

Arbitrage in Combinatorial Exchanges^{*}

Andrew Gilpin and Tuomas Sandholm

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract. Combinatorial exchanges are trading mechanisms that allow agents to specify preferences over bundles of goods. When agents' preferences exhibit complementarity and/or substitutability, this additional expressiveness can lead to more efficient allocations than is possible using traditional exchanges. In the context of combinatorial exchanges, this paper examines *arbitrage*, a risk-free profit opportunity. We show that some combinatorial exchanges allow agents to perform arbitrage and thus extract a positive payment from the market while contributing nothing, something that is not possible in traditional exchanges. We analyze the extent to which arbitrage is possible and computationally feasible in combinatorial exchanges. We show that the surplus-maximizing combinatorial exchange with free disposal is resistant to arbitrage, but without free disposal arbitrage is possible. For volume-maximizing and liquidity-maximizing combinatorial exchanges, we show that arbitrage is sometimes possible and we propose an improved combinatorial exchange that achieves the same economic objective but eliminates a particularly undesirable form of arbitrage. We show that the computational complexity of detecting winning arbitraging bids is \mathcal{NP} -complete and that the ability for an agent to submit arbitraging bids depends on the type of feedback in the exchange. We also show that a variant of combinatorial exchanges in which arbitrage is impossible becomes susceptible to arbitrage if certain side constraints are placed on the allocation or if an approximating clearing algorithm is used.

1 Introduction

A combinatorial exchange (CE) is a trading mechanism where agents are able to specify preferences over bundles of good. This additional expressiveness can lead to more economically desirable outcomes than is possible in traditional exchanges. However, CEs also present problems not found in traditional markets that need to be studied before CEs can safely be put to use.

Despite the theoretical advantages inherent in CEs, they have been of little practical use thus far. CEs have been proposed by the Federal Communications Commission (FCC) as a method of efficiently trading wireless spectrum licenses.

^{*} This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, ITR IIS-0121678, ITR IIS-0427858, and a Sloan Fellowship.

The FCC is currently in the process of designing a CE suitable for their purposes, but is still probably several years away from implementing a CE. BondConnect, a combinatorial exchange for bond trading, was deployed, but is no longer in use. For a related description of CEs in the context of financial markets see [1]. In this infant stage of CE development, there are still many challenges for researchers.

One of the challenges faced by implementers of CEs is that most forms of the clearing problem are \mathcal{NP} -complete [2]. The first special-purpose algorithm for clearing CEs was presented in [3] and clearing time for different variants of this problem was experimentally studied in [2]. Another challenge in CEs is that of preference elicitation [4], in which the market maker solicits information from participating agents so as to achieve the desired economic objective while minimizing communication and/or computational costs. Yet another challenge is mechanism design for CEs [5, 6], in which the rules governing the exchange are carefully designed so as to achieve various economic goals.

In this paper, we study an additional problem that arises in CEs: arbitrage. Typically, agents participating in a CE possess an endowment of money or some subset of the goods (or both) and wish to exchange goods and money with other agents so as to increase their well-being (i.e. utility). This paper examines how well an agent that initially has no endowment can do in various combinatorial marketplaces. Put simply, where are the free lunches in CEs?

The computational complexity of arbitrage has previously been studied for various frictional markets [7], including frictional markets in foreign exchange [8] and frictional markets in securities [9]. A frictional market is defined as a market in which assets are traded in integer numbers and when there is a maximum limit on the number of assets that can be traded at a fixed price. Via a reduction from 3SAT, it has been shown that it is strongly \mathcal{NP} -hard to compute an arbitrage opportunity in frictional markets. This result has important consequences because a fundamental assumption in many theories of finance is the arbitrage-free assumption (for example, see [10]). This assumption is based on the belief that if an arbitrage opportunity ever existed, it would disappear in an arbitrarily short period of time. If generating an arbitrage opportunity is a computationally hard problem, then this assumption may not hold in practice. In this paper, we show that many problems related to arbitrage in CEs are also hard and depend heavily on the amount of information available to participating agents. Thus, any theories about CEs should be cautious when making assumptions regarding the existence of arbitrage opportunities in CEs.

2 Combinatorial Exchanges

In this section, we formally define our model of a combinatorial exchange, and define the various different types of exchanges that we treat in this paper. We also present hardness results for the clearing problems of each type of exchange.

