Is Shapley Cost Sharing Optimal?*
(For the special issue in honor of Lloyd Shapley)

Shahar Dobzinski† Aranyak Mehta‡ Tim Roughgarden§
Mukund Sundararajan¶

August 29, 2016

Abstract

A general approach to the design of budget-balanced cost-sharing mechanisms is to use the Shapley value, applied to the given cost function, to define payments from the players to the mechanism. Is the corresponding Shapley value mechanism "optimal" in some sense? We consider the objective of minimizing worst-case inefficiency subject to a revenue constraint, and prove results in three different regimes.

1. For the public excludable good problem, the Shapley value mechanism minimizes the worst-case efficiency loss over all truthful, deterministic, and budget-balanced mechanisms that satisfy equal treatment. This result follows from a characterization of the Shapley value mechanism as the unique one that satisfies two additional technical conditions.

2. For the same problem, even over a much more general mechanism design space that allows for randomization and approximate budget-balance and does not impose equal treatment, the worst-case efficiency loss of the Shapley value mechanism is within a constant factor of the minimum possible.

*The results in Sections 3–5 appeared, in preliminary form, in the extended abstract [8]. The results in Section 6 appeared in the PhD thesis of the fourth author [27].

†Department of Applied Math and Computer Science, the Weizmann Institute of Science, Rohovot 76100, Israel. This work was done while the author was visiting Stanford University, and 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: `dobzin@gmail.com`.

‡Google, Inc., Mountain View, CA. Email: `aranyak@google.com`.

§Corresponding author. Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. This work was supported in part by NSF awards CCF-0448664 and CCF-1524062, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship. Email: `tim@cs.stanford.edu`.

¶Google, Inc., Mountain View, CA. This work was done while a PhD student at Stanford University and supported by NSF Award CCF-0448664 and a Stanford Graduate Fellowship. Email: `mukunds@google.com`

3. For no-deficit mechanisms that need not satisfy approximate budget-balance, we prove a general positive result: for every monotone cost function, a suitable blend of the VCG and Shapley value mechanisms is no-deficit and enjoys good approximate efficiency guarantees.

Keywords: Shapley value, cost-sharing mechanisms, approximate efficiency

1 Introduction

In a cost-sharing mechanism design problem, several participants with unknown preferences vie to receive some good or service, and each possible outcome has a known cost. Formally, we consider problems defined by a set U of players and a cost function $C : 2^U \to \mathbb{R}^+$ that describes the cost incurred by the mechanism as a function of the outcome (i.e., of the set S of "winners"). We assume that each player i has a private nonnegative value v_i for winning.

For example, in the *public excludable good* problem (e.g. [7, 22]), the problem is to determine whether or not to finance a public good and, if so, who is allowed to use it.[1] This problem corresponds to the cost function C with $C(\emptyset) = 0$ and $C(S) = 1$ for every $S \neq \emptyset$. Many other cost functions have been considered in the cost-sharing literature (Section 1.2), and most of them include public excludable good problems as a special case.

A (direct-revelation) *cost-sharing mechanism* is a protocol that decides, as a function of players' bids, which players win and at what prices. For example, the general Vickrey-Clarke-Groves (VCG) mechanism specializes to the following procedure for a public excludable good problem.

VCG Mechanism (Public Excludable Good)

1. Accept a bid b_i from each player i.
2. Choose the outcome $S := U$ if $\sum_{i \in U} b_i > 1$, and $S := \emptyset$ otherwise.
3. Charge each winner i the minimum bid for which she would still win (holding others' bids fixed), namely $\max\{0, 1 - \sum_{j \in U \setminus \{i\}} b_j\}$.

It is well known that the VCG mechanism is *truthful*, meaning that for every player it is a dominant strategy to set her bid equal to her private value for winning. By design, the VCG mechanism is also *efficient*, meaning that it always selects the set $S \subseteq U$ of winners that maximizes the total value to the winners less the cost incurred, that is, the *social welfare* $\sum_{i \in S} v_i - C(S)$. One drawback of the VCG mechanism is that its revenue can be far from the cost incurred. For example, in a public excludable good problem in which all of the players have valuations larger than $\frac{1}{|U|-1}$, the VCG mechanism obtains zero revenue (while the cost is 1).

[1]In 1959, the citizens of Palo Alto, Portola Valley, and Los Altos Hills voted on whether or not to finance a new park. The measure only passed in Palo Alto, and to this day entrance to Foothills Park is restricted to Palo Alto residents.

