Is Shapley Cost Sharing Optimal?

Shahar Dobzinski[1⋆], Aranyak Mehta[2], Tim Roughgarden[3⋆⋆], and Mukund Sundararajan[3⋆⋆⋆]

[1] The School of Computer Science and Engineering, the Hebrew University of Jerusalem
[2] Google, Inc., Mountain View, CA
[3] Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305

Abstract. We study the best guarantees of efficiency approximation achievable by cost-sharing mechanisms. Our main result is the first quantitative lower bound that applies to all truthful cost-sharing mechanisms, including randomized mechanisms that are only truthful in expectation, and only β-budget-balanced in expectation. Our lower bound is optimal up to constant factors and applies even to the simple and central special case of the public excludable good problem. We also give a stronger lower bound for a subclass of deterministic cost-sharing mechanisms, which is driven by a new characterization of the Shapley value mechanism. Finally, we show a separation between the best-possible efficiency guarantees achievable by deterministic and randomized cost-sharing mechanisms.

1 Introduction

1.1 Approximation in Algorithmic Mechanism Design

Algorithmic mechanism design studies the possibilities and impossibilities of optimization with incomplete information by incentive-compatible mechanisms. The main positive result in the area is, of course, the *VCG mechanisms* [19, 3, 8], a family of truthful, direct-revelation mechanisms that maximize objective functions of the form

$$\max_{o \in \Omega} \sum_i w_i v_i(o) - C(o), \tag{1}$$

where Ω is the outcome space, v_i is a valuation private to a self-interested player i, and the w_i's and $C(o)$'s are known real-valued constants. In other

⋆ This work was done while the author was visiting Stanford University. Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, and by grants from the Israel Science Foundation and the USA-Israel Bi-national Science Foundation. Email: `shahard@cs.huji.ac.il`.

⋆⋆ Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship. Email: `tim@cs.stanford.edu`.

⋆⋆⋆ Supported by NSF Award CCF-0448664 and a Stanford Graduate Fellowship. Email: `mukunds@cs.stanford.edu`.

words, affine maximization of private data is always possible by compensating the self-interested participants appropriately.

For many central applications, VCG mechanisms are irrelevant or infeasible, and research has focused on the design and analysis of truthful approximation mechanisms (see e.g. [10]). For example, for some optimization problems different from affine maximization, *no* truthful mechanism can achieve full optimality, even with unbounded computational power (e.g. [16, 17]). Another common reason for designing truthful approximation mechanisms is the exponential communication and/or computation required by the VCG mechanism for some affine maximization problems, such as welfare maximization in combinatorial auctions (see e.g. [1]). This paper is motivated by a different flaw with truthful welfare-maximizing mechanisms: no such mechanism achieves non-trivial worst-case revenue guarantees, even if unbounded computation is allowed. Precisely, when the outcome-dependent constant $C(o)$ in (1) represents the production costs for outcome o, then no truthful and individually rational mechanism that maximizes the welfare $\sum_i v_i(o) - C(o)$ guarantees that the revenue obtained is at least a constant fraction of the incurred cost. This impossibility result applies even to extremely simple single-parameter settings [6, 7, 17]. An important research goal, to which this paper contributes, is to quantify the minimum efficiency loss required to recover non-trivial budget-balance guarantees.

1.2 Randomization in Algorithmic Mechanism Design

A related issue is quantifying the power of randomization in the design of truthful approximation mechanisms. Recall that a randomized mechanism is *truthful in expectation* if truthful revelation is a dominant strategy for a player that wants to maximize its expected payoff, and is *universally truthful* if it is a distribution over truthful deterministic mechanisms. (The second condition effectively assumes that players can predict the outcome of the mechanism's internal randomization and therefore is stronger than the first.) For non-affine problems, universally truthful mechanisms are provably more powerful than deterministic ones [16]. We show, for the first time, an analogous separation between the best-possible performance of deterministic and randomized revenue-constrained mechanisms.

