the cost of serving any subset of the users. We will assume that this function is *submodular*; i.e.,

$$\forall S, T \subseteq U, \quad \text{cost}(S) + \text{cost}(T) \geq \text{cost}(S \cup T) + \text{cost}(S \cap T).$$

Following is an equivalent definition that makes it clear that such cost functions satisfy a natural economies-of-scale condition. The marginal cost of including a new user can only be smaller if a superset is being served:

$$\forall S \subset T \subset U,\ i \notin T, \quad \text{cost}(S + i) - \text{cost}(S) \geq \text{cost}(T + i) - \text{cost}(T).$$

[1] See section 8 for an implication that is easier to derive from our formulation.

We consider the following game. The service provider picks a *mechanism* for deciding the set $S \subseteq U$ of users that get the service and the individual cost shares of users in S so as to retrieve the cost of serving them, $\text{cost}(S)$. It obtains from the users their utilities for receiving the service. Thus each user's strategy is simply the utility he reports. The service provider is not allowed to charge a user more than his reported utility (otherwise, the user will refuse the service).

We will say that the service provider's mechanism is *strategyproof* if the dominant strategy of each user is to report his true utility. It is *group strategyproof* if the above holds despite collusions among users. Let us make this precise. Consider a coalition C of users. Let $\boldsymbol{u}$ and $\boldsymbol{u}'$ be two vectors of bids (we will think of the former as agents' true values and $\boldsymbol{u}'$ as strategically chosen bids). Assume that $\boldsymbol{u}_j = \boldsymbol{u}'_j$ for all $j \notin C$. Let $(\boldsymbol{q}, \boldsymbol{x})$ and $(\boldsymbol{q}', \boldsymbol{x}')$ denote the users served and costs at $\boldsymbol{u}$ and $\boldsymbol{u}'$, respectively. Now, group strategyproofness requires that if the inequality

$$\boldsymbol{u}'_i \boldsymbol{q}_i - \boldsymbol{x}_i \geq \boldsymbol{u}'_i \boldsymbol{q}'_i - \boldsymbol{x}'_i$$

holds for all $i \in C$, then it must hold with equality for all $i \in C$ as well; i.e., if no member of C is made worse off by misreporting of their utility values, then no member of C is made better off either. Moulin [23] showed that it is sufficient for the service provider to pick a cross-monotone cost sharing method in order to obtain a group strategyproof mechanism. His procedure for obtaining such a mechanism from the method is recapitulated below.

A *cost sharing method*, ξ, specifies how to distribute, for any set $S \subseteq U$, $\text{cost}(S)$ among the users in S. It satisfies the following:

1. Users will not be paid for receiving service; i.e.,

$$\forall S \subseteq U,\ i \in S, \quad \xi(S, i) \geq 0.$$

2. Budget balance

$$\forall S \subseteq U, \quad \sum_{i \in S} \xi(S, i) = \text{cost}(S).$$

3. Users not being served will not be charged; i.e.,

$$\forall S \subseteq U,\ i \notin S, \quad \xi(S, i) = 0.$$

For any set $S \subseteq U$, the slice of ξ at S, i.e., $\xi(S, \cdot)$, will be denoted by ξ^S. Thus, $\xi^S : S \to \mathbf{R}^+$ specifies the cost shares of users in S, assuming that S is the set being served. We will say that ξ is *cross-monotonic* if it satisfies the following economies of scale condition:

$$\forall S \subset T \subseteq U,\ \forall i \in S, \quad \xi^S(i) \geq \xi^T(i);$$

i.e., the cost share of a user can only be smaller if a superset is being served.

Let ξ be a cross-monotone cost sharing method. Consider the following mechanism. Initialize $S \leftarrow U$. If for each user $i \in S$, his cost share $\xi(S, i)$ is at most his utility, HALT. Else, drop users whose utilities are smaller than their cost shares, update S, and repeat.

THEOREM 1 (Moulin [23]). *If ξ is a cross-monotone cost sharing method, then the mechanism specified above is group strategyproof.*

A cost allocation for set $S \subseteq U$, $\alpha : S \to \mathbf{R}^+$ is said to satisfy the *coalition participation constraint* if it satisfies the following:

1. Budget balance, i.e.,

$$\sum_{i \in S} \alpha(i) = \text{cost}(S).$$

2. No subset $S' \subset S$ is charged more than the stand-alone cost of serving S'; i.e.,

$$\forall S' \subset S, \quad \sum_{i \in S'} \alpha(i) \leq \text{cost}(S').$$

The core is usually defined as the set of all cost allocations for U (i.e., the grand coalition) satisfying the coalition participation constraint. In this paper, we will define the *core* to consist of all cost allocations satisfying the coalition participation constraint for all sets $S \subseteq U$, rather than only U. We will say that *a cost sharing method* ξ *is in the core* if for all $S \subseteq U$, ξ^S is in the core.

3. Equitable cost sharing methods. Let $\boldsymbol{q}$ and $\boldsymbol{r}$ be n-dimensional vectors with nonnegative coordinates. We will denote by $\boldsymbol{q}_{INC}$ the vector obtained by sorting the components of $\boldsymbol{q}$ in increasing order. Thus $\boldsymbol{q}_{INC}(i) \leq \boldsymbol{q}_{INC}(i+1)$ for $1 \leq i \leq n-1$. Define a partial order as follows: say that $\boldsymbol{q}$ *max-min dominates* $\boldsymbol{r}$ if $\boldsymbol{q}_{INC}$ is lexicographically larger than $\boldsymbol{r}_{INC}$, i.e., if there is an i such that $\boldsymbol{q}_{INC}(i) > \boldsymbol{r}_{INC}(i)$ and $\boldsymbol{q}_{INC}(j) = \boldsymbol{r}_{INC}(j)$ for $j < i$. Clearly, $\boldsymbol{q}_{INC} = \boldsymbol{r}_{INC}$ may hold even though $\boldsymbol{q} \neq \boldsymbol{r}$.

