

Limitations of cross-monotonic cost sharing schemes

Nicole Immorlica*

Mohammad Mahdian*

Vahab S. Mirrokni*

Abstract

A cost sharing scheme is a set of rules defining how to share the cost of a service (often computed by solving a combinatorial optimization problem) amongst serviced customers. A cost sharing scheme is cross-monotonic if it satisfies the property that everyone is better off when the set of people who receive the service expands. Cross-monotonic cost sharing schemes are used to define group-strategyproof mechanisms. In this paper, we investigate the limitations imposed by the cross-monotonicity property on cost-sharing schemes for several combinatorial optimization games including edge cover, vertex cover, set cover, metric facility location, maximum flow, arborescence packing, and maximum matching. We develop a novel technique based on the probabilistic method for proving upper bounds on the budget-balance factor of cross-monotonic cost sharing schemes, deriving tight or nearly-tight bounds for each game that we study. For the set cover game, which generalizes many of the above games, we show that no cross-monotonic cost sharing scheme can recover more than a $O(\frac{1}{n})$ fraction of the total cost, respectively, and thus we can not hope to use a set-cover cost sharing scheme as a black box for the cost sharing schemes of covering games. For the vertex cover game, we show no cross-monotonic cost sharing scheme can recover more than a $O(n^{-1/3})$, demonstrating that cross-monotonicity is strictly harder to achieve than the core property (vertex cover games have a solution in the core that is 1/2-budget balanced). For the facility location game, we show that there is no cross-monotonic cost sharing scheme that recovers more than a third of the total cost. This result together with a recent 1/3-budget-balanced cross-monotonic cost sharing scheme of Pál and Tardos [16] closes the gap for the facility location game. Finally, we study the implications of our results on the existence of group-strategyproof mechanisms. We observe that the definition of group-strategyproofness does not exclude trivial mechanisms that recover all the cost. However, with extra assumptions, we show that group-strategyproof mechanisms give rise to cross-monotonic cost

sharing schemes and therefore our upper bounds hold.

1 Introduction

Consider a situation where a group of customers (which we call *agents*) wish to buy a service such as connectivity to a network. The total cost of this service is a function of the group of customers that are serviced: a group of customers in distant towns might incur a larger cost than a group of customers in the same town. The service provider must develop a pricing policy, or *cost-sharing scheme*, that, given any group of customers, divides the cost of the service amongst them. For example, one plausible cost-sharing scheme divides the cost of the service evenly amongst the customers. However, in the case of network connectivity, this scheme seems to undercharge distant customers with high connection costs and overcharge other customers. Developing a fair and economically viable cost-sharing scheme is a central problem in cooperative game theory. One commonly explored condition is that of *cross-monotonicity* [14, 15]. Intuitively, cross-monotonicity requires that the price charged to any individual in a group decreases as the group expands. Thus customers have an economic incentive to promote the service.

For the important class of services with submodular cost functions, various cross-monotonic cost-sharing schemes were studied by Moulin and Shenker [14] and further by Jain and Vazirani [10]. For submodular cost functions, there are cross-monotonic cost-sharing schemes that are *budget-balanced*, i.e., the sum of prices charged to the agents covers the full cost of serving them. There are many other interesting classes of cost functions that arise from NP-hard optimization problems. For example, the cost of providing the service for a set of agents S could be expressed as the cost of building the cheapest Steiner tree that covers the elements of S , or the minimum cost of opening facilities and connecting each member of S to an open facility. These two games, and many others of practical import, are instances of covering problems. For such problems, it is usually impossible for a cross-monotonic cost sharing scheme to be budget-balanced. Moreover, even if a budget-balanced cross-monotonic cost sharing scheme exists, it might be hard to compute. Therefore, it is natural to consider cost sharing schemes that are *approximately budget balanced*, i.e., they recover only a fraction of the cost of the service. Such

*Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139, USA. Email: {nickle,mahdian,mirrokn}@theory.csail.mit.edu. The first author was supported in part by an NSF fellowship. The second author was supported by a Microsoft fellowship. The third author was supported in part by NSF contracts ITR-0121495 and CCR-0098018

schemes have been studied by Kent and Skorin-Kapov [11], Feigenbaum et al. [4], Jain and Vazirani [8], and Pál and Tardos [16]. It is easy to show that if there is an α -budget balanced cross-monotonic cost-sharing scheme for the fractional set cover, then for any special case of the set cover problem of integrality gap at most μ , there is an $\alpha\mu$ -budget balanced cross-monotonic cost-sharing scheme. For example, if we could get a constant-factor for fractional set cover, it would have implied a constant-factor for metric facility location and generalized Steiner tree games. Unfortunately, our result shows that no cross-monotonic cost-sharing scheme for fractional set cover with a reasonable budget-balance factor exists, and thus this approach for designing cross-monotonic cost-sharing schemes fails to recover much of the cost. This raises the natural question of whether it is possible to design well budget-balanced schemes for these combinatorial optimization games.

We can derive simple bounds on the budget-balance factor of combinatorial optimization games using the integrality gaps of the “natural” LP-relaxations. The cross-monotonicity of a cost sharing scheme implies that for every set of agents the cost shares form an allocation in the core of the game (see Section 2 for definitions). Therefore, the best budget-balance factor achievable by a cross-monotonic cost sharing scheme cannot be better than that of a cost sharing in the core. This, together with the folklore theorem that the best budget-balance factor for a cost sharing in the core of integer covering games is equal to the integrality gap of the “natural” LP-relaxation of the problem, gives us upper bounds for the best cross-monotonic cost sharing scheme for various combinatorial optimization problems. For example, this argument implies that cross-monotonic cost sharing schemes for metric facility location, vertex cover, and set cover games cannot recover more than a $\frac{1}{1.463}$, $\frac{1}{2}$, and $\frac{1}{\ln n}$ fraction of the total cost, respectively. Prior to this work, this was the only method known for upper bounding the cross-monotonic cost sharing schemes. In this paper, we show stronger upper bounds for several combinatorial optimization problems using a novel technique based on the probabilistic method that will be explained in Section 3.1. In particular, we prove that the best budget-balance factor achievable for the facility location game is $1/3$. This matches an upper bound recently given by Pál and Tardos [16]. Also, for the vertex cover and set cover games, we show that no cross-monotonic cost sharing scheme can recover more than an $O(n^{-1/3})$ and $O(\frac{1}{n})$ fraction of the total cost, respectively. Previously, Devanur et al. [3] give a strategyproof $\Omega(1/\log(n))$ -budget balanced cost-sharing mechanism in the core for the set cover game, but their underlying cost-sharing scheme is not cross-monotonic. We also apply this technique to several other cost or profit sharing problems including edge cover, maximum flow, maximum matching, and arborescence packing.

