Pseudonyms in cost-sharing games*

Paolo Penna† Florian Schoppmann‡ Riccardo Silvestri§ Peter Widmayer¶

Abstract

This work initiates the study of cost-sharing mechanisms that, in addition to the usual incentive compatibility conditions, make it disadvantageous for the users to employ pseudonyms. We show that this is possible only if all serviced users pay the *same* price, which implies that such mechanisms do not exist even for certain subadditive cost functions. In practice, a user can increase her utility by lying in one way (misreport her willingness to pay) or another (misreport her identity). We prove also results for approximately budget-balanced mechanisms. Finally, we consider mechanisms that rely on some kind of "reputation" associated to the pseudonyms and show that they are provably better.

First draft: February 2009; **This version:** October 2009.

1 Introduction

Incentives play a crucial role in distributed systems of almost any sort. Typically users want to get resources (a service) without contributing or contributing very little. For instance, file-sharing users in peer-to-peer systems are only interested in downloading data, even though uploading is essential for the system to survive (but maybe costly for some of the users). Some successful systems, like BitTorrent [IUKB+, Coh03], have already incorporated incentive compatibility considerations in their design (users can download data only if they upload some content to others).

One can regard such systems as *cost-sharing mechanisms* in which the overall cost must be recovered from the users in a reasonable and *incentive compatible* manner; the mechanism determines which users get the service and at what price. The study of these mechanisms is a very important topic in economics and in cooperative game theory, in which individuals (users) can coordinate their strategies (i.e., they can form coalitions and collude). Cost-sharing mechanisms should guarantee several essential properties: Users recover the overall cost and are not charged more than necessary (*budget-balance*), even coalitions of users cannot benefit from misreporting their willingness to pay

*Work supported by the European Project IST-15964 "Algorithmic Principles for Building Efficient Overlay Computers" (AEOLUS), by the European Project COST 295 (DYNAMO), and by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). Part of this work has been done while the first author was visiting ETH Zürich. Research done while the second author was at the University of Paderborn, Germany.

†Dipartimento di Informatica ed Applicazioni, Università di Salerno, via S. Allende 2, Baronissi (SA), Italy. Email: penna@dia.unisa.it

‡Department of Computer Science, Stanford University, Stanford, CA, 94305, USA. Email: fschopp@stanford.edu

§Dipartimento di Informatica, Università di Roma "La Sapienza", via Salaria 113, 00198 Roma, Italy. Email: silvestri@di.uniroma1.it

¶Institut für Theoretische Informatik, ETH Zürich, CAB H 15 Universitätstrasse 6, 8092 Zürich, Switzerland. Email: widmayer@inf.ethz.ch.

(*groupstrategyproofness*), no user is excluded a priori (*consumer sovereignty*), no user is charged more than her willingness to pay (*voluntary participation*), and no user receives money (*no positive transfer*). To get a feel of the difficulties of obtaining such mechanisms, consider the following:

Example 1 (identical prices) *There are three players, and the overall cost depends on how many users get the service: Servicing three users costs* 2 *while servicing one or two users costs only* 1*. Consider a mechanism that iteratively drops all users who cannot afford an equal fraction of the total cost (three users pay* $2/3$ *each, two users pay* $1/2$ *each, and one user pays* 1*). Unfortunately, this mechanism is* not even strategyproof[1]*: For valuations* $(0.6, 0.6, 1)$ *the first two users are dropped, so their utility is zero (no service and nothing to pay). If the first user misreports her valuation to 1 she gets the service for price* 0.5 *(now only the second user is dropped), and her utility (valuation minus price to pay* $= 0.6 - 0.5 = 0.1$*) is strictly better than before.*

As Example 1 shows, a groupstrategyproof and budget-balanced mechanism has to charge the serviced users *different* prices in general. In this work, we consider a new form of manipulation that, to the best of our knowledge, has not been considered in the cost-sharing literature before: In most of the Internet applications a user can easily create several virtual identities [Dou02, FR01, YSM04, FPCS04, CF05]. The typical scenario is that there is a universe of possible names (e.g., all strings of up to 40 characters), and only a subset of these correspond to actual users. Each user can replace her name with some *pseudonym*, that is, another name in the universe of names which does not correspond to any other user (e.g., an email account that has not been taken by anybody else). In this new scenario, each user can manipulate the mechanism in *two* ways: by misreporting her name and/or by misreporting her willingness to pay for the service. Ideally, we would like resistance against both kinds of manipulations.

Example 2 (manipulation via pseudonyms) *Consider a different mechanism for Example 1: Order players alphabetically by their name. The first two players bidding at least* $1/2$ *only have to pay* $1/2$*. Otherwise, the price is always 1. This mechanism is group-strategyproof [BMST07]; however, it can be manipulated via pseudonyms. For instance, in the situation*

Names	*Alice*	*Bob*	*Cindy*
Valuations	0.6	0.6	1
Prices	0.5	0.5	1

there is an obvious incentive for Cindy to use a pseudonym "Adam":

Names	*Alice*	*Bob*	*Adam*
Valuations	0.6	0.6	1
Prices	0.5	1	0.5

Now Bob is dropped and the remaining two users get the service for a price of 0.5 *and thus Cindy's utility is better off.*

A natural idea to discourage pseudonyms is to *randomize* the mechanism by ordering the players randomly. In this way users can be motivated to use their own name, provided they trust the fact

[1] A mechanism is *strategyproof* if no single player can benefit from miserporting her willingness to pay (i.e., group-strategyproof for maximum coalition size 1).

that the mechanism picks a truly random shuffle of the (reported) names *and* they cannot guess the randomness before they bid. Unfortunately, the notion of group-strategyproofness is problematic for randomized mechanisms [GH05] (see also Appendix A). Moreover, randomness is generally considered a scarce resource and derandomization of mechanisms is especially difficult because we can run the mechanism ("the auction") only once [AFG+05].

1.1 Our contribution and related work

In our model, the overall cost depends only on the number of users that get the service and the mechanism is deterministic. Call a mechanism *renameproof* if, in the scenario above, no user has an incentive to change her name (i.e., create a pseudonym and obtain the service for a better price). We first prove that for *all* budget-balanced groupstrategyproof mechanisms the following equivalence holds:

$$\text{renameproof} \iff \text{identical prices} \tag{1}$$

By "identical prices" we mean that all serviced users pay the same fraction of the total cost (identical prices imply renameproofness and our contribution is to prove that the converse holds). In some cases, this result gives a *characterization* of such mechanisms, while in other cases it implies that these mechanisms *do not exist*! This is because, even in our restricted scenario, budget-balance and groupstrategyproofness can be achieved only with mechanisms that use *different* prices [BMST07].

We actually prove a more general version of Equivalence (1) which applies also to *approximately* budget-balanced ones. In proving this, we exhibit an intriguing connection with a problem of hypergraph coloring, a linear algebra result by Gottlieb [Got66] and, ultimately, with the Ramsey Theorem (the impossibility of coloring large hypergraphs with a constant number of colors without creating monochromatic components).

The notion of renameproof mechanisms leads naturally to a weaker condition which we call *reputationproof* because it captures the idea that names have some reputation associated, and a user cannot create a pseudonym with a better reputation than that of her true name [CF05]. In some sense, we can regard reputation as a way to "derandomize" some of the mechanisms that are renameproof in expectation. We show that reputation does help because there exist cost functions that admit budget-balanced, groupstrategyproof, and and reputationproof mechanisms, while no budget-balanced mechanism can be simultaneously groupstrategyproof and renameproof (all budget-balanced mechanisms can be manipulated in one way or another). We find it interesting that certain new constructions of mechanisms [BMST07], which were originally introduced to overcome the limitations of identical prices, are reputationproof (though not renameproof).