1.3 Our Results

Our main result is the first quantitative lower bound on efficiency loss that applies to all truthful and budget-balanced mechanisms. Our lower bound applies even in the special case of a single-parameter *public excludable good* problem, where the outcome set Ω is the subsets of the participants (the "winners") and the cost $C(o)$ is zero for the empty set and 1 otherwise. The public excludable good problem occupies a central position in the economic cost-sharing literature [5, 4]. It is also a special case of nearly all of the cost-sharing problems that have been studied in the theoretical computer science literature, including fixed-tree multicast, uncapacitated facility location, and vertex cover cost-sharing problems (see [2]). Naturally, our lower bound carries over to all of these more general classes of cost-sharing problems. Previous lower bounds for approximate efficiency in cost-sharing mechanisms applied only to subclasses of deterministic mechanisms (to Moulin mechanisms in [18] and to acyclic mechanisms in [11]).

Precisely, we prove the following. Call a truthful and individually rational mechanism for a public excludable good problem *β-budget-balanced* if its revenue is always at least a $1/\beta$ fraction of and no more than the incurred cost. We show that *every β-budget-balanced truthful mechanism is $\Omega(\log k/\beta)$-approximate* in the sense of [18], where k is the number of participants. Our lower bound applies even to randomized mechanisms that are only truthful in expectation, and only β-budget-balanced in expectation. Our lower bound is optimal up to constant factors for all $\beta = O(\sqrt{\log k})$, with the nearly matching upper bound provided by a scaled version of the Shapley value mechanism [15, 18]. All of our lower bounds apply to both the social cost approximation measure introduced in [18] and to the additive efficiency loss measure studied earlier by Moulin and Shenker [15].

We also give stronger results for a subclass of deterministic cost-sharing mechanisms. Specifically, we show that the *Shapley value mechanism is optimal* among all deterministic, symmetric, and budget-balanced cost-sharing mechanisms for public excludable good problems. (A similar result of Moulin and Shenker [15] proves only that the Shapley value mechanism is an optimal Moulin mechanism [14].) Here, "symmetric" means that players that submit equal bids are given the same allocations and prices. This proof is based on a new characterization of the Shapley value mechanism, which improves upon a previous characterization of Deb and Razzolini [5].

Finally, we give the first separation between the power of deterministic and randomized cost-sharing mechanisms: we prove a lower bound on the approximation factor of all deterministic mechanisms for the 2-player public excludable good problem, and exhibit a universally truthful randomized mechanism that possesses a strictly better approximation guarantee.

2 Preliminaries

There is a population U of k players and a public cost function C defined on all subsets of U. We always assume that $C(\emptyset) = 0$ and that C is nondecreasing (i.e., $S \subseteq T$ implies that $C(S) \leq C(T)$). Player i has a private value v_i for service. We focus on direct revelation mechanisms; such mechanisms accept a bid b_i from each player i and determine an allocation $S \subseteq U$ and payments p_i for the players.

We discuss only mechanisms that satisfy the following standard assumptions: *individual rationality*, meaning that $p_i = 0$ if $i \notin S$ and $p_i \leq b_i$ if $i \in S$; and *no positive transfers*, meaning that prices are always nonnegative. We also assume that players have quasilinear utilities, meaning that each player i aims to maximize $u_i(S, p_i) = v_i x_i - p_i$, where $x_i = 1$ if $i \in S$ and $x_i = 0$ if $i \notin S$.

A mechanism is *strategyproof*, or *truthful*, if no player can ever strictly increase its utility by misreporting its valuation. Formally, truthfulness means that for every player i, every bid vector b with $b_i = v_i$, and every bid vector b' with $b_j = b'_j$ for all $j \neq i$, $u_i(S, p_i) \geq u_i(S', p'_i)$, where (S, p) and (S', p') denote the outputs of the mechanism for the bid vectors b and b', respectively. When discussing truthful mechanisms, we typically assume that players bid their valuations and conflate the (unknown) valuation profile v with the (known) bid vector b.