2.1 Model

We assume the presence of an exchange administrator (i.e., mediator) that runs the centralized clearing algorithm. Our model for a CE consists of $M =$

$\{1, 2, \dots, m\}$ goods, determined by an exchange administrator. Goods can be discrete (e.g. automobiles) or continuous (e.g. oil). In this paper, we do not make any distinction among these two possibilities; all of our results apply to markets with discrete or continuous goods.

Definition 1. A combinatorial bid is a tuple $B_j = \langle (\lambda_j^1, \lambda_j^2, \dots, \lambda_j^m), p_j \rangle$ with $\lambda_j^k \in \mathbb{R}$ and $p_j \in \mathbb{R}$. λ_j^k is the amount of good k demanded by bid j (negative values indicate supply). Similarly, p_j indicates the amount the agent will pay for the bundle (a negative value indicates that the agent requires a payment).

The following special cases of combinatorial bids will be discussed later in the paper.

Definition 2. A combinatorial bid $B_j = \langle (\lambda_j^1, \lambda_j^2, \dots, \lambda_j^m), p_j \rangle$ with either $\lambda_j^i \geq 0$ for each i , or $\lambda_j^i \leq 0$ for each i , is called a pure combinatorial bid (i.e., it is purely supplying or purely demanding). Otherwise, it is called a mixed combinatorial bid (i.e., it is both supplying and demanding).

Definition 3. A combinatorial bid $B_j = \langle (\lambda_j^1, \lambda_j^2, \dots, \lambda_j^m), p_j \rangle$ is called a single-item bid if only one of λ_j^i is non-zero and the rest are zero.

Definition 4. A combinatorial bid $B_j = \langle (\lambda_j^1, \dots, \lambda_j^m), p_j \rangle$ with $\lambda_j^i \in \{-1, 0, 1\}$ is called a single-unit bid. A CE consisting of only single-unit bids is called a single-unit CE.

Given a collection of combinatorial bids, the exchange administrator is faced with the problem of clearing the exchange. There are several variations of this problem that we treat in this paper: the *surplus-maximizing CE clearing problem*, the *liquidity-maximizing CE clearing problem*, and the *volume-maximizing CE clearing problem*.

Definition 5. Given a set $\mathcal{B} = \{B_1, B_2, \dots, B_n\}$ of combinatorial bids, the general CE clearing problem is to label the bids as winning or losing so as to maximize an objective function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ subject to the constraint that demand does not exceed supply:

$$\begin{aligned} \max \quad & f(x) \\ \text{such that} \quad & \sum_{j=1}^n \lambda_j^i x_j \leq 0 \quad i = 1, 2, \dots, m \\ & x_j \in \{0, 1\} \quad j = 1, 2, \dots, n \end{aligned}$$

In the situation where good i is not freely disposable (i.e. buyers are not willing (or unable) to take extra units, and sellers are not willing (or unable) to keep any units of their winning bids), then the inequality in the first constraint is replaced with an equality.

Definition 6. The surplus-maximizing CE clearing problem is the general CE clearing problem with $f(x) = \sum_{j=1}^n p_j x_j$.

Definition 7. *The liquidity-maximizing CE clearing problem is the general CE clearing problem with $f(x) = \sum_{j=1}^n x_j$.*

Definition 8. *The volume-maximizing CE clearing problem is the general CE clearing problem with $f(x) = \sum_{i=1}^m \sum_{j=1}^n |\lambda_j^i| x_j$.*

2.2 Complexity of Clearing

This paper is primarily concerned with computational complexity results related to CEs. An important starting point is the computational complexity of the various clearing problems. We begin with the following result from [2].

Proposition 1. *[2] The decision version of the surplus-maximizing CE clearing problem with free disposal is \mathcal{NP} -complete.*

This was proven in [2] by noting that clearing a combinatorial auction with free disposal, which is a special case of a surplus-maximizing CE with free disposal, is \mathcal{NP} -complete [11] since it is equivalent to WEIGHTED SET PACKING. This proof only applies to the surplus-maximizing CE with free disposal clearing problem, but as the following two propositions show, the clearing problem is also \mathcal{NP} -complete for the liquidity-maximizing and volume-maximizing versions of the CE.