1.2 Connections with prior work

The notion of renameproof mechanism is weaker than that of *falsenameproof* mechanism [YSM04] where users can submit *multiple* bids to the mechanism, each one under a different name (pseudonym).[2] In the context of combinatorial auctions, where budget-balance is ignored, groupstrategyproofness and falsenemeproofness are independent notions [YSM04]; several constructions of mechanisms that are both strategyproof and falsenameproof are known (see e.g. [YSM04, Yok03]).

[2] For instance, a falsenameproof mechanism for combinatorial auctions guarantees that a user interested in a bundle of items cannot obtain the bundle for a lower price by submitting several bids. Each bid is made under a different identity and for some of the objects for sale.

It is common opinion that manipulations via pseudonyms arise because a single user can gain a large influence on the game by "voting" (bidding) many times [Dou02, CF05, FR01, FPCS04, LMSW06, YSM04]. Indeed, most of the research focuses on solutions that make it impossible (or very difficult) to vote more than once [Con07, WC08, NSSP04]. While this is enough in combinatorial auctions [Con07], manipulations are still possible by declaring false names in cost-sharing games where the mechanism must guarantee budget-balance.

Cost-sharing mechanisms are usually studied in more general settings where costs do not depend only on the number of users. The main technique for obtaining groupstrategyproof mechanisms is due to Moulin [Mou99]. These mechanisms are budget-balanced for arbitrary submodular[3] costs [Mou99]. When costs are not submodular, Moulin mechanisms can achieve only *approximate* budget-balance [JV01, IMM05, KLSvZ08]. An alternative method called *two-price* mechanisms has recently been presented in [BMST07]. The authors showed that, for some problems where the cost depends only on the number of serviced users, groupstrategyproof mechanisms using identical prices cannot be budget-balanced, while their two-price mechanisms are both groupstrategyproof and budget-balanced. This holds for subadditive costs that depend only on the number of users [BMST07]. In [Jua08] it is shown that the so-called *sequential* mechanisms are groupstrategyproof and budget-balanced for supermodular costs. Only partial characterizations of general budget-balanced and groupstrategyproof mechanisms are known [IMM05, PV04, PV06, Jua07, Jua08]. Acyclic mechanisms [MRS07] can achieve budget-balance for all non-decreasing cost-functions by considering a weaker version of groupstrategyproofness (see also [BMS07]).

Finally, [DMRS08] studies the *public excludable good* problem and gives sufficient conditions for which every mechanism must use identical prices. Since this problem is a special case of those studied here, we obtain an alternative axiomatic characterization of identical prices.

Road map The formal model and the definition of renameproof mechanism are given in Section 2. Equivalence (1) and its relaxation to approximately budget-balanced mechanisms are proved in Section 3. Reputationproof mechanisms are defined and analyzed in Section 4. The appendix provides additional material: Appendix A concerns with randomized mechanisms obtained as probability distributions over deterministic ones; Appendix B deals with the design of reputationproof mechanisms based on Moulin mechanisms.

2 The formal model

We consider names as integers taken from a universe $\mathcal{N} = \{1, \ldots, n\}$. A user can register herself under one or more names and two users cannot share the same name. A user who first registered under name i can thus create a number of additional names which we call her *pseudonyms*, while the first name she created is her *true name*. The set of all users (the true names) is private knowledge, that is, the mechanism does not know if a name i is a true name or a pseudonym, and it cannot distinguish if two names have been created by the same user. We assume further that a user cannot use two names simultaneously and thus can only make a *single bid*. For instance, the system may be able to detect that two names correspond to the same IP address and thus to the same user. Moreover, every negative result proved under the single bid assumption also holds when users can make multiple bids.

[3] A cost function is submodular if $C(A \cup B) \leq C(A) + C(B) - C(A \cap B)$.

A *mechanism* for a cost-sharing game is a pair (S, P) defined as follows. The input to the mechanism is an n-dimensional *bid vector* v where the i^{th} coordinate v_i is either $\perp$, indicating that no user submitted her bid using the name i, or it is equal to the bid submitted under the name i. The mechanism outputs the names of the winners and their prices: the user who submitted her bid using a name $i \in S(v)$ receives the service at the price $P(v, i)$; all other users do not get served and do not pay. When user i bids using the name j and her valuation for the service is v_i^*, she derives a *utility* equal to

$$S(v, j) \cdot v_i^* - P(v, j),$$

where $S(v, j)$ is equal to 1 if $j \in S(v)$, and 0 otherwise. We consider only mechanisms (S, P) that satisfy the following standard requirements: *voluntary participation* meaning that $P(v, j) = 0$ if $j \notin S(v)$ and $P(v, j) \leq v_j$ otherwise; *no positive transfer* meaning that prices $P(v, j)$ are always nonnegative; *consumer sovereignty*, meaning that every user can get the service if bidding sufficiently high, regardless of the bids of the other users.

There is a *symmetric cost function* C whose value depends only on the number of users that get the service, that is, $C(S) = C(T)$ whenever S and T are two subsets of names of the same size. Following [BMST07], we speak of a *symmetric cost-sharing game*. A mechanism (S, P) is α*-budget-balanced* if for every bid vector v, it holds that

$$C(S(v)) \leq \sum_{i \in S(v)} P(v, i) \leq \alpha \cdot C(S(v))$$

and it is *budget-balanced* if this condition holds for $\alpha = 1$. We sometimes write C_s in place of $C(S)$, where $s = |S|$.

We require that, in the scenario in which pseudonyms are not used, users cannot improve their utilities by misreporting their valuations. A mechanism (S, P) is *groupstrategyproof* if no group of users can raise the utility of some of its members without lowering the utility of some other member. Mechanism (S, P) is *strategyproof* if this condition is required to hold for G of size one only.

The next desideratum is that, when a user reports truthfully her valuation, she cannot improve her utility when bidding with a pseudonym. For any bid vector b, we let $U(b)$ be the set of names that have been used to submit the bids, that is, those i's such that $b_i \neq \perp$. We also let $b_{i \to j}$ be the vector obtained from b by exchanging b_i with b_j, where $i \in U(b)$ and $j \notin U(b)$.

Definition 3 (renameproof mechanism) *A cost-sharing mechanism* (S, P) *is* renameproof *if the following holds. For any bid vector* v *and for any* i *and* j *such that* $i \in U(v)$ *and* $j \notin U(v)$

$$S(v, i) \cdot v_i - P(v, i) \geq S(v_{i \to j}, j) \cdot v_i - P(v_{i \to j}, j).$$

We think of v as the bid vector in which user i reports truthfully her valuation and her name. Thus, the above definition says that no user can improve her utility by using a pseudonym in place of her name, no matter if the other users report truthfully their names and valuations. Symmetry implies that, in a renameproof mechanism, the utility of an agent must be constant over all names. Note that this does not imply that the price must be the same for all serviced users (see Example 9 below). We decide to present the definition in this form since this will naturally lead to a weaker condition used by mechanisms based on "reputation" in Section 4.