A second approach to designing a cost-sharing mechanism is to insist on *budget balance*, meaning that the sum of players' payments equals the cost of the outcome chosen. For a symmetric problem like a public excludable good problem, perhaps the most natural approach is to require equal payments from the winners, and subject to this choose as many winners as possible. The Shapley value mechanism implements this idea. For the special case of a public excludable good, the Shapley value mechanism chooses the largest set S of players such that $b_j \geq 1/|S|$ for all $j \in S$. The mechanism can be described more procedurally as follows.[2]

Shapley Value Mechanism (Public Excludable Good)

1. Accept a bid b_i from each player i.
2. Initialize $S := U$.
3. If $b_i \geq 1/|S|$ for every $i \in S$, then halt with winners S, and charge each player $i \in S$ the price $p_i = 1/|S|$.
4. Let $i^* \in S$ be a player with $b_{i^*} < 1/|S|$.
5. Set $S := S \setminus \{i^*\}$ and return to Step 3.

The Shapley value mechanism is also truthful—overbidding can only cause a player to win when she would prefer to lose, and vice versa for underbidding. By design, it is budget-balanced. It is not efficient, however.

Example 1.1 (Inefficiency of Shapley Value Mechanism) Consider a public excludable good problem with k players, where the valuation of player i is $\frac{1}{i} - \delta$ for small $\delta > 0$. By induction, the Shapley value mechanism will remove player $k+1-i$ in its ith iteration, terminating with the empty outcome, which has zero social welfare. The welfare-maximizing outcome is to choose the full set $S = U$ of winners. This results in social welfare approaching $\mathcal{H}_k - 1$ as $\delta \to 0$, where $\mathcal{H}_k = \sum_{i=1}^{k} \frac{1}{i}$ denotes the kth Harmonic number (which lies between $\ln k$ and $\ln k + 1$).

Thus, the VCG mechanism sacrifices budget-balance in the service of efficiency, while the Shapley value mechanism makes the opposite trade-off. This trade-off between efficiency and budget-balance is fundamental: no truthful mechanism can be both [11, 24]. This impossibility result raises the issue of understanding the feasible trade-offs between the two objectives.

This paper is motivated primarily by the following question:

[2] We call this mechanism the Shapley value mechanism (following [22]) because the prices charged to the winning set S correspond to the Shapley value applied to the cost function C restricted to S (since C is symmetric, the Shapley values are equal). The Shapley value mechanism can be defined analogously for arbitrary cost-sharing problems (see [22]).

Does the Shapley value mechanism have the best-possible efficiency guarantee for a budget-balanced mechanism, for public excludable goods and for more general cost-sharing problems?

1.1 Summary of Results

Formalizing the question above requires specifying a performance measure that we want to optimize and the design space of mechanisms that we are willing to consider. Throughout this paper, we consider the objective of minimizing inefficiency, in the worst case over all valuation profiles. (See Section 2 for formal definitions.) We prove results for three different choices of the mechanism design space.

1. Our strongest optimality result is for the space of truthful, deterministic, and budget-balanced mechanisms for public excludable good problems that satisfy equal treatment (two players with equal bids receive the same allocations and prices). Here, the Shapley value mechanism minimizes the worst-case efficiency loss (Corollary 3.5). This result follows from a characterization of the Shapley value mechanism as the unique one that satisfies two additional technical conditions (Theorem 3.4).

2. Our second regime continues to focus on public excludable good problems but considers a much more general mechanism design space that allows for randomization and approximate budget-balance, and with no imposition of equal treatment. We show that randomized mechanisms can have strictly smaller worst-case efficiency loss than any deterministic mechanism (Propositions 4.1 and 4.2), including the Shapley value mechanism. Our main result here is that, even with respect to this large design space, the Shapley value mechanism still has worst-case efficiency loss within a constant factor of the minimum possible (Theorem 5.1).

3. Finally, we relax the approximate budget-balance constraint to a no-deficit condition, requiring that a mechanism's revenue is at least the cost incurred. Here, we give a quite general positive result (Theorem 6.1): for every monotone cost function, a "hybrid" mechanism that blends the VCG and Shapley value mechanisms is no-deficit and enjoys good approximate efficiency guarantees.

1.2 Related Work

Moulin [20] considers truthful mechanisms for public excludable good problems, and subsequent work on the problem includes Deb and Razzolini [6, 7], Moulin and Shenker [22], and Massó et al. [16]. The idea of comparing different cost-sharing mechanisms using the worst-case efficiency loss measure is from Moulin and Shenker [22]. The notion of loss in [22] is additive; Roughgarden and Sundararajan [26] develop an analogous theory for relative approximation guarantees. Subsequent works that give approximate efficiency guarantees for specific cost-sharing mechanisms include Chawla et al. [5], Roughgarden and Sundararajan [25], Mehta et al. [17], Brenner and Schäfer [3, 4], Bleischwitz et al. [1], Bleischwitz and Schoppmann [2], Gupta et al. [12], Moulin [21], and Juarez [15].