An equitable cost sharing method is parameterized by n strictly increasing, continuous, and unbounded functions from $\mathbf{R}^+$ to $\mathbf{R}^+$, $f_1, \ldots, f_n$ satisfying further that $f_i(0) = 0$. These will be called *equalizing functions*. The *equitable cost sharing method* corresponding to this set of equalizing functions, say ξ, is defined as follows. We will specify ξ^S for each set $S \subseteq U$. Without loss of generality assume that S consists of users $1, \ldots, s$. Let α be a cost allocation for S that lies in the core. Let $\boldsymbol{t}(\alpha)$ denote the s-dimensional vector whose ith component is $f_i^{-1}(\alpha(i))$. We will show in Theorem 6 that there is a unique cost allocation for set S in the core, say β, such that $\boldsymbol{t}(\beta)$ max-min dominates $\boldsymbol{t}(\alpha)$ for all other allocations, α, for S in the core. We will define $\xi^S = \beta$.

If all n equalizing functions are picked to be the identity function, the resulting cost sharing method will be the egalitarian method of Dutta and Ray (strictly speaking, this is not the way they defined their method; see section 5.2). The egalitarian method maximizes the minimum cost shares, i.e., ensures that no one underpays, subject to core constraints—in that sense, it tries to make the cost shares of individual users as equal as possible. Our generalization optimizes the max-min objective function relative to the equalizing functions $f_1, \ldots, f_n$, which encode the particular fairness criterion chosen. As shown in Theorem 10, equitable methods (including the egalitarian method) optimize the min-max objective as well, ensuring that no one overpays.

4. The algorithm. We now present a primal–dual-type algorithm for obtaining ξ^S for any set S. First we give some definitions. Let $x : S \to \mathbf{R}^+$ be a function assigning costs to users in S. Set $A \subseteq S$ will be said to be *tight* if $\sum_{i \in A} x_i = \text{cost}(A)$. It will be said to be *overtight* if $\sum_{i \in A} x_i > \text{cost}(A)$. We will say that x is *feasible* if no subset of S is overtight. Note that we have not imposed the condition that $\sum_{i \in S} x_i = \text{cost}(S)$. The algorithm will utilize the following properties.

LEMMA 2. *Let* cost *be submodular and x be feasible for S. If $A, B \subseteq S$ are both tight, then $A \cup B$ is also tight.*

Proof. By submodularity,

$$\operatorname{cost}(A \cup B) \leq \operatorname{cost}(A) + \operatorname{cost}(B) - \operatorname{cost}(A \cap B).$$

Since x is feasible for S,

$$\sum_{i \in (A \cap B)} x_i \leq \operatorname{cost}(A \cap B).$$

Combining this with the fact that A and B are both tight, we get

$$\operatorname{cost}(A \cup B) \leq \sum_{i \in A} x_i + \sum_{i \in B} x_i - \sum_{i \in (A \cap B)} x_i = \sum_{i \in (A \cup B)} x_i.$$

Therefore, $A \cup B$ must also be tight. ◻

COROLLARY 3. *If* cost *is submodular and x is feasible for S, then there is a unique maximal tight set. It is given by $\{i \in S \mid i$ belongs to some tight set$\}$.*

For each set $S \subseteq U$, the algorithm below computes a cost allocation for S.

ALGORITHM 1. We will associate a notion of time with our algorithm. Initially, the time t is set to zero. As the algorithm proceeds, we raise cost shares of users in S in proportion to their respective functions f_i; thus, at time t, the cost share of a user i is $f_i(t)$. Whenever a set $A \subseteq S$ goes tight, the cost shares of all users in A are frozen at the current value. The cost shares of the remaining users keep increasing with time as before. The algorithm terminates when the cost shares of all users are frozen. For each user $i \in S$, define $\xi^S(i)$ to be i's cost share at termination.

By Corollary 3, at any time, there is a unique maximal tight set. This tight set can be found in polynomial time using a submodular function minimization algorithm [14, 26] as follows. The difference of a submodular function and a modular function is a submodular function. Hence, the following function, defined on 2^S, is submodular:

$$\operatorname{cost}'(A) = \operatorname{cost}(A) - \sum_{i \in A} f_i(t_i)$$

for $A \subseteq S$, where t_i's are fixed for each element $i \in S$. For an element i that is already frozen, fix t_i to be the time at which i froze. For an element i that is not yet frozen, let $t_i = t$. Now we will do a binary search on t to find the smallest time at which there is a set $A \subseteq S$ such that $\operatorname{cost}'(A)$ is a small negative number. At that value of time, the set whose cost$'$ is minimum will clearly be the maximal set to go tight next.

REMARK 4. *Observe that $\boldsymbol{t}(\xi^S)$ is precisely the vector of times at which individual elements went tight ($\boldsymbol{t}$ is defined at the beginning of section 3).*

LEMMA 5. *For each set $S \subseteq U$, the cost allocation given by ξ^S lies in the core.*

Proof. Clearly, ξ^S is feasible for S and no subset of S is overtight. Furthermore, by Corollary 3, at termination, set S must be tight. ◻

THEOREM 6. *For any set $S \subseteq U$, the cost allocation, ξ^S, found by Algorithm 1 is such that $\boldsymbol{t}(\xi^S)$ max-min dominates $\boldsymbol{t}(\alpha)$ for all other cost allocations, α, for S in the core.*

Proof. Let α be an allocation for set S that lies in the core. Suppose that $\boldsymbol{t}(\xi^S)$ does not max-min dominate $\boldsymbol{t}(\alpha)$. Then we will show that ξ^S and α are in fact the same allocation, hence proving the theorem.

Let $A_1 \subset A_2 \subset \cdots \subset S$ be the sequence of maximal sets that go tight when the algorithm is run on set S. We will show by induction on i that all users in A_i must have the same cost allocation in α and ξ^S. Observe that all elements in $A_i - A_{i-1}$ go tight at the same time, and hence the components corresponding to them in $\boldsymbol{t}(\xi^S)$ are identical.

Clearly,

$$\sum_{i \in A_1} \alpha(i) \leq \sum_{i \in A_1} \xi^S(i) = \mathrm{cost}(A_1).$$

If this inequality is strict, there exists $i \in A_1$ such that $\alpha(i) < \xi^S(i)$. Since users $i \in A_1$ give rise to the smallest entries of $\boldsymbol{t}_{INC}(\xi^S)$, $\boldsymbol{t}(\xi^S)$ max-min dominates $\boldsymbol{t}(\alpha)$, leading to a contradiction. Therefore, this inequality must hold with equality. If for some user $i \in A_1$, $\alpha(i) > \xi^S(i)$, then for some other user $j \in A_1$, $\alpha(j) < \xi^S(j)$, and again $\boldsymbol{t}(\xi^S)$ max-min dominates $\boldsymbol{t}(\alpha)$, leading to a contradiction. Therefore,

$$\forall i \in A_1, \quad \alpha(i) = \xi^S(i).$$

The idea for the induction step is the same as for the basis. □

REMARK 7. *Observe that the precise manner of "dual" increase in the algorithm was essential for proving Theorem* 6.