In Section 4 we use the following standard fact about truthful mechanisms (see e.g. [12]).

Proposition 1. *Let M be a truthful, individually rational cost-sharing mechanism with the player set U. Then for every $i \in U$ and bid vector b_{-i} for players other than i, there is a* threshold *$t_i(b_{-i})$ such that: (i) if i bids more than $t_i(b_{-i})$, then it receives service at price $t_i(b_{-i})$; (ii) if i bids less than $t_i(b_{-i})$, then it does not receive service.*

A randomized mechanism is, by definition, a probability distribution over deterministic mechanisms. Such a mechanism is *universally truthful* if every mechanism in its support is truthful. Such a mechanism is *truthful in expectation* if no player can ever strictly increase its *expected* utility by misreporting its valuation. Every universally truthful mechanism is truthful in expectation, but the converse need not hold.

We study two kinds of objectives for cost-sharing mechanisms, one for the revenue of the mechanism, and one for its economic efficiency. First, for a parameter $\beta \geq 1$, a mechanism is *β-budget-balanced* if it always recovers at least a $1/\beta$ fraction of and at most the cost incurred. We say that a mechanism is *budget-balanced* if it is 1-budget-balanced.

We measure the efficiency (loss) achieved by a cost-sharing mechanism via the *social cost* objective. The social cost of an outcome S with respect to a cost function C and valuation profile v is, by definition, the service cost $C(S)$ plus the excluded value $v(U \setminus S) = \sum_{i \notin S} v_i$. This objective function is ordinally equivalent to the more standard welfare objective, which is the difference between the value served $\sum_{i \in S} v_i$ and the cost $C(S)$. Moreover, it is, in a precise sense, the "minimal perturbation" of the welfare objective function that admits non-trivial relative approximation guarantees; see [18] for details and additional justification for studying this objective. A cost-sharing mechanism is *α-approximate* if, assuming truthful bids, it is an α-approximation algorithm for the social cost objective. We state all of our lower bounds in terms of this approximation measure, but our proofs immediately yield comparable lower bounds for the additive efficiency loss measure adopted by Moulin and Shenker [15].

For a *public excludable problem*, in which $C(S) = 1$ for every non-empty S, the optimal solution is either U (for valuation profiles v with $v(U) \geq 1$) or $\emptyset$ (otherwise).

We conclude this section by describing a central mechanism [18, 5, 4, 15] for the public excludable good problem. Following [15], we call this the *Shapley value mechanism*. Given a set of bids, the mechanism serves the largest set $S \subseteq U$ such that for each player $i \in S$, $b_i \geq 1/|S|$. (Such sets are closed under union, and hence there is a unique largest such set.) Every player in S pays $1/|S|$ and the other players pay 0; the price that a player in S pays is precisely its Shapley value in the set S with respect to the function $C(\cdot)$. The mechanism is obviously budget-balanced; it is also truthful [15] and $\mathcal{H}_k$-approximate [18], where $k = |U|$ and $\mathcal{H}_k$ is the kth harmonic number. (Recall that $\mathcal{H}_k \approx \ln k$.) We recall here the example that shows that the result is tight.

Example 1. Let ϵ be a small positive number. Consider the truthful bid vector $1-\epsilon, 1/2-\epsilon, 1/3-\epsilon \ldots 1/k-\epsilon$. The solution which optimizes social cost serves all the players and has social cost 1. On the other hand, the Shapley value mechanism serves no players and has social cost $\mathcal{H}_k - k\epsilon$. Since ϵ can be arbitrarily small, the Shapley value mechanism is no better than $\mathcal{H}_k$-approximate.

This paper investigates whether or not there are truthful budget-balanced mechanisms that outperform the Shapley value mechanism.