Proposition 2. *The decision version of the liquidity-maximizing CE clearing problem with free disposal is \mathcal{NP} -complete.*

Proposition 3. *The decision version of the volume-maximizing CE clearing problem with free disposal is \mathcal{NP} -complete.*

So far, we have shown that the clearing problems as defined in Definitions 6-8 are \mathcal{NP} -complete for the case where there is free disposal. Proposition 4 shows that even finding a *non-trivial* feasible allocation is \mathcal{NP} -complete for the case without free disposal.

Definition 9. *Let $\mathcal{B} = \{B_1, B_2, \dots, B_n\}$ be the bids in a CE. Let x_j be the binary decision variable for B_j in some feasible allocation. We say that the allocation is a *non-trivial* feasible allocation if there is at least one bid B_j in the allocation such that $x_j \neq 0$ (i.e., at least one item is traded).*

Proposition 4. *[2] Without free disposal, finding a non-trivial feasible allocation for each of the problems in Definitions 6-8 is \mathcal{NP} -complete, even if there is only one item.*

To summarize Propositions 1-4, each of the problems defined in Definitions 6-8 with free disposal is \mathcal{NP} -complete. Without free disposal, even finding a non-trivial feasible allocation is \mathcal{NP} -complete.

In the liquidity-maximizing and volume-maximizing CE clearing problems, it is possible for a negative surplus to occur. Some exchange administrators may wish to constrain the allocations so that the allocation is *budget-nonnegative*.

Definition 10. Let $x = \{x_1, \dots, x_n\}$ be the allocation given by an exchange clearing algorithm. If $\sum_{j=1}^n p_j x_j \geq 0$ then the allocation is said to be budget-nonnegative.

The surplus-maximizing CE clearing problem always results in a budget-nonnegative allocation. It can easily be shown that that each of the clearing problems remains \mathcal{NP} -complete if the budget-nonnegative constraint is added.

3 Existence of Arbitrage in Combinatorial Exchanges

In this section we formally define what arbitrage is in the context of a CE. We discuss types of exchanges for which arbitrage is provably impossible, and discuss other exchanges for which arbitrage is sometimes possible and sometimes impossible. We also present a new CE clearing method that curtails some types of arbitrage.

Definition 11. A combinatorial bid $B_j = \langle (\lambda_j^1, \lambda_j^2, \dots, \lambda_j^m), p_j \rangle$ is an arbitraging bid if the following two conditions hold:

1. $\lambda_j^i \geq 0$ for all i (i.e. the bid is not supplying any goods);
2. $p_j < 0$ (i.e. the bid, if accepted, gives a positive payoff to the agent who placed the bid).

The way we have defined arbitrage in a CE reinforces the point that arbitrage is *risk-free* to the agent submitting the bid. It is possible for an agent to receive a positive payment from an exchange without providing any net supply of goods to the exchange while not submitting any arbitraging bids as defined above. This could be accomplished, for example, by submitting a set of bids, none of which is arbitraging, but in aggregate the sum of the prices of the bids is negative and the sum of the quantities of the items is nonnegative. However, this is not a risk-free situation for the agent, and so we do not consider it as arbitrage.

Before discussing the possibility and impossibility of arbitrage in combinatorial exchanges, it is worth first noting that arbitrage is never possible in a combinatorial auction or combinatorial reverse auctions.

3.1 Impossibility of Arbitrage

An arbitraging bid, if accepted, results in a positive payoff to the agent that submitted it, and the agent does not supply any goods. This is indeed a very profitable situation for the agent, possibly at the expense of other agents in the exchange. However, as the following theorem shows, in some variants it is impossible for an agent to place a winning arbitraging bid.

Theorem 1. (*Main Impossibility Result*) In a surplus-maximizing CE with free disposal, an arbitraging bid will never be in an optimal allocation.