3 Renameproof mechanisms and (non-) identical prices

A general approach to design (approximately) budget-balanced mechanisms is to define a suitable *cost-sharing method* for the cost function of the problem. That is, a function ξ which (approximately) divides the overall cost among the serviced users: $\xi(S,i)$ is the price associated to user i when the subset S is served, and

$$1 \leq \sum_{i \in S} \xi(S,i)/C(S) \leq \alpha,$$

for $\alpha \geq 1$ and for every subset S of users. Such a function is an α-*budget-balanced* cost-sharing method (or simply budget-balanced if $\alpha = 1$). The main technique to construct groupstrategyproof cost-sharing mechanisms is due to Moulin [Mou99]:

Example 4 (Moulin mechanism M_ξ) *Initially we set S as all users in $U(b)$. At each round we remove all users in S whose bid is less than the price $\xi(S,i)$ offered by the mechanism. We iterate this step until all users in the current set S accept the offered price, or no user is left. We service the final set S obtained in this way and charge each user $i \in S$ an amount $\xi(S,i)$.*

We first observe that renameproofness by itself is not a problem since the *average-cost* mechanism [MS01] which charges all users the same price is budget-balanced and renameproof. However, this mechanism is (group) strategyproof only when the average cost does not increase with the number of users (see Example 1). Therefore, it is natural to ask if there exist other mechanisms that are renameproof. To answer this question, we consider a general class of mechanisms that define a unique cost-sharing method and that are known to include *all* groupstrategyproof mechanisms [Mou99]. Such mechanisms are termed *separable*:

Definition 5 (separable mechanism) *A cost-sharing mechanism (S,P) is* separable *if it induces a unique cost-sharing method $\xi = \xi_{(S,P)}$. That is, there exists a cost-sharing method $\xi = \xi_{(S,P)}$ such that $P(v,i) = \xi(S(v),i)$ for all bids v and for all i.*

Our first observation is that, in order to be renameproof, the price assigned to a user should be independent from her name:

Definition 6 *A cost-sharing method ξ is* name independent *if the price associated to a user does not depend on her name. That is, for every $R \subset \mathcal{N}$ and for any two $i, j \notin R$, it holds that $\xi(R \cup \{i\}, i) = \xi(R \cup \{j\}, j)$.*

Theorem 7 *Every separable renameproof mechanism induces a name independent cost-sharing method.*

Proof. Let (S,P) be a separable renameproof mechanism with cost-sharing method ξ. By contradiction, suppose ξ is not name-independent, i.e., there is a set $R \subset \mathcal{N}$ and $i, j \notin R$ with $\xi(R \cup \{i\}, i) > \xi(R \cup \{j\}, j)$. Recall that due to consumer sovereignty, there is some value ℓ so that every user bidding at least ℓ gets the service, independent of all other users' bids. Consider the vector v defined by $v_k = \ell$ if $k \in R \cup \{i\}$ and $v_k = \perp$ if $k \notin R$. Then,

$$S(v,i) \cdot v_i - P(v,i) = v_i - \xi(R \cup \{i\}, i) < v_i - \xi(R \cup \{j\}, j) = S(v_{i \to j}, j) \cdot v_i - P(v_{i \to j}, j)$$

which contradicts the fact that the mechanism is renameproof (see Definition 3). □

Because of the previous result, we will focus on name independent cost-sharing methods. The main idea is to regard each name-independent cost-sharing method as a "fractional coloring" of the complete hypergraph on n nodes.

Definition 8 *An* α-balanced (n, s)-coloring *is a function* x *assigning a nonnegative weight* x_R *to every* $(s-1)$*-subset* [4] $R \subset \mathcal{N}$ *such that, for every* s*-subset* $S \subseteq \mathcal{N}$*, it holds that*

$$1 \leq \sum_{R \in subsets(S)} x_R \leq \alpha \tag{2}$$

where $subsets(S)$ *denotes the family of all* $(s-1)$*-subsets of* S*.*

The connection between a name-independent ξ and (n, s)-colorings can be made explicit by considering

$$x_R^{(\xi,s)} := \xi(R \cup \{i\}, i)/C_s \tag{3}$$

and by observing that $x^{(\xi,s)}$ must be α-balanced if ξ is α-budget-balanced. With this equivalent representation, we can argue about the existence of cost-sharing mechanisms using non-identical prices by studying "non-uniform" (n, s)-colorings.

Example 9 (three users) *When servicing three users, Definition 8 boils down to the problem of assigning weights to the edges of the complete graph over* n *nodes so that every triangle has total weight in between* 1 *and* α*. For* $n = 4$*, one possible way is as follows:*

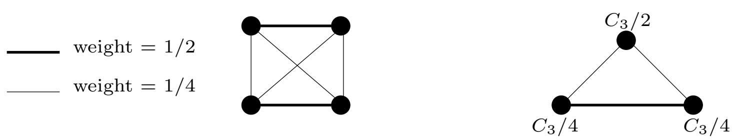

where on the right we show the corresponding prices (cost-sharing scheme): the names of the three serviced users correspond to a triangle as above. The triangle tells us that one user pays $1/2$ *of the total price and the other two pay* $1/4$ *each. To see that this scheme is name independent, observe that the price of a user depends only on the edge connecting the other two vertices of the triangle (that is the names of the other two users). Let us extend the scheme by assigning identical prices if two or one user get the service. For the cost function* $C_1 = C_2 = 1$ *and* $C_3 = 2$*, used in Example 1, the cost-sharing method corresponds to a "non-Moulin" two-price mechanism in [BMST07]. This mechanism is budget-balanced groupstrategyproof and* renameproof *for* $n = 4$ *and for three users.*

Despite the above example, our main result below says that one cannot avoid identical prices as soon as the number of possible names is not very small:

Theorem 10 *Any* α*-budget-balanced renameproof separable mechanism must charge each serviced user a price which is at least* $\frac{C_s}{s}\left(2^{s-1}(1-\alpha)+\alpha\right)$ *and at most* $\frac{C_s}{s}\left(2^{s-1}(\alpha-1)+1\right)$*, unless the number* s *of serviced users is more than half the number of names (i.e, unless* $s > n/2$*). In particular, if the mechanism is budget-balanced (i.e.,* $\alpha = 1$*) then all serviced users are charged the same price* C_s/s*.*

[4] An r-subset is a subset of cardinality r.

Before proving the theorem, we show its implications. Notice that every groupstrategyproof mechanism must be separable [Mou99]. Since we have just shown that (a) every budget-balanced renameproof separable mechanism is an average-cost mechanism, but since we also know (recall Example 1) that (b) average-cost mechanisms are not groupstrategyproof if costs are not submodular, we obtain the following characterization and impossibility result:

Corollary 11 *If the number of possible names is at least twice the number of users, then the following holds. If costs are submodular, then the Moulin mechanism charging all players equally is the only renameproof, budget-balanced and groupstrategyproof mechanism (up to welfare-equivalence).*[5] *When costs are not submodular, there is no mechanism that is at the same time renameproof, groupstrategyproof and budget-balanced.*

We note that the bounds of Theorem 10 become rather weak as s grows (and α is fixed). It will be clear from the proof that such bounds are tight. In Section 3.2 we show a different kind of bounds that essentially rule out constructions of mechanisms based on different prices.

3.1 Proof of Theorem 10

We prove bounds on α-balanced (n, s)-colorings. Let us consider the following *incidence* matrix A with $\binom{n}{s}$ rows and $\binom{n}{s-1}$ columns (originally defined in [Got66]). For any s-subset (row) S and any $(s-1)$-subset (column) R, define:

$$A_{S,R} = \begin{cases} 1 & \text{if } R \subset S, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

By definition

$$A_S \cdot x = \sum_{R \in subsets(S)} x_R$$

for any s-subset S. Thus the set of all α-balanced (n, s)-colorings is given by the following polytope:

$$\mathbf{polytope} := \{x \in \mathbb{R}^q | \ \mathbf{1} \leq A \cdot x \leq \boldsymbol{\alpha}\}$$

where $q = \binom{n}{s-1}$ and $\boldsymbol{\alpha} = (\alpha, \ldots, \alpha)$ for every real α. Intuitively, each x_R corresponds to the price $\xi(S, i)/C_s$ for $R = S \setminus \{i\}$, which is independent of i.

The crucial point is that the minimum/maximum value of any component of x in **polytope** can be found by looking only at vectors with a "nice structure":

Lemma 12 *For every x in* **polytope** *and for every $(s-1)$-subset $R \subset \mathcal{N}$ there exists $\overline{x}$ in* **polytope** *such that $x_R = \overline{x}_R$ and all the values $\overline{x}_Q$ depend only on the size of $R \cap Q$. That is, $\overline{x}_Q = \overline{x}_{Q'}$ for all $(s-1)$-subsets Q and Q' such that $|Q \cap R| = |Q' \cap R|$.*

Proof. Consider the set Π_R of all permutations $\pi : \mathcal{N} \to \mathcal{N}$ satisfying $\pi(R) = R$. Then, for every π in Π_R, the vector x^π defined by $x^\pi_Q = x_{\pi(Q)}$ for all $(s-1)$-subsets Q satisfies $x^\pi_R = x_R$ by definition.