Moulin and Shenker [22] and Roughgarden and Sundararajan [26] prove that the Shapley value mechanism has the minimum-possible worst-case inefficiency of any budget-balanced "Moulin mechanism," meaning a cost-sharing mechanism derived from a cross-monotonic cost-sharing method (see [22]). Brenner and Schäfer [3] give additional negative results for the inevitable inefficiency of Moulin mechanisms. None of the optimality results in this paper (Corollary 3.5 or Theorem 5.1) are confined to Moulin mechanisms; for example, we never require group-strategyproofness (a property of every Moulin mechanism [22]).

Our Theorem 3.4 is similar to the characterization result of Deb and Razzolini [7], 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. Less directly related are characterizations of group-strategyproof (rather than only truthful) cost-sharing mechanisms that satisfy various conditions, including those of Moulin and Shenker [22], Immorlica, Mahdian, and Mirrokni [13], Pountourakis and Vidali [23], and Juarez [14]. Subsequent to the preliminary publication of our characterization result [8], Massó et al. [16] extended our Corollary 3.5 by relaxing the budget-balance requirement to the no-deficit condition, and replacing the equal treatment condition with different fairness conditions. Also after the publication of [8], our lower bound in Theorem 5.1 was extended to Bayesian incentive-compatible mechanisms by Fu et al. [9].

Finally, our Theorem 6.1 can be viewed as a simplification of a comparable result proved by Georgiou and Swamy [10], who also focus on computationally efficient mechanisms for various cost-sharing problems. Prior to [10], Bleischwitz et al. [1] gave an analogous result (with a different mechanism) for all subadditive cost functions (where $C(S \cup T) \leq C(S) + C(T)$ for all $S, T \subseteq U$).

2 Preliminaries

2.1 Cost-Sharing Problems and Mechanisms

We consider cost-sharing problems with 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 monotone (i.e., $S \subseteq T$ implies that $C(S) \leq C(T)$). 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 a payment p_i for each player $i \in U$. Player i has a private value v_i for being included in the chosen set S. We 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 $i \notin S$.

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 payments from the bidders to the mechanism are always nonnegative.

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

Section 3 uses the following standard fact about deterministic truthful mechanisms.

Proposition 2.1 *Let M be a deterministic, truthful, and individually rational cost-sharing mechanism with player set U. Then for every player $i \in U$ and bid vector $\mathbf{b}_{-i}$ for players other than i, there is a* threshold $t_i(\mathbf{b}_{-i}) \in \mathbb{R}^+ \cup \{+\infty\}$ *such that: (i) if player i bids more than $t_i(\mathbf{b}_{-i})$, then she is included in the output set S, at the price $t_i(\mathbf{b}_{-i})$; (ii) if player i bids less than $t_i(\mathbf{b}_{-i})$, then she is excluded from S.*

For example, in the VCG mechanism for a general cost-sharing problem—which chooses the welfare-maximizing outcome and charges each winner the minimum bid at which she would continue to win—the threshold $t_i(\mathbf{b}_{-i})$ equals the difference between the maximum reported social welfare that can be attained by the other players when player i is not and is included in the winner set S. (This is nonnegative under our assumption that the cost function C is monotone.) In the Shapley value mechanism (Section 1), for a public excludable good problem, the threshold $t_i(\mathbf{b}_{-i})$ is defined by

$$t_i(\mathbf{b}_{-i}) = \frac{1}{f_i(\mathbf{b}_{-i}) + 1}, \tag{1}$$

where $f_i(\mathbf{b}_{-i})$ denotes the size of the largest subset S of $U \setminus \{i\}$ such that $b_j \geq 1/(|S|+1)$ for all $j \in S$. This is precisely the set of other players that win in the Shapley value mechanism when player i also bids high enough to win.

2.2 Randomized Mechanisms

A randomized mechanism is, by definition, a probability distribution over deterministic (and possibly non-truthful) 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 her expected utility by misreporting her valuation. Every universally truthful mechanism is truthful in expectation, but the converse does not hold.

By a *randomized threshold mechanism*, we mean a mechanism that, for each player i, chooses a random threshold $t_i(\mathbf{b}_{-i})$ (cf., Proposition 2.1) from a distribution that is independent of b_i. Thresholds for different players need not be stochastically independent. Every randomized threshold mechanism is universally truthful.

Section 5 uses the following known result.

Proposition 2.2 ([18]) *For every truthful-in-expectation cost-sharing mechanism M, there is a universally truthful cost-sharing mechanism M' such that, for every bid vector $\mathbf{b}$, the expected revenues of M and M' are equal.*

2.3 Social Cost Minimization and Budget-Balance

We study two kinds of objectives for cost-sharing mechanisms, one for revenue and one for economic efficiency. First, for a parameter $\beta \geq 1$, a mechanism is *β-budget-balanced* if the sum $\sum_{i \in S} p_i$ of the prices charged lies between $C(S)/\beta$ and $C(S)$, where S is the chosen set of winners. We say that a mechanism is *budget-balanced* if it is 1-budget-balanced. Section 6 also considers *no-deficit* mechanisms, where the sum of the payments is always at least as large as the cost incurred.