Proof. Let $x^* = \{x_1^*, x_2^*, \dots, x_n^*\} \in \{0, 1\}^n$ be the optimal (surplus-maximizing) allocation. By way of contradiction, say that $x_k^* = 1$ and B_k is an arbitraging bid. Consider what happens when we change the allocation to have $x_k^* = 0$ (i.e. we remove the arbitraging bid from the allocation). For some $i \in \{1, 2, \dots, m\}$, consider the constraint $\sum_{j=1}^n \lambda_j^i x_j \leq 0$ in the clearing problem. Since $\lambda_k^i \geq 0$, we have $\sum_{j=1}^{k-1} \lambda_j^i x_j^* + \sum_{j=k+1}^n \lambda_j^i x_j^* \leq \sum_{j=1}^n \lambda_j^i x_j^* \leq 0$. So the allocation remains feasible. Since $p_k < 0$, subtracting p_k from the objective value results in a strictly higher surplus. Thus the arbitraging bid x_k can be discarded while raising surplus. This contradicts the optimality of the allocation x^* . \square

3.2 Possibility of Arbitrage

In contrast to the impossibility results above, the following example shows that an arbitraging bid may be accepted in surplus-maximizing CEs without free disposal and in liquidity-maximizing and volume-maximizing CEs with or without free disposal.

Example 1. Consider the following CE:

$$M = \{1, 2\}, \mathcal{B} = \{B_1, B_2\}, B_1 = \langle (-1, 0), -8 \rangle, B_2 = \langle (1, -1), 10 \rangle.$$

(In the first bid, an agent is offering to sell one unit of good 1 for \$8. In the second bid, an agent wishes to buy one unit of good 1 and sell one unit of good 2, and is offering to pay \$10.)

In any of the CEs defined in Definitions 6-8 with no free disposal for good 2, the exchange administrator is unable to clear the market since there is no available demand for the second good. Now consider what happens if an agent submits the bid $B_3 = \langle (0, 1), -1 \rangle$ indicating that she will accept one unit of good 2, and is asking for \$1. This is clearly an arbitraging bid, but since it satisfies the no free disposal constraint that demand equals supply, it allows the market to clear. So the arbitraging agent receives one unit of good 2, and also receives \$1.

Now consider a liquidity-maximizing exchange with free disposal. Considering only bids B_1 and B_2 the exchange will clear with a trading volume of 2 (both bids are accepted). When we add bid B_3 , the exchange will clear with a trading volume of 3. Again, the arbitraging agent is receiving a positive payment yet is not supplying anything. This same example can be used to show that arbitrage is possible in volume-maximizing exchanges with free disposal.

An argument could be made that arbitrage in an exchange without free disposal is unlikely to occur for the very reason that ownership (or disposal) of the good entails some costs. However, in some situations the presence of free disposal may be different among different agents. If one agent is able to dispose of the good for free, then she would be able to gain from arbitraging bids.

The above example simply showed that an arbitrage opportunity existed for certain instances of CEs. As the following example illustrates, even if arbitrage is possible in a given variant, there may be no arbitrage opportunities available to an arbitraging agent in a given instance.

Example 2. Consider the surplus-maximizing CE without free disposal and with $M = \{1, 2\}$ and $\mathcal{B} = \{B_1, B_2, B_3\}$ with $B_1 = \langle(-1, 0), -8\rangle$, $B_2 = \langle(1, -1), 10\rangle$, and $B_3 = \langle(0, 1), 2\rangle$. The exchange administrator faces the following optimization problem:

$$\begin{aligned} \max \quad & -8x_1 + 10x_2 + 2x_3 \\ \text{such that} \quad & -x_1 + x_2 = 0 \\ & -x_2 + x_3 = 0 \\ & x_1, x_2, x_3 \in \{0, 1\} \end{aligned}$$

The optimal solution to this problem is to accept all three bids which results in a surplus of 4. Consider an agent trying to place an arbitraging bid in this exchange. The agent can create an arbitraging bid to either accept item 1, item 2, or both items. Consider the problem faced by the agent if she tries to arbitrage by accepting item 1:

$$\begin{aligned} \text{find} \quad & p \\ \text{such that} \quad & -8x_1 + 10x_2 + 2x_3 - p \geq 4 \\ & -x_1 + x_2 = 1 \\ & -x_2 + x_3 = 0 \\ & x_1, x_2, x_3 \in \{0, 1\} \end{aligned}$$

Routine computation shows that this problem is infeasible. Similar computations show the same for when the agent tries to compute an arbitraging bid accepts item 2 or accepts both items. Thus there is no arbitrage opportunity currently available in this exchange.