[5] Two mechanisms are welfare-equivalent if they produce the same utilities.

We next show that $x^\pi \in$ **polytope**. For any s-subset S we have that $\pi(Q) \subseteq \pi(S) \Leftrightarrow Q \subseteq S$. Hence, we have the following identities:

$$\begin{aligned}\sum_{Q \text{ is } (s-1)\text{-subset}} A_{S,Q} \cdot x^\pi_Q &= \sum_{Q \text{ is } (s-1)\text{-subset}} A_{S,Q} \cdot x_{\pi(Q)} = \sum_{Q \text{ is } (s-1)\text{-subset}} A_{\pi(S),\pi(Q)} \cdot x_{\pi(Q)} \\ &= \sum_{Q \text{ is } (s-1)\text{-subset}} A_{\pi(S),Q} \cdot x_Q\end{aligned}$$

where the last equality follows from the observation that each $\pi(Q)$ is an $(s-1)$-subset and π is a bijection. Now observe that the first and the last summations above are $A_S \cdot x^\pi$ and $A_{\pi(S)} \cdot x$, respectively. Since x is in **polytope** we have $A_{\pi(S)} \cdot x \in [1, \alpha]$ and thus x^π is also in **polytope**. Now we can define $\overline{x}$ as the following convex combination of elements x^π with $\pi \in \Pi_R$:

$$\overline{x} := \frac{1}{|\Pi_R|} \cdot \sum_{\pi \in \Pi_R} x^\pi.$$

Since each x^π is in **polytope** and since the latter is a convex polytope, we have that $\overline{x}$ is also in **polytope**. Moreover, $\overline{x}_R = x_R$ because $x^\pi_R = x_R$ for all $\pi \in \Pi_R$. Finally, due to symmetry, we have $\overline{x}_Q = \overline{x}_{Q'}$ for all $(s-1)$-subsets Q, Q' with $|R \cap Q| = |R \cap Q'|$. □

Lemma 13 *For any x in* **polytope** *and for any $(s-1)$-subset R, it holds that* $\frac{2^{s-1}(1-\alpha)+\alpha}{s} \le x_R \le \frac{2^{s-1}(\alpha-1)+1}{s}$.

Proof. Let $x \in$ **polytope** and R be any $(s-1)$-subset. Because of Lemma 12 it is enough to prove the bound for the vector $\overline{x}$ because $x_R = \overline{x}_R$. Let p_k denote the unique value such that $x_Q = p_k$ for all Q such that $|R \cap Q| = k$. We are interested in the value of $p_{s-1} = \overline{x}_R$. We first prove that, for every s-subset S with $k := |S \cap R|$ it holds that

$$\sum_{Q \text{ is } (s-1)\text{-subset}} A_{S,Q} \cdot \overline{x}_Q = k \cdot p_{k-1} + (s-k) \cdot p_k. \tag{5}$$

Observe that $Q \subset S$ implies that $|Q \cap R|$ is either $k-1$ or k because Q consists of all but an excluded element from S: either one in $S \cap R$ or one in $S \setminus R$, respectively. Therefore, the number of $(s-1)$-subsets Q with $|R \cap Q| = k-1$ is precisely k, while the number of $(s-1)$-subsets Q with $|R \cap Q| = k$ is precisely $s-k$. This proves (5).

Now let $b_k := k \cdot p_{k-1} + (s-k) \cdot p_k$ for $k = 0, \ldots, s-1$. Since $2s \le n$, for every $k = 0, \ldots, s-1$ there exists S such that $|S \cap R| = k$. From (5) and from the fact that $\overline{x} \in$ **polytope** we have that $b_k \in [1, \alpha]$. Notice that

$$p_0 = \frac{b_0}{s} \qquad \text{and} \qquad p_k = \frac{b_k - k \cdot p_{k-1}}{s-k} \qquad \text{for } k = 1, \ldots, s-1. \tag{6}$$

By eliminating the recursion from (6) we have in particular

$$\begin{aligned} p_{s-1} &= b_{s-1} - (s-1) \cdot p_{s-2} \\ &= b_{s-1} - \frac{s-1}{2} \cdot \left(b_{s-2} - \frac{s-2}{3} \cdot \left(b_{s-3} - \frac{s-3}{4} \cdot \left(\cdots \left(b_1 - \frac{1}{s} \cdot b_0\right)\right)\right)\right) \\ &= \sum_{i=0}^{s-1} (-1)^{s-i-1} \cdot b_i \cdot \frac{(s-1)\cdots(i+1)}{2\cdots(s-i)} = \sum_{i=0}^{s-1} \frac{(-1)^{s-i-1} \cdot b_i}{s-i} \cdot \binom{s-1}{s-i-1} \end{aligned}$$

To derive upper and lower bounds on $p_{s-1} = \overline{x}_R = x_R$ we show that its maximum and minimum are achieved for "alternating" sequences of the form $b_i = \alpha, b_{i+1} = 1, b_{i+2} = \alpha, b_{i+3} = 1$, etc. Note that p_{s-1} is increasing is b_{s-1}, decreasing in b_{s-2}, and so on. So p_{s-1} is maximized for $b_{s-1} = \alpha$, $b_{s-2} = 1$, $b_{s-3} = \alpha$, $b_{s-4} = 1$, and so on, i.e., for $b_i = (1 + \alpha + (\alpha - 1)(-1)^{s-i})/2$. By plugging in these values we obtain

$$\begin{aligned} p_{s-1} &\geq \sum_{i=0}^{s-1} \frac{(1+\alpha)\cdot(-1)^{s-i-1} - (\alpha-1)}{2\cdot(s-i)} \cdot \binom{s-1}{s-i-1} \\ &= -\frac{1+\alpha}{2} \cdot \sum_{k=0}^{s-1} \frac{(-1)^{k+1}}{k+1}\binom{s-1}{k} + \frac{1-\alpha}{2} \cdot \sum_{k=0}^{s-1} \frac{1}{k+1}\binom{s-1}{k} \\ &= \frac{1+\alpha}{2s} + \frac{(1-\alpha)\cdot(2^s-1)}{2s} = \frac{2^{s-1}(1-\alpha)+\alpha}{s}, \end{aligned}$$

where in the second equality we have simply substituted $s - i - 1$ with k, and in the last equality we have applied the following well-known identity (see e.g. [GR07, Equation 0.155])

$$\sum_{k=0}^{m} \frac{\lambda^{k+1}}{k+1}\binom{m}{k} = \frac{(\lambda+1)^{m+1} - 1}{m+1}, \qquad \text{for every } \lambda \in \mathbb{R}.$$

We prove the upper bound in a similar way: p_{s-1} is maximized for $b_{s-1} = \alpha$, $b_{s-2} = 1$, $b_{s-3} = \alpha$, $b_{s-4} = 1$, etc. That is, for $b_k = (1 + \alpha + (\alpha - 1) \cdot (-1)^{s-k-1})/2$ and thus

$$\begin{aligned} p_{s-1} &\leq \sum_{i=0}^{s-1} \frac{(1+\alpha)\cdot(-1)^{s-i-1} + \alpha - 1}{2\cdot(s-i)} \cdot \binom{s-1}{s-i-1} \\ &= \frac{1+\alpha}{2} \cdot \sum_{i=0}^{s-1} \frac{(-1)^{i}}{i+1}\binom{s-1}{i} + \frac{\alpha-1}{2} \cdot \sum_{i=0}^{s-1} \frac{1}{i+1}\binom{s-1}{i} \\ &= \frac{1+\alpha}{2s} + \frac{(\alpha-1)\cdot(2^s-1)}{2s} = \frac{2^{s-1}(\alpha-1)+1}{s} \end{aligned}$$

□

From this lemma, Theorem 7, and using the relation in (3) we obtain Theorem 10.