For a definition of equitable cost sharing method, see the beginning of section 3.

COROLLARY 8. *The cost sharing method, ξ, found by Algorithm* 1 *is the equitable cost sharing method for equalizing functions $f_1, \dots, f_n$.*

THEOREM 9. *The cost sharing method ξ is cross-monotonic.*

Proof. Suppose that $S \subset T \subseteq U$. Let us call the two runs of the algorithm S-run and T-run, respectively. It suffices to prove that at each time t, the tight set in the T-run is a superset of the tight set in the S-run, because then each user $i \in S$ can be frozen only at an earlier time in the T-run and hence can have only a smaller cost share under the T-run.

Consider time t, and let A and B be the tight sets in the S- and T-runs, respectively. Let x_i denote the cost share of $i \in S$ at time t under the S-run, and let x'_i denote the cost share of $i \in T$ at time t under the T-run.

By submodularity,

$$\mathrm{cost}(A \cup B) + \mathrm{cost}(A \cap B) \leq \mathrm{cost}(A) + \mathrm{cost}(B).$$

Since x is feasible for S, we have

$$\sum_{i \in (A \cap B)} x_i \leq \mathrm{cost}(A \cap B).$$

Using the additional fact that A and B are tight in the S- and T-runs, respectively, we get

$$\mathrm{cost}(A \cup B) + \sum_{i \in (A \cap B)} x_i \leq \sum_{i \in A} x_i + \sum_{i \in B} x'_i.$$

Therefore,

$$\mathrm{cost}(A \cup B) \leq \sum_{i \in (A - B)} x_i + \sum_{i \in B} x'_i.$$

Observe that at time t, the users in $A - B$ are frozen in the S-run but not in the T-run. Therefore for each $i \in A - B$, $x_i \leq x'_i$. Therefore,

$$\mathrm{cost}(A \cup B) \leq \sum_{i \in (A \cup B)} x'_i.$$

Therefore, $A \cup B$ is tight at time t in the T-run. Hence $A \subseteq B$, and the theorem follows. ∎

5. Alternative characterizations of equitable methods.

In this section, we will give two alternative characterizations of equitable methods. The proofs of both characterizations appeal to Algorithm 1, and we do not know of direct proofs. These alternative characterizations are as basic as the definition of equitable methods and could have been taken as alternative definitions of these methods.

5.1. Min-max domination.

Let $\boldsymbol{q}$ and $\boldsymbol{r}$ be n-dimensional vectors with nonnegative coordinates. We will denote by $\boldsymbol{q}_{DEC}$ the vector obtained by sorting the components of $\boldsymbol{q}$ in decreasing order and will say that $\boldsymbol{q}$ *min-max dominates* $\boldsymbol{r}$ if $\boldsymbol{q}_{DEC}$ is lexicographically smaller than $\boldsymbol{r}_{DEC}$, i.e., if there is an i such that $\boldsymbol{q}_{DEC}(i) < \boldsymbol{r}_{DEC}(i)$ and $\boldsymbol{q}_{DEC}(j) = \boldsymbol{r}_{DEC}(j)$ for $j < i$.

THEOREM 10. *For any set $S \subseteq U$, the cost allocation, ξ^S, found by Algorithm 1 is such that $\boldsymbol{t}(\xi^S)$ min-max dominates $\boldsymbol{t}(\alpha)$ for all other cost allocations, α, for S in the core.*

Proof. The proof is similar to that of Theorem 6. Let α be an allocation for set S that lies in the core. Suppose that $\boldsymbol{t}(\xi^S)$ does not min-max dominate $\boldsymbol{t}(\alpha)$. Then we will show that ξ^S and α are in fact the same allocation, hence proving the theorem.

Let $S = A_1 \supset A_2 \supset \cdots \supset \emptyset$ be the *reverse* order in which sets go tight when the algorithm is run on set S. We will show by induction on i that all users in $A_i - A_{i+1}$ must have the same cost allocation in α and ξ^S. Observe that all elements in $A_i - A_{i+1}$ go tight at the same time, and hence the components corresponding to them in $\boldsymbol{t}(\xi^S)$ are identical.

Clearly,

$$\sum_{i \in A_2} \alpha(i) \leq \sum_{i \in A_2} \xi^S(i) = \mathrm{cost}(A_2).$$

Therefore,

$$\sum_{i \in A_1 - A_2} \alpha(i) \geq \sum_{i \in A_1 - A_2} \xi^S(i).$$

If this inequality is strict, there exists $i \in A_1 - A_2$ such that $\alpha(i) > \xi^S(i)$. Since users $i \in A_1 - A_2$ give rise to the largest entries in $\boldsymbol{t}_{INC}(\xi^S)$, $\boldsymbol{t}(\xi^S)$ min-max dominates $\boldsymbol{t}(\alpha)$, leading to a contradiction. Therefore, this inequality must hold with equality. If for some user $i \in A_1 - A_2$, $\alpha(i) < \xi^S(i)$, then for some other user $j \in A_1 - A_2$, $\alpha(j) > \xi^S(j)$, and again $\boldsymbol{t}(\xi^S)$ min-max dominates $\boldsymbol{t}(\alpha)$, leading to a contradiction. Therefore,

$$\forall i \in A_1 - A_2, \quad \alpha(i) = \xi^S(i).$$

The idea for the induction step is the same as for the basis. ∎

5.2. *L*-domination. The next characterization is along the lines of Dutta and Ray's definition of the egalitarian method which uses the notion of Lorentz orderings. Following is their definition when applied to the case of a submodular cost function (for the complete definition, which involves a recursive construct, see [6]).

Let $\alpha_1 \leq \cdots \leq \alpha_n$ and $\beta_1 \leq \cdots \leq \beta_n$ be such that $\alpha_1 + \cdots + \alpha_n = \beta_1 + \cdots + \beta_n$. We will say that $(\alpha_1, \ldots, \alpha_n)$ *Lorentz dominates* $(\beta_1, \ldots, \beta_n)$ if for $1 \leq k \leq n$, we have

$$\alpha_1 + \cdots + \alpha_k \geq \beta_1 + \cdots + \beta_k,$$

and the inequality is strict for at least one k.