3 A Lower Bound on Cost-Sharing Mechanisms

In this section we prove that every $O(1)$-budget-balanced cost-sharing mechanism for the public excludable good problem is $\Omega(\log k)$-approximate. This lower bound applies even to randomized mechanisms, and even to mechanisms that are only truthful in expectation.

Theorem 1. *Every cost-sharing mechanism for the public excludable good problem that is truthful in expectation and β-budget-balanced in expectation is $\Omega((\log k)/\beta)$-approximate, where k is the number of players.*

Proof. Fix values for k and $\beta \geq 1$. The plan of the proof is to define a distribution over valuation profiles such that the sum of the valuations is likely to be large but every mechanism is likely to produce the empty allocation. Let $a_1, \ldots, a_k$ be i.i.d. draws from the distribution with density $1/z^2$ on $[1, k]$ and remaining mass $(1/k)$ at zero. Set $v_i = a_i/4k\beta$ for each i and $V = \sum_{i=1}^{k} v_i$. We first note that V is likely to be $\Omega((\log k)/\beta)$. To see why, we have $\mathbf{E}[V] = k\mathbf{E}[v_i] = (\ln k)/4\beta$, $\mathbf{Var}[V] = k\mathbf{Var}[v_i] \leq k\mathbf{E}[v_i^2] = 1/(16\beta^2)$, and $\sigma[V] = 1/4\beta$. By Chebyshev's Inequality, V is at least $(\ln k - 2)/4\beta = \Omega(\log k/\beta)$ with probability at least 3/4.

Let M be a mechanism that is truthful in expectation and β-budget-balanced in expectation, meaning that for every bid vector, the expected revenue of M is at least a β fraction of its expected cost. For a public excludable good problem, the expected cost equals 1 minus the probability that no player is served. We can finish the proof by showing that the expected revenue of M, over both the random choice of valuation profile and the internal coin flips of the mechanism, is at most $1/4\beta$: if true, the expected cost of M is at most 1/4, so no player is served with probability at least 3/4. By the Union Bound, the probability that no player is served and also the sum of the valuations is $\Omega((\log k)/\beta)$ is at least 1/2. Thus, there is a valuation profile for which the optimal social cost is 1 but the expected social cost of M is $\Omega((\log k)/\beta)$.

We next apply a transformation of Mehta and Vazirani [12], originally developed for digital goods auctions, to assist in upper bounding the revenue obtained by M. Given a bid vector b, a *randomized threshold mechanism* chooses a random threshold $t_i(b_{-i})$ for each player i (cf., Proposition 1) from a distribution that is independent of b_i. Such mechanisms are truthful in the universal sense. By Mehta and Vazirani [12], there is a randomized threshold mechanism M' that has the same expected revenue as M on every bid vector.

To upper bound the expected revenue of M', consider a single truthful player i with (random) valuation v_i. Every fixed threshold t extracts expected revenue $t \cdot \mathbf{Pr}[v_i \geq t] \leq 1/4k\beta$ from the player. By the Principle of Deferred Decisions, a randomized threshold that is independent of v_i also obtains expected revenue at most $1/4k\beta$ from player i. Linearity of expectation implies that the expected revenue of M', and hence of M, is at most $1/4\beta$, completing the proof.

Scaling the prices of the Shapley value mechanism down by a $\beta \geq 1$ factor gives a β-budget-balanced, $O(\beta + (\log k)/\beta)$-approximate mechanism [18]. Thus, the lower bound in Theorem 1 is optimal up to constant factors for all $\beta = O(\sqrt{\log k})$.

4 Deterministic, Symmetric Mechanisms: Characterizations and Lower Bounds

In this section we prove a lower bound on the social cost approximation factor of every deterministic, budget-balanced cost-sharing mechanism that satisfies the "equal treatment" property. We derive this lower bound from a new characterization of the Shapley value mechanism, discussed next.