Cross-monotonic cost sharing schemes are mainly used to obtain group-strategyproof mechanisms using a method developed by Moulin and Shenker [15, 14]. In fact, almost all known group-strategyproof mechanisms are Moulin-Shenker mechanisms. Therefore, one might hope that our negative results on cross-monotonic cost sharing schemes might imply similar negative results for group-strategyproof mechanisms. However, we observe that for almost any problem there are trivial group-strategyproof mechanisms that recover all the cost. These mechanisms completely ignore the structure of the problem and can therefore be unfair and inefficient. This suggests that new conditions should be added to the definition of group-strategyproofness to exclude such mechanisms. We study a few such conditions, and prove that with the extra assumptions of *no free riders* and *upper continuity*, group-strategyproof mechanisms give rise to cross-monotonic cost sharing schemes, and hence our upper bounds hold for group-strategyproof mechanisms with these extra assumptions. We also consider subsidy-freeness [13], which is a stronger fairness condition and prove the equivalence of budget-balanced cross-monotone cost sharing schemes and budget-balanced group-strategyproof mechanisms with this property.

The rest of this paper is organized as follows. In Section 2, we present the definitions of cross-monotonic cost sharing schemes and group-strategyproof mechanisms. Section 3 contains a description of our upper bound techniques, our upper bounds for the covering game and the facility location game, and the statement of some other results that we have been able to prove using this technique. In Section 4 we present several trivial group-strategyproof mechanisms and study some of the axioms that can be added to the definition of group-strategyproof mechanisms to eliminate such trivial mechanisms.

2 Definitions

Let \mathcal{A} denote a set of n users who are interested in a service. The cost of providing service to a set $S \subseteq \mathcal{A}$ of users is denoted by $C(S)$. A *cost allocation* for a set $S \subseteq \mathcal{A}$ is a function $\psi : S \mapsto \mathbb{R}^+ \cup \{0\}$, that for each user $i \in S$, specifies the share $\psi(i)$ of i of the total cost of servicing S . A *cost-sharing scheme* is a collection of cost allocations for every $S \subseteq \mathcal{A}$. More formally, a cost sharing scheme is a function $\xi : \mathcal{A} \times 2^{\mathcal{A}} \mapsto \mathbb{R}^+ \cup \{0\}$, such that for every $S \subseteq \mathcal{A}$ and every $i \notin S$, $\xi(i, S) = 0$. Intuitively, we think of $\xi(i, S)$ as the share of i of the total cost if S is the set of agents receiving the service.

Ideally, we want cost sharing schemes (and cost allocations) to be *budget-balanced*, i.e., for every $S \subseteq \mathcal{A}$, $\sum_{i \in S} \xi(i, S) = C(S)$. However, it is not always possible to achieve budget balance in combination with other properties, or even if it is possible, it might be computationally hard to compute the cost shares. Therefore, we relax this notion

to the notion of α -budget balance (for some $\alpha \leq 1$), which means that for every $S \subseteq \mathcal{A}$, $\alpha C(S) \leq \sum_{i \in S} \xi(i, S) \leq C(S)$.

In addition to budget balance, we usually require cost allocation and cost sharing schemes to satisfy additional properties. One property that is extensively studied in the cooperative game theory literature is the property of being in the *core* (see, for example, Bondareva [1] and Shapley [18]), which intuitively says that no subset of users should be overcharged for the service.

DEFINITION 2.1. *A cost allocation ψ for a set $S \subseteq \mathcal{A}$ is in the α -core if and only if it is α -budget balanced and for every $T \subseteq S$, $\sum_{i \in T} \psi(i) \leq C(T)$. A cost sharing scheme ξ is in the α -core if and only if for every S , $\xi(\cdot, S)$ is in the α -core.*

Another property, which was studied by Moulin [15] and Moulin and Shenker [14] in order to design group-strategyproof mechanisms (see the definition below), and has recently received considerable attention in the computer science literature (see, for example, [11, 10, 8, 16]), is *cross-monotonicity*. This property captures the notion that users should not be penalized as the serviced set grows. Namely,

DEFINITION 2.2. *A cost sharing scheme ξ is cross-monotone if for all $S, T \subseteq \mathcal{A}$ and $i \in S$, $\xi(i, S) \geq \xi(i, S \cup T)$.*

It is a simple exercise to show that every α -budget-balanced cross-monotonic cost sharing scheme is in the α -core, but the converse need not hold. Therefore, cross-monotonicity is a strictly stronger requirement than being in the core.

The main application of cross-monotonic cost sharing schemes is in the design of cost sharing mechanisms, defined in the following setting: Each user i has a *willingness to pay* $u_i \in \mathbb{R}^+ \cup \{0\}$ for the service, i.e., she is willing to pay up to u_i dollars to get the service. We further assume that the utility (or happiness) of user i is given by $u_i q_i - x_i$, where q_i is an indicator variable which indicates whether she has received the service or not, and x_i is the amount she has to pay¹. A *cost sharing mechanism* (also known as a *social choice function*) is an algorithm that elicits a bid $b_i \in \mathbb{R}^+ \cup \{0\}$ from each agent, and based on these bids, decides which agents should receive the service and how much each of them has to pay. More formally, a cost sharing mechanism is a function that associates to each vector b of non-negative bids a set $Q(b) \subseteq \mathcal{A}$ of agents to be serviced, and a vector $x(b) \in \mathbb{R}^n$ of non-negative payments. When there is no ambiguity, we write Q and x instead of $Q(b)$

and $x(b)$, respectively. Throughout this paper, we assume that a mechanism does not charge an agent who does not receive the service (i.e., $x_i = 0$ for $i \notin Q$), does not charge an agent who receives the service more than her bid (i.e., $x_i \leq b_i$ for $i \in Q$), and for each agent i , there is some bid b_i^* such that if i bids b_i^* , she will get the service, no matter what others bid². Furthermore, we would like the mechanisms to be approximately budget balanced. We call a mechanism α -budget balanced if the total amount the mechanism charges the agents is between $\alpha C(Q)$ and $C(Q)$ (i.e., $\alpha C(Q) \leq \sum_{i \in Q} x_i \leq C(Q)$).