Following several previous works (see Section 1.2), we measure the inefficiency of a cost-sharing mechanism via the *social cost* objective. The social cost $\pi(S)$ of an outcome S with respect to a cost function C and valuation profile $\mathbf{v}$ is, by definition, the cost $C(S)$ of the outcome plus the excluded value $\sum_{i \notin S} v_i$:

$$\pi(S) = C(S) + \sum_{i \notin S} v_i. \tag{2}$$

In Example 1.1, the minimum-possible social cost (with $S = U$) is 1, while the social cost of the outcome of the Shapley value mechanism (with $S = \emptyset$) is $\approx \mathcal{H}_k$.

The social cost objective function is ordinally equivalent to the negative of the social welfare (i.e., the negative of $\sum_{i \in S} v_i - C(S)$). It is also, in a precise sense, the "minimal perturbation" of the welfare objective function that admits non-trivial relative approximation guarantees; see [26] for details and additional justification for studying this objective.[3]

A (randomized) cost-sharing mechanism is *α-approximate* if, assuming truthful bids, the (expected) social cost of its outcome is at most $\alpha \geq 1$ times that of a, optimal (i.e., social cost-minimizing) outcome. For example, the VCG mechanism is 1-approximate by design. Example 1.1 shows that the Shapley value mechanism is not α-approximate for public excludable good problems for any $\alpha < \mathcal{H}_k$. Roughgarden and Sundararajan [26] proved that the Shapley value mechanism is $\mathcal{H}_k$-approximate for public excludable good problems, and more generally for all monotone submodular cost functions.

3 Deterministic Symmetric Mechanisms: Characterization and Lower Bound

This section proves a lower bound on the social cost approximation factor of every deterministic and 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.

This section considers mechanisms that satisfy the following symmetry property.

[3] We state all of our results in terms of relative approximations to the optimal social cost, but our proofs immediately yield analogous results for the additive efficiency loss measure proposed by Moulin and Shenker [22].

Definition 3.1 (Equal Treatment) 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.

Next is a technical condition. Proposition 2.1 does not specify the behavior of a truthful mechanism when a player bids exactly her threshold $t_i(\mathbf{b}_{-i})$. There are two valid possibilities, each of which yields zero utility to a truthful player: the player does not win (at price 0), or wins and is charged her bid. The following condition breaks ties in favor of the second outcome.

Definition 3.2 (Upper Semi-Continuity) A mechanism satisfies *upper semi-continuity* if and only if the following condition holds for every player i and bids $\mathbf{b}_{-i}$ of the other players: if player i wins with every bid larger than b_i, then it also wins with the bid b_i.

We stress that while our characterization result (Theorem 3.4) relies on upper semi-continuity, our lower bound (Corollary 3.5) does not depend on it. The same comment applies to the final condition.

Definition 3.3 (Consumer Sovereignty) A mechanism satisfies *consumer sovereignty* if and only if, for all players i and bids $\mathbf{b}_{-i}$ of the other players, there exists a bid b_i such that player i wins when the bid profile is $(b_i, \mathbf{b}_{-i})$.

By definition, the Shapley value mechanism (Section 1) satisfies equal treatment, upper semi-continuity, and consumer sovereignty. Our characterization result states that no other deterministic and budget-balanced cost-sharing mechanism satisfies these three conditions.

Theorem 3.4 (Characterization) *A deterministic, truthful, and budget-balanced cost-sharing mechanism for public excludable good problems satisfies equal treatment, consumer sovereignty, and upper semi-continuity if and only if it is the Shapley value mechanism.*

Proof: Fix such a mechanism M. We first note that all thresholds $t_i(\mathbf{b}_{-i})$ 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 $\mathbf{b}_{-i}$ by the other players, the threshold function t_i has the same value as that for the Shapley value mechanism (1). We prove this by downward induction on the number of coordinates of $\mathbf{b}_{-i}$ that are equal to 1.

For the base case, fix i and suppose that $\mathbf{b}_{-i}$ 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 win. By equal treatment and budget-balance, all players pay $1/k$. Thus, $t_i(\mathbf{b}_{-i}) = 1/k$ when $\mathbf{b}_{-i}$ is the all-ones vector, as for the Shapley value mechanism.