Proposition 1 does not specify the behavior of a truthful mechanism when a player bids exactly its threshold $t_i(b_{-i})$. There are two valid possibilities, each of which yields zero utility to a truthful player: the player is not served (at price 0), or is served and charged its bid. The following technical condition breaks ties in favor of the second outcome.

Definition 1. *A mechanism satisfies* upper semi-continuity *if and only if the following condition holds for every player i and bids b_{-i} of the other players: if player i receives service at every bid larger than b_i, then it also receives service at bid b_i.*

We stress that while our characterization result (Theorem 2) relies on this condition, our lower bound (Corollary 1) does not depend on it.

Our results concern mechanisms satisfying the following symmetry property.

Definition 2. *A mechanism satisfies* equal treatment *if and only if every two players i and j that submit the same bid receive the same allocation and price.*

The Shapley value mechanism (Section 2) satisfies equal treatment and upper semi-continuity. It uses the same threshold function for each player, namely:

$$\forall\ b_{-i}: \qquad t(b_{-i}) = \frac{1}{f(b_{-i}) + 1}. \tag{2}$$

Here, $f(b_{-i})$ is the size of the largest subset S of $U \setminus \{i\}$ such that $b_j \geq 1/(|S|+1)$ for all $j \in S$. Intuitively, this is precisely the set of other players that the Shapley value mechanism services if player i pays its share and also receives service.

Our characterization theorem is the following.

Theorem 2. *A deterministic and budget-balanced cost-sharing mechanism satisfies equal treatment, consumer sovereignty, and upper-semicontinuity if and only if it is the Shapley value mechanism.*

Proof. Fix such a mechanism M. We first note that all thresholds $t_i(b_{-i})$ induced by M must lie in $[0, 1]$: every threshold is finite by consumer sovereignty, and is at most 1 by the budget-balance condition. We proceed to show that for all players i and bids b_{-i} by the other players, the threshold function t_i has the same value as that for the Shapley value mechanism. We prove this by downward induction on the number of coordinates of b_{-i} that are equal to 1.

For the base case, fix i and suppose that b_{-i} is the all-ones vector. Suppose that $b_i = 1$. Since all thresholds are in $[0, 1]$ and M is upper semi-continuous, all players are served. By equal treatment and budget-balance, all players pay $1/k$. Thus, $t_i(b_{-i}) = 1/k$ when b_{-i} is the all-ones vector, as for the Shapley value mechanism.

For the inductive step, fix a player i and a bid vector b_{-i} that is not the all-ones vector. Set $b_i = 1$ and consider the bid vector $b = (b_i, b_{-i})$. Let S denote the set of players j with $b_j = 1$. Let $R \supseteq S$ denote the output of the Shapley value mechanism for the bid vector b — the largest set of players such that $b_j \geq 1/|R|$ for all $j \in R$.

As in the base case, consumer sovereignty, budget-balance, and equal treatment imply that M serves all of the players of S at a common price p. For a player j outside S, b_{-j} has one more bid of 1 than b_{-i} (corresponding to player i), and the inductive hypothesis implies that its threshold is that of the Shapley value mechanism for the same bid vector b. For players of $R \setminus S$, this threshold is $1/|R|$. For a player outside R, this threshold is some value strictly greater than its bid. Since $b_j \geq 1/|R|$ for all $j \in R$ and M is upper semicontinuous, it serves precisely the set R when given the bid vector b. This generates revenue $|S|p + (|R| - |S|)/|R|$. Budget-balance dictates that the common threshold p for all players of S, and in particular the value of $t_i(b_{-i})$, equals $1/|R|$. This agrees with player i's threshold for the bids b_{-i} in the Shapley value mechanism, and the proof is complete.

Theorem 2 implies that the Shapley value mechanism is the optimal deterministic, budget-balanced mechanism that satisfies the equal treatment property.