Dutta and Ray [6] showed that if the underlying cost function is submodular, there is a core cost allocation for S, say $\boldsymbol{q}$, such that $\boldsymbol{q}_{INC}$ Lorentz dominates $\boldsymbol{r}_{INC}$ for all other cost allocations, $\boldsymbol{r}$, for S in the core. $\boldsymbol{q}$ is the egalitarian cost allocation for S.

Note that the definition of Lorentz ordering may compare cost shares of different users, since $\boldsymbol{q}_{INC}(i)$ and $\boldsymbol{r}_{INC}(i)$ may be cost shares of different users. In our setting, the equalizing functions for different users may be very different, thus making such comparisons meaningless. We give below an ordering that takes this into consideration. This ordering is not a generalization of Lorentz ordering—if all f_i's are the identity function, it does not necessarily reduce to the Lorentz ordering, but it does preserve the property established in Lemma 11.

Let V be the set of all nonnegative s-dimensional vectors $\boldsymbol{q}$ such that $f_1(\boldsymbol{q}(1)) + \cdots + f_s(\boldsymbol{q}(s)) = \text{cost}(S)$. Let $V_c \subseteq V$ be the set of vectors $\boldsymbol{q}$ such that $(f_1(\boldsymbol{q}(1)), \ldots, f_s(\boldsymbol{q}(s)))$ forms a cost allocation for S lying in the core. For $\boldsymbol{q}, \boldsymbol{r} \in V$, say that $\boldsymbol{q}$ *L-dominates* $\boldsymbol{r}$ if there exists a permutation π such that

1. $\boldsymbol{q}(\pi(1)) \leq \cdots \leq \boldsymbol{q}(\pi(s))$,
2. for $1 \leq i \leq s$, we have

$$\sum_{k=1}^{i} f_{\pi(k)}(\boldsymbol{q}(\pi(k))) \geq \sum_{k=1}^{i} f_{\pi(k)}(\boldsymbol{r}(\pi(k))),$$

   and the inequality is strict for at least one i.

It is easy to construct examples showing that the relation defined above is not necessarily transitive. However, it is *acyclic* in the following sense: if $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_k \in V$, then it cannot be the case that $\boldsymbol{q}_i$ L-dominates $\boldsymbol{q}_{i+1}$, for $1 \leq i \leq k$, and $\boldsymbol{q}_k$ L-dominates $\boldsymbol{q}_1$. The definition of L-dominates requires that vectors $\boldsymbol{q}$ and $\boldsymbol{r}$ be ordered according to the *same* permutation—it is for this reason that this notion does not generalize the notion of Lorentz ordering.

LEMMA 11. *Let $\boldsymbol{q}, \boldsymbol{r} \in V$, and suppose that $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$. Then $\boldsymbol{q}_{INC}$ is lexicographically larger than $\boldsymbol{r}_{INC}$.*

Proof. Without loss of generality, assume that the permutation π showing that $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$ is the identity permutation. Let i be the smallest index, $1 \leq i \leq s$, such that $\sum_{k=1}^{i} f_k(\boldsymbol{q}(k)) > \sum_{k=1}^{i} f_k(\boldsymbol{r}(k))$. Clearly, $\boldsymbol{q}(k) = \boldsymbol{r}(k)$, for $k < i$, and $\boldsymbol{q}(i) > \boldsymbol{r}(i)$. Therefore, $\boldsymbol{r}(1) \leq \cdots \leq \boldsymbol{r}(i)$. By assumption, $\boldsymbol{q} = \boldsymbol{q}_{INC}$. On the other hand, the first i components of $\boldsymbol{r}_{INC}$ can be only smaller than the corresponding components of $\boldsymbol{r}$. Hence $\boldsymbol{q}_{INC}$ is lexicographically larger than $\boldsymbol{r}_{INC}$. ∎

Since lexicographic domination is acyclic, we get the following corollary.

COROLLARY 12. *The relation of L-domination is acyclic.*

We will need the following technical lemma for the main result.

LEMMA 13. *Let $M, a_1, \ldots, a_l, m, b_1, \ldots, b_l$ be nonnegative real numbers satisfying*

- $M \geq m$,
- $M + a_1 + \cdots + a_l \geq m + b_1 + \cdots + b_l$.

Then, there is a permutation π *over* $[1, \ldots, l]$ *such that for* $1 \leq i \leq l$,

$$M + \sum_{k=1}^{i} a_{\pi(k)} \geq m + \sum_{k=1}^{i} b_{\pi(k)}.$$

Proof. Let us show how to construct π. Let us first determine $\pi(1)$. We claim that there exists k, $1 \leq k \leq l$, such that

$$M + a_k \geq m + b_k.$$

Suppose not. Then, for $1 \leq k \leq l$,

$$M + a_k < m + b_k.$$

Adding these l inequalities, we get

$$(l-1)M + (M + a_1 + \cdots + a_l) < (l-1)m + (m + b_1 + \cdots + b_l).$$

But this contradicts the assumptions made on these numbers, proving existence of k satisfying the inequality above. Set $\pi(1) = k$. The idea for constructing the rest of π is the same. ∎

THEOREM 14. *Let* $\boldsymbol{q} \in V_c$ *correspond to the equitable cost allocation* χ^S, *and let* $\boldsymbol{r} \in V_c$ *be any other vector. Then* $\boldsymbol{q}$ *L-dominates* $\boldsymbol{r}$.

Proof. We need to construct permutation π that shows that $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$. Let $A_1 \subset A_2 \subset \cdots \subset S$ be the sequence in which sets go tight when Algorithm 1 computes cost shares for S. In the simple case that each of $A_{i+1} - A_i$ is a singleton, let π be the order in which elements go tight. Then, for $1 \leq i \leq s$,

$$\sum_{k=1}^{i} f_{\pi(k)}(\boldsymbol{q}(\pi(k))) = \mathrm{cost}(A_i) \geq \sum_{k=1}^{i} f_{\pi(k)}(\boldsymbol{r}(\pi(k))).$$

The inequality holds because the cost allocation corresponding to $\boldsymbol{r}$ lies in the core. Since $\boldsymbol{q} \neq \boldsymbol{r}$, one of these inequalities must be strict, thereby showing that $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$.