Unlike surplus-maximizing exchanges, the objective function in liquidity-maximizing and volume-maximizing exchanges does not depend on the prices of bids. As we will see below, this difference has important consequences in the existence of arbitrage opportunity.

One aspect of allocations that has been ignored until now is that of *uniqueness*. In general, there can be multiple *optimal* allocations. This is especially likely to occur in liquidity-maximizing and volume-maximizing exchanges. Since bid prices are irrelevant in these exchanges, bids on the same bundles are treated equally regardless of their prices. The following theorem gives a sufficient condition for the existence of an arbitrage opportunity.

Theorem 2. (*Main Possibility Result*) *In liquidity- and volume-maximizing exchanges, if there is at least one pure demand bid in an optimal allocation, then there always exists an arbitrage opportunity where the arbitraging bid can be accepted in an equally optimal allocation. This holds with and without free disposal and holds even if the exchange is constrained to be budget-nonnegative.*

Proof. Let $B_j \in \mathcal{B}$ be a pure demand bid that is currently in an optimal allocation. So $\lambda_j^i \geq 0$ for all $i \in M$. We can create an arbitraging bid B_k with $\lambda_k^i = \lambda_j^i$ for all $i \in M$ that can be interchanged with B_j in the optimal allocation. Thus we can set p_j to any value we choose. If the exchange is constrained to be budget-nonnegative, then we can set the price equal to the negative of the current surplus. \square

In addition to the possibility of arbitrage that is present in the above CE variants, it is easy to see that in any CE that uses a non-optimal (*i.e.* approximating) clearing algorithm, arbitrage is always a possibility.

3.3 Curtailing Arbitrage Opportunities

Theorem 2 shows that in very common circumstances an arbitrage opportunity exists. What allows this to happen is the fact that in liquidity-maximizing and volume-maximizing exchanges, bids on the same bundle are treated equally regardless of their prices. A natural question arises: can we define liquidity-maximizing and volume-maximizing exchanges in which this type of arbitrage is impossible? We answer that question in the affirmative for the variant with free disposal by showing that a modification of the liquidity-maximizing and volume-maximizing exchanges results in a new exchange that achieves the same economic objective, but also eliminates the possibility of this type of arbitrage.

Definition 12. Let $f(x)$ be defined as in Definitions 6-8. Let x^* be the allocation determined using $f(x)$ as the objective function and the optimization problem defined as in Definition 5. Let $z = f(x^*)$ be the objective value. Now consider the following optimization problem:

$$\begin{aligned} \max \quad & \sum_{j=1}^n p_j x_j \\ \text{such that} \quad & \sum_{j=1}^n \lambda_j^i x_j \leq 0 \quad i = 1, 2, \dots, m \\ & f(x) \geq z \\ & x_j \in \{0, 1\} \quad j = 1, 2, \dots, n \end{aligned}$$

This is the surplus-constrained CE clearing problem. As with the general CE clearing problem, in the situation where good i is not freely disposable (*i.e.* buyers are not willing (or unable) to take extra units, and sellers are not willing (or unable) to keep any units of their winning bids), then the inequality in the first constraint is replaced with an equality.

Theorem 3 shows that in a surplus-constrained CE, the type of situation described in Theorem 2 can never occur.

Theorem 3. In a surplus-constrained CE with free disposal, an arbitraging bid will never be in an optimal allocation when a non-arbitraging bid on the same bundle is not in the optimal allocation. This holds with and without free disposal.

Proof. Assume, by way of contradiction, that there is an arbitraging bid that is in an optimal allocation and there is a non-arbitraging bid on the same bundle that is not in the optimal allocation. Let B_k be the arbitraging bid and let B_j be the non-arbitraging bid on the same bundle. Replacing B_k with B_j in the allocation results in a strictly higher surplus. Since the bids are on the same bundle, the new allocation is feasible and it achieves both the same level of volume and liquidity. \square

Although Theorem 3 does not rule out the possibility of arbitrage, it does perform a form of tie-breaking that eliminates nondeterminism in the clearing problem.

4 Detecting Arbitrating Bids

As discussed in the previous section, the presence of arbitrating bids can lead to less desirable outcomes. The exchange administrator may be interested in detecting arbitrating bids so as to remove them from consideration or to identify agents attempting arbitrage. Also, identifying arbitrage opportunities can be used by the exchange administrator as an analysis tool to indicate segments of the market where demand is weak. If a bundle of goods is identified as allowing for an arbitrage opportunity, this is a signal to the exchange administrator that there is insufficient demand on the items in the bundle by agents with a positive valuation of the goods. Furthermore, identifying arbitrage opportunities can lead to a new form of feedback to agents participating in the exchange. Instead of just providing price feedback, agents can be told which areas of the market are currently lacking competition.

Exchanges can be parameterized based on whether or not bids are allowed to be deleted. Clearly, allowing for the deletion of winning bids can change the current allocation.¹ From the perspective of the clearing problem objective, the exchange administrator may do better by not removing arbitrating bids. (The exchange administrator certainly does no worse by leaving them all in.) For this reason, we advocate not allowing for the deletion of bids.

Finding an arbitrating bid is trivial: simply check each bid to see if it satisfies the conditions specified in Definition 11. However, locating an arbitrating bid in the input does not indicate whether or not the arbitrating bid will be in an optimal allocation, and thus does not indicate anything about the strength or weakness of demand in the market. In order to compute the outcome of placing a given arbitrating bid, we need to take into account all of the other submitted bids. In fact, as Theorem 4 below indicates, determining if a given arbitrating bid is winning is as hard as clearing the market.

Before presenting Theorem 4, consider a CE where the exchange administrator continuously solves the market-clearing problem so as to give feedback to the agents participating. In this case, when determining if a submitted bid is a winning arbitrating bid, the exchange administrator already knows the current optimal allocation. A natural question arises: is it easier to determine if the new bid is a winning arbitrating bid when the current optimal allocation is known? Clearly this problem is no harder than answering the same question when the current optimal allocation is unknown. As the following theorem shows, this problem is hard even if the current optimal allocation is known.²

¹ Unless certain assumptions are made about the clearing algorithm, allowing for the deletion of losing bids can also change the current allocation.

² In practice, however, it is likely that knowing the current optimal allocation will allow the exchange administrator to solve the new problem more quickly.

Theorem 4. *Given a set of combinatorial bids and a corresponding optimal allocation, determining if there is a better allocation that contains a newly submitted arbitraging bid is \mathcal{NP} -complete for any of the clearing problems defined in Definitions 6-8, with and without free disposal.*

Proof. We show that a polynomial-time algorithm for the problem implies the existence of a polynomial-time algorithm for SUBSET SUM. For simplicity, we only prove this for the liquidity-maximizing CE. The proof is similar for the other types of exchanges.

Assume, by way of contradiction, that we have a polynomial-time algorithm \mathcal{A} for the problem. Given a collection C of finite subsets of some domain X and an integer K , we first create an item for each element in X . We also create a single-item supply bid for each item. (Recall that an arbitraging bid is never a supply bid.) Given just these single-item bids, an optimal allocation will be to accept all of the bids. Thus we have the current optimal allocation. Now examine the first set $c \in C$. In the reduction, we make a pure demand bid with one unit of demand for each item in c . Using our algorithm \mathcal{A} , we can determine if this new bid (which can be considered an arbitraging bid) is in a new optimal allocation. Thus we have a new optimal allocation considering this bid. We can repeat this process for each additional $c \in C$. Since, by assumption, \mathcal{A} runs in polynomial-time, we can solve SUBSET SUM in polynomial time. This is true if and only if $\mathcal{P} = \mathcal{NP}$. \square

5 Generating Arbitraging Bids

In the previous section we showed that it is hard for the exchange administrator to determine if a given arbitraging bid is guaranteed to be winning (i.e., in an optimal allocation). In this section, we discuss the problem faced by an agent that is attempting to arbitrage (i.e., to generate an arbitraging bid).

CEs can be parameterized according to the type of feedback given to participating agents. We define four types of feedback: NONE, OWN-WINNING, ALL-WINNING, and ALL. With feedback NONE, the agents are not told which (if any) of their submitted bids are currently in the optimal allocation, nor are they told any information about the current optimal allocation. (This mechanism corresponds to a sealed-bid mechanism.) In this variant, it is possible for an agent to place arbitraging bids, but it is impossible for an agent to place an arbitraging bid that she knows for sure will be in the optimal allocation.³ For this assurance, the agent needs to know something about the other bids.

³ An agent determined to find an arbitrage opportunity could place a large number of arbitraging bids in the hope that some of them would be accepted. The exchange administrator could prevent this type of strategy from occurring (for example) by charging a small payment for each bid placed. However, this could result in a smaller revenue (and thus a less efficient outcome) for the exchange administrator. A reasonable policy for the exchange administrator to use is to limit the number of arbitraging bids an agent can place to a small constant, possibly even zero.

Similarly, when the exchange is using feedback OWN-WINNING, then agents are told which of their own bids are currently winning. With this amount of information, it is generally not possible for an agent to compute an arbitraging bid that is guaranteed to be in an optimal allocation.

When the exchange is using feedback ALL-WINNING, agents are told what the current winning bids are in the current optimal allocation. In general, this is not enough information to compute an arbitraging bid that is guaranteed to be in an optimal allocation. For example, in Example 1, initially there were no winning bids. With feedback ALL-WINNING, there would not be any information available to the arbitraging agent, and thus she could not compute an arbitraging bid that is guaranteed to be winning.

When the exchange is using feedback ALL, then all agents are told about all of the bids that have been submitted, as well as what the winning bids are. In this variant, it is clear that an agent has enough information to find an arbitraging bid that is guaranteed to be in an optimal allocation, if such a bid exists.

We know from Theorem 2 that if a pure demand bid is in the optimal allocation, then an allocation exists that has the same objective value and contains an arbitraging bid in the liquidity-maximizing and volume-maximizing CEs. So, for example, if the exchange only contains pure demand and pure supply bids, an arbitraging bid can be found in polynomial time by simply scanning the winning bids and finding a demand bid. By Theorem 3, we know that this method will not necessarily find an arbitraging bid that is guaranteed to be in an optimal allocation when our surplus-constrained version of the liquidity-maximizing or volume-maximizing CE is used.

We can generalize this idea to a method for generating arbitraging bids that is guaranteed to create an arbitrage opportunity for an agent exactly when an arbitrage opportunity exists. If all bids have integer item quantities, an agent wishing to arbitrage can perform the following. For each item being supplied in each combinatorial bid, create a single-item unit demand bid for each unit of quantity in the bid with a small ask price, say p . For a small enough value of $|p|$, this method is guaranteed to find an arbitrage opportunity, if such an opportunity exists.

There are three drawbacks to this method: it does not compute an optimal value of p , it potentially requires submitting a huge number of bids, and it does not work when item quantities are allowed to be fractional. The following method addresses each of these issues. In the more general setting where item quantities can be fractional, an arbitraging agent needs to compute the correct item quantities for an arbitraging bid. It is easy to show that the following optimization problem is guaranteed to compute the appropriate item quantities y_i for an arbitraging bid, if one exists.

$$\begin{aligned}
& \max && f(x) + p \\
& \text{such that} && \sum_{j=1}^n \lambda_j^i x_j + y_i \leq 0 && i = 1, 2, \dots, m \\
& && x_j \in \{0, 1\} && j = 1, 2, \dots, n \\
& && y_i \leq 0 && i = 1, 2, \dots, m
\end{aligned}$$

If p is defined appropriately for the type of objective function in use, solving this problem also computes the most profitable arbitraging bid.

Although these approaches will find an arbitraging bid, they are unlikely to be used in practice by agents participating in the exchange. The main drawback to these methods for an arbitraging agent is that they require the arbitraging agent to know every bid that has been submitted. This level of transparency is unlikely to be present in exchanges (e.g. due to privacy concerns). Even with this level of feedback, an exchange administrator can discourage this type of arbitraging strategy, for example, by charging a small price for each bid submitted. As discussed previously, this arbitraging bid strategy is most likely to be used by the exchange administrator to locate segments of the market in which demand is weak and as a form of feedback to bidders to indicate which segments of the market currently have low demand, so the bidders can be guided into bidding on certain items with weak demand.

6 Side Constraints

In this paper thus far we have mainly considered *unconstrained allocations*. The only constraints that we have placed on allocations is that the demand of each item must meet the supply, or that the surplus is nonnegative. However, exchange administrators may wish to place other constraints on the allocation. We refer to these additional constraints as *side constraints*. There are