The main property that we want a mechanism to satisfy is incentive compatibility. We want our mechanism to encourage participants to submit their true willingness to pay as their bid. Agents should not be able to benefit from lying about the prices they are willing to pay. Ideally, not even a group of users should be able to benefit by cooperatively lying, thus discouraging complicated bidding strategies. More precisely, we look for mechanisms, called *group strategyproof mechanisms* which satisfy the following additional property. Let $S \subseteq \mathcal{A}$ be a coalition of users, and u, u' be two vectors of non-negative bids satisfying $u_i = u'_i$ for every $i \notin S$ (we think of u as the true willingness to pay of users, and u' as a vector of strategically chosen bids). Let (Q, x) and (Q', x') denote the outputs of the mechanism when the bids are u and u' , respectively. We say that the mechanism is *group strategyproof* if for every such S, u, u' , if the inequality $u_i q'_i - x'_i \geq u_i q_i - x_i$ holds for all $i \in S$, then it holds with equality for every $i \in S$. In other words, there should not be any coalition S and vector u' of bids such that if members of S announce u' instead of u (their true willingness to pay) as their bids, then every member of the coalition S is at least as happy as in the truthful scenario, and at least one person is happier.

3 Upper bounds for cross-monotonic cost sharing schemes

In this section we present the main idea behind our upper bound technique and prove several upper bounds for the games defined based on edge cover, vertex cover, and facility location. In Section 3.1 we explain the technique with a simple example of the edge cover game. Sections 3.2 and 3.3 contain the proofs of the upper bounds for the vertex cover and facility location games. Finally, in Section 3.4 we state (without proof) several other upper bounds that can be proved using our technique.

3.1 A simple example: the edge cover game In this section, we explain our technique using the edge cover game as a guiding example. The edge cover cost function is

¹As noted by Moulin and Shenker [14], this assumption is without loss of generality

²For a discussion about these properties see Moulin [15] and Moulin and Shenker [14].

defined as follows.

DEFINITION 3.1. *Let $G = (V, E)$ be a graph with no isolated vertices. The set of agents in the edge cover game on G is the set of vertices of G . Given a subset S of vertices, the cost of S is the minimum size of a set $F \subseteq E$ of edges such that for every $v \in S$, at least one of the edges incident to v is in F . Such a set F is called an edge cover for S .*

It is easy to see that for every set S , one can obtain a minimum edge cover of S by taking a maximum matching on S and adding one edge for every vertex that is not covered by the maximum matching (see [2]). Using this fact, we can give a cost-sharing scheme that is in the $\frac{2}{3}$ -core of the game: charge each vertex that is covered by the maximum matching $\frac{1}{3}$, and other vertices $\frac{2}{3}$. Since there is no edge between two vertices that are not covered by the maximum matching, this cost-sharing scheme satisfies the core property. Furthermore, it is easy to see that the sum of the cost shares is always equal to $\frac{2}{3}$ times the edge cover for S . Therefore, there is a cost-sharing scheme satisfying the core property with a budget-balance factor of $\frac{2}{3}$. In fact, Goemans [5] showed that for every graph there is a cost sharing scheme in the $\frac{3}{4}$ -core. However, in the following, we show that no cross-monotonic cost-sharing scheme can achieve a budget-balance factor better than $\frac{1}{2}$.

THEOREM 3.1. *For every $\epsilon > 0$, there is no $(\frac{1}{2} + \epsilon)$ -budget balanced cross-monotonic cost sharing scheme for the edge cover problem.*

Here's the high-level idea of the proof: We assume, for contradiction, that there is a cross-monotonic cost sharing scheme that always recovers at least a $(\frac{1}{2} + \epsilon)$ fraction of the total cost. We explicitly construct a graph G (or in general the set of agents \mathcal{A} and the structure based on which the cost function is defined), and look at the cost-sharing scheme on this graph. For edge cover, this graph is simply a complete bipartite graph $K_{n,n}$, with n large enough. Then, we need to argue that there is a set S of agents such that the total cost shares of the elements of S is less than $\frac{1}{2} + \epsilon$ times the size of the minimum edge-cover for S . This is done using the probabilistic method: we pick a subset S at random from a certain distribution and show that in expectation, the ratio of the recovered cost to the cost of S is low. Therefore, there is a manifestation of S for which this ratio is low. In the edge-cover example, we pick one vertex v of G uniformly at random and let S be the union of v and the set of vertices adjacent to v . We now need to bound the expected value of the sum of cost shares of the elements of S . We do this by using cross-monotonicity and bounding the cost share of each vertex $u \in S$ by the cost share of u in a substructure T_u of S . Bounding the expected cost share of u in T_u is done by showing that for every substructure T , every $u \in T$ has

the same probability of occurring in a structure S in which $T_u = T$. This implies that the expected cost share of u in T_u (where the expectation is over the choice of S) is at most the cost of T_u divided by the number of agents in T_u . Summing up these values for all u gives us the desired contradiction.