In the general case, we will order elements of A_1 first, the elements of $A_2 - A_1$ next, and so on. This ensures that

$$\sum_{k \in A_i} f_k(\boldsymbol{q}(k)) = \mathrm{cost}(A_i) \geq \sum_{k \in A_i} f_k(\boldsymbol{r}(k)).$$

Next, let us specify the precise order given to elements of $A_{i+1} - A_i = \{j_1, \ldots, j_l\}$. For this, we will use Lemma 13 with

$$M = \sum_{k \in A_i} f_k(\boldsymbol{q}(k)) = \mathrm{cost}(A_i), \quad m = \sum_{k \in A_i} f_k(\boldsymbol{r}(k))$$

and for $1 \leq k \leq l$,

$$a_k = f_{j_k}(\boldsymbol{q}(j_k)) \quad \text{and} \quad b_k = f_{j_k}(\boldsymbol{r}(j_k)). \qquad ∎$$

Observe that the proof given above uses the fact that the functions f_j are *strictly* increasing.

6. The opportunity egalitarian method. For each user $i \in U$, let $G_i : \mathbf{R}^+ \to [0,1]$ be the cumulative probability distribution function from which i's utility is drawn; assume that these distributions are independent. Assume that G_i is monotonically increasing.

Let α be any cost allocation for $S \subseteq U$ that lies in the core. User i will accept the service only if his utility turns out to be at least $\alpha(i)$, his cost share. The probability of this event is $1 - G_i(\alpha(i))$. Let $\boldsymbol{p}(\alpha)$ denote the vector whose ith component is this probability. We will say that a cost sharing method ξ is the *opportunity egalitarian method* for cumulative distribution functions $G_1, \ldots, G_n$ if for each set $S \subseteq U$, ξ^S is the cost allocation in the core such that

1. $\boldsymbol{p}(\xi^S)$ max-min dominates $\boldsymbol{p}(\alpha)$ and
2. $\boldsymbol{p}(\xi^S)$ min-max dominates $\boldsymbol{p}(\alpha)$

for all other cost allocations, α, for S in the core. Theorem 15 shows that there is a unique such cost allocation. Observe that ξ is attempting to equalize the probabilities of users receiving the service, subject to core constraints. The two characterizations, min-max and max-min, ensure that both extremes are avoided.

Let $f_i : [0,1] \to \mathbf{R}^+$ denote the inverse of G_i. Let ξ be the equitable cost sharing method for functions $f_1, \ldots, f_n$.

THEOREM 15. *ξ is the opportunity egalitarian method for probability density functions $G_1, \ldots, G_n$.*

Proof. Consider any set $S \subseteq U$, and let α be a cost allocation for S lying in the core. Clearly, for $i \in S$,

$$\boldsymbol{t}(\alpha)(i) = f_i^{-1}(\alpha(i)) = G_i(\alpha(i)) = 1 - \boldsymbol{p}(\alpha)(i);$$

i.e., this is the probability that user i does not accept service under cost allocation α. Since ξ is the equitable method for $f_1, \ldots, f_n$, $\boldsymbol{t}(\xi^S)$ max-min dominates $\boldsymbol{t}(\alpha)$ for all cost allocations, α, for S in the core. Equivalently, $\boldsymbol{p}(\xi^S)$ min-max dominates $\boldsymbol{p}(\alpha)$ for all core cost allocations for S. Its uniqueness follows from Theorem 6. By Theorem 10, $\boldsymbol{p}(\xi^S)$ max-min dominates $\boldsymbol{p}(\alpha)$ for all other cost allocations, α, for S in the core. Hence, ξ is the opportunity egalitarian method for cumulative distribution functions $G_1, \ldots, G_n$. $\square$

By Theorem 9, the opportunity egalitarian method is cross-monotone.

7. Maximizing acceptance probability. As in the last section, for each user $i \in U$, let $G_i : \mathbf{R}^+ \to [0,1]$ be the cumulative probability distribution function from which i's utility is drawn; assume that these distributions are independent. Let g_i be the corresponding probability density function, i.e., $g_i(x) = \frac{\partial G_i(x)}{\partial x}$. Assume that $g_i(0) = 0$. We further assume that G_i satisfies the *strict monotone hazard rate condition*, i.e.,

$$\frac{\partial}{\partial x}\left[\frac{g_i(x)}{1 - G_i(x)}\right] > 0.$$

If g_i represents the failure probability of a component as a function of time, the above condition says that the failure rate, conditioned on the component still being intact, strictly increases with age. The monotone hazard rate condition is satisfied by most standard probability distributions and is a standard assumption in the Bayesian mechanism design literature [8].

Let $\lambda_i : \mathbf{R}^+ \to \mathbf{R}^+$ be the function

$$\lambda_i(x) = \frac{g_i(x)}{1 - G_i(x)}.$$

Observe that under the assumption of strict monotone rate hazard rate condition, λ_i is an invertible function. Let f_i be the inverse of this function, and let χ be the equitable method corresponding to equalizing functions $f_1, f_2, \ldots, f_n$.

Let ξ be any cost sharing method in the core. Let $P(\xi^S)$ denote the probability that all users in S accept service when it is offered at cost shares ξ^S. Clearly,

$$P(\xi^S) = \prod_{i \in S} [1 - G_i(\xi_i^S)].$$

THEOREM 16. *Let χ be the equitable cost sharing method defined above, and ξ be any cost sharing method in the core. Then, for each set of users $S \subseteq U$, $P(\chi^S) \geq P(\xi^S)$.*

This is a generalization of Mutuswami's result [24], which deals with the case that G_i's are i.i.d. (on the other hand, he requires only monotone hazard rate condition—not necessarily *strict*). In this case, χ is the Dutta–Ray egalitarian method. Mutuswami's proof uses the original definition of Dutta and Ray, in terms of Lorentz orderings, and a theorem of Hardy, Littlewood, and Polya [10], giving a condition that is equivalent to Lorentz domination. To prove the generalization, we prove a characterization for L-domination, in Lemma 18, in the style of the Hardy–Littlewood–Polya theorem. The proof of Theorem 16 is given below.

By Theorem 16, χ *simultaneously* maximizes the probability of all users accepting service, for each set S of users, among all cost sharing methods in the core. One may be led to believe that if the mechanism of Theorem 1 is run with χ, then the expected size of the set served is maximized, over all cost sharing methods in the core. However, this is not true, as shown in the example below.