Corollary 1. *Every deterministic, budget-balanced cost-sharing mechanism that satisfies equal treatment is at least $\mathcal{H}_k$-approximate.*

We briefly sketch the proof. Let M be such a mechanism. If M fails to satisfy consumer sovereignty, then we can find a player i and bids b_{-i} such that $t_i(b_{-i}) = +\infty$. Letting the valuation of player i tend to infinity shows that the mechanism fails to achieve a finite social cost approximation factor.

Suppose that M also satisfies consumer sovereignty. The proof of Theorem 2 shows that the outcome of the mechanism agrees with that of the Shapley value mechanism except on the measure-zero set of bid vectors for which there is at least one bid equal to $1/i$ for some $i \in \{1, \dots, k\}$. As in Example 1, bid vectors of the form $1 - \epsilon$, $\frac{1}{2} - \epsilon, \dots, \frac{1}{k} - \epsilon$ for small $\epsilon > 0$ show that M is no better than $\mathcal{H}_k$-approximate.

Remark 1. Other characterizations of the Shapley value mechanism are known. See Moulin and Shenker [15] and Immorlica, Mahdian, and Mirrokni [9] for related characterizations of *groupstrategyproof* mechanisms that satisfy various properties. (A groupstrategyproof mechanism is robust to coordinated false bids when there are no side payments between players. The Shapley value mechanism satisfies this strong incentive-compatibility condition.) Our Theorem 2 is incomparable to these results because we work with the much richer class of truthful, not necessarily groupstrategyproof, mechanisms. Our characterization is more similar to that of Deb and Razzolini [5], who also show that the Shapley value mechanism is the only one that satisfies certain conditions. We weaken their stand-alone condition to consumer sovereignty and do not require the voluntary non-participation condition. Also, our proof is arguably simpler.

An interesting research problem is to characterize the class of mechanisms obtained after dropping the (admittedly strong) equal treatment condition. There are several mechanisms that satisfy the remaining conditions and appear hard to characterize (e.g. [9, Example 4.1]).

5 The Power of Randomization

Theorem 1 shows that the best-possible approximation guarantee of a randomized cost-sharing mechanism cannot be more than a constant factor smaller than that of the (deterministic) Shapley value mechanism. We now show that randomized mechanisms are in fact strictly more powerful than deterministic ones, even in the two-player public excludable good problem.

Proposition 2. *Let M be a deterministic budget-balanced cost-sharing mechanism for the 2-player public excludable good problem. Then, M is at least* 1.5*-approximate.*

Proof. Consider the bid vector with $b_1 = b_2 = 1$. Every mechanism that provides an approximation ratio better than 2 must serve both players. Suppose this is the case and player 1 pays p while player 2 pays $1 - p$. Without loss of generality, assume that $p \leq 0.5$. By Proposition 1, player 2's threshold function satisfies $t_2(1) = 1 - p$.

Now suppose $b_1 = 1$ and $b_2 = 1 - p - \epsilon$ for small $\epsilon > 0$. The optimal social cost is 1, with both players served. Since $t_2(1) = 1 - p$, player 2 is not served by M. Whether or not player 1 is served, the incurred social cost is $1 + 1 - p - \epsilon \geq 1.5 - \epsilon$.

There is a randomized mechanism with strictly better approximate efficiency.

Proposition 3. *There is a universally truthful, budget-balanced,* 1.25*-approximate randomized mechanism for the two-player public excludable good problem.*

Proof. The mechanism starts by selecting $\gamma \in [0,1]$ uniformly at random. Then, players 1 and 2 are offered service at prices γ and $1-\gamma$, respectively. A player who refuses is not served. If both players accept, then both are served at their respective prices. If exactly one player accepts, it is served (at price 1) if and only if its bid is at least 1.

The mechanism is clearly universally truthful and budget-balanced with probability 1. To bound its expected social cost, assume truthful bids with $v_1 \geq v_2$ and define $x = v_1 + v_2 - 1$. If $x < 0$ then, with probability 1, neither player is served and this is optimal. If $v_2 \geq 1$, then both players are served with probability 1, which again is optimal.