Proof of Theorem 3.1. Assume that there is a $(\frac{1}{2} + \epsilon)$ -budget-balanced cross-monotonic cost sharing scheme ξ . Let G be the complete bipartite graph $K_{n,n}$, where n will be fixed later, and consider ξ on G . For every $v \in V(G)$, we let S_v be the union of v and the set of vertices adjacent to v (i.e., vertices of the other part). We pick a set S of agents by picking v uniformly at random from $V(G)$ and letting $S = S_v$. By the definition of the edge cover problem,

$$(3.1) \quad C(S) = n \quad \text{for every } S.$$

On the other hand,

$$(3.2) \quad \begin{aligned} \mathbb{E}_S \left[\sum_{i \in S} \xi(i, S) \right] &= \mathbb{E}_v \left[\xi(v, S_v) \right] \\ &\quad + \mathbb{E}_v \left[\sum_{u \in S_v \setminus \{v\}} \xi(u, S_v) \right] \\ &\leq 1 + \mathbb{E}_v \left[\sum_{u \in S_v \setminus \{v\}} \xi(u, \{u, v\}) \right], \end{aligned}$$

where the last inequality follows from the facts that for every vertex u and every set S , $\xi(u, S) \leq 1$, and that for every $v \in V(G)$ and $u \in S_v \setminus \{v\}$, $\xi(u, S_v) \leq \xi(u, \{u, v\})$. Both of these facts are consequences of the cross-monotonicity of ξ . By the definition of expected values, we have

$$(3.3) \quad \mathbb{E}_v \left[\sum_{u \in S_v \setminus \{v\}} \xi(u, \{u, v\}) \right] = n \mathbb{E}_{v,u} \left[\xi(u, \{u, v\}) \right],$$

where the second expectation is over the choice of v from $V(G)$ and u in $S_v \setminus \{v\}$. However, choosing a vertex v and then a neighbor u of v at random is equivalent to choosing a random edge e in G at random, and letting u be a random endpoint of e and v be the other one. By the budget-balance condition, the sum of the cost shares of the endpoints of e is at most one. Therefore, for every e , if u is a random endpoint of e and v is the other endpoint, $\mathbb{E}[\xi(u, \{u, v\})] \leq \frac{1}{2}$. Thus, the right-hand side of Equation 3.3 is at most $\frac{n}{2}$. Therefore, by Equations 3.1 and 3.2, we have

$$\mathbb{E}_S \left[\frac{\sum_{i \in S} \xi(i, S)}{C(S)} \right] \leq \frac{1 + \frac{n}{2}}{n} < \frac{1}{2} + \epsilon$$

for $n > 1/\epsilon$. Therefore, there is a set S satisfying $\frac{\sum_{i \in S} \xi(i, S)}{C(S)} < \frac{1}{2} + \epsilon$, which is a contradiction with the assumption that ξ is $(\frac{1}{2} + \epsilon)$ -budget balanced. \square

It is not difficult to see that the cost-sharing scheme ξ satisfying $\xi(i, S) = \frac{1}{2}$ for every $i \in S$ is cross-monotonic

and $\frac{1}{2}$ -budget balanced. Therefore, the bound given in the above theorem is tight. Also, one can think of the edge-cover problem as a special case of the set cover problem in which the size of each set is 2. It is not difficult to generalize the above result to the special case of set cover in which the size of each set is k , and prove that for k constant, no cross-monotonic cost-sharing scheme for this problem can recover more than a $\frac{1}{k}$ fraction of the cost. Using similar argument the next theorem shows that for the general case of the set cover game, no cross-monotonic cost-sharing scheme can recover more than a $O(\frac{1}{n})$ of the total cost.

THEOREM 3.2. *There is no cross-monotonic cost-sharing scheme ξ for the set cover game such that for every set $S \subseteq \mathcal{A}$, ξ recovers more than a $O(\frac{1}{|S|})$ fraction of the cost of S .*

Proof. Assume that there is such a cross-monotonic cost sharing scheme ξ . Consider the following instance of the set cover problem. Let \mathcal{A} be a set of n^2 agents that can be partitioned as $\mathcal{A} = A_1 + A_2 + \dots + A_n$, where A_i 's are disjoint sets each of size n . Define \mathcal{C} as the collection of all sets $S \subset \mathcal{A}$ such that $|S \cap A_i| = 1$ for every $i = 1, \dots, n$. An alternative way to look at this is that \mathcal{A} and \mathcal{C} are sets of vertices and edges of an n -uniform n -partite complete hypergraph.

We pick a random set S of agents in the above instance as follows: Pick a random i from $\{1, \dots, n\}$, and for every $j \neq i$, pick an agent a_j uniformly at random from A_j . Let $T = \{a_j : j \neq i\}$ and $S = A_i \cup T$. The cost of the optimal set cover solution on S is always at least n , since no set in \mathcal{C} contains two distinct elements of A_i , and therefore each element of A_i must be covered with a distinct set in \mathcal{C} .

We now bound the average recovered cost over the random choice of S .

$$\begin{aligned}
\mathbb{E}_S \left[\sum_{x \in S} \xi(x, S) \right] &= \mathbb{E} \left[\sum_{x \in A_i} \xi(x, S) \right] \\
&\quad + \mathbb{E} \left[\sum_{j \neq i} \xi(a_j, S) \right] \\
&\leq \mathbb{E} \left[\sum_{x \in A_i} \xi(x, \{x\} \cup T) \right] \\
(3.4) \quad &\quad + \mathbb{E} \left[\sum_{j \neq i} \xi(a_j, T) \right]
\end{aligned}$$

Since all elements of T can be covered by one set, the second term in the above expression is at most 1. We write the first term as $n\mathbb{E}_{S,x} [\xi(x, \{x\} \cup T)]$ where the expectation is over the random choice of S and the random choice of x from A_i . As in the proof of Theorem 3.5, the expected value of $\xi(x, \{x\} \cup T)$ in this experiment is equal to the expected value of $\frac{1}{n} \sum_{j=1}^n \xi(a_j, \{a_1, \dots, a_n\})$ in an experiment that

consists of choosing an agent a_j from each A_j uniformly at random. By the budget-balance property, we always have $\sum_{j=1}^n \xi(a_j, \{a_1, \dots, a_n\}) \leq C(\{a_1, \dots, a_n\}) = 1$. Therefore, the first term in the left-hand side of the inequality (3.4) is at most one. This means that the expected total cost share recovered from the set S is at most two. Therefore, the ratio of recovered cost to total cost of S is at most $2/n < 4/|S|$. \square

3.2 The vertex cover game The vertex cover game is defined on a graph $G = (V, E)$. The set of agents is the set of edges of G , and the cost of serving a set $S \subseteq E$ is equal to the minimum size of a set A of vertices such that for each $e \in S$, at least one of the endpoints of e is in A . Such a set is called a *vertex cover* for the set S . It is well-known that the integrality gap of the LP relaxation of vertex cover is 2, and therefore no allocation in core can recover more than half the cost of the solution in the worst case. We show in the following theorem that if we require the cost-sharing scheme to be cross-monotonic, then no constant-factor budget balanced scheme exists.