For the inductive step, fix a player i and a bid vector $\mathbf{b}_{-i}$ that is not the all-ones vector. Set $b_i = 1$ and consider the bid vector $\mathbf{b} = (b_i, \mathbf{b}_{-i})$. 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 $\mathbf{b}$—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 allocates to all of the players of S at a common price p. For a player j outside S,

$\mathbf{b}_{-j}$ has one more bid of 1 than $\mathbf{b}_{-i}$ (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 semi-continuous, it chooses precisely the winner set R when the bid vector is $\mathbf{b}$. 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 $t_i(\mathbf{b}_{-i})$, equals $1/|R|$. This agrees with player i's threshold for the bids $\mathbf{b}_{-i}$ in the Shapley value mechanism, and the proof is complete. ■

Theorem 3.4 implies that the Shapley value mechanism is the optimal deterministic, budget-balanced mechanism for public excludable good problems that satisfies the equal treatment property.

Corollary 3.5 (Lower Bound for Deterministic Symmetric Mechanisms) *No deterministic and budget-balanced cost-sharing mechanism for public excludable good problems that satisfies equal treatment is better than $\mathcal{H}_k$-approximate, where k is the number of players.*

Proof: Let M be such a mechanism. If M fails to satisfy consumer sovereignty, then we can find a player i and bids $\mathbf{b}_{-i}$ such that $t_i(\mathbf{b}_{-i}) = +\infty$. 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 satisfies consumer sovereignty. The proof of Theorem 3.4 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 $i \in \{1, \ldots, k\}$. As in Example 1.1, setting players' valuations to $v_i = \frac{1}{i} - \delta$ for each i, for arbitrarily small $\delta > 0$, shows that M is no better than $\mathcal{H}_k$-approximate. ■

Corollary 3.5 immediately applies also to deterministic and budget-balanced cost-sharing mechanisms for more general families of problems (Section 1), as long as the mechanism satisfies equal treatment for the special case of public excludable good problems.

4 The Power of Randomization

Section 5 generalizes the lower bound in Corollary 3.5 to wider classes of mechanisms, including randomized mechanisms. But does randomization ever help in cost-sharing mechanism design? This section proves the first efficiency separation between deterministic and randomized budget-balanced mechanisms. It suffices to consider two-player public excludable good problems.

Proposition 4.1 (Lower Bound for Deterministic Mechanisms) *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 a social cost approximation ratio better than 2 must allocate to 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 2.1, 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 winning. Since $t_2(1) = 1-p$, player 2 does not win in M. Whether or not player 1 wins, the incurred social cost is $1 + 1 - p - \epsilon \geq 1.5 - \epsilon$. ■

There is a randomized mechanism with strictly better approximate efficiency.

Proposition 4.2 (Upper Bound for Randomized Mechanisms) *There is a universally truthful, budget-balanced, and* 1.25*-approximate mechanism for the two-player public excludable good problem.*

Proof: We use a simple modification of the Shapley value mechanism from Section 1. In the first iteration, the mechanism selects $\gamma \in [0,1]$ uniformly at random. If players 1 and 2's bids are at least γ and $1-\gamma$, respectively, then the mechanism halts with $S = \{1,2\}$, $p_1 = \gamma$, and $p_2 = 1-\gamma$. If one of the players did not bid high enough, then this player is removed and, in the second iteration, the remaining player is asked to pay the full cost 1.

This mechanism is universally truthful (as is easily verified), and it is clearly budget-balanced. 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, the empty outcome is chosen, and this outcome is welfare-maximizing. If $v_2 \geq 1$, then the outcome $S = \{1,2\}$ is chosen with probability 1, and this is again 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 win and the incurred social cost is 1. Otherwise, neither player wins and the social cost 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 win with probability v_2, and only player 1 wins 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 is 1.25. ■

Remark 4.3 (Lower Bound for Universally Truthful Mechanisms) For two-player public excludable good problems, no universally truthful and budget-balanced cost-sharing mechanism is better than 1.25-approximate. To see this, consider choosing one of the valuation profiles $(1, \frac{1}{2} - \delta)$ or $(\frac{1}{2} - \delta, 1)$, uniformly at random, where $\delta > 0$ is arbitrarily small. The optimal social cost is 1 with probability 1.