Example. Let $U = \{a, b\}$ and the cost function be $\text{cost}(a, b) = 10$, $\text{cost}(a) = 8$, $\text{cost}(b) = 6$, $\text{cost}(\emptyset) = 0$. Suppose the utilities of a and b are picked from the uniform distribution over the interval $[0, 20]$. In this case, χ will be the egalitarian method:

$$\chi(\{a, b\}, a) = \chi(\{a, b\}, b) = 5, \quad \chi(\{a\}, a) = 8, \quad \chi(\{b\}, b) = 6.$$

The expected size of set picked by the mechanism is 1.45. However, using the following cross-monotonic cost sharing method, ξ, the expected size of set picked is 1.45125:

$$\xi(\{a, b\}, a) = 5.5, \quad \xi(\{a, b\}, b) = 4.5, \quad \xi(\{a\}, a) = 8, \quad \xi(\{b\}, b) = 6.$$

For completeness, we first state the following theorem.

THEOREM 17 (Hardy, Littlewood, and Polya [10]). *Let $\alpha_1 \leq \cdots \leq \alpha_n$ and $\beta_1 \leq \cdots \leq \beta_n$ be such that $\alpha_1 + \cdots + \alpha_n = \beta_1 + \cdots + \beta_n$. The following two statements are equivalent.*

1. *$(\alpha_1, \ldots, \alpha_n)$ Lorentz dominates $(\beta_1, \ldots, \beta_n)$.*
2. *$(\alpha_1, \ldots, \alpha_n)$ can be obtained from $(\beta_1, \ldots, \beta_n)$ by applying the following transformations to $(\beta_1, \ldots, \beta_n)$ a finite number of times:*
   (a) *Find $i < j$ such that $\beta_i < \beta_j$ and the next step can be performed.*
   (b) *Increase β_i and decrease β_j by a small $\epsilon > 0$.*

Following is an analogous fact for L-domination. As defined in section 5.2, let V be the set of all nonnegative s-dimensional vectors $\boldsymbol{q}$ such that $f_1(\boldsymbol{q}(1)) + \cdots + f_s(\boldsymbol{q}(s)) = \text{cost}(S)$.

LEMMA 18. *Let $\boldsymbol{q}, \boldsymbol{r} \in V$. The following two statements are equivalent.*

1. *$\boldsymbol{q}$ L-dominates $\boldsymbol{r}$, with π being the identity permutation.*
2. *$\boldsymbol{q}$ can be obtained from $\boldsymbol{r}$ by applying the following transformations to $\boldsymbol{r}$ a nonzero number of times:*

(a) *Find* $i < j$ *such that* $\boldsymbol{r}(i) < \boldsymbol{r}(j)$ *and the next step can be performed.*

(b) *Increase* $\boldsymbol{r}(i)$ *and decrease* $\boldsymbol{r}(j)$ *by small positive amounts such that*

- $f_i(\boldsymbol{r}(i)) + f_j(\boldsymbol{r}(j))$ *is preserved,*
- $\boldsymbol{r}(i) \le \boldsymbol{r}(j)$ *is preserved.*

Proof. We will prove the forward direction only; the reverse direction is straightforward. Let i be the smallest index such that $\boldsymbol{q}(i) \neq \boldsymbol{r}(i)$ (if such an index i does not exist, then $\boldsymbol{q}$ and $\boldsymbol{r}$ are identical). Since $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$, $\boldsymbol{q}(i) > \boldsymbol{r}(i)$. By definition,

$$f_1(\boldsymbol{q}(1)) + \cdots + f_s(\boldsymbol{q}(s)) = f_1(\boldsymbol{r}(1)) + \cdots + f_s(\boldsymbol{r}(s)).$$

Let j be the smallest index $> i$ such that

$$f_1(\boldsymbol{q}(1)) + \cdots + f_j(\boldsymbol{q}(j)) = f_1(\boldsymbol{r}(1)) + \cdots + f_j(\boldsymbol{r}(j)).$$

By the choice of j, it must be the case that

$$f_1(\boldsymbol{q}(1)) + \cdots + f_{j-1}(\boldsymbol{q}(j-1)) > f_1(\boldsymbol{r}(1)) + \cdots + f_{j-1}(\boldsymbol{r}(j-1));$$

therefore $\boldsymbol{q}(j) < \boldsymbol{r}(j)$. Since $i < j$, $\boldsymbol{q}(i) \le \boldsymbol{q}(j)$. Now we have $\boldsymbol{r}(i) < \boldsymbol{q}(i) \le \boldsymbol{q}(j) < \boldsymbol{r}(j)$. Increase $\boldsymbol{r}(i)$ and decrease $\boldsymbol{r}(j)$ as much as possible so that the following are preserved:

1. $f_i(\boldsymbol{r}(i)) + f_j(\boldsymbol{r}(j))$.
2. $\boldsymbol{r}(i) \le \boldsymbol{r}(j)$.
3. For $1 \le k \le s$,

$$f_1(\boldsymbol{q}(1)) + \cdots + f_k(\boldsymbol{q}(k)) \ge f_1(\boldsymbol{r}(1)) + \cdots + f_k(\boldsymbol{r}(k)).$$

Clearly, $\boldsymbol{r}(i)$ must increase by a positive amount, and $\boldsymbol{r}(j)$ must decrease by a positive amount. Therefore, in the limit, $\boldsymbol{q}$ can be obtained from $\boldsymbol{r}$ by applying such steps. ∎

Proof of Theorem 16. Consider any set of users $S \subseteq U$. Let $\boldsymbol{q}, \boldsymbol{r} \in V_c$ correspond to χ^S and ξ^S, respectively. By Theorem 14, $\boldsymbol{q}$ L-dominates $\boldsymbol{r}$. Now, $\boldsymbol{q}$ can be obtained from $\boldsymbol{r}$ by steps specified in Lemma 18. Finally, by Lemma 19 given below, each of these steps will only increase the probability that all users accept service. Hence $P(\chi^S) \ge P(\xi^S)$. ∎

LEMMA 19. *Let* $f_1, \ldots, f_s$ *be the equalizing functions specified above. Let* $x_i \ge 0$, $1 \le i \le s$, *and define* $\boldsymbol{v} = (f_1(x_1), \ldots, f_s(x_s))$. *Define* $P(\boldsymbol{v}) = \prod_{i=1}^{s}[1 - G_i(f_i(x_i))]$. *Suppose there exist* i *and* j *such that* $x_i < x_j$. *Then,*

$$dP(\boldsymbol{v}) = \sum_{k=1}^{s} \frac{\partial P(\boldsymbol{v})}{\partial f_k(x_k)} df_k(x_k) \ge 0,$$

where $0 < df_i(x_i) = -df_j(x_j)$ *and* $df_k(x_k) = 0$ *if* $k \neq i, j$.