3.2 More bounds for approximate budget-balance

The bounds of Theorem 10 apply no matter which subset of names, S, is serviced. In the worst case the bound can be strengthened. That is, for *some* S the prices need to be almost identical:

Theorem 14 *For any α, for any s, and for any δ, there exists $N = N(\alpha, s, \delta)$ such that the following holds for all $n \geq N$. For every α-budget-balanced name-independent cost-sharing method ξ there exists a subset S of s users such that their prices satisfy*

$$|\xi(S, i) - \xi(S, i')| \leq \delta$$

for every i and i' in S.

Proof. We make use of the following result (see e.g. [Juk01]):

Ramsey Theorem (1930). *Let c, r, l be given positive integers, $l \geq r$. Then there is a number $n = R_c(r; l)$ with the following property. If r-subsets of an n-set are colored with c colors, then there is an l-set all of whose r-subsets have the same color.*

The function $x^{(\xi,s)}$ is an α-balanced (n, s)-coloring. Consider the (n, s)-coloring x obtained by rounding each $x_R^{(\xi,s)}$ to the closest power of δ', for suitable $\delta' > 0$. Since every $x_R^{(\xi,s)}$ is in the interval $[0, \alpha]$, the function x takes at most $c = \lceil \log_{\delta'} \alpha \rceil$ different values. Moreover, if $x_R = x_{R'}$ then $\left| x_R^{(\xi,s)} - x_{R'}^{(\xi,s)} \right| \leq \delta'$. The Ramsey Theorem says that, for $n \geq R_c(s-1; s)$, there is an s-subset S such that, for any two $(s-1)$- subsets R and R' of S, it holds that $x_R = x_{R'}$. Hence, for any i and i' in S, we have

$$|\xi(S, i) - \xi(S, i')| = \left| C_s \cdot x_R^{(\xi,s)} - C_s \cdot x_{R'}^{(\xi,s)} \right| \leq C_s \delta'$$

where R and R' are the subsets of S obtained by removing i and i', respectively. The theorem thus follows by setting $\delta' := \delta / C_s$ and $N(\alpha, s, \delta) := R_c(s-1; s)$, where $c = \lceil \log \alpha / \log \delta' \rceil$. □

3.3 Two applications

We present two applications of our results. The first one is the *public excludable good* problem, which plays a central role in the cost-sharing literature and arises as a special case of many optimization problems [DMRS08]. It corresponds to the simple cost function $C_0 = 0$ and $C_s = 1$ for all $s \geq 1$. Corollary 11 implies that the Shapley value mechanism[6] is *the only* deterministic one that meets budget-balance, groupstrategyproofness, and renameproofness. This gives an alternative axiomatic characterization of the Shapley value to the one in [DMRS08, Theorem 2].

The second application concerns certain *bin packing* or *scheduling* problems, in which users (items or jobs) are homogeneous [BM06, BMST07]. For instance, costs are of the form $C_s = \lceil s/m \rceil$, where m can be interpreted as the capacity of the bins (bins are identical and the cost is the number of bins needed) or the number of machines (the cost is the makespan of the computing facility offering the service). For these problems, Bleischwitz *et al* [BMST07] proved that budget-balanced groupstrategyproof mechanisms must use *non-identical* prices. Corollary 11 implies that deterministic budget-balanced mechanisms which are both groupstrategyproof and renameproof *do not exist* (these two notions are *mutually exclusive*). Theorem 14 essentially rules out constructions with non-identical prices, and identical prices can only achieve 2-approximate budget-balance.

4 Reputation helps

In this section we consider the scenario in which the mechanism has some additional information about the users. For instance, there might be *time stamps* for the times at which new identities have been created associated to users. Similarly, we might employ *reputation functions* that make it impossible for a user to obtain a better ranking [CF05]. An abstraction of these two scenarios is to consider the case in which each user can replace her name only with a "larger" one.

[6] The Shapley value mechanism services the largest set S such that all users in S bid at least $1/|S|$.

Definition 15 (reputationproof mechanism) *A cost-sharing mechanism is* reputationproof *if it satisfies the condition of Definition 3 limited to $j > i$.*

We show a *sufficient* condition for obtaining mechanisms that are both groupstrategyproof and reputationproof. Here the cost-sharing method can be represented by an $n \times n$ matrix $\xi = \{\xi_i^s\}$, where each ξ^s is a vector of s prices. This is an example of the kind of the cost-sharing schemes we use here:

$$\xi = \begin{pmatrix} 1/2 & 1/2 & 1/2 & 1/2 \\ 1/2 & 1/2 & 1 & \\ 1/2 & 1/2 & & \\ 1 & & & \end{pmatrix} \tag{7}$$

where the scheme is regarded as a triangular matrix in which we leave empty the (irrelevant) values ξ_i^s with $i > s$. Users are ranked according to their (reported) names and the highest price is offered to the user with largest (reported) name.

Definition 16 *Let $\xi = \{\xi_j^s\}$ be a cost-sharing scheme containing only two values low and $\mathsf{high} > \mathsf{low}$, and say that ξ^s* offers low prices *if it is equal to $(\mathsf{low}, \ldots, \mathsf{low})$. Such a cost-sharing scheme ξ is a* same-two-price *scheme if the following holds for every s. If ξ^s does not offer low prices, then ξ^{s-1} offers low prices and $\xi^s = (\mathsf{low}, \ldots, \mathsf{low}, \mathsf{high})$.*

The main idea in [BMST07] is to drop some "indifferent" users who bid the low price:

Definition 17 *Let ξ be a* same-two-price *scheme and call* indifferent *those users bidding exactly the low price. The corresponding* same-two-price mechanism *is as follows:*

1. *Drop all users who do not even bid the low price;*
2. *If the resulting subset has size s such that ξ^s offers low prices, then service this set and charge all these users the low price;*
3. *Otherwise do the following:*
   (a) *If there are indifferent users, then drop the one with the last name and charge all others the low price;*
   (b) *Otherwise, drop the last user if she is not willing to pay the high price and charge all others the same price;*
   (c) *If no user has been dropped in Steps (3a-3b) then charge the low price to all but the last user who is charged the high price.*

Theorem 18 *Every same-two-price mechanism is both groupstrategyproof and reputationproof.*

Proof. The same-two-price mechanisms is a special case of two-price mechanisms in [BMST07] in which we have

1. All high and low prices (cost shares in [BMST07, Definition 5]) are the same.
2. The number of users paying the high price is at most one.

Together, these two conditions trivially satisfy the definition of valid two-price cost-sharing forms [BMST07, Definition 6]. Since the two-price mechanisms in [BMST07] are groupstrategyproof for all valid two-price cost-sharing form, the same holds for our special case of same-two-price mechanisms.

Now we prove that the same-two-price mechanism is also reputationproof. By contradiction, suppose there exists v, i and $j > i$ such that

$$S(v,i) \cdot v_i - P(v,i) < S(v_{i\to j}, j) \cdot v_i - P(v_{i\to j}, j). \tag{8}$$

An invariant of the mechanism is that the utility of every user whose valuation is at most the low price is 0 because she is either excluded or she is serviced paying the low price (this case arises only if her valuation is low). So, we have to consider only two cases:

($v_i > \mathsf{high}$) Any user bidding more than v_i is serviced by the same-two-price mechanism. Hence, the above inequality implies $P(v,i) = \mathsf{high} < P(v_{i\to j}, j) = \mathsf{low}$. Since in v, i is charged the high price, it means that the mechanism enters Step 3a and there are no indifferent users. So the mechanism enters Step 3b and i is the largest name. Now observe that the same holds for vector $v_{i\to j}$ because $j > i$. Indeed, also in this case we enter Step 3a and, because there are no indifferent users, we enter Step 3b. Here, the user with the highest name is "j" because of $j > i$. Hence, $P(v_{i\to j}, j) = \mathsf{high}$, thus a contradiction.