THEOREM 3.3. *For every $\epsilon > 0$, there is no cross-monotonic cost sharing scheme for vertex cover that on every set S of n agents, recovers at least a $(2 + \epsilon)n^{-1/3}$ fraction of the cost of S .*

Proof. Assume, for contradiction, that such a scheme ξ exists. We let G be a complete graph on $m + 2\ell$ vertices, where m and ℓ ($m < \ell$) are numbers that will be fixed later, and consider the cost-sharing scheme ξ on G . We show that there is some set S of edges of G for which ξ recovers at most a $|S|^{-1/3}$ fraction of the cost. We do this by picking S randomly from a distribution described below, and showing that the above statement holds in expectation, and therefore there should be a particular S satisfying the above statement.

Let π be a permutation of the $m + 2\ell$ vertices. Let A be the set of the first m vertices, B be the set of the next ℓ vertices, and C be the set of the remaining ℓ vertices. We denote the i 'th vertices of B and C (based on the ordering given by π) by b_i and c_i . Let S_π denote the set of all $m\ell$ edges between A and B , union the set of edges $b_i c_i$ for $i = 1, \dots, \ell$. We pick S by picking the permutation π uniformly at random and letting $S = S_\pi$. See Figure 1 for an example.

If we denote the set of edges between A and B by T , we have

$$(3.5) \quad \mathbb{E} \left[\sum_{e \in T} \xi(e, S) \right] \leq \mathbb{E} \left[\sum_{e \in T} \xi(e, T) \right] \leq m,$$

where the first inequality follows from the cross-monotonicity of ξ and the second inequality is implied by the budget balance assumption and the fact that the cost of the minimum vertex cover in T is m . We also let T_i be




Figure 1: The structure S in the vertex cover game

the set of all $m + 1$ edges in S that have b_i as an endpoint (see Figure 1). Equation 3.5 and the cross-monotonicity of ξ imply the following.

$$\begin{aligned}
 E_S \left[\sum_{i \in S} \xi(i, S) \right] &= E \left[\sum_{e \in T} \xi(e, S) \right] \\
 &\quad + \sum_{i=1}^{\ell} E \left[\xi(b_i c_i, S) \right] \\
 (3.6) \quad &\leq m + \sum_{i=1}^{\ell} E \left[\xi(b_i c_i, T_i) \right],
 \end{aligned}$$

We now need to analyze the expectation of $\xi(b_i c_i, T_i)$ over the random choice of π . Notice that the only elements of π that are important in $\xi(b_i c_i, T_i)$ are the first m elements and the $m + i$ 'th and $m + \ell + i$ 'th elements (b_i and c_i). Therefore, the expectation of $\xi(b_i c_i, T_i)$ over the choice of π is equal to the expectation of $\xi(v_{m+2} v_{m+1}, \{v_1 v_{m+1}, v_2 v_{m+1}, \dots, v_m v_{m+1}, v_{m+2} v_{m+1}\})$ over the random choice of an ordered list v_1, v_2, \dots, v_{m+2} of $m + 2$ different vertices of G . However, in this experiment it is clear by symmetry that the expected cost share of $v_i v_{m+1}$ is the same for $i = 1, \dots, m, m + 2$, and therefore by the budget balance condition each of these expected cost shares is at most $\frac{1}{m+1}$. This, together with Equation 3.6 imply the following.

$$(3.7) \quad E_S \left[\sum_{i \in S} \xi(i, S) \right] \leq m + \frac{\ell}{m+1}.$$

On the other hand, the size of the minimum vertex cover in S is always ℓ . Therefore, the expected value of the ratio of $\sum_{i \in S} \xi(i, S)$ to $C(S)$ is at most $\frac{m}{\ell} + \frac{1}{m+1}$. Thus, there is a set S for which this ratio is at most $\frac{m}{\ell} + \frac{1}{m+1}$. Taking $m = \sqrt{\ell}$, we see that the allocation on S recovers at most a $\frac{2}{\sqrt{\ell}} < (2 + \epsilon)|S|^{-1/3}$ fraction of the cost. \square

We can show the following positive result for cross-monotonic cost sharing schemes for the vertex cover. We do not know the right bound for the budget-balance factor of the vertex cover game. The proof of the following theorem is omitted here.

THEOREM 3.4. *For the vertex cover game, the cost sharing scheme that charges the edge uv in the set S an amount equal to $\min(1/\deg_S(u), 1/\deg_S(v))$ is cross-monotonic and $\frac{1}{2\sqrt{n}}$ -budget balanced.*

Proof. It is clear that this scheme is cross-monotone. We only need to verify the budget-balance factor. Consider a set S of n agents (i.e., edges), and the graph $G[S]$ induced on this set of edges. We prove that the total cost share of the agents in S is at least $\frac{1}{2\sqrt{n}}$ times the cost of a vertex cover for S .

Divide the set of vertices into two subsets L and H , where L is the set vertices of degree less than \sqrt{n} in $G[S]$ and H is the rest of vertices ($H = V(G) - L$). As a vertex cover solution, select H and both endpoints of all edges (u, v) such that $u, v \in L$. We show that the cost shares of the edges in S sum to at least a $\frac{1}{2\sqrt{n}}$ fraction of the cost of this solution. First consider any edge e between vertices in L . The cost share of e is at least $\frac{1}{\sqrt{n}}$, thus its cost share covers the cost of picking both its endpoints. Now consider the vertices in H . Since the degree of each vertex $v \in H$ is greater than or equal to \sqrt{n} , the sum of the cost shares of the edges adjacent to v is at least $\frac{1}{n} \sqrt{n} = \frac{1}{\sqrt{n}}$. Each edge is included in at most two such summations, and thus the sum of the cost shares of edges adjacent to vertices in H is at least a $\frac{1}{2\sqrt{n}}$ fraction of the cost of H . Therefore, the sum of the cost shares of the agents in S is at least $\frac{1}{2\sqrt{n}}$ times the cost of the optimal vertex cover for S . \square

3.3 The metric facility location game Given a set of cities, facilities with opening costs, and metric connection costs between cities and facilities, the facility location problem seeks to open a subset of facilities and connect each city to a facility in a manner that minimizes the total cost. In the facility location game, each city is an agent. The cost of a subset of agents is the cost of the minimum facility location solution for that subset; a cross-monotonic cost-sharing scheme tries to share this cost among the agents. In this section, we prove that any cross-monotonic cost-sharing scheme for facility location is at best $\frac{1}{3}$ -budget-balanced. This matches the budget-balance factor of the scheme given by Pál and Tardos [16].