Proof.

$$dP(\boldsymbol{v}) = \sum_{k=1}^{s} \frac{\partial P(\boldsymbol{v})}{\partial f_k(x_k)} df_k(x_k)$$

$$= \sum_{k=1}^{s} \frac{P(\boldsymbol{v})}{1 - G(f_k(x_k))} \times (-g(f_k(x_k))) \times df_k(x_k)$$

$$= P(\boldsymbol{v}) \left[\frac{g_j(f_j(x_j))}{1 - G_j(f_j(x_j))} - \frac{g_i(f_i(x_i))}{1 - G_i(f_i(x_i))} \right] dx$$

$$= P(\boldsymbol{v})[\lambda_j(f_j(x_j)) - \lambda_i(f_i(x_i))]dx$$

$$= P(\boldsymbol{v})[x_j - x_i]dx \geq 0,$$

where $dx = df_i(x_i) = -df_j(x_j)$.

The second-to-last step uses the fact that f_i is the inverse of λ_i and the last step follows from the assumption that $x_i < x_j$. □

8. Discussion. Submodular cost functions admit a rich class of cross-monotone cost sharing methods. Of these, equitable methods capture a large subclass, though not all, as evidenced by the example below. Our definition as well as Hokari's [11] appears quite natural, and despite differences, they lead to the same class of cost sharing methods. This raises the following question: how to impose an economic criterion on cross-monotone methods so as to obtain precisely the class of equitable methods.

Example. Consider the following cross-monotone cost sharing method (the costs of individual sets are simply the sum of cost shares of their elements).
$\xi(\{a, b, c, d\}, a) = 2$,
$\xi(\{a, b, c, d\}, b) = 2$, $\xi(\{a, b, c, d\}, c) = 1$, $\xi(\{a, b, c, d\}, d) = 1$,
$\xi(\{a, b, c\}, a) = 2$, $\xi(\{a, b, c\}, b) = 3$, $\xi(\{a, b, c\}, c) = 1$,
$\xi(\{a, b, d\}, a) = 3$, $\xi(\{a, b, d\}, b) = 2$, $\xi(\{a, b, d\}, d) = 1$,
$\xi(\{a, c, d\}, a) = 2$, $\xi(\{a, c, d\}, c) = 2$, $\xi(\{a, c, d\}, d) = 2$,
$\xi(\{b, c, d\}, b) = 2$, $\xi(\{b, c, d\}, c) = 2$, $\xi(\{b, c, d\}, d) = 2$,
$\xi(\{a, b\}, a) = 3$, $\xi(\{a, b\}, b) = 3$,
$\xi(\{a, c\}, a) = 3$, $\xi(\{a, c\}, c) = 3$,
$\xi(\{a, d\}, a) = 3$, $\xi(\{a, d\}, d) = 3$,
$\xi(\{b, c\}, b) = 3$, $\xi(\{b, c\}, c) = 3$,
$\xi(\{b, d\}, b) = 3$, $\xi(\{b, d\}, d) = 3$,
$\xi(\{c, d\}, c) = 3$, $\xi(\{c, d\}, d) = 3$,
$\xi(\{a\}, a) = 5$, $\xi(\{b\}, b) = 5$, $\xi(\{c\}, c) = 5$, $\xi(\{d\}, d) = 5$.
We will show, by contradiction, that this is not an equitable cost sharing method. First run Algorithm 1 on $S = \{a, b, c\}$. In order to produce the above method, S must be the first set to go tight. Suppose this happens at time t. Clearly, none of the proper subsets of S is tight at time t. Therefore, $f_a(t) = 2$, and $f_b(t) = 3$. Next, run Algorithm 1 on set $S' = \{a, b, d\}$. The first set to go tight must be the entire set S', at time t', say. Therefore, $f_a(t') = 3$, and $f_b(t') = 2$. But then at least one of f_a or f_b is not monotonically increasing.

Observe that Algorithm 1 is very explicitly trying to impose equality among users, as specified by the equalizing functions. Of course, the precise notion of "equality" or "fairness" imposed depends on these functions. An exciting research direction is to characterize the notions of fairness captured by particular choices of the equalizing functions.

This very explicit seeking of "equality" is also the chief difference between our definition and Hokari's definition and algorithm. Hokari defines *sequential monotone path* cost sharing methods by partitioning the set of users, ordering the partitions sequentially, and applying monotone path methods within each partition, with an incremental method applied on the ordered list of partitions. In our setting it is easy to see that the resulting method is also equitable, and so this operation does not lead to new cost sharing methods.

Considering the fact that for submodular cost functions, our definition is a natural generalization of the Dutta–Ray egalitarian method, one wonders whether there is a similar generalization for nonsubmodular cost functions as well. L-orderings, presented in section 5.2, may lead to such a definition. At present it is not clear how to generalize Algorithm 1 to this setting—if the cost function is not submodular, then every element may be in a tight set, without the entire set being tight. Algorithm 1 will halt at this point, and the resulting cost allocation will not be budget balanced.

Next, assume that the cost function is nondecreasing and submodular. Let σ be a permutation on $1, \ldots, n$. The *incremental cost sharing method*, ξ_σ, corresponding to permutation σ is defined as follows. Let $S \subseteq U$, $|S| = k$, and let $i_1, \ldots, i_k$ be the users in S ordered according to σ. Then, $\xi_\sigma(S, i_1) = \text{cost}(i_1)$, and for $2 \leq j \leq k$, $\xi_\sigma(S, i_j) = \text{cost}(\{i_1, \ldots, i_j\}) - \text{cost}(\{i_1, \ldots, i_{j-1}\})$.