Consider first a deterministic mechanism M with threshold functions t_1 and t_2 (Proposition 2.1). Since M is budget-balanced on the bid vector $(1,1)$, $t_1(1) + t_2(1) = 1$. Since either $t_1(1) \geq \frac{1}{2}$ or $t_2(1) \geq \frac{1}{2}$, the expected social cost of M is at least $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (1 + \frac{1}{2} - \delta)$, which tends to $5/4$ as $\delta \to 0$. Since every universally truthful mechanism is just a distribution over deterministic truthful mechanism, the same computation applies to them. Since the expected cost of a universally truthful and budget-balanced mechanism M' on this input distribution is (arbitrarily close to) $\frac{5}{4}$ times the optimal social cost, there exists a valuation profile ($(1, \frac{1}{2} - \delta)$ or $(\frac{1}{2} - \delta, 1)$) on which the expected social cost incurred by M' is at least $\frac{5}{4}$ times the optimal social cost.[4]

5 A General Lower Bound

This section proves a lower bound that is stronger than that in Corollary 3.5 in several respects: it applies even to randomized (truthful-in-expectation) mechanisms, to mechanisms that are only approximately budget-balanced, and to mechanisms that need not satisfy the equal treatment condition. Quantitatively, the lower bound is weaker than that of Corollary 3.5 by a constant factor (independent of the number k of players). Propositions 4.1 and 4.2 imply that this constant-factor loss is inevitable when comparing the Shapley value mechanism to randomized mechanisms, even assuming universal truthfulness and exact budget-balance.

Theorem 5.1 (Lower Bound for Randomized Mechanisms) *There is a constant $c > 0$ such that the following holds: no truthful-in-expectation and β-budget-balanced mechanism for public excludable good problems is better than $c \cdot \mathcal{H}_k/\beta$-approximate, where k is the number of players.*

Proof: Fix values for k and $\beta \geq 1$. Since no mechanism can be better than 1-approximate, we can assume that k is sufficiently large (otherwise take $c = \beta/\mathcal{H}_k$). 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 independent 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 at least a constant fraction of $(\ln k)/\beta$. To see why, we have

$$\mathbf{E}[V] = k\mathbf{E}[v_i] = \frac{\ln k}{4\beta}$$

and

$$\mathbf{Var}[V] = k\mathbf{Var}[v_i] \leq k\mathbf{E}[v_i^2] = \frac{1}{16\beta^2},$$

and hence

$$\sigma[V] \leq \frac{1}{4\beta}.$$

By Chebyshev's inequality, which states that

$$\Pr[|X - \mathbf{E}[X]| \geq \gamma \cdot \sigma[X]] \leq \frac{1}{\gamma^2}$$

[4]In computer science, this method of proving limitations on randomized algorithms via a suitable distribution over inputs is sometimes called "Yao's minimax principle" (see e.g. [19]).

for all $\gamma > 0$ (e.g. [19]), we have that V is at least $(\ln k - 2)/4\beta$ probability at least 3/4. For sufficiently large k, $(\ln k - 2)/4\beta$ is at least $\mathcal{H}_k/8\beta$.

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 wins. We claim 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 $1/4\beta$. To see why the claim implies the theorem, note that this would imply that the expected cost of M is at most 1/4, and so with probability at least 3/4, M chooses the empty allocation. Conditioned on the event that $\sum_{i \in U} v_i \geq \mathcal{H}_k/8\beta$, the probability that M chooses the empty allocation is at least 1/2. Thus, there exists a valuation profile $\mathbf{v}$ with $\sum_{i \in U} v_i \geq \mathcal{H}_k/8\beta$ such that, with probability at least 1/2 over the internal randomness of M, M chooses the empty allocation. The expected social cost of M on this valuation profile is at least $\mathcal{H}_k/16\beta$, while the optimal social cost is at most 1.

To prove the claim and upper bound the expected revenue of M with respect to this distribution over valuation profiles, first assume that M is a truthful deterministic mechanism. For every fixed threshold $t = t_i(\mathbf{b}_{-i})$ that arises in the mechanism (Proposition 2.1), the expected (over v_i) revenue extracted from player i is $t \cdot \Pr[v_i \geq t] \leq 1/4k\beta$. By the linearity of expectation, the expected (over $\mathbf{v}$) revenue of every deterministic truthful mechanism is at most $1/4\beta$. Since a universally truthful mechanism is just a distribution over deterministic truthful mechanisms, the expected revenue of every such mechanism is at most $1/4\beta$. Finally, Proposition 2.2 states that for every truthful-in-expectation mechanism M, there is a universally truthful mechanism M' such that M and M' have the same expected revenue (over the mechanisms' internal randomness) on every bid profile. We conclude that the expected revenue of every truthful-in-expectation mechanism is at most $1/4\beta$, which completes the proof. ■

The worst-case lower bound in Theorem 5.1 applies immediately to approximately budget-balanced cost-sharing mechanisms for families of cost-sharing problems that include public excludable good problems as a special case.

Scaling the prices of the Shapley value mechanism (Section 1) down by a $\beta \geq 1$ factor yields a truthful mechanism that is β-budget-balanced and $(\frac{\mathcal{H}_k}{\beta} + \beta)$-approximate [26]. This fact shows that the linear degradation in β of the lower bound in Theorem 5.1 is necessary, at least up to $\beta \approx \sqrt{\mathcal{H}_k}$.