We start by giving an example on which the scheme of Pál and Tardos [16] recovers only a third of the cost³. This example will be used as the randomly chosen structure in our proof. The proof of the following lemma is simple and is omitted here.

LEMMA 3.1. *Let \mathcal{I} be an instance of the facility location problem consisting of $m+k$ cities $c_1, \dots, c_m, c'_1, \dots, c'_k$ and*

³This example also shows that the dual computed by the Jain-Vazirani facility location algorithm [9] can be a factor 3 away from the optimal dual.




Figure 2: Upper bound for facility location game

m facilities f_1, \dots, f_m each of opening cost 3. For every i and j , the connection costs between f_i and c_i and between f_i and c'_j are all 1, and other connection costs are obtained by the triangle inequality. See Figure 2(a). Then if $m = \omega(k)$ and k tends to infinity, the optimal solution for \mathcal{I} has cost $3m + o(m)$.

THEOREM 3.5. Any cross-monotonic cost-sharing scheme for the facility location game is at most $1/3$ -budget balanced.

Proof. Consider the following instance of the facility location problem. There are k sets A_1, \dots, A_k of m cities each, where $m = \omega(k)$ and $k = \omega(1)$. For every subset B of cities containing exactly one city from each A_i ($|B \cap A_i| = 1$ for all i), there is a facility f_B with connection cost 1 to each city in B . The remaining connection costs are defined by extending the metric, i.e., the cost of connecting city i to facility f_B for $i \notin B$ is 3. The facility opening costs are all 3.

We pick a random set S of cities in the above instance as follows: Pick a random i from $\{1, \dots, k\}$, and for every $j \neq i$, pick a city a_j uniformly at random from A_j . Let $T = \{a_j : j \neq i\}$ and $S = A_i \cup T$. See Figure 2(b) for an example. It is easy to see that the set S induces an instance of the facility location problem almost identical to the instance \mathcal{I} in Lemma 3.1 (the only difference is that here we have more facilities, but it is easy to see that the only relevant facilities are the ones that are present in \mathcal{I}). Therefore, the cost of the optimal solution on S is $3m + o(m)$.

We show that for any cross-monotonic cost-sharing scheme ξ , the average recovered cost over the choice of S is at most $m + o(m)$ and thus conclude that there is some S whose recovered cost is $m + o(m)$. As in the previous proofs, we start bounding the expected total cost share by using the linearity of expectations and cross-monotonicity:

$$\begin{aligned} \mathbb{E}_S \left[\sum_{c \in S} \xi(c, S) \right] &= \mathbb{E} \left[\sum_{c \in A_i} \xi(c, S) \right] \\ &\quad + \mathbb{E} \left[\sum_{j \neq i} \xi(a_j, S) \right] \end{aligned}$$

$$\begin{aligned} &\leq \mathbb{E} \left[\sum_{c \in A_i} \xi(c, \{c\} \cup T) \right] \\ &\quad + \mathbb{E} \left[\sum_{j \neq i} \xi(a_j, T) \right] \end{aligned}$$

Notice the set T has a facility location solution of cost $3+k-1$ and thus by the budget balance condition the second term in the above expression is at most $k+2$. The first term in the above expression can be written as $m \mathbb{E}_{S,c} [\xi(c, \{c\} \cup T)]$ where the expectation is over the random choice of S and the random choice of c from A_i . However, it can be seen easily that this is equivalent to the following random experiment: From each A_j , pick a city a_j uniformly at random. Then pick i from $\{1, \dots, k\}$ uniformly at random and let $c = a_i$ and $T = \{a_j : j \neq i\}$. From this description it is clear that the expected value of $\xi(c, \{c\} \cup T)$ is equal to $\frac{1}{k} \sum_{j=1}^k \xi(a_j, \{a_1, \dots, a_k\})$. This, by the budget balance property and the fact that $\{a_1, \dots, a_k\}$ has a solution of cost $k+3$ cannot be more than $\frac{k+3}{k}$. Therefore,

$$(3.8) \quad \mathbb{E}_S \left[\sum_{c \in S} \xi(c, S) \right] \leq m \left(\frac{k+3}{k} \right) + (k+2) = m + o(m),$$

when $m = \omega(k)$ and $k = \omega(1)$. Therefore, the expected value of the ratio of recovered cost to total cost tends to $1/3$. \square

3.4 Other combinatorial optimization problems In this section we state upper bounds for three other combinatorial optimization games (in particular, the ones considered by Deng et al. [2]). These problems are maximization problems, therefore instead of cost sharing, we need to design *profit sharing* schemes. Definitions of profit sharing schemes and their properties are similar to the ones for cost sharing schemes (usually with the direction of inequalities reversed).

The first example is the maximum flow game. In the maximum flow game, we are given a directed graph $G = (V, E)$ with a source s and a sink t . Agents are directed edges

of G . Given a subset of edges, S , the profit of S is the size of maximum flow from s to t on subgraph of G induced on the edges of S . It is known that the core of maximum flow game is nonempty [2]. The story is different for cross-monotonic profit sharing schemes.

THEOREM 3.6. *There is no $o(n)$ -budget balanced profit sharing scheme for the maximum flow game where n is the number of agents.*

Proof. Let G be a graph consisting of three nodes: s , u , and t . There are $n - 1$ edges from s to u , and $n - 1$ edges from u to t . Let E_{su} and E_{ut} denote the set of edges from s to u and from u to t , respectively. See Figure 3. We pick a random set S of n agents as follows: With probability $1/2$, pick a random edge e from s to u , and let $S = \{e\} \cup E_{ut}$. With probability $1/2$, pick a random edge e from u to t , and let $S = \{e\} \cup E_{su}$. For example the set S could contain the thick edges in Figure 3.