($\mathsf{low} < v_i \leq \mathsf{high}$) If the mechanisms does not enter Step 3b, then i is served for the low price and thus gets the highest possible utility (no user gets served for a better price). If the mechanisms enters Step 3b for v, then the same arguments as in the previous case show that it also enters Step 3b for $v_{i\to j}$, and j is the largest name. For $v_i < \mathsf{high}$, the user in question is excluded in both cases, while for $v_i = \mathsf{high}$ she get served for the high price in both cases. This contradicts (8).

□

Example 19 (two-machine scheduling) *The cost of scheduling s identical jobs on two identical machines is the makespan, that is, $C_s = \lceil s/2 \rceil$. No budget-balanced mechanism which is groupstrategyproof can be renameproof (see Section 3). In contrast, there exists a budget-balanced groupstrategyproof mechanism which is also reputationproof: The same-two-price mechanism corresponding to the cost-sharing methods of the form (7).*

Unfortunately, not all two-price mechanisms from [BMST07] are reputationproof. The argument is similar to the discussion in Appendix B about Moulin mechanisms. Concerning the notion of reputationproof mechanism, we note that BitThief [LMSW06] is a sophisticated client that free rides on BitTorrent using a weakness of the protocol: newcomers are allowed to download data "for free" because they are supposed to have nothing to upload yet. BitTorrent resembles a mechanism that is *not* reputationproof because it offers a better price to users that have "no reputation".

5 Concluding remarks and open questions

Despite the fact that in symmetric cost-sharing games all users play the same role, in most of the cases one must employ *non-identical* prices in order to get budget-balanced and groupstrategyproof mechanisms [BMST07]. In sharp contrast, we have shown that in order to make it disadvantageous for users to use pseudonyms, one must use *identical* prices, thus implying that groupstrategyproofness and renameproofness can be achieved only *separately*. The results apply also to *randomized* mechanisms in which using the "true" name is a dominant strategy for all coin tosses. The notion of reputationproof mechanism captures in a natural way the use of "reputation" to overcome these difficulties.

It would be very interesting to characterize the class of "priority-based" mechanisms (sequential [Jua08], acyclic [MRS07], two-price [BMST07]) that are also reputationproof. Another important issue is to consider multiple bids (falsenameproofness [YSM04, Yok03]) and consider non-symmetric ("combinatorial") cost-sharing games [JV01, IMM05, RS06].

An interesting issue concerns the "social cost" of pseudonyms in cost-sharing games. For instance, we might consider the economic *efficiency loss* [RS06] caused by the use of pseudonyms (namely, the efficiency of arbitrary mechanisms versus those that use identical prices) and thus quantify how much "reputation" helps.

Acknowledgements. We wish to thank Roger Wattenhofer for several useful discussions, and to Carmine Ventre for comments on an earlier version of this work. We are grateful to an anonymous referee for many insightful comments and, in particular, for suggesting the use of negative bids to simplify some of the proofs.

References

[AFG+05] G. Aggarwal, A. Fiat, A.V. Goldberg, J.D. Hartline, N. Immorlica, and M. Sudan. Derandomization of auctions. In *Proc. of the 37th ACM Symposium on Theory of Computing (STOC)*, pages 619–625, 2005.

[BM06] Y. Bleischwitz and B. Monien. Fair cost-sharing methods for scheduling jobs on parallel machines. In *Proc. of the 6th International Conference on Algorithms and Complexity (CIAC)*, volume 3998 of *LNCS*, pages 175–186, 2006.

[BMS07] Y. Bleischwitz, B. Monien, and F. Schoppmann. To be or not to be (served). In *Proc. of the 3rd Workshop on Internet & Network Economics (WINE)*, volume 4858 of *LNCS*, pages 515–528, 2007.

[BMST07] Y. Bleischwitz, B. Monien, F. Schoppmann, and K. Tiemann. The power of two prices: Beyond cross-monotonicity. In *Proc. of the 32nd Int. Symp. on Mathematical Foundations of Computer Science (MFCS)*, volume 4708 of *LNCS*, pages 657–668, 2007.

[CF05] A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *Proc. of the 3rd Workshop on Economics of Peer-to-Peer Systems (P2PECON)*, pages 128–132, 2005.

[Coh03] B. Cohen. Incentives Build Robustness in BitTorrent. In *Proc. of the 1st Workshop on Economics of Peer-to-Peer Systems (P2PECON)*, 2003.

[Con07] V. Conitzer. Limited verification of identities to induce false-name-proofness. In *Proc. of the 11th conference on Theoretical aspects of rationality and knowledge (TARK)*, pages 102–111, New York, NY, USA, 2007. ACM.

[DMRS08] S. Dobzinski, A. Mehta, T. Roughgarden, and M. Sundararajan. Is Shapley Cost Sharing Optimal? In *Proc. of the 1st Int. Symp. on Algorithmic Game Theory (SAGT)*, volume 4997 of *LNCS*, pages 327–336, 2008.

[Dou02] J. R. Douceur. The sybil attack. In *Proc. of the 1st Int. Workshop on Peer-to-Peer Systems (IPTPS)*, pages 251–260, 2002.

[FPCS04] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. In *Proc. of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems (PINS)*, pages 228–236, 2004.

[FR01] E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

[GH05] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proc. of the 16th ACM Symposium on Discrete Algorithms (SODA)*, pages 620–629, 2005.

[Got66] D.H. Gottlieb. A certain class of incidence matrices. *Proc. Amer. Math. Soc.*, pages 1233–1237, 1966.

[GR07] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series, and Products*. Academic-Press, 7th edition, 2007.

[IMM05] N. Immorlica, M. Mahdian, and V. S. Mirrokni. Limitations of cross-monotonic cost sharing schemes. In *Proc. of the 16th ACM Symposium on Discrete Algorithms (SODA)*, pages 602–611, 2005.

[IUKB$^+$] M. Izal, G. Urvoy-Keller, E. Biersack, P. Felber, A. Hamra, and L. Garces-Erice.

[Jua07] R. Juarez. Collusion-proof cost sharing. Technical report, Mimeo University of Hawaii, 2007.

[Jua08] R. Juarez. Group strategyproof cost sharing: the role of indifferences. Technical report, Mimeo University of Hawaii, 2008.

[Juk01] Stasys Jukna. *Extremal Combinatorics With Applications in Computer Science*. Springer-Verlag, 2001.

[JV01] K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. of the 33rd ACM Symposium on Theory of Computing (STOC)*, pages 364–372, 2001.

[KLSvZ08] J. Könemann, S. Leonardi, G. Schäfer, and S. H. M. van Zwam. A group-strategyproof cost sharing mechanism for the steiner forest game. *SIAM J. Comput.*, 37(5):1319–1341, 2008.

[LMSW06] T. Locher, P. Moor, S. Schmid, and R. Wattenhofer. Free Riding in BitTorrent is Cheap. In *Proc. of the 5th Workshop on Hot Topics in Networks (HotNets), Irvine, California, USA*, November 2006.

[Mou99] H. Moulin. Incremental cost sharing: characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.

[MRS07] A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin Mechanisms. In *Proc. of the 8th ACM Conference on Electronic Commerce (EC)*, pages 1–10, 2007.

[MS01] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[NSSP04] J. Newsome, E. Shi, D. Song, and A. Perrig. The sybil attack in sensor networks: analysis & defenses. In *Proc. of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 259–268, 2004.