6 A General Upper Bound

This section notes that a blend of the VCG mechanism (for general cost functions) and the Shapley value mechanism (for a public excludable good) is truthful and approximately efficient even for very general cost functions.[5] This hybrid mechanism need not be approxi-

[5]Why not just run the Shapley value mechanism, which remains well defined for arbitrary cost functions? (Each iteration, each remaining player is asked to pay her Shapley value with respect to the cost function

mately budget-balanced, but it is no-deficit, meaning that the revenue obtained is always at least the cost incurred. This result can be viewed as a simplification of one in Georgiou and Swamy [10].

Consider an arbitrary player set U and monotone cost function C.

Hybrid Mechanism (Monotone Cost Functions)

1. Accept a bid b_i from each player i.
2. Let
$$S^* \in \text{argmax}_{S \subseteq U} \left[\sum_{i \in S} b_i - C(S) \right]$$
denote a welfare-maximizing outcome.
3. Initialize $S := S^*$.
4. If $b_i \geq C(S^*)/|S|$ for every $i \in S$, then halt with winners S.
5. Let $i^* \in S$ be a player with $b_{i^*} < C(S^*)/|S|$.
6. Set $S := S \setminus \{i^*\}$ and return to Step 4.
7. Charge each winner $i \in S$ a payment equal to the minimum bid at which i would continue to win (holding $\mathbf{b}_{-i}$ fixed).

That is, the hybrid mechanism computes the welfare-maximizing allocation S^* (according to the reported bids) and then uses the Shapley value (i.e., equal-share) mechanism to share the cost $C(S^*)$ among the largest subset of S^* that is willing to pay it. The final payment by a winner can be viewed as the maximum of the payment from the "VCG phase" and from the "cost-sharing phase."

Theorem 6.1 (Properties of the Hybrid Mechanism) *For every monotone cost function, the hybrid mechanism is truthful, satisfies the no positive transfers, individual rationality and no-deficit conditions, and is $\mathcal{H}_k$-approximate.*

Proof: The individual rationality and no positive transfers properties are immediate from the definition of the payment rule. For truthfulness, the key observation is that the allocation rule of the hybrid mechanism is monotone, meaning that a winner would continue to win at any higher bid. To see this, fix a player i and bids $\mathbf{b}_{-i}$ of the other players. Fix a winning bid b_i and another bid $\bar{b}_i$ with $\bar{b}_i > b_i$. Let S^* and $\bar{S}^*$ be the welfare-maximizing outcomes for the bid vectors $(b_i, \mathbf{b}_{-i})$ and $(\bar{b}_i, \mathbf{b}_{-i})$, respectively. Because player i wins with bid b_i, it is in S^*. After raising player i's bid from b_i to $\bar{b}_i$, S^* remains the welfare-maximizing outcome (i.e., $\bar{S}^* = S^*$). Thus player i graduates to the second phase and participates in

restricted to the remaining player set.) The problem is that this mechanism need not be truthful (without making an additional submodularity assumption) [22].

same cost-sharing problem whether she bids b_i or $\bar{b}_i$. Since the Shapley value mechanism for a public excludable good has a monotone allocation rule, if player i wins with the bid b_i, she also wins with the bid $\bar{b}_i$. Because the allocation rule is monotone and each player is charged the minimum bid at which she would continue to win, the mechanism is truthful.

For the no-deficit condition, observe that the mechanism either halts with $S = \emptyset$ (in which case there are no costs and no payments), or else it terminates with a winning set $S \subseteq S^*$ and revenue at least $C(S^*)$. Since C is monotone, $C(S) \leq C(S^*)$ and hence the mechanism's revenue is at least its cost.

Finally, for the social cost approximation guarantee, by definition the set S^* of players that graduate to the cost-sharing phase minimizes the social cost $C(S^*) + \sum_{i \notin S^*} v_i$ (this is equivalent to maximizing social welfare). For the public excludable good problem considered in the second phase, the additional additive efficiency loss caused by the Shapley mechanism is $\mathcal{H}_k - 1$, scaled by the cost $C(S^*)$ of producing the good [22, 26]. Thus the social cost of the final winner set S satisfies

$$C(S) + \sum_{i \notin S} v_i \leq C(S^*) + \sum_{i \notin S^*} v_i + (\mathcal{H}_k - 1) \cdot C(S^*) \leq \mathcal{H}_k \cdot \left(C(S^*) + \sum_{i \notin S^*} v_i \right),$$

completing the proof. ■

Remark 6.2 (Generalizations) The hybrid mechanism can be generalized to work with surrogates of the VCG and Shapley value mechanisms; computational efficiency is one motivation for using surrogates in place of these mechanisms. The VCG mechanism can be replaced by any truthful mechanism — ideally, an approximately welfare-maximizing one — for which the output set is invariant under increases of a winner's bid. The Shapley value mechanism can be replaced by any truthful and no-deficit mechanism, which ideally excludes as little valuation as possible. For instance, for subadditive cost functions (with $C(S \cup T) \leq C(S) + C(T)$ for all $S, T \subseteq U$), the mechanism from Bleischwitz et al. [1] is a good surrogate for the Shapley value mechanism. Georgiou and Swamy [10] formalize this composition technique and use it to identify computationally efficient implementations of no-deficit mechanisms with good worst-case inefficiency for families of cost functions derived from facility location, network design, and covering problems.