Figure 3: The graph G for the maximum flow game

Assume ξ is an $o(n)$ -budget balanced cross-monotonic profit-sharing scheme for G . We have that $E \equiv \mathbb{E}_S[\sum_{a \in S} \xi(a, S)]$ is

$$\begin{aligned}
E &\geq \frac{1}{2} \mathbb{E}_{e \leftarrow E_{su}} \left[\sum_{a \in E_{ut}} \xi(a, \{e\} \cup E_{ut}) \right] \\
&\quad + \frac{1}{2} \mathbb{E}_{e \leftarrow E_{ut}} \left[\sum_{a \in E_{su}} \xi(a, \{e\} \cup E_{su}) \right] \\
&\geq \frac{1}{2} \mathbb{E}_{e \leftarrow E_{su}} \left[\sum_{a \in E_{ut}} \xi(a, \{a, e\}) \right] \\
&\quad + \frac{1}{2} \mathbb{E}_{e \leftarrow E_{ut}} \left[\sum_{a \in E_{su}} \xi(a, \{a, e\}) \right] \\
&= \frac{n-1}{2} \mathbb{E}_{a \leftarrow E_{su}, b \leftarrow E_{ut}} \left[\xi(a, \{a, b\}) + \xi(b, \{a, b\}) \right] \\
&\geq \frac{n-1}{2}.
\end{aligned}$$

On the the hand, the profit of every set S picked using the above procedure is one. Therefore, the expected ratio of the total profit shares to the profit of S is at least $(n - 1)/2$. \square

The second problem is the problem of packing the maximum number of arborescences in a digraph. An r -arborescence is a spanning tree rooted at r in which all edges are directed away from r . In the maximum r -arborescence game, we are given a directed graph $G = (V, E)$ with a root r . Agents are directed edges of G . Given a subset of edges, S , the value of S is the maximum number of edge-disjoint r -arborescences on the subgraph induced by S . One can think of the profit of S as the maximum bandwidth for broadcasting messages from r to all vertices of the graph. It is known that the core of this game is nonempty [2].

THEOREM 3.7. *There is no $o(n)$ -budget balanced profit sharing scheme for maximum r -arborescence game.*

The same construction as the one used in the above proof gives us a proof for Theorem 3.7.

Finally, we consider the maximum matching game, in which the agents are vertices of a graph G , and the profit of a subset of vertices S is the size of maximum matching in $G[S]$. One can show that there is a 2-budget balance profit allocation in the core of this game.

THEOREM 3.8. *There is no $o(n)$ -budget balanced profit sharing scheme for the maximum matching game.*

Proof. We use the same construction that was used in the proof of Theorem 3.1. Let G be a complete bipartite graph with $n - 1$ vertices in each part (here we use $n - 1$ instead of n so that the size of S becomes n), and pick S by picking a random vertex in G and all vertices in the other part. Using an argument essentially the same as the one used in the proof of Theorem 3.1, the expected total profit share of the elements of S is at least $(n - 1)/2$. On the other hand, the profit of S is always one. Thus, there is an S on which the ratio between the total profit-share and the profit of S is at least $(n - 1)/2$. \square

4 Group-strategyproof mechanisms

A main motivation behind cross-monotonic cost-sharing schemes is that they can be used to define group-strategyproof mechanisms [14]. In the previous section, we proved that for certain games every cross-monotonic cost sharing scheme is poorly budget balanced. A natural question to ask is whether all group-strategyproof mechanisms for these games are so poorly budget balanced. Towards this aim, one might hope to show that every group-strategyproof mechanism corresponds to a cross-monotonic cost sharing scheme. In fact, given a group-strategyproof mechanism \mathcal{M} , it is possible to define a corresponding cost-sharing scheme $\xi_{\mathcal{M}}$ as follows: Consider the scenario where the agents in a set S bid a sufficiently large value and others zero, and let $\xi_{\mathcal{M}}(i, S)$ be the payment charged by the mechanism to the agent i in this scenario. Unfortunately, this scheme is not

necessarily budget-balanced or cross-monotonic. In fact, the following simple examples show that for every cost function, there is a group-strategyproof mechanism recovering all the cost.

EXAMPLE 4.1. *Single Payment Mechanism: Arbitrarily order the agents from 1 to n . Then, find the first agent i in this order whose bid is at least $C(\{i, \dots, n\})$. The set that will receive the service is $Q = \{i, \dots, n\}$, and the total cost of servicing this set is paid by the agent i . Other agents pay nothing.*

A less unfair mechanism that still recovers all the cost and works for every subadditive function is the following.

EXAMPLE 4.2. *Arbitrarily order the agents from 1 to n . Initialize the set of serviced agents S to the empty set and the amount of money M that is already charged to the agents to 0. For i from 1 to n , do the following: if the bid of i is at least $\min(C(\{i\}), C(S \cup \{i, \dots, n\}) - M)$, then include i in S , and charge her $\min(C(\{i\}), C(S \cup \{i, \dots, n\}) - M)$ (therefore, M will be increased by this quantity).*

Intuitively, the mechanisms in Examples 4.1 and 4.2 are neither fair nor efficient. They always place the burden of the entire service cost on a small subset of agents while servicing others for free. Agents who receive the service for free, called *free riders*, increase the cost of the solution but do not contribute any payment. We consider constraining our mechanism to rule out free riders (This is also known as the *no free lunch* property). With this added constraint, we can prove that there is no budget-balanced group strategyproof mechanism for vertex cover. statement is

THEOREM 4.1. *There is no budget-balanced group strategyproof mechanism without free riders for the vertex cover game.*

Proof of Theorem 4.1 is omitted here. Although there is no budget-balanced group strategyproof mechanism for the vertex cover game, the following theorem shows that for every $\epsilon > 0$, there are $(1 - \epsilon)$ -budget balance group strategyproof mechanisms without free riders for any non-decreasing cost function.