The most interesting case is when $x, v_1, v_2 \in [0,1]$. The optimal social cost in this case is 1. The mechanism selects a γ such that $v_1 \geq \gamma$ and $v_2 \geq 1-\gamma$ with probability x. In this event, both players are served and the incurred social cost is 1. Otherwise, neither player is served and the social is $1+x$. The expected approximation ratio obtained by the algorithm for this valuation profile is $x \cdot 1 + (1-x) \cdot (1+x)$. Choosing $x = 0.5$ maximizes this ratio, at which point the ratio is 1.25.

Finally, if $v_1 \geq 1$ but $v_2 < 1$, both players are served with probability v_2, and the mechanism serves only player 1 otherwise. The optimal social cost is again 1 and the expected social cost incurred by the mechanism is $v_2 \cdot 1 + (1-v_2)(1+v_2)$. This quantity is maximized when $v_2 = 0.5$, at which point the expected social cost (and hence the expected approximation ratio) is 1.25.

Unfortunately, universally truthful mechanisms cannot help further.

Proposition 4. *Let M be a universally truthful, budget-balanced cost-sharing mechanism for the two-player public excludable good problem. Then M is no better than* 1.25*-approximate.*

Proof. By Yao's Minimax Principle (e.g. [13]), we only need to exhibit a distribution over valuation profiles so that the approximate efficiency of every deterministic budget-balanced mechanism is large.

Let M be a deterministic, budget-balanced truthful mechanism. Let t_1 and t_2 denote the threshold functions for M in the sense of Proposition 1. Since M is budget-balanced on the bid vector $(1,1)$, $t_1(1) + t_2(1) = 1$. Fix $\epsilon > 0$ and randomize uniformly between the profiles $v_1 = 1, v_2 = (1/2) - \epsilon$ and $v_1 = (1/2) - \epsilon, v_2 = 1$. The optimal social cost is 1 for both of these profiles. Since either $t_1(1) \geq 1/2$ or $t_2(1) \geq 1/2$, the expected social cost of M is at least $(1/2) \cdot 1 + (1/2) \cdot (1 + (1/2) - \epsilon)$, which tends to $5/4$ as $\epsilon \to 0$.

References

1. L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 11, pages 267–299. Cambridge, 2007.

2. J. Brenner and G. Schäfer. Cost sharing methods for makespan and completion time scheduling. In *Proceedings of the 24th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 4393 of *Lecture Notes in Computer Science*, pages 670–681, 2007.
3. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
4. R. Deb and L. Razzolini. Auction-like mechanisms for pricing excludable public goods. *Journal of Economic Theory*, 88(2):340–368, 1999.
5. R. Deb and L. Razzolini. Voluntary cost sharing for an excludable public project. *Mathematical Social Sciences*, 37:123–138, 1999.
6. J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Hardness results for multicast cost sharing. *Theoretical Computer Science*, 304(1-3):215–236, 2003.
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. T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
9. N. Immorlica, M. Mahdian, and V. S. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms*, 2008. To appear.
10. R. Lavi. Computationally efficient approximation mechanisms. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 12, pages 301–329. Cambridge, 2007.
11. A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin mechanisms. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*, pages 1–10, 2007. Full version available from `http://theory.stanford.edu/~tim/`.
12. A. Mehta and V. V. Vazirani. Randomized truthful auctions of digital goods are randomizations over truthful auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 120–124, 2004.
13. R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge, 1995.
14. H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
15. H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
16. N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, 2001.
17. K. Roberts. The characterization of implementable choice rules. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.
18. T. Roughgarden and M. Sundararajan. New trade-offs in cost-sharing mechanisms. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 79–88, 2006. Full version available from `http://theory.stanford.edu/~tim/`.
19. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.