In [17] we showed that the class of cross-monotone methods for a nondecreasing submodular cost function form a polytope and that incremental cost sharing methods form corner points of this polytope. We also gave an example of a cross-monotone method that is not in the convex hull of these corner points (corresponding to incremental cost sharing methods) and left the open problem of characterizing the rest of the corner points. We show below that this example is in fact an equitable method. It is easy to see that all incremental methods are also equitable. Do equitable methods capture all corner points of this polytope? Are equitable methods closed under convex combinations? Since not all cost sharing methods are equitable, the answers to both these questions cannot be "Yes."

Example. Consider the following cost sharing method:
$\xi(\{a,b,c\},a) = 2$, $\xi(\{a,b,c\},b) = 3$, $\xi(\{a,b,c\},c) = 4$,
$\xi(\{a,b\},a) = 4$, $\xi(\{a,b\},b) = 3$,
$\xi(\{a,c\},a) = 3$, $\xi(\{a,c\},c) = 4$,
$\xi(\{b,c\},b) = 3$, $\xi(\{b,c\},c) = 4$,
$\xi(\{a\},a) = 4$, $\xi(\{b\},b) = 4$, $\xi(\{c\},c) = 4$.
In [17] we showed that this cross-monotone method is not a convex combination of incremental cost sharing methods. However, it is an equitable cost sharing method corresponding to equalizing functions f_a, f_b, f_c satisfying

$$f_a(2) = 0, \quad f_b(1) = 0, \quad f_b(2) = 3, \quad f_b(3) = 3, \quad f_b(4) = 4, \quad \text{and } f_c(1) = 4.$$

Finally, the example given in section 7 raises the question of identifying the cost sharing method that maximizes the expected size of the set served.

Acknowledgments. We wish to thank Bhaskar Dutta, Suresh Mutuswami, and Christos Papadimitriou for valuable comments and for providing references to the literature. Thanks also go to the anonymous STOC referee for valuable suggestions on the paper.

REFERENCES

[1] Y. Afek, Y. Mansour, and Z. Ostfeld, *Convergence complexity of optimistic rate based flow control algorithms*, in Proceedings of the 28th Annual ACM Symposium on Theory of Computing, 1996.

[2] J. Arin, J. Kupers, and D. Vermeulen, *An Axiomatic Approach to Egalitarianism in TU Games*, Discussion paper, Department of Mathematics, Maastricht University, Maastricht, The Netherlands, 2000.

[3] J. Arin, J. Kupers, and D. Vermeulen, *Some Characterization of Egalitarian Solutions on Classes of TU Games*, mimeo, Maastricht University, Maastricht, The Netherlands, 2000.
[4] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, Upper Saddle River, NJ, 1987.
[5] B. Dutta, *The egalitarian solution and reduced game properties in convex games*, Internat. J. Game Theory, 19 (1990), pp. 153–169.
[6] B. Dutta and D. Ray, *A concept of egalitarianism under participation constraints*, Econometrica, 57 (1989), pp. 615–635.
[7] B. Dutta and D. Ray, *Constrained egalitarian allocations*, Games Econom. Behav., 3 (1991), pp. 403–422.
[8] D. Fudenberg and J. Tirole, *Game Theory*, MIT Press, Cambridge, MA, 1991.
[9] A. Goel, A. Meyerson, and S. Plotkin, *Approximate majorization and fair online load balancing*, in Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, 2001, pp. 384–390.
[10] G. Hardy, H. Littlewood, and G. Polya, *Inequalities*, Cambridge University Press, Cambridge, UK, 1934.
[11] T. Hokari, *Monotone-path Dutta-Ray solutions on convex games*, Soc. Choice Welf., 19 (2002), pp. 825–844.
[12] T. Hokari, *Weighted Dutta-Ray Solutions on Convex Games*, mimeo, University of Rochester, Rochester, NY, 1998.
[13] J. L. Hougard, B. Peleg, and L. P. Osterdal, *The Dutta-Ray Solution on the Class of Convex Games: A Generalisation and Its Monotonicity Properties*, mimeo, University of Copenhagen, Copenhagen, Denmark, 2000.
[14] S. Iwata, L. Fleischer, and S. Fujishige, *A strongly polynomial-time algorithm for minimizing submodular functions*, J. ACM, 48 (2001), pp. 761–777.
[15] J. Jaffe, *Bottleneck flow control*, IEEE Trans. Comm., 29 (1981), pp. 954–962.
[16] J. Y. Jaffray and P. Mongin, *Constrained egalitarianism in a simple redistribution model*, Theory and Decision, 54 (2003), pp. 33–56.
[17] K. Jain and V. V. Vazirani, *Applications of approximation algorithms to cooperative games*, in Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, 2001, pp. 364–372.
[18] K. Jain and V. V. Vazirani, *Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation*, J. ACM, 48 (2001), pp. 274–296.
[19] K. Jain and V. V. Vazirani, *Eisenberg-Gale markets: Algorithms and structural properties*, in Proceedings of the 39th Annual ACM Symposium on Theory of Computing, 2007, pp. 364–373.
[20] J. Kleinberg, Y. Rabani, and E. Tardos, *Fairness in routing and load balancing*, in Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science, 1999, p. 568.
[21] F. Klijn, M. Slikker, S. Tijs, and J. Zarzuelo, *The egalitarian solution for convex games: Some characterizations*, Math. Social Sci., 40 (2000), pp. 111-121.
[22] N. Megiddo, *Optimal flows in networks with multiple sources and sinks*, Math. Programming, 7 (1974), pp. 97–107.
[23] H. Moulin, *Incremental cost sharing: Characterization by coalition strategy-proofness*, Soc. Choice Welf., 16 (1999), pp. 279–320.
[24] S. Mutuswami, *Strategyproof Mechanism for Cost Sharing*, Economics Discussion Papers 520, University of Essex, Colchester, UK, 2000.
[25] A. M. Odlyzko, *Internet pricing and the history of communications*, Computer Networks, 36 (2001), pp. 493–517.
[26] A. Schrijver, *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, J. Combin. Theory Ser. B, 80 (2000), pp. 346–355.
[27] H. Varian, *Pricing electronic journals*, D-Lib Magazine, June (1996); available online from http://www.dlib.org/dlib/june96/06varian.html.
[28] H. Varian, *Differential pricing and efficiency*, First Monday, 1 (2) (1996).
[29] H. Varian, *Pricing information goods*, in Proceedings of Research Libraries Group Symposium on "Scholarship in the New Information Environment," Harvard Law School, Cambridge, MA, 1995.
[30] V. V. Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.