[PV04] P. Penna and C. Ventre. More powerful and simpler cost-sharing methods. In *Proc. of the 2nd International Workshop on Approximation and Online Algorithms (WAOA)*, volume 3351 of *LNCS*, pages 97–110, 2004.

[PV06] P. Penna and C. Ventre. The algorithmic structure of group strategyproof budget-balanced cost-sharing mechanisms. In *Proc. of the 23rd Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3884 of *LNCS*, pages 337–348, 2006.

[RS06] T. Roughgarden and M. Sundararajan. New trade-offs in cost-sharing mechanisms. In *Proc. of the 38th ACM Symposium on Theory of Computing (STOC)*, pages 79–88, 2006.

[WC08] L. Wagman and V. Conitzer. Optimal false-name-proof voting rules with costly voting. In *Proc. of the 23rd AAAI Conference on Artificial Intelligence*, pages 190–195, 2008.

[Yok03] M. Yokoo. Characterization of Strategy/False-name Proof Combinatorial Auction Protocols: Price-oriented, Rationing-free Protocol. In *Proc. of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, volume 18, pages 733–742, 2003.

[YSM04] M. Yokoo, Y. Sakurai, and S. Matsubara. The Effect of False-name Bids in Combinatorial Auctions: New Fraud in Internet Auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

A Randomization does not preserve groupstrategyproofness

It has been observed in [GH05] that a randomized mechanism which is obtained as a probability distribution of two groupstrategyproof (deterministic) mechanisms is no longer groupstrategyproof. This is also true for cost-sharing games, as shown by the following example. Consider the following cost-sharing method

$$\xi = \begin{pmatrix} 1 & 1 & 2 \\ 1 & 2 & \\ 1 & & \end{pmatrix}$$

This is in fact a "valid" cost-sharing method as defined in [BMS07] and it specifies a two-price groupstrategyproof cost-sharing mechanism. The mechanism breaks ties in a way to ensure groupstrategyproofness, but when there are no ties it will (essentially) act as a Moulin mechanism. In particular, for valuations $v = (2 - \varepsilon, 2 - \varepsilon, 2 - \varepsilon)$ it will only serve the first user for a price 1 and her utility will be $1 - \varepsilon$. The utility of the other two is zero. Now consider a randomized version in which we randomize over all 3! orders (of the names). Then the expected utility of each user is $\frac{1-\varepsilon}{3}$. However, if the three users form a coalition in which they bid $b = (2 + \varepsilon, 2 + \varepsilon, 2 + \varepsilon)$, each of them gets an expected utility equal to $\frac{2}{3} - \varepsilon$. That is, the randomized mechanisms is not groupstrategyproof in expectation.

B Reputationproof mechanisms based on Moulin mechanisms

To demonstrate that also renameproofness imposes limitations on the existing constructions of mechanisms, we consider a simple instantiation of the Moulin mechanism [MS01]:

Definition 20 (rank-based Moulin mechanism M_ξ.) *Initially we set S as all users in $U(b)$. At each round the mechanism offers a price ξ_i^s to the user with the i^{th} name in S, where $s = |S|$. Moreover, we remove from S the user with the largest name among those whose bid is less than the offered price. We iterate this step until all users in the current set S accept the offered price, or no user is left. We service the final set S obtained in this way and charge each user according to ξ (that is, the user with the i^{th} name in S pays ξ_i^s).*

A rule of thumb for obtaining reputationproof mechanisms is that the mechanism should not favor "newcomers" or users that have a "worse reputation". In order to guarantee groupstrategyproofness, we require that each user is offered a price which never decreases if another user is dropped. These two conditions are formally stated here:

Definition 21 *A rank-based cost-sharing method ξ is* rank monotonic *if the price offered to the i^{th} user is at most the price offered to the $(i+1)^{th}$ user (i.e., $\xi_i^s \leq \xi_{i+1}^s$ for all $1 \leq i < s \leq n$). A rank-based cost-sharing method ξ is* cross monotonic *if the price offered to the i^{th} user does not decrease if one user is dropped (i.e., $\xi_i^{s+1} \leq \xi_i^s$ for all $1 \leq i \leq s < n$, and $\xi_{i+1}^{s+1} \leq \xi_i^s$ for all $1 \leq i \leq s < n$).*

It is quite simple to show that these two conditions are necessary for obtaining a rank-based Moulin mechanism which is reputationproof and groupstrategyproof, respectively. Quite surprisingly, the rank monotonicity condition does not ensure that the mechanism is reputationproof.

Consider the following simple example:

$$\xi = \begin{pmatrix} 1 & 2 \\ 3 & \end{pmatrix}. \quad (9)$$

It turns out that user 1 can obtain a better price by submitting her valuation under the pseudonym 3. Indeed, for a vector $v = (3, 1, \perp)$, the mechanism drops user 2 and services user 1 for a price of 3. On the contrary, for $v_{1\to 3} = (\perp, 1, 3)$, the mechanism services both users since, in the first round, user 2 is offered the lower price (1 instead of 2) and accepts.

The above example can be generalized as follows:

Definition 22 (offers three prices) *A rank-based cost-sharing method* ξ offers three prices *if it is both rank and cross monotonic, and there are* $\xi_i^s < \xi_{i+1}^s < \xi_1^{s-i}$*, with* $1 \le i < s \le n$*.*

We stress that the above definition is somewhat orthogonal to the *two-prices* mechanisms in [BMST07] in which the serviced users are offered at most two different prices.

Theorem 23 *If* ξ *offers three prices then the corresponding rank-based Moulin mechanism* M_ξ *is not reputationproof.*

Proof. Since ξ offers three prices we consider i as in Definition 22 and we define

$$\mathsf{low} := \xi_i^s, \quad \mathsf{medium} := \xi_{i+1}^s, \quad \mathsf{high} := \xi_1^{s-i}.$$

We show that there exists a vector v such that the following two things happen. On input v, the mechanism drops several users and user 1 is charged an amount high. If user 1 misreport her name to $i+2$, then none is dropped and this user is charged an amount $\mathsf{medium} < \mathsf{high}$.

The vector v is as follows:

$$v := (\underbrace{\mathsf{high}, \mathsf{low}, \ldots, \mathsf{low}}_{i+1}, \perp, \underbrace{\xi_{i+2}^s, \xi_{i+3}^s, \ldots, \xi_s^s}_{s-i-1}, \underbrace{\perp, \ldots, \perp}_{n-s-1}).$$

Since there are s users that bid, the prices offered at the first round and the bids are

Names	1	2	$\cdots$	$i+1$	$i+3$	$i+4$	$\cdots$	s
Bids	high	low	$\cdots$	low	ξ_{i+2}^s	ξ_{i+3}^s	$\cdots$	ξ_s^s
Prices	ξ_1^s	ξ_2^s	$\cdots$	ξ_{i+1}^s	ξ_{i+2}^s	ξ_{i+3}^s	$\cdots$	ξ_s^s

Then, the $(i+1)^{th}$ user is dropped since she bids less than the offered price ($\mathsf{low} < \xi_{i+1}^s$). By cross monotonicity, at the second round the i^{th} user is offered a price $\xi_i^{s-1} \ge \xi_{i+1}^s = \mathsf{high}$. Since prices never decrease, the i^{th} user is eventually dropped. Similarly, all other users bidding low are also dropped. Since there are i users bidding low, the 1^{st} user is offered a final price which is at least $\xi_1^{s-i} = \mathsf{high}$.

Now consider the execution of the mechanism when the 1^{st} user misreports her name to $i+2$. This corresponds to the vector $v_{1\to i+2}$, and the prices offered at the first round and the bids are

Names	2	3	$\cdots$	$i+1$	$i+2$	$i+3$	$\cdots$	s
Bids	low	low	$\cdots$	low	high	ξ_{i+1}^s	$\cdots$	ξ_s^s
Prices	ξ_1^s	ξ_2^s	$\cdots$	ξ_i^s	ξ_{i+1}^s	ξ_{i+2}^s	$\cdots$	ξ_s^s

and thus no user is dropped ($\mathsf{low} = \xi_i^s$ and $\mathsf{high} > \mathsf{medium} = \xi_{i+1}^s$). In particular, user 1 is now bidding with name $i+2$ and thus pays the price $\xi_{i+1}^s = \mathsf{medium} < \mathsf{high}$. This show that the mechanism is not weakly renameproof. □

Although the above result puts some restriction on ξ, it does *not* say that that serviced users must be charged at most two prices as in the mechanisms in [BMST07]. Nevertheless, it is possible to show that, if the average cost C_s/s is monotone non-decreasing in s, then every non-uniform cross-monotonic ξ offers three prices. This argument essentially proves the following:

Theorem 24 *For the case of monotone non-decreasing average cost the only rank-based Moulin mechanism which is budget-balanced and reputationproof is the trivial mechanism. This mechanism is groupstrategyproof only if the average cost is constant.*

Proof. Assume that ξ does not offer three prices. We show that the average cost cannot be monotone non-decreasing (i.e., it decreases in at least in one case). In particular, we prove the following implication:

$$\xi_i^s < \xi_{i+1}^s \Rightarrow C_s/s < C_{s-i}/(s-i).$$

So, the average cost decreases when the number of serviced users increases from $s-i$ to $s > s-i$.

Since ξ does not offer three prices, then it must be $\xi_1^{s-i} \geq \xi_{i+1}^s$. Notice also that cross monotonicity and rank monotonicity imply that, for all j with $i+1 \leq j \leq s$, it holds that $\xi_{i+1}^s \leq \xi_j^s \leq \xi_{j-1}^{s-1} \leq \cdots \leq \xi_1^{s-j+1} \leq \xi_1^{s-i}$. We conclude that $\xi_1^{s-i} = \xi_{i+1}^s$ and thus none of these inequality is strict. In particular, it must must be $\xi_j^s = \xi_{j-i}^{s-i}$, for $i+1 \leq j \leq s$, thus yielding

$$\sum_{j=i+1}^{s} \xi_j^s = \sum_{j=1}^{s-i} \xi_j^s = C_{s-i}.$$

From this and by rank monotonicity we obtain

$$\frac{C_s - C_{s-i}}{i} = \frac{\sum_{j=1}^{i} \xi_j^s}{i} \leq \xi_i^s < \xi_{i+1}^s \leq \frac{\sum_{j=i+1}^{s} \xi_j^s}{s-i} = \frac{C_{s-i}}{s-i}.$$

With simple algebraic manipulations, we obtain

$$\begin{aligned} \frac{C_s - C_{s-i}}{i} < \frac{C_{s-i}}{s-i} \quad &\Leftrightarrow \\ s \cdot C_s - s \cdot C_{s-i} - i \cdot C_s + i \cdot C_{s-i} < i \cdot C_{s-i} \quad &\Leftrightarrow \\ (s-i)C_s < s \cdot C_{s-i} \quad &\Leftrightarrow \\ C_s/s < C_{s-i}/(s-i) \quad & \end{aligned}$$

which contradicts our hypothesis that the average cost is non-decreasing. □

B.1 The case of almost constant marginal cost

We prove an impossibility result for *almost constant* marginal cost functions. That is, the case in which the cost function satisfies

$$C_{s+1} - C_s = m$$

for all $s \geq 1$ and C_1 is any nonnegative number. We show that the only alternative to the trivial mechanism charging all users the same price is to have cost-sharing methods of the following type:

Definition 25 (marginal-and-uniform) *A rank-based cost-sharing method* ξ *is* marginal-and-uniform *if it is of the following form:*

$$\begin{pmatrix}
\xi_1^n & \cdots & & & & & & \cdots & \xi_n^n \\
m & \cdots & \cdots & m & a_i & \cdots & \cdots & a_i & \\
\vdots & & & & & & & & \\
 & & \iddots & \iddots & & & \iddots & & \\
m & m & a_i & \cdots & \cdots & a_i & & & \\
m & a_i & \cdots & \cdots & a_i & & & & \\
a_i & \cdots & \cdots & a_i & & & & & \\
\vdots & & & & & & & & \\
 & & \iddots & & & & & & \\
a_2 & a_2 & & & & & & & \\
a_1 & & & & & & & &
\end{pmatrix} \tag{10}$$

where a_s *is the average cost* C_s/s*.*

First of all, we have:

Lemma 26 *Unless* ξ *offers three prices, the following holds. Assume almost constant marginal cost, that is,* $m = C_{s+1} - C_s$ *for all* $s \geq 1$*. If all rows of* ξ *from* 1 *up to row* i *are uniform, and row* $i+1$ *is not uniform, then* ξ *must be like in (10).*

Proof. We consider an arbitrary row $j = i + p$, with $p \geq 0$. We prove by induction on p that ξ^j is as in (10), that is

$$\underbrace{\xi_1^j = \cdots = \xi_p^j}_{\text{all equal to } m} < \underbrace{\xi_{p+1}^j = \cdots = \xi_j^j}_{\text{all equal to } a_i}.$$

The base case is for $p = 0$, which follows from the hypothesis that ξ^i is uniform, i.e., all of its elements are equal to a_i.

As for the inductive step, we observe that any row $j > i$ is not uniform, that is, we have a strict inequality $\xi_q^j < \xi_{q+1}^j$, for some $1 \leq q \leq j$. Since ξ does not offers three prices, this strict inequality implies the following:

$$\xi_1^j \leq \cdots \leq \xi_q^j < \underbrace{\xi_{q+1}^j}_{\text{equal to } \xi_1^{j-q}} \leq \cdots \leq \underbrace{\xi_j^j}_{\text{equal to } \xi_{j-q}^{j-q}}. \tag{11}$$

All we have to do is to prove that it must be $q = p$, which implies that there is exactly one strict inequality in ξ^j. Moreover, since $j - q = j - p = i$, we have that $\xi^{j-q} = \xi^{j-p} = \xi^i$ which is identically equal to a_i. Finally, it cannot be $\xi_1^j < m$ since otherwise the sum of all these elements is *strictly less* than $m \cdot q + C_q = C_j$, which contradicts the fact that this sum must be C_j.

We first show that it cannot be $q < p$. Indeed, this would imply that

$$\xi_1^j \leq \cdots \leq \xi_q^j < \xi_1^{j-q} \leq \cdots \leq \xi_p^{j-q} < \xi_{p+1}^{j-q} \leq \cdots \leq \xi_{j-q}^{j-q}. \tag{12}$$

By inductive hypothesis, the row ξ^{j-q} satisfies (10), thus implying that

$$\xi_1^j \le \cdots \le \xi_q^j < \underbrace{m = \cdots = m}_{j-q-i} < \underbrace{a_i = \cdots = a_i}_{i}. \tag{13}$$

Hence, the sum of all the elements of ξ^j is *strictly less* than $m(j-i) + i \cdot a_i = m(j-i) + C_i$. Since $C_j = C_i + m(j-i)$, we have a contradiction.

Next, we show that it cannot be $q > p$ either. Indeed, in this case we would have

$$\xi_1^j \le \cdots \le \xi_q^j \le \xi_{q+1}^j \le \cdots \le \xi_p^j < \underbrace{a_i = \cdots = a_i}_{j-q}. \tag{14}$$

By cross monotonicity of ξ, we have $\xi_a^j \le \xi_a^{j-1}$. In particular, we have that all elements from 1 up to p must be at most m. Therefore, the sum of all elements of ξ^j is at most

$$mp + a_i(j-q) < mp + a_i(j-p) = m(j-i) + a_i \cdot i = m(j-i) + C_i = C_j,$$

thus a contradiction to the fact that this sum is equal to the cost C_j.

This concludes the proof. □