References

[1] Y. Bleischwitz, B. Monien, and F. Schoppmann. To be or not to be (served). In *Proceedings of the Third Annual International Workshop on Internet and Network Economics (WINE)*, volume 4858 of *Lecture Notes in Computer Science*, pages 515–528, 2007.

[2] Y. Bleischwitz and F. Schoppmann. Group-strategyproof cost sharing for metric fault tolerant facility location. In *Proceedings of the First International Symposium on Algorithmic Game Theory (SAGT)*, pages 350–361, 2008.

[3] J. Brenner and G. Schäfer. Group-strategyproof cost sharing mechanisms for makespan and other scheduling problems. *Theoretical Computer Science*, 401:96–106, 2008.

[4] J. Brenner and G. Schäfer. Singleton acyclic mechanisms and their applications to scheduling problems. In *Proceedings of the First International Symposium on Algorithmic Game Theory (SAGT)*, pages 315–326, 2008.

[5] S. Chawla, T. Roughgarden, and M. Sundararajan. Optimal cost-sharing mechanisms for network design. In *Proceedings of the Second Annual International Workshop on Internet and Network Economics (WINE)*, volume 4286 of *Lecture Notes in Computer Science*, pages 112–123, 2006.

[6] R. Deb and L. Razzolini. Auction-like mechanisms for pricing excludable public goods. *Journal of Economic Theory*, 88(2):340–368, 1999.

[7] R. Deb and L. Razzolini. Voluntary cost sharing for an excludable public project. *Mathematical Social Sciences*, 37:123–138, 1999.

[8] S. Dobzinski, A. Mehta, T. Roughgarden, and M. Sundararajan. Is Shapley cost-sharing optimal? In *Proceedings of the First International Symposium on Algorithmic Game Theory (SAGT)*, pages 327–336, 2008.

[9] H. Fu, B. Lucier, B. Sivan, and V. Syrgkanis. Cost-recovering Bayesian algorithmic mechanism design. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 453–470, 2013.

[10] K. Georgiou and C. Swamy. Black-box reductions for cost-sharing mechanism design. *Games and Economic Behavior*. To appear.

[11] J. Green, E. Kohlberg, and J. J. Laffont. Partial equilibrium approach to the free rider problem. *Journal of Public Economics*, 6:375–394, 1976.

[12] A. Gupta, J. Könemann, S. Leonardi, R. Ravi, and G. Schäfer. Efficient cost-sharing mechanisms for prize-collecting problems. *Mathematical Programming*, 152(1):147–188, 2015.

[13] N. Immorlica, M. Mahdian, and V. S. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms*, 4(2), 2008.

[14] R. Juarez. Group strategyproof cost sharing: The role of indifferences. *Games and Economic Behavior*, 82:218–234, 2013.

[15] R. Juarez. Optimal groupstrategyproof cost sharing. Working paper, 2015.

[16] J. Massó, A. Nicolò, A. Sen, T. Sharma, and L. Ülkü. On cost sharing in the provision of a binary and excludable public good. *Journal of Economic Theory*, 155:30–49, 2015.

[17] A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin mechanisms. *Games and Economic Behavior*, 67(1):125–155, 2009.

[18] 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.

[19] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge, 1995.

[20] H. Moulin. Serial cost sharing of excludable public goods. *Review of Economic Studies*, 61:305–325, 1994.

[21] H. Moulin. The price of anarchy of serial, average, and incremental cost sharing. *Economic Theory*, 36(3):379–405, 2008.

[22] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[23] E. Pountourakis and A. Vidali. A complete characterization of group-strategyproof mechanisms of cost-sharing. *Algorithmica*, 63(4):831–860, 2012.

[24] K. Roberts. The characterization of implementable choice rules. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

[25] T. Roughgarden and M. Sundararajan. Optimal efficiency guarantees for network design mechanisms. In *Proceedings of the 12th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 4513 of *Lecture Notes in Computer Science*, pages 469–483, 2007.

[26] T. Roughgarden and M. Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM*, 56(4), 2009. Article 23.

[27] M. Sundararajan. *Trade-Offs in Cost-Sharing*. PhD thesis, Stanford University, 2009.