On the Approximability of Combinatorial Exchange Problems

Moshe Babaioff* Patrick Briest† Piotr Krysta‡

Abstract

In a combinatorial exchange a set of indivisible products is traded between buyers and sellers which are interested in bundles (multi-sets of products). Although combinatorial exchanges are a natural and important generalization of combinatorial auctions, their approximability has not been studied. We investigate the computational approximability of several social goals and show that the problems of surplus maximization and volume maximization (subject to positive surplus) are inapproximable even with free disposal and even if each agent's bundle is of size at most 3. Similar results based on communication complexity are shown for agents with general valuation functions.

In light of the negative results for surplus maximization we consider a complementary goal of social cost minimization and present tight approximation results for several social cost minimization problems. Considering the more general supply chain problem we prove that social cost minimization remains inapproximable even with bundles of size 3, yet becomes polynomial time solvable for agents trading bundles of size 1 or 2. This yields a complete characterization of the approximability of supply chain and combinatorial exchange problems based on the size of traded bundles.

Finally, we point out that economic considerations prevent any social cost approximation when agents hold private information about their valuations that needs to be elicited via some truthful mechanism.

1 Introduction

Following the emergence of the Internet as the world's foremost market place, much interest has been paid to problems naturally arising in a context where large scale economic problems need to be solved efficiently by computers. Many of these problems' essential difficulties can be captured by the class of (combinatorial) auction problems, which have in turn received a lot of attention from both practitioners and theoreticians in computer science. A major drawback of this model of abstraction, however, is the fact that it implicitly assumes a monopolistic market structure. While practitioners have therefore turned to more general (and more complex) supply chain models, these have not been subject to a rigorous theoretical investigation.

In this paper we consider the computational aspects of combinatorial exchanges and their extension to general supply chain formation problems. The combinatorial exchange model is a generalization of combinatorial auctions, which departs from the assumption that a monopolist seller holds a set of products which are of no actual value to him. Instead, it is assumed that apart from the set of buyers interested in purchasing bundles of items from the auctioneer, he also has access to a number of sellers offering to supply different bundles of products at a certain price. Thus, the problem we are faced with is to simultaneously run both a forward and reverse auction, which should return a feasible trade leaving a reasonable profit margin for the auctioneer. In

*School of Information, University of California at Berkeley, California. E-mail: moshe@ischool.berkeley.edu. Supported by NSF ITR Award ANI-0331659.

†Dept. of Computer Science, University of Liverpool, UK. E-mail: patrick@csc.liv.ac.uk. Supported by DFG grant Kr 2332/1-2 within Emmy Noether program.

‡Dept. of Computer Science, University of Liverpool, UK. E-mail: p.krysta@csc.liv.ac.uk. Supported by DFG grant Kr 2332/1-2 within Emmy Noether program.

the more general supply chain model, we drop the assumption that each trader is either a seller or a buyer, but may in fact offer some bundle of products under the condition that he is supplied with a different bundle and some appropriate side payment in return.

Combinatorial auctions have drawn much recent research attention (see for example the book by Cramton, Shoham and Steinberg [5]), and their computational and communication hardness are well understood. Lehmann, O'Callaghan and Shoham [8] and Sandholm [17] have shown that combinatorial auctions are hard to approximate within $min(m^{1/2-\varepsilon}, n^{1-\varepsilon})$ for any $\varepsilon > 0$, unless P=NP, where m and n denote the numbers of products and bidders, respectively. Nisan and Segal [12] present communication lower bounds for combinatorial auctions. They show that if 2 buyers have general valuations then exponential communication (in m) is required to find an allocation with maximum surplus (or $2^{1-\epsilon}$-approximation). Additionally, Nisan [11] shows that for n buyers distinguishing the case that the surplus is 1 and the case that the surplus is n requires exponential communication in m, assuming that $m^{1/2-\epsilon} > n$.

Clearly, all these hardness results carry to the combinatorial exchange case, as has been observed before. It turns out, however, that exchanges are even essentially more difficult than combinatorial auctions even in quite restricted cases. On the other hand, they nevertheless allow a number of positive results when the problem formulation is chosen carefully. We proceed by presenting a formal problem definition before stating our results in more detail.

1.1 Preliminaries

We first introduce formally the combinatorial exchange scenario, which most of the paper will be focused on. At the end of the section, we briefly explain the more general supply chain scenario. Assume that we are given a set of *agents* $\mathcal{A} = \mathcal{S} \cup \mathcal{B}$, where the collections $\mathcal{S}$, $\mathcal{B}$ are disjoint sets of *sellers* and *buyers*, respectively, with $|\mathcal{A}| = n$. We are interested in trades that include *indivisible* products $\mathcal{U}$, where $|\mathcal{U}| = m$. Each seller i is offering a bundle $q_i = (q_i^1, \ldots, q_i^m)$ of products at some price $v_i \in \mathbb{R}_+$. Buyer j is requesting to buy the bundle $q_j \in \mathbb{N}^m$ at price $v_j \in \mathbb{R}_+$. By $q_i^e, q_j^e \in \mathbb{N}$ we refer to the number of copies of product $e \in \mathcal{U}$ offered by seller i or requested by buyer j, respectively. For agent $k \in \mathcal{A}$, her bundle is a *set* if for all $e \in \mathcal{U}$, $q_k^e \in \{0, 1\}$.

A *feasible trade* $T = (S, B)$, $S \subseteq \mathcal{S}$ and $B \subseteq \mathcal{B}$, is a selection of sellers and buyers, such that

$$\sum_{i \in S} q_i^e \geq \sum_{j \in B} q_j^e \text{ for all } e \in \mathcal{U},$$

i.e., the supply provided by sellers in S is sufficient to satisfy the requests of all buyers in B. Note, that we assume free disposal here, i.e., supply and demand do not have to match exactly. We say that trade $T = (S, B)$ has *surplus* [*volume*]

$$\text{sur}(T) = \sum_{j \in B} v_j - \sum_{i \in S} v_i \qquad \text{vol}(T) = \sum_{e \in \mathcal{U}} \sum_{j \in B} q_j^e.$$

The *combinatorial exchange surplus problem* is the problem of finding a trade which maximizes the surplus. The *combinatorial exchange volume problem* is the problem of finding a trade which maximizes the volume, subject to positive surplus. By *combinatorial exchange positive surplus problem* we refer to the problem variation in which we simply want to find any feasible trade with strictly positive surplus.

Following [16], as another objective we define the *social cost* of trade T as

$$\text{cost}(T) = \sum_{j \notin B} v_j + \sum_{i \in S} v_i,$$

i.e., the cost is the sum of valuations of the trading sellers and non-trading buyers. Note that for every instance, the sum of the social cost and the surplus is constant, $\text{cost}(T) + \text{sur}(T) = \sum_{j \in \mathcal{B}} v_j$. Thus, a trade maximizes the surplus if and only if it minimizes the social cost. We consider the problem of minimizing the social cost subject to non-negative surplus, which we call the *combinatorial exchange social cost problem*. From a computational perspective, the social cost objective is preferable, because it allows us to derive approximation algorithms and express their approximation ratio in terms of a multiplicative factor, which is generally difficult if we are faced with any mixed sign objective and in fact turns out to be impossible in the case of combinatorial exchanges.

A natural generalization of the combinatorial exchange problem is obtained if we allow agents that are both sellers and buyers simultaneously and confront our algorithm with offers of the form "given bundle A I will supply bundle B for an additional payment of x". Formally, in the *supply chain problem* we are given a set $\mathcal{A}$ of n agents. Agent k is represented by $(\{\delta_k^e\}_{e \in \mathcal{U}}, v_k)$, where $\delta_k^e \in \mathbb{Z}$ denotes the number of copies of product e requested or supplied by agent k (modelled as $\delta_k \geq 0$ or $\delta_k < 0$) and $v_k \in \mathbb{R}$ is the additional payment offered or requested (modelled as $v_k \geq 0$ or $v_k < 0$). The objectives of surplus or volume maximization generalize naturally to the supply chain scenario. However, we need to adapt the notion of social cost to fit our generalized type of agents. We let $\mathcal{A}^+ = \{k \,|\, v_k > 0\}$ and $\mathcal{A}^- = \mathcal{A} \backslash \mathcal{A}^+$. Thus, $\mathcal{A}^+$ is the set of agents that have a positive utility for being included in the trade, agents in $\mathcal{A}^-$ incur a cost when included for which they need to be repaid. We can then naturally define the social cost of a trade $T \subseteq \mathcal{A}$ as

$$\text{cost}(T) = \sum_{k \in \mathcal{A}^+ \backslash (T \cap \mathcal{A}^+)} v_k - \sum_{k \in T \cap \mathcal{A}^-} v_k.$$

Finally, let us introduce some more notation that will come in handy. For a given instance $I = (\mathcal{S}, \mathcal{B}, v)$ of the combinatorial exchange problem we let $T^\star_{\text{sur}}(I)$ refer to the surplus maximizing trade. The trade computed by some algorithm A on the same instance is denoted as $A(I)$. Analogous notation is used in the context of volume maximization or cost minimization, respectively. We use $\mathcal{F}$ to denote a family of instances which satisfy some specified additional constraints (as maximal bundle size).