THEOREM 4.2. *For any non-decreasing cost function C and $\epsilon > 0$, there is a $(1 - \epsilon)$ -budget-balanced group-strategyproof mechanism for C with no free-riders.*

Proof Sketch. The idea is to charge every agent a small participation fee δ that depends on ϵ . Let T be the subset of bidders who bid *strictly* greater than δ . Other bidders will not be considered by the mechanism. Now, we run a variant of the single payment mechanism for bidders in T : Arbitrarily order the agents (indexed by numbers $1, \dots, |T|$),

and find the first agent i in this order whose bid is at least $C(\{i, \dots, |T|\}) - n\delta$. The set $Q = \{i, \dots, |T|\}$ receives the service, i pays $C(Q) - n\delta$, and other agents in Q pays δ . \square

In the mechanism given in the above proof, there are scenarios in which an agent does not receive the service, but would receive service if she increased her bid by any positive amount, however small. In fact, we can show that a cross-monotonic cost sharing scheme can be derived from any mechanism with no free riders for which such situations do not occur. More precisely, we call a mechanism \mathcal{M} *upper continuous* if for every player i , if i gets the service for every bid value greater than x holding other bids fixed, then i gets the service if she bids x . Upper continuity by itself is not difficult to satisfy. In fact, both mechanisms in Examples 4.1 and 4.2 are upper continuous. However, the following theorem shows that mechanisms satisfying upper continuity and no-free-rider conditions are as hard as cross-monotonic cost sharing schemes.

THEOREM 4.3. *The cost function C has an upper-continuous α -budget-balanced group-strategyproof mechanism with no free riders if and only if it has an α -budget-balanced cross-monotonic cost-sharing scheme.*

In the proof of Theorem 4.3 we do not use coalitions of size greater than 2. Thus, with the assumptions of no-free-rider and upper-continuity, coalitions of size 2 are as strong as coalitions of arbitrary size. We do not know if this equivalence holds without these assumptions. Theorem 4.2 shows that the no-free-rider property is not strong enough to guarantee that the mechanism is fair, since there are almost budget-balanced group-strategyproof mechanisms satisfying this property in which a few users pay a large portion of the total cost. A stronger fairness property is the subsidy-freeness property considered by Moulin⁴ [13]. This condition says that the total charge of the mechanism to all users in a set S is at most the cost of the set $c(S)$. In the following theorem, we show that group-strategyproof mechanisms with this property are equivalent to cross-monotonic cost sharing schemes.

THEOREM 4.4. *There exists a budget-balanced group-strategyproof mechanism satisfying the subsidy-freeness property for a cost function C if and only if this cost function has a budget-balanced cross-monotonic cost-sharing scheme.*

The proof of this theorem is omitted here. We do not know if this theorem holds for budget-balance factors other than 1.

⁴See Devanur et al. [3] for a discussion of strategyproof (but not group-strategyproof) mechanisms satisfying subsidy-freeness for set cover and facility location problems.

5 Conclusion

In this paper, we studied upper bounds for the budget-balance factor of cross-monotonic cost-sharing schemes for a variety of combinatorial optimization games. Our techniques are quite general and may prove applicable to variety of other interesting games. For example, the facility location game restricted to a tree always has a budget-balanced cost sharing in the core [6], but we do not have a tight lower and upper bound for the budget-balance factor of a cross-monotonic cost sharing scheme. For the facility location on the line, we have the upper bound of $\frac{6}{7}$. A more challenging open question is that of cross-monotonic cost-sharing schemes for the Steiner tree game. There are 1/2-budget-balanced scheme for the Steiner tree and Steiner forest [11, 12], but the best upper bound we know for both problems is 8/9 (based on the core property).

A main motivation behind cross-monotonic cost-sharing schemes is the development of group-strategyproof mechanisms. As mentioned in this paper, almost any cost function (including all those for which we derived upper bounds on cost sharing schemes) has a trivial group-strategyproof mechanism. Several different sets of axioms, some of which explored in this paper, can be added to the mechanisms to rule out these trivial ones. An interesting open question in this area is whether α -budget-balanced group-strategy mechanisms with subsidy-freeness are equivalent to α -budget-balanced cross-monotone cost-sharing schemes. Theorem 4.4 shows that for $\alpha = 1$ the answer to this question is positive.

It is a standard economic result that a strategyproof mechanism can not be both efficient (i.e., return a solution that maximizes social welfare) and budget-balanced [7, 17]. It would be interesting to explore the possible budget-balance factor of group-strategyproof mechanisms that are in some sense close to efficient.

Acknowledgments. We would like to thank Michel Goemans for helpful discussions. Also, we would like to thank Martin Pál for introducing the problem and for helpful discussions.

References

- [1] O.N. Bondareva. Some applications of linear programming to cooperative games. *Problemy Kibernetiki*, 1963.
- [2] X. Deng, T. Ibaraki, and H. Nagamochi. Algorithms and complexity in combinatorial optimization games. In *SODA*, 1997.
- [3] N. Devanur, M. Mihail, and V. Vazirani. Strategyproof cost sharing mechanisms for set cover and facility location problems. *Proceedings of ACM Conference on Electronic Commerce*, 2003.
- [4] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmission. *Journal of Computer and System Sciences*, 63:21–41, 2001.
- [5] M. Goemans. Personal communication.
- [6] M.X. Goemans and M. Skutella. Cooperative facility location games. *SODA*, 2000.
- [7] J. Green, E. Kohlberg, and J. J. Laffont. Partial equilibrium approach to the free rider problem. *Journal of Public Economics*, 6:375–394, 1976.
- [8] K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. *Proceedings STOC*, 2001.
- [9] K. Jain and V.V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM*, 48:274–296, 2001.
- [10] K. Jain and V.V. Vazirani. Equitable cost allocations via primal-dual-type algorithms. *Proceedings of STOC*, 2002.
- [11] K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on MST’s. In *Operational Research Proceedings KOI*, pages 43–48, 1996.
- [12] J. Koenemann, S. Leonardi, and G. Schaefer. A group-strategyproof mechanism for steiner forests. In *Proceedings of SODA*, 2005.
- [13] H. Moulin. *Axioms of cooperative decision making*, volume 1, chapter Cost sharing games and the core. Cambridge University Press, 1988.
- [14] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *to appear in Economic Theory*, 1997.
- [15] Hervé Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16:279–320, 1999.
- [16] M. Pál and E. Tardos. Strategy proof mechanisms via primal-dual algorithms. *FOCS’03*, 2003.
- [17] K. Roberts. The characterization of implementable choice rules. In J.J. Laffont, editor, *Aggregation and Revelation of Preferences*. Studies in Public Economics, Amsterdam, North Holland, 1979.
- [18] Lloyd S.Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.