1.2 Our Results

Although combinatorial exchanges model many interesting trade scenarios and understanding their computational hardness is clearly an important problem, we are not aware of any paper that addresses this important issue. The only result that we know of is the rather immediate reduction from combinatorial auctions to exchanges without free disposal, presented by Sandholm et al. [19], which shows that without free disposal (i.e., if the total supply of each product must exactly match the demand), it is NP-complete to find a feasible allocation and, thus, the combinatorial exchange surplus problem without free disposal cannot be approximated (Corollary 2.1 in [19]). This result heavily relies on the assumption of no free disposal. Is the problem as hard when we allow for free disposal? What if severe restrictions are imposed on the bundles of the agents? We show that the combinatorial exchange surplus and volume maximization problems are inapproximable[1], even with free disposal, unless $P = NP$. They remain inapproximable even on instances in which each agent's bundle is a set of size at most 3, and when restricting to a single product (multi-unit exchange).

Our inapproximability result for bounded size sets is based on a reduction to a family of combinatorial exchange instances with sets of size at most 3 that has the following property. Once the quantity in which each product is bought is decided, the packing and covering problems defined by these quantities (finding

[1] Inapproximable within any factor that is a polynomially computable function of n and m, and even if we allow for an additive term that is a polynomially computable function of n and m.

the trading buyers and sellers, respectively) are polynomial time solvable. This implies that combinatorial exchange has inherent hardness that does not come from the hardness of packing and covering. Finally, we prove that the problem remains inapproximable even when we restrict ourselves to instances with large *packing to covering factor*, i.e., in cases where there is a large gain from the optimal trade, and derive complementary inapproximability results for exchanges with sub-exponential communication based on the communication lower bounds of [11, 12].

Sandholm and Suri [18] present an anytime algorithm for combinatorial auctions which also extends to combinatorial exchanges. Parkes et al. [13] have recently presented ICE, an iterative combinatorial exchange. Our hardness results have implications for the worst case performance of these algorithms and emphasize the need for algorithms with provable performance guarantees.

We show that, focusing on the social cost objective rather than surplus or trade volume maximization, such algorithmic results can be obtained. More formally, while the surplus objective turns out to be inapproximable within any factor, we show that the social cost can be approximated within a factor of H_k, where k is the maximal size of any bundle and that this is essentially tight when agents are bidding for sets (not for multi-sets), as it is hard to get $(1 - o(1))H_m$ approximation, unless NP $\subseteq$ DTIME($n^{\mathcal{O}(\log \log n)}$). Similarly, for the relevant case of a single type of product (multi-unit exchange) we show that social cost minimization is NP-hard yet there is an FPTAS.

In light of these positive results, we ask whether the more general supply chain social cost problem allows approximate social cost minimization, as well. It turns out that this is not the case and, in fact, social cost minimization in the supply chain scenario is inapproximable, even with bundles of size at most 3. This is interesting, as it is the first formal result separating combinatorial exchanges from the more general supply chain scenario.

We then consider the special case in which agents are restricted to sets of size 1 or 2. While this case is polynomially solvable for combinatorial auctions, it was unknown whether this is true for combinatorial exchanges. We prove here that it holds, in fact, even for the supply chain scenario. Our algorithm works by reducing the problem to a weighted b-matching problem through a number of transformations and reformulation in terms of the social cost objective. This also shows that our hardness result for supply chains with sets of size 3 is tight.

Finally, although the computational problem of social cost minimization is much easier than surplus maximization, we show that there is another problem that arises when trying to construct mechanisms with this goal in mind, even when computation is not a problem at all. We move to consider the case that agents have privately known values for their bundles and we need to elicit this information from them. A classical result by Myerson and Satterthwaite [10] shows that for bilateral trade, any mechanism that is truthful, individually rational and budget-balanced must sometimes be inefficient. This implies that a truthful, individually rational and budget-balanced mechanism cannot always minimize the social cost. We extend this result and show that for any constant $\alpha \geq 1$, α-approximation for social cost is impossible. Circumventing the Myerson and Satterthwaite impossibility result [10] for exchanges has been the subject of several papers. Parkes et al. [14] enforce budget-balance as a hard constraint, and explore payment rules that are fairly efficient and fairly truthful. On the other hand, Babaioff and Walsh [2] consider the problem of supply chains and present a truthful and budget-balanced mechanism with efficiency which depends on the size of the efficient trade.

The rest of the paper is organized as follows. Section 2 presents our inapproximability results for surplus and volume maximization. Section 3 presents approximation algorithms as well as inapproximability results for social cost minimization. Section 4 considers the case of agents with sets of size at most 2. Finally, Section 5 shows that economic considerations prevent mechanisms that achieve social cost approximation. Several of the proofs have been moved to the Appendix due to space limitations.

2 Surplus and Volume Inapproximability

We start by showing a number of strong inapproximability results for both surplus and volume maximization in combinatorial exchanges. Section 2.1 derives results for the single-minded setting under standard complexity theoretic assumptions. Section 2.2 presents similar results for general agents in terms of communication complexity.

2.1 Computational Hardness

We show that both the combinatorial exchange surplus and volume problems do not allow polynomial time approximation algorithms with any reasonable approximation guarantee. This is formalized in the following definition.

Definition 1 *Let $\alpha, \beta : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be any polynomial-time computable functions. The combinatorial exchange surplus problem is* approximable *(for family of instances $\mathcal{F}$) if for some α, β there exists a poly-time approximation algorithm A, such that* $\text{sur}(T^\star_{\text{sur}}(I)) \leq \alpha(n,m) \cdot \text{sur}(A(I)) + \beta(n,m)$ *holds on every problem instance I ($I \in \mathcal{F}$). The combinatorial exchange surplus problem is* inapproximable *if it is not approximable.*

Approximability in the volume maximization and social cost minimization case is defined analogously. We are interested in a very natural restricted class of families of instances.

Definition 2 *The family of instances $\mathcal{F}$ is* rational*, if all values v_i are rational numbers. $\mathcal{F}$ is* scalable *if it is closed under scaling of agents' valuations, i.e., if instance $I = (\mathcal{S}, \mathcal{B}, v) \in \mathcal{F}$ then $I' = (\mathcal{S}, \mathcal{B}, \gamma \cdot v) \in \mathcal{F}$ for any rational $\gamma > 0$.*

It turns out that for the families of instances defined above inapproximability can be derived from the fact that detecting any trade with positive surplus or volume is hard in itself. The proof of Lemma 1 is found in Appendix A.1.

Lemma 1 *If the combinatorial exchange positive surplus problem for a rational and scalable family $\mathcal{F}$ is NP-hard then the combinatorial exchange surplus (volume) problem for the family $\mathcal{F}$ is inapproximable, unless $P = NP$.*

Theorem 1 *The combinatorial exchange surplus (volume) problem is inapproximable, unless $P = NP$. It remains inapproximable unless $P = NP$ for the following families of instances $\mathcal{F}$:*

- *instances with only sets ($q_k^e \in \{0,1\}$) of size at most 3 ($\sum_{e \in \mathcal{U}} q_k^e \leq 3$ for every agent k).*
- *instances with only one type of product (multi-unit exchange, $|\mathcal{U}| = 1$).*

Theorem 1 follows from Lemmas 2 and 3. The proof of Lemma 3 is straightforward and omitted due to space limitations.

Lemma 2 *The combinatorial exchange positive surplus problem for the family $\mathcal{F}$ of instances with only sets ($q_k^e \in \{0,1\}\ \forall k \in \mathcal{A}, e \in \mathcal{U}$) of size at most 3 ($\sum_{e\in\mathcal{U}} q_k^e \leq 3\ \forall k \in \mathcal{A}$) is NP-hard.*

Proof: We show a reduction from the set packing problem, which is known to be NP-hard [7] even for sets of size at most 3. Let an unweighted set packing instance $S_1, \ldots, S_n \subseteq \mathcal{U}$, $|\mathcal{U}| = m$, be given and assume that $|S_j| \leq 3$ for all j. We ask whether a collection $P \subseteq \{1, \ldots, n\}$ of sets with $S_i \cap S_j = \emptyset$ for all $i, j \in P$ and $\bigcup_{i\in P} S_i = \mathcal{U}$ exists.

We construct an instance of the combinatorial exchange problem on the extended ground set $\mathcal{U}' = \mathcal{U} \cup \{e_{j,1}, e_{j,2}, e_{j,3} \,|\, j = 1, \ldots, n\}$.

Depending on its size we define a number of buyers for each set S_j. If S_j has size 1 and consists of some single element e, we define buyers j_1, j_2 with $q_{j_i}^e = q_{j_i}^{e_{j,i}} = 1$, $q_{j_i}^{e'} = 0$ else and $v_{j_i} = 2(2m + 3n + 1)$. If S_j has size 2 and consists of elements e_1 and e_2, we define buyers j_1, j_2 with $q_{j_i}^{e_1} = q_{j_i}^{e_2} = q_{j_i}^{e_{j,i}} = 1$, $q_{j_i}^{e'} = 0$ else and $v_{j_i} = 3(2m + 3n + 1)$. If S_j has size 3 and consists of elements e_1, e_2 and e_3, we define 3 buyers j_1, j_2 and j_3 with $q_{j_1}^{e_1} = q_{j_1}^{e_2} = q_{j_1}^{e_{j,1}} = 1$, $q_{j_2}^{e_2} = q_{j_2}^{e_3} = q_{j_2}^{e_{j,2}} = 1$, $q_{j_3}^{e_1} = q_{j_3}^{e_3} = q_{j_3}^{e_{j,3}} = 1$, $q_{j_i}^{e'} = 0$ else and $v_{j_i} = 3(2m + 3n + 1)$. Thus, the union of buyers belonging to a single set S_j requests 2 copies of each element in the set. Additionally, each of them requests a distinct item from the set of associated items $\{e_{j,1}, e_{j,2}, e_{j,3}\}$.

For ease of notation let $\mathcal{U} = \{e_0, \ldots, e_{n-1}\}$. We define sellers indexed by $i = 1, \ldots, n$, where $q_i^{e_j} = 1$ if $j \in \{i, i+1 \mod n\}$, $q_i^{e_j} = 0$ else. Seller i's valuation is given as $v_i = 2(2m + 3n)$. Additionally, we have a seller j for each set S_j with $q_j^{e_{j,1}} = q_j^{e_{j,2}} = 1$, $q_j^e = 0$ for all other $e \in \mathcal{U}'$ and $v_j = 2(2m + 3n)$, if $|S_j| = 1$ or $|S_j| = 2$, $q_j^{e_{j,1}} = q_j^{e_{j,2}} = q_j^{e_{j,3}} = 1$, $q_j^e = 0$ for all other $e \in \mathcal{U}'$ and $v_j = 3(2m + 3n)$, if $|S_j| = 3$.

Now let P be a packing, such that $\bigcup_{i\in P} S_i = \mathcal{U}$. We let $P = P_1 \cup P_2 \cup P_3$, where P_k contains all sets of size k from P. By selecting all buyers j_i with $j \in P$, all sellers i and sellers j supplying products $e_{j,k}$, $j \in P$, we obtain a feasible trade with surplus

$$
\begin{aligned}
&\sum_{j\in P_1} 2 \cdot 2(2m + 3n + 1) + \sum_{j\in P_2} 2 \cdot 3(2m + 3n + 1) + \sum_{j\in P_3} 3 \cdot 3(2m + 3n + 1) \\
&\qquad - \sum_{j\in P_1\cup P_2} 2(2m + 3n) - \sum_{j\in P_3} 3(2m + 3n) - |\mathcal{U}| \cdot 2(2m + 3n) \\
= \;& (4|P_1| + 6|P_2| + 9|P_3|)(2m + 3n + 1) - (2|\mathcal{U}| + 2|P_1 \cup P_2| + 3|P_3|)(2m + 3n) \\
= \;& (2|\mathcal{U}| + 2|P_1 \cup P_2| + 3|P_3|) > 0
\end{aligned}
$$

where we use the fact that $2|P_1| + 4|P_2| + 6|P_3| = 2|\mathcal{U}|$. Thus, we have shown that there exists a trade with positive revenue whenever the desired packing exists. We continue by arguing that any trade with positive revenue implies the existence of a corresponding packing.

Consider an arbitrary trade $T = (S, B)$ and assume that sellers in S supply a total number s of products, while buyers request only a total number of $t < s$ of products. Observe that in our instance every seller charges a price of exactly $2m + 3n$ for each product she supplies. Similarly, buyers offer exactly $2m + 3n + 1$ per product they request. Hence, we can bound the surplus from trade T by

$$t(2m + 3n + 1) - s(2m + 3n) \leq t - (2m + 3n) < 0,$$

where we use the fact that $s \geq t + 1$ and obviously $t < |\mathcal{U}'| = 2m + 3n$.

Thus, in any trade with positive revenue sellers supply exactly the number of products requested by the selected buyers. It is straightforward to argue that this implies the following two properties of the returned trade. First, whenever some buyer j_i corresponding to set S_j is selected for the trade, then this is true for all buyers corresponding to that set. Secondly, the trade must include all products $e_0, \ldots, e_{n-1}$. This, however, implies that selecting sets S_j corresponding to selected buyers j_i yields the desired packing. $\square$

Lemma 3 *The combinatorial exchange positive surplus problem for the family $\mathcal{F}$ of instances with only a single type of good (multi-unit exchange problem) is NP-hard.*

Finally, we prove inapproximability even in the cases where there are relatively large gains from the trade.

Definition 3 *For instance I we define the* packing to covering factor $f(I)$ *as the maximal value of* $\sum_{i \in B^*} v_i / \sum_{j \in S^*} v_j$ *over all surplus maximizing trades $T^\star_{sur}(I) = (S^*, B^*)$.*

We prove inapproximability with a packing to covering factor corresponding to the set packing lower bound.

Theorem 2 *There exists a function $\gamma(n,m) = \Omega(\min\{n^{1-\varepsilon}, m^{1/2-\varepsilon}\})$ such that the combinatorial exchange surplus (volume) problem is inapproximable for the family $\mathcal{F}$ of instances which satisfy $f(I) \geq \gamma(n,m)\ \forall I \in \mathcal{F}$, unless P=NP.*

2.2 Communication Lower Bounds

We next consider the problem of achieving approximation to the combinatorial exchange surplus (volume) problem when agents have general valuations (buyers are not single-minded). We show that the two goals cannot be approximated unless exponential communication in m is used. The inapproximability results hold even in the case of a single seller holding a set, and buyers that have general monotone valuations over sets (not multi-sets).

Let $G = (1, \ldots, 1)$ be the bundle with one item of each product. Assume that buyer i has a monotone valuation function $v_i : 2^m \rightarrow \mathbb{R}_+$. As we assume that a single seller offers G, the goal of a communication protocol P is to find a partition of the items to the buyers such that the surplus (volume) is maximized. Similar to Definition 1 we define inapproximability of a communication protocol in Appendix A.2 (with α, β using only m as their argument). Based on a result from Nisan and Segal [12] we show the following.

Theorem 3 *The combinatorial exchange surplus (volume) problem is not approximable in less than $\binom{m}{m/2}$ bits. This holds even with only a single seller and two buyers which have valuations over sets.*

Moreover, even if there is a large gain from trade the problem remains inapproximable. Based on a lower bound of Nisan [11] we prove that when $f(I) \geq n$ and $n < m^{1/2-\epsilon}$, any approximation requires exponential communication in m .

Theorem 4 *The combinatorial exchange surplus (volume) problem for the family of instances with packing to covering factor at least n is not approximable with less than $e^{m/(2n^2)-5log\ n}$ bits. The lower bound holds for randomized and nondeterministic protocols.*

3 Approximating Social Cost

3.1 The Combinatorial Exchange Social Cost Problem

In the following section we present an algorithm that achieves a logarithmic approximation ratio for the objective of minimizing social cost of the trade. We additionally present a matching lower bound for the case without multi-sets. Our algorithm is based on the well known greedy approximation algorithm for the multi-set multi-cover problem [15]. For the remainder of this section, let $k_i = \sum_{e \in \mathcal{U}} q_i^e$ for all $i \in \mathcal{S}, \mathcal{B}$ and define $k = \max_i k_i$.

1. Let $\Delta_e = \sum_{j \in \mathcal{B}} q_j^e$ for all products $e \in \mathcal{U}$.
2. Apply the greedy approximation algorithm to the following multi-set multi-cover problem:

$$\begin{aligned} \text{min.} \quad & \sum_{i \in \mathcal{S}} v_i x_i + \sum_{j \in \mathcal{B}} v_j y_j \\ \text{s.t.} \quad & \sum_{i \in \mathcal{S}} q_i^e x_i + \sum_{j \in \mathcal{B}} q_j^e y_j \geq \Delta_e \quad \forall e \in \mathcal{U} \\ & x_i, y_j \in \{0, 1\} \end{aligned}$$

3. Let $S = \{i \mid x_i = 1\}$, $B = \{j \mid y_j = 0\}$. If $\sum_{j \in B} v_j \geq \sum_{i \in S} v_i$ return trade $T = (S, B)$, else return $T = (\emptyset, \emptyset)$.

Figure 1: Approximating optimal social cost by algorithm COVER.

It is known that the greedy approximation algorithm approximates the covering integer program (CIP) in Figure 1 within H_k, where H_k denotes the k'th harmonic number. Essentially, Theorem 5 follows from the observation that this CIP is an exact formulation of the problem of finding a trade of minimal social cost. We briefly mention that the greedy approximation algorithm we apply is inherently monotone and, thus, yields a truthful exchange mechanism if combined with an appropriate (critical value based) payment scheme. However, we point out in Section 5 that there are other reasons that prevent us from obtaining reasonable truthful mechanisms.

Theorem 5 *Algorithm* COVER *is a H_k-approximation algorithm for the combinatorial exchange social cost problem.*

Proof: Consider the problem of social cost minimization in combinatorial exchanges. We can write this problem as the following integer linear program, where variables x_i, x_j indicate which sellers and buyers are selected for the trade and constraints (2) ensure feasibility of the trade:

$$\text{min.} \quad \sum_{j \in \mathcal{B}} v_j (1 - x_j) + \sum_{i \in \mathcal{S}} v_i x_i \tag{1}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{S}} q_i^e x_i \geq \sum_{j \in \mathcal{B}} q_j^e x_j \quad \forall e \in \mathcal{U} \tag{2}$$

$$x_i, x_j \in \{0, 1\} \tag{3}$$

Defining $\Delta_e = \sum_{j\in\mathcal{B}} q_j^e$ as in the algorithm we can rewrite $\sum_{j\in\mathcal{B}} q_j^e x_j$ as $\Delta_e - \sum_{j\in\mathcal{B}} q_j^e(1-x_j)$. Thus, constraints (2) become

$$\sum_{i\in\mathcal{S}} q_i^e x_i \geq \Delta_e - \sum_{j\in\mathcal{B}} q_j^e(1-x_j).$$

Substituting new variable y_j for $1-x_j$ for all buyers $j\in\mathcal{B}$ we obtain exactly the covering integer program defined in algorithm COVER.

Let now $T=(S,B)$ as in the algorithm. By the fact that the greedy algorithm for multi-set multi-covering [15] has approximation ratio H_k we immediately obtain that $\text{cost}(T) \leq H_k \cdot \text{cost}(T^\star_{\text{cost}}(I))$. However, it might be the case that trade T has negative surplus. We fix this problem by returning the empty trade whenever surplus is indeed negative. Obviously, this does not violate our approximation guarantee, since in this case the social cost of the empty trade is below the cost of the trade that we have originally computed. □

The following theorem states that the approximation ratio of algorithm COVER is essentially best possible, as parameter k is trivially upper bounded by the number m of distinct goods whenever we do not allow multi-sets. The proof of Theorem 6 is found in Appendix B.1.

Theorem 6 *The combinatorial exchange social cost problem cannot be approximated in polynomial time better than within $(1-o(1))\ln m$, unless NP $\subseteq$ DTIME($n^{\mathcal{O}(\log\log n)}$).* [2]

We finally mention that algorithm COVER can in fact be viewed as a generic reduction of social cost minimization in combinatorial exchanges to social cost minimization in reverse combinatorial auctions. Loosely speaking, we first allocate to all buyers in the exchange scenario their desired bundles at their offered price and then run a reverse auction algorithm considering all exchange participants as sellers. If one of the original sellers is selected, we buy her offered bundle. If one of the original buyers is selected, we buy the previously allocated bundle back from her. Setting demand for every product as done above, we ensure that the auction algorithm achieves sufficient supply for all buyers not returning their bundles.

3.2 The Multi-Unit Exchange Social Cost Problem

We briefly consider the special case of multi-unit exchanges, in which all sellers and buyers want to trade a number of copies of the same good. By Theorem 1 we are unlikely to achieve in polynomial time any approximation ratio for surplus or volume maximization. Thus, we again reformulate the problem in terms of a covering problem and aim to minimize social cost. Algorithm MINKNAPSACK (see Appendix B.2) is based on the dynamic programming based FPTAS for the min-knapsack (or reverse multi-unit auction) problem. For a monotone version of the algorithm tailored to our setting see [3]. Similar to the proof of Theorem 6 a simple reduction to the partition problem (which is known to be NP-hard [7]) yields optimality of our algorithm's approximation guarantee.

Theorem 7 *Algorithm* MINKNAPSACK *is an FPTAS for the combinatorial exchange social cost problem. Furthermore, social cost minimization in multi-unit exchanges is NP-hard.*

[2] Note that we can replace this assumption by $P \neq NP$ if we relax the lower bound to $\Omega(\ln m)$ by [1].

3.3 The Supply Chain Social Cost Problem

The objectives of surplus or volume maximization generalize naturally to the supply chain scenario. Thus, all the hardness results presented in Section 2.1 hold for this more general model, as well. However, we show that the situation is even worse and that for the supply chain scenario even (approximate) social cost minimization is out of reach. Theorem 8 states that no polynomial time algorithm can achieve any approximation guarantee for the supply chain social cost problem.

Theorem 8 *The supply chain social cost problem is inapproximable, unless P=NP.*

Proof: We show a reduction from the decision version of the set packing problem, which is known to be NP-hard [7]. Let sets $S_1, \ldots, S_n$ over ground set $\mathcal{U}$, $|\mathcal{U}| = m$, and integer r be given. We want to decide whether there exist r non-intersecting sets. Towards a contradiction, assume that we are given an algorithm A for supply chain social cost minimization, such that

$$cost(T^{\star}_{cost}(I)) \leq \alpha(n, m) \cdot cost(A(I)) + \beta(n, m)$$

holds on every problem instance I. We define our supply chain social cost instance over ground set $\mathcal{U} \cup \{e^*\}$ as follows. For each set S_j we create a corresponding agent j who requests set S_j, supplies one copy of e^* and has valuation $v_j = -1$, formally $\delta^e_j = 1$ if $e \in S_j$, $\delta^e_j = -1$ if $e = e^*$ and $\delta^e_j = 0$ else. Additionally, we define agent $n+1$ who supplies $\mathcal{U}$ (i.e., $\delta^e_{n+1} = -1$ for all $e \in \mathcal{U}$), requests r copies of e^* ($\delta^{e^*}_{n+1} = r$) and offers to pay $v_{n+1} = \alpha(n, m)r + \beta(n, m) + 1$. Thus, agents corresponding to sets from the packing instance form set $\mathcal{A}^-$, agent $n+1$ forms set $\mathcal{A}^+$.

Now assume that there are r non-intersecting sets $\mathcal{S}' \subseteq \{1, \ldots, m\}$ and consider trade $T = \mathcal{S}' \cup \{n + 1\}$. This trade is feasible, since agents in $\mathcal{S}'$ request a subset of the goods in $\mathcal{U}$, which are supplied by agent $n+1$. On the other hand, agents in $\mathcal{S}'$ supply r copies of product e^* as requested by $n + 1$. Furthermore, trade T has social cost r. Now observe that the empty trade hast cost $v_{n+1} = \alpha(n, m)r + \beta(n, m) + 1$. Thus, if there exist r non intersecting sets algorithm A must return a non-empty trade.

Assume then that in every selection $\mathcal{S}'$ of at least r sets there are sets with a non-empty intersection. We claim that no feasible trade other than the empty trade exists. To see this, note that every non-empty trade must contain agent $n + 1$, as no other agent can supply products from $\mathcal{U}$. Furthermore, no feasible trade can select agents corresponding to intersecting sets, since agent $n+1$ supplies only a single copy of each product in $\mathcal{U}$. Finally, every feasible trade must contain at least r of the agents $1, \ldots, m$, since we need a supply of at least r copies of product e^* in order to select agent $n + 1$. Hence, no feasible non-empty trade exists under the above assumption. □

A natural question to ask is whether similar hardness can be shown for supply chain instances with bounded size bundles, where we define the size of agent j's bundle as $\sum_e |\delta^e_j|$. Theorem 9 states that the problem is inapproximable even with bundles of size 3.

Theorem 9 *The supply chain social cost problem with bundles of size at most* 3 *is inapproximable, unless P=NP.*

We omit the formal proof, which is based on a reduction of set packing with sets of size 3 and the fact that we can transform the problem instances constructed in the proof of Theorem 8 into instances with only bounded size bundles, see Figure 2. This hardness result is tight, as we will see in the following section that bundles of size 2 allow exact polynomial time algorithms.

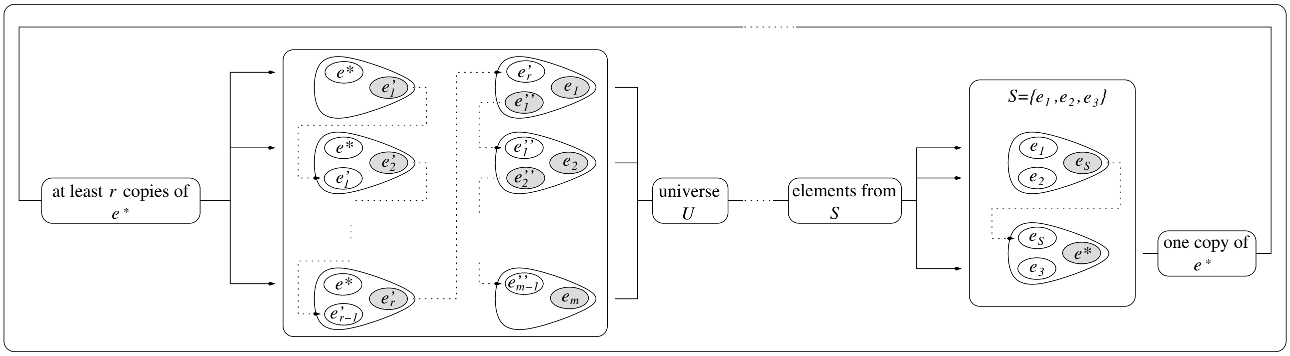

Figure 2: Gadgets for the proof of Theorem 9 for agents with sets of bounded size. Agents are depicted as rounded triangles, supplied and requested products as shaded or non-shaded ellipses, respectively. Agent $n+1$ from the proof of Theorem 8 is simulated by the agents in the box to the left, each of which has valuation $(2\alpha(n,m)r+\beta(n,m)+1)/(m+r)$. Agent j corresponding to set S_j is simulated by the two agents depicted in the box to the right, both with valuation -1. Again, every feasible non-empty trade must contain agents corresponding to at least r non-intersecting sets.

4 Supply Chains with Sets of Size 2

It is well known that combinatorial auctions become solvable in polynomial time when the participating bidders are restricted to bid on sets of size at most 2, since in this case the problem can be formulated in terms of a weighted matching problem. For combinatorial exchanges or even supply chains, on the other hand, it is not immediate whether the problem reduces to matching in this case. As we shall see, the social cost objective turns out to be the key in obtaining such a problem formulation.

Theorem 10 *The supply chain surplus problem with sets of size* 1 *and* 2 *can be solved in polynomial time.*

Theorem 10 is a direct consequence of Lemmas 4, 5 and 6, which describe the reduction of the supply chain surplus problem to weighted b-matching.

Lemma 4 *The supply chain surplus problem with bundles of size at most* k *reduces to the combinatorial exchange surplus problem with bundles of size at most* k.

Proof: Let an instance of the supply chain surplus problem with agents $\mathcal{A}$ be given and let $\alpha = \sum_{k\in\mathcal{A}^+} v_k + 1$, thus, $\mathrm{sur}(T^\star) < \alpha$.

Consider agent $k \in \mathcal{A}$ and let $S_k = \{e \in \mathcal{U} \,|\, \delta_k^e < 0\}$, $B_k = \{e \in \mathcal{U} \,|\, \delta_k^e > 0\}$ denote the sets of products offered and requested, respectively. Assume that agent k is both offering and requesting products, i.e, $|S_k|, |B_k| \geq 1$. We define a new product e_k, seller s_k and buyer b_k corresponding to agent k as follows. Seller s_k supplies product e_k and the products supplied by k at price α, formally $q_{s_k}^{e_k} = 1$, $q_{s_k}^e = -\delta_k^e$ for all $e \in S_k$ and $v_{s_k} = \alpha$. Buyer b_k requests product e_k and the products requested by k at price $\alpha + v_k$, formally $q_{b_k}^{e_k} = 1$, $q_{b_k}^e = \delta_k^e$ for all $e \in B_k$ and $v_{s_k} = \alpha + v_k$.

Consider any solution to the resulting instance of the combinatorial exchange surplus problem. If buyer b_k is selected for the trade, then so is seller s_k, since no other agent can supply product e_k. If seller s_k is selected, then so is buyer b_k, since otherwise surplus is bounded by $\mathrm{sur}(T^\star) - \alpha < 0$. Finally, observe that bundle sizes have not increased. □

Lemma 5 *The combinatorial exchange surplus problem with sets of size* 1 *and* 2 *reduces to the combinatorial exchange surplus problem with sets of size* 2*.*

Proof: First, for each buyer with a set of size 1, define a new product requested by this buyer and supplied by a newly defined seller at no cost. This leaves us with a problem instance in which each buyer has a set of size 2 and sellers have sets of size 1 or 2. For each seller with a set of size 1 now simply define a new product which she supplies at no additional cost. □

Lemma 6 *The combinatorial exchange surplus problem with sets of size* 2 *can be solved in polynomial time.*

Proof: Let an instance of the combinatorial exchange surplus problem with sellers $\mathcal{S}$ and buyers $\mathcal{B}$ with sets of size 2 be given. Using that surplus maximization is equivalent to minimization of social cost, we formulate the problem as an ILP as in Figure 1. Let $\Delta_e = |\{j \in \mathcal{B} \,|\, q_j^e = 1\}|$ as before, $\Gamma_e = |\{j \in \mathcal{S} \,|\, q_j^e = 1\}|$ and $V = \sum_{i\in\mathcal{S}} v_i + \sum_{j\in\mathcal{B}} v_j$. Defining $\lambda_i = 1 - x_i$ and $\mu_j = 1 - y_j$ we obtain

$$\sum_{i\in\mathcal{S}} v_i x_i + \sum_{j\in\mathcal{B}} v_j y_j = V - \sum_{i\in\mathcal{S}} v_i\lambda_i - \sum_{j\in\mathcal{B}} v_j\mu_j$$

and can rewrite the ILP as:

$$\text{max.} \quad \sum_{i\in\mathcal{S}} v_i\lambda_i + \sum_{j\in\mathcal{B}} v_j\mu_j \tag{4}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{S}} q_i^e\lambda_i + \sum_{j\in\mathcal{B}} q_j^e\mu_j \le \Gamma_e \quad \forall e \in \mathcal{U} \tag{5}$$

$$\lambda_i, \mu_j \in \{0,1\} \tag{6}$$

Now observe that the above is the ILP formulation of the maximum weighted b-matching problem with weights v_i, v_j on the edges and degree constraints Γ_e for vertices $e \in \mathcal{U}$, which is known to be solvable in polynomial time [4]. □

5 Non-Existence of Mechanisms for Social Cost Approximation

In this section we discuss the existence of truthful mechanisms for the combinatorial exchange problem. A mechanism consists of some algorithm A that outputs a trade $T = (S, B)$ and additional payments $(p_i)_{i\in\mathcal{S}}$, $(q_j)_{j\in\mathcal{B}}$ determining the payments given to sellers and collected from buyers, respectively. For ease of notation we denote input as $v = (v_1, \ldots, v_{|\mathcal{S}|})$ and $w = (w_1, \ldots, w_{|\mathcal{B}|})$. Trade and payments are referred to as $A(v, w)$, $p(v, w)$ and $q(v, w)$, respectively. For both sellers and buyers $i \in \mathcal{B}, \mathcal{S}$ we denote by $i \in A(v, w)$ the situation that i is selected for the trade. A mechanism is *normalized* and satisfies *voluntary participation* (VP), if selected buyers never pay more than their declared valuation, selected sellers are never paid less than their valuation and payments to and from non-selected agents are 0. Furthermore, a mechanism is *budget-balanced* (BB) if the sum of payments is always non-negative ($\sum_{i\in S} p_i \le \sum_{j\in B} q_j$), α-approximate *cost-efficient* if it computes trades that are α-approximate with respect to social cost and *truthful* if it is a dominant strategy for every agent to declare their true valuations. A classical result by Myerson and Satterthwaite [10] shows that no truthful 1-approximately cost-efficient mechanism can satisfy both (VP) and (BB). This result extends to approximately cost-efficient mechanisms. The proof of Theorem 11 is found in Appendix C.

Theorem 11 *Fix some $\alpha \ge 1$. Let $M = (A, p, q)$ be a truthful and α-approximately cost-efficient combinatorial exchange mechanism satisfying VP. Then M is not budget-balanced.*

References

[1] N. Alon, D. Moshkovitz, and M. Safra. Algorithmic Construction of Sets for k-Restrictions. *ACM Transactions on Algorithms*, 2:153–177, April 2006.

[2] M. Babaioff and W. E. Walsh. Incentive-Compatible, Budget-Balanced, yet Highly Efficient Auctions for Supply Chain Formation. *Decision Support Systems*, 39:123–149, March 2005.

[3] P. Briest. Frugality and Truthfulness in Approximate Mechanism Design. Diploma Thesis, University of Dortmund, 2004.

[4] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. John Wiley & Sons, 1998.

[5] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[6] U. Feige. A Threshold of ln n for Approximating Set Cover. *J. of the ACM*, 45(4):634–652, 1998.

[7] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. Freeman, 1979.

[8] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. *J. of the ACM*, 49(5):1–26, 2002.

[9] A. Mu'alem and N. Nisan. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. In *Proc. of AAAI/IAAI*, pages 379–384, 2002.

[10] R. B. Myerson and M. A. Satterthwaite. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29:265–281, 1983.

[11] N. Nisan. The Communication Complexity of Approximate Set Packing and Covering. In *Proc. of ICALP*, pages 868–875, 2002.

[12] N. Nisan and I. Segal. The Communication Requirements of Efficient Allocations and Supporting Prices. *Journal of Economic Theory*, 2006.

[13] D. C. Parkes, R. Cavallo, N. Elprin, A. Juda, S. Lahaie, B. Lubin, L. Michael, J. Shneidman, and H. Sultan. ICE: An Iterative Combinatorial Exchange. In *Proc. of EC*, pages 249–258, 2005.

[14] D. C. Parkes, J. Kalagnanam, and M. Eso. Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In *Proc. of IJCAI*, pages 1161–1168, 2001.

[15] S. Rajagopalan and V. V. Vazirani. Primal-Dual RNC Approximation Algorithms for (multi)Set (multi)Cover and Covering Integer Programs. *SIAM J. of Computing*, 28(2):525–540, 1998.

[16] T. Roughgarden and M. Sundararajan. New trade-offs in cost-sharing mechanisms. In *Proc. of STOC*, pages 79–88, 2006.

[17] T. Sandholm. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

[18] T. Sandholm and S. Suri. Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations. In *Proc. of AAAI/IAAI*, pages 90–97, 2000.

[19] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Winner Determination in Combinatorial Auction Generalizations. In *Proc. of AAMAS*, pages 69–76. ACM Press, 2002.

A Proofs from Section 2

A.1 Proofs from Section 2.1

Lemma 1 *If the combinatorial exchange positive surplus problem for a rational and scalable family $\mathcal{F}$ is NP-hard then the combinatorial exchange surplus (volume) problem for the family $\mathcal{F}$ is inapproximable, unless $P = NP$.*

Proof: The hardness of the volume problem is immediate, as we require a trade with positive surplus. We next consider the surplus maximization problem. Assume that for some α, β there exists a poly-time approximation algorithm A for the combinatorial exchange surplus problem such that

$$\mathit{sur}(T^\star_{\mathit{sur}}(I)) \leq \alpha(n,m) \cdot \mathit{sur}(A(I)) + \beta(n,m)$$

holds on every problem instance $I \in \mathcal{F}$ with n agents. Given instance I in which agent i has value v_i we let $r = 1/2^\ell$, where ℓ denotes the maximum number of bits used to encode any valuation in I. We define instance I' to be I after multiplying each v_i by $(1+\beta(n))r^{-1}$ (observe that the number of agents in I' is still n and the number of products is still m, thus we define $\beta = \beta(n,m)$ and $\alpha = \alpha(n,m)$). We denote $OPT = \mathit{sur}(T^\star_{\mathit{sur}}(I))$ and $OPT' = \mathit{sur}(T^\star_{\mathit{sur}}(I'))$. Note that $OPT' = (1+\beta)r^{-1}OPT$, thus $OPT = 0$ if and only if $OPT' = 0$. We apply A on I' and by our assumption

$$OPT' \leq \alpha \cdot \mathit{sur}(A(I')) + \beta.$$

To conclude the proof we show that if $\mathit{sur}(A(I')) = 0$ then $OPT = 0$ and if $\mathit{sur}(A(I')) > 0$ then $OPT > 0$. This contradicts the hardness of deciding if OPT is positive or 0 (unless $P = NP$).

As $OPT' \geq \mathit{sur}(A(I')) \geq 0$, if $OPT' = OPT = 0$ then $\mathit{sur}(A(I')) = 0$. On the other hand, if $OPT > 0$ then $Opt \geq r$ and, thus, $OPT' = (1+\beta)r^{-1}OPT \geq 1+\beta$. We conclude that if $OPT > 0$ then

$$\alpha \cdot \mathit{sur}(A(I')) + \beta \geq OPT' = (1+\beta)OPT \geq OPT + \beta$$

or equivalently $\alpha \cdot \mathit{sur}(A(I')) \geq OPT > 0$, thus, $\mathit{sur}(A(I')) > 0$. □

Lemma 3 *The combinatorial exchange positive surplus problem for the family $\mathcal{F}$ of instances with only a single type of good (multi-unit exchange problem) is NP-hard.*

Proof: We use a simple reduction from the knapsack problem, which is known to be NP-hard [7]. Given n objects with weights $w_1, \ldots, w_n$ and integral utilities $u_1, \ldots, u_n$ we want to decide whether there is a selection of objects with total utility u and weight at most W.

For each object i our multi-unit exchange instance has a buyer i requesting $r_i = w_i$ copies of the traded item and valuation $v_i = u_i$. A single seller supplies $s_1 = W$ copies of the item at price $u - 1$. Clearly, there exists a trade with positive surplus if and only if the desired selection of objects exists. □

As deciding if the packing to covering factor is 1 or some $\gamma > 1$ is equivalent to solving the combinatorial exchanges positive surplus problem, using Lemma 1 we derive the following.

Lemma 7 *If for a rational and scalable family $\mathcal{F}$ it is NP-hard to decide whether the packing to covering factor is at most 1 or at least $\gamma(n,m)$ for $\gamma(m,n) > 1$, then the combinatorial exchange surplus (volume) problem for the family $\mathcal{F}$ of instances with $f(I) \geq \gamma(n,m)$ is inapproximable, unless $P = NP$.*

Lemma 8 *There exists a function $\gamma(n, m) = \Omega(\min\{n^{1-\varepsilon}, m^{1/2-\varepsilon}\})$ such that for combinatorial exchange with n agents and m products it is NP-hard to decide whether the packing to covering factor is at most 1 or at least $\gamma(n, m)$.*

Proof: Consider the set packing problem with n sets $S_1, \ldots, S_n$ over some ground set of size m. For functions $a(n, m)$ and $b(n, m)$ we define $\mathcal{S}_{\leq a(n,m)}$ and $\mathcal{S}_{\geq b(n,m)}$ as the families of instances with optimal packings of size at most $a(n, m)$ or at least $b(n, m)$, respectively. It is known that there exist polynomial time computable functions $a(n, m)$, $b(n, m)$ with $b(n, m) = \gamma(n, m) \cdot a(n, m)$ for some $\gamma(n, m) = \Omega(\min\{n^{1-\varepsilon}, m^{1/2-\varepsilon}\})$, such that given instance $I \in \mathcal{S}_{\leq a(n,m)} \cup \mathcal{S}_{\geq b(n,m)}$ it is NP-hard to decide whether $I \in \mathcal{S}_{\leq a(n,m)}$ or $I \in \mathcal{S}_{\geq b(n,m)}$.

We encode every set as a buyer requesting the respective set at price 1 in the exchange scenario. Furthermore, we have a single seller supplying the complete ground set at price $a(n, m)$. Clearly, the packing to covering factor of the resulting CE instance is at most 1 if $I \in \mathcal{S}_{\leq a(n,m)}$ and $b(n, m)/a(n, m) = \gamma(n, m)$ whenever $I \in \mathcal{S}_{\geq b(n,m)}$. $\square$

Theorem 2 is a direct result from Lemma 7 and Lemma 8.

A.2 Proofs from Section 2.2

We next prove a communication lower bounds for combinatorial exchange surplus and volume problems, we present inapproximability results under sub-exponential communication. The lower bounds are obtained in Yao's standard model of communication complexity. Our first lower bound is based on a lower bound of Nisan and Segal [12] for combinatorial auctions (with a single seller that has no value for the products she holds). Nisan and Segal show that in the combinatorial auctions problem with general valuations, any protocol that always finds the optimal allocation communicates at least $\binom{m}{m/2}$ bits. We show that for combinatorial exchange this implies that achieving *any* approximation to the surplus requires $\binom{m}{m/2}$ bits. This holds even with only two buyers and a single seller. This is in contrast to the combinatorial auctions case, for which a trivial n-approximation exists, when n is the number of buyers, as all items can be assigned to the buyer with the highest valuation for the bundle of all items. We formally define inapproximability for the communication model. Inapproximability means that we need exponential communication in m if we would like the protocol to always get an approximate solution. Our results hold even for a single seller offering a set of items and buyers that have valuations for sets (not multi-sets). For communication protocol P, let *sur*$(P(I))$ be the surplus of the trade chosen by P on instance I.

Definition 4 *The combinatorial exchange surplus problem is* approximable *in t bits if for some $\alpha, \beta : \mathbb{N} \to \mathbb{N}$ there exists a communication protocol P that uses at most t bits, such that* $\mathrm{sur}(T^{\star}_{\mathrm{sur}}(I)) \leq \alpha(m) \cdot \mathrm{sur}(P(I)) + \beta(m)$ *holds on every problem instance I.*

Similar proof to the one of Lemma 1 shows the following.

Lemma 9 *If the combinatorial exchange positive surplus problem for a rational and scalable family $\mathcal{F}$ is non approximable in t bits the combinatorial exchange surplus (volume) problem for the family $\mathcal{F}$ is not approximable in t bits.*

Let $G = (1, \ldots, 1)$ be the bundle with one item of each product. Assume that buyer i has a monotone valuation function $v_i : 2^m \rightarrow \{0, 1\}$. Such a valuation is called a 0/1 valuation, as the value of any bundle is either 0 or 1. Buyer i's valuation of bundle $S_i \subseteq G$ is $v_i(S_i) \in \{0, 1\}$. For combinatorial auctions with a supply G, the optimal allocation problem is the problem of finding a partition of G to S and $S^c = G \setminus S$ that maximizes $v_1(S) + v_2(S^c)$.

Theorem 12 ([12]) *Every protocol that finds the optimal allocation for every pair of 0/1 valuations v_1, v_2 such that $v_1(G) = v_2(G) = 1$, must use at least $\binom{m}{m/2}$ bits of total communication in the worst case.*

We derive the following for the combinatorial exchange problem.

Lemma 10 *Consider a combinatorial exchange with a single seller offering G for the price of 1, and two buyers with $0/1$ valuations v_1, v_2 such that $v_1(G) = v_2(G) = 1$. Every protocol for the combinatorial exchange positive surplus problem must use at least $\binom{m}{m/2}$ bits of total communication in the worst case.*

Proof: We use the straightforward reduction from the optimal allocation problem for combinatorial auctions, to the problem of revenue maximization in combinatorial exchange. Given a pair of 0/1 valuations v_1, v_2 such that $v_1(G) = v_2(G) = 1$, we consider the combinatorial exchange problem with a single seller offering G for the price of 1, and two buyers with valuations v_1, v_2. Clearly, there is a positive revenue exchange if and only if there is an allocation of value 2 for the optimal allocation problem. As the later requires communication of at least $\binom{m}{m/2}$ bits in the worst case (by Theorem 12), so does the former. □

A direct corollary of the above lemma and Lemma 9 is the following.

Theorem 3 *The combinatorial exchange surplus (volume) problem is not approximable in less than $\binom{m}{m/2}$ bits. This holds even with only a single seller and two buyers.*

Note that the theorem holds with or without free disposal.

We next strengthen the result and show that even if there is a very large gain from trade that can be realized, no approximation can be obtained unless exponential communication is used. We use a reduction that is based on the following result by Nisan [11] for combinatorial auctions with n agents and m items.

Theorem 13 (Derived from the proof of Theorem 3 in [11]) *Any n-agent protocol for CA that always distinguishes the case that the surplus is 1 and the case that the surplus is n requires $e^{m/(2n^2)-5log\ n}$ bits of communication. The lower bound holds for randomized and nondeterministic protocols. It also holds even if all buyers have 0/1 valuations.*

In particular, as long as $n < m^{1/2-\epsilon}$, the communication complexity is exponential in m. Thus we show that if we only consider instances for which $f(I) \geq n$ (and $n < m^{1/2-\epsilon}$), the communication complexity of any approximation is exponential in m.

From Theorem 13 we derive the following for the combinatorial exchange problem.

Lemma 11 *Any n-agent protocol for the combinatorial exchange positive surplus problem for the family of instances with packing to covering factor $f(I) \geq n$ requires $e^{m/(2n^2)-5log\ n}$ bits of communication. The lower bound holds for randomized and nondeterministic protocols.*

Proof: The proof of Nisan [11] for the above theorem shows that $e^{m/(2n^2)-5log\ n}$ bits of communication are required to decide if the surplus is 1 or n in combinatorial auctions with buyers having 0/1 valuations. Given such an instance of combinatorial auctions, we construct and instance of combinatorial exchange by keeping the same set of buyers and adding a seller with the set $G = (1, \ldots, 1)$ and value of 1. Clearly for such an instance $f(I) \geq n$ and any positive value trade trivially corresponds to a solution for the combinatorial auctions with value n. □

A direct corollary of the above lemma and Lemma 9 is the following.

Theorem 14 *The combinatorial exchange surplus (volume) problem for the family of instances with packing to covering factor $f(I) \geq n$ is not approximable with less than $e^{m/(2n^2)-5log\ n}$ bits. The lower bound holds for randomized and nondeterministic protocols.*

Corollary 1 *For any $\epsilon > 0$, the combinatorial exchange surplus (volume) problem with packing to covering factor $f(I) \geq n$ when $n < m^{1/2-\epsilon}$ is not approximable by a protocol that uses sub-exponential communication in m.*

Note that even if $f(I)$ is infinity the surplus (volume) problem cannot be approximated better than CA can be approximated, as combinatorial auctions reduce to combinatorial exchange with a single seller selling G for 0.

B Proofs from Section 3

B.1 Proofs from Section 3.1

Theorem 6 *The combinatorial exchange social cost problem cannot be approximated in polynomial time better than within $(1 - o(1)) \ln m$, unless NP $\subseteq$ DTIME($n^{\mathcal{O}(\log \log n)}$).* [3]

Proof: It is known that the set cover problem cannot be approximated below a threshold of $(1-o(1)) \ln m$ (m being the size of the ground set) in polynomial time, unless NP $\subseteq$ DTIME($n^{\mathcal{O}(\log \log n)}$) [6]. Let $S_1, \ldots, S_n$ be a collection of subsets of $\mathcal{U} = \{1, \ldots, m\}$. The set cover problem asks for a selection of as few as possible of the subsets covering $\mathcal{U}$.

Assume now that we are given algorithm A with approximation ratio α for social cost minimization in combinatorial exchanges. We show that this implies an α-approximation for set cover, as well. Consider the following combinatorial exchange instance. For each S_i we have a seller i with $q_i^e = 1$ if $e \in S_i$ ($q_i^e = 0$ else) and $v_i = 1$. Additionally, there is a single buyer with $q_b = (1, \ldots, 1)$ and $v_b = (\alpha+1)n$, i.e., buyer b is offering to buy the whole ground set $\mathcal{U}$ at price $(\alpha+1)$ times the number of subsets in the set cover instance.

We assume w.l.o.g. that $\mathcal{U}$ can be covered by selecting all the subsets. Thus, selecting buyer b and all the sellers defines a feasible trade of social cost n. On the other hand, not selecting buyer b results in the empty trade and causes cost $(\alpha + 1)n$. Thus, the approximation ratio of algorithm A guarantees that a non-empty trade must be returned.

[3] Note that we can replace this assumption by $P \neq NP$ if we relax the lower bound to $\Omega(\ln m)$ by [1].

However, there is a one-to-one correspondence between non-empty trades and feasible covers of ground set $\mathcal{U}$. Thus, if we denote by opt the cost of the minimum cover of $\mathcal{U}$, this is also the social cost of the optimal trade. It follows that the trade returned by A has social cost at most $\alpha \cdot opt$ and we obtain a cover of identical cost.

Hence, every α-approximate algorithm for social cost minimization immediately yields the same approximation ratio for set cover and the known hardness results for set cover can be applied. $\square$

B.2 MINKNAPSACK Algorithm missing from Section 3.2

1. Let $\Delta = \sum_{j \in \mathcal{B}} q_j$.
2. Apply the dynamic programming based FPTAS to the following min-knapsack problem:

$$\begin{aligned} \text{min.} \quad & \sum_{i \in \mathcal{S}} v_i x_i + \sum_{j \in \mathcal{B}} v_j y_j \\ \text{s.t.} \quad & \sum_{i \in \mathcal{S}} q_i x_i + \sum_{j \in \mathcal{B}} q_j y_j \geq \Delta \\ & x_i, y_j \in \{0,1\} \end{aligned}$$

3. Let $S = \{i \,|\, x_i = 1\}$, $B = \{j \,|\, y_j = 0\}$. If $\sum_{j \in B} v_j \geq \sum_{i \in S} v_i$ return trade $T = (S, B)$, else return $T = (\emptyset, \emptyset)$.

Figure 3: Approximating social cost in multi-unit exchanges by algorithm MINKNAPSACK.

C Proofs from Section 5

In this section we discuss the existence of truthful mechanisms for the combinatorial exchange problem. A mechanism consist of some algorithm A that outputs a trade $T = (S, B)$ and additional payments $(p_i)_{i \in \mathcal{S}}$, $(q_j)_{j \in \mathcal{B}}$ determining the payments given to sellers and collected from buyers, respectively. For ease of notation we denote input as $v = (v_1, \ldots, v_{|\mathcal{S}|})$ and $w = (w_1, \ldots, w_{|\mathcal{B}|})$. Trade and payments are referred to as $A(v, w)$, $p(v, w)$ and $q(v, w)$, respectively. For both sellers and buyers $i \in \mathcal{B}, \mathcal{S}$ we denote by $i \in A(v, w)$ the situation that i is selected for the trade. We are interested in mechanisms with the following properties:

- *Voluntary Participation (VP):* For all sellers $i \in S$, buyers $j \in B$ we have that $p_i \geq v_i$ and $q_j \leq w_j$, respectively.
- *Normalization:* Payments to and from sellers and buyers that are not selected for the trade are 0, i.e., $p_i(v, w) = q_j(v, w) = 0$ whenever $i, j \notin A(v, w)$.
- *Budget-Balance (BB):* Let $T = (S, B)$ be the computed trade. It holds that $\sum_{i \in S} p_i \leq \sum_{j \in B} q_j$.

- *α-approximate Cost-Efficiency:* The trade computed by algorithm A is an α-approximation w.r.t. social cost.
- *Truthfulness:* For all sellers and buyers let v_i^*, w_j^* denote their *true* valuations. The *utility* of seller i declaring valuation v_i to the mechanism is defined as

  $$u_i(v_i) = p_i((v_i, v_{-i}), w) - v_i^*$$

  if $i \in S$ and $u_i(v_i) = 0$ else. Analogously, buyer j has utility $u_j(v_j) = v_j^* - q_j(v, (w_j, w_{-j}))$ if $j \in B$, $u_j(v_j) = 0$ else. A mechanism (A, p, q) is truthful, if

  $$u_i(v_i^*) \geq u_i(v_i) \quad \forall \quad v_i \neq v_i^* \quad \text{and} \quad u_j(w_j^*) \geq u_j(w_j) \quad \forall \quad w_j \neq w_j^*$$

  holds for all sellers $i \in \mathcal{S}$ and buyers $j \in \mathcal{B}$.

A classical result by Myerson and Satterthwaite [10] shows that no truthful 1-approximately cost-efficient mechanism can satisfy both (VP) and (BB). In the remainder of this section, we extend this result to general α-approximate cost-efficiency.

Theorem 11 *Fix some $\alpha \geq 1$. Let $M = (A, p, q)$ be a truthful and α-approximately cost-efficient combinatorial exchange mechanism satisfying VP. Then M is not budget-balanced.*

The theorem follows from the characterization of truthful combinatorial auctions in [9], which generalizes naturally to the exchange scenario.

C.1 Characterization of Truthfulness

Definition 5 *Algorithm A for combinatorial exchanges is said to be* monotone*, if*

- *$i \in A((v_i, v_{-i}), w)$ implies that $i \in A((v_i', v_{-i}), w)$ for any $v_i' \leq v_i$ for every seller i.*
- *$j \in A(v, (w_j, w_{-j}))$ implies that $j \in A(v, (w_j', w_{-j}))$ for any $w_j' \geq w_j$ for every buyer j.*

The following is a straightforward observation.

Lemma 12 *Let A be a monotone algorithm for combinatorial exchanges. Then for each seller $i \in \mathcal{S}$ and for each buyer $j \in \mathcal{B}$ there exist* critical values *θ_i and θ_j (independent of v_i and w_j, respectively), such that*

- *$i \in A((v_i, v_{-i}), w)$ iff $v_i \leq (<)\theta_i$.*
- *$j \in A(v, (w_j, w_{-j}))$ iff $w_j \geq (>)\theta_j$.*

Theorem 15 extends the characterization of truthful auctions in the known single-minded case to the exchange scenario.

Theorem 15 *A normalized mechanism $M = (A, p, q)$ satisfying (VP) for combinatorial exchanges is truthful if and only if algorithm A is monotone and p, q are based on critical values.*

Proof: "$\Leftarrow$": Consider mechanism $M = (A, p, q)$, where A is monotone algorithm and p, q the corresponding critical value payment schemes. Fix seller i with true valuation v_i^* and consider the cases $i \in A((v_i^*, v_{-i}), w)$ and $i \notin A((v_i^*, v_{-i}), w)$, respectively.

If $i \in A((v_i^*, v_{-i}), w)$, then seller i is paid her critical value θ_i and has utility $\theta_i - v_i^* \geq (>)0$. Declaring $v_i \neq v_i^*$, the utility remains unchanged for any $v_i \leq (<)\theta_i$ and is 0 for any $v_i > (\geq)\theta_i$. If $i \notin A((v_i^*, v_{-i}), w)$, then we have that $v_i^* > (\geq)\theta_i$ and seller i's utility is 0. Declaring $v_i \neq v_i^*$, the utility remains unchanged for any $v_i > (\geq)\theta_i$ and is $\theta_i - v_i^* < (\leq)0$ for any $v_i \leq (<)\theta_i$. Similar arguments apply to buyers, as well.

"$\Rightarrow$": Consider $M = (A, p, q)$ and assume that algorithm A is not monotone w.r.t. to some seller i. Thus, there exist valuations $v_i^1 < v_i^2$, such that $i \notin A((v_i^1, v_{-i}), w)$ and $i \in A((v_i^2, v_{-i}), w)$. By the (VP) condition we know that $p_i((v_i^2, v_{-i}), w) \geq v_i^2$. Assume that seller i's true valuation is v_i^1. It follows that

$$u_i(v_i^2) = p_i((v_i^2, v_{-i}), w) - v_i^1 \geq v_i^2 - v_i^1 > 0 = u_i(v_i^1),$$

and, thus, M is not a truthful mechanism. A similar argument applies if A is not monotone w.r.t. some buyer.

Finally, assume that A is monotone, but p is not based on its critical values. Fix again seller i, let the declaration of all other buyers and sellers be given and $p_i \neq \theta_i$. First observe that $p_i < \theta_i$ is impossible due to the (VP) condition, since for declarations between p_i and θ_i seller i would be selected for the trade but paid less than her actual declaration. Let then $\theta_i < p_i$. If seller i has true valuation $\theta_i < v_i^* < p_i$, then she has utility $u_i(v_i^*) = 0$ when truthfully declaring v_i^*, but can increase it to $u_i(v_i) = p_i - v_i^* > 0$ with any false declaration $v_i \leq (<)\theta_i$. It follows that M is not a truthful mechanism. Again buyers can be treated analogously. $\square$

We consider the situation that we have just one seller and one buyer interested in trading a single good. Let us assume that $M = (A, p, q)$ is truthful α-approximately cost-efficient mechanism. Let us denote the bid of the seller by v, the buyer's bid by w. By Lemma 12 (Appendix C) A defines critical values $\theta_s(w)$ and $\theta_b(v)$, such that both participants are selected for the trade whenever $v \leq \theta_s(w)$ and $w \geq \theta_b(v)$. Furthermore, in this case the critical values are exactly the prices at which the good is traded.

Lemma 13 *Let A be α-approximately cost-efficient monotone algorithm for combinatorial exchanges. Consider a single seller and buyer trading one good as described above and let $\theta_s(w)$, $\theta_b(v)$ be their critical values defined by A. It holds that $\theta_s(w) \geq \alpha^{-1} \cdot w$ and $\theta_b(v) \leq \alpha \cdot v$.*

Proof: The mechanism always returns either no trade or the trade that selects both the seller and buyer resulting in total cost w or v, respectively. Since our mechanism is α-approximately cost efficient, we know that it must choose the second alternative whenever $w > \alpha \cdot v$. Rearranging for v and w, respectively, yields the claim. $\square$

This basically concludes the proof of Theorem 11. Fix any declarations v and w with $w > \alpha^2 v$. Then $\theta_b(v) - \theta_s(w) \leq \alpha \cdot v - \alpha^{-1} \cdot w < 0$, and, thus, the mechanism is not budget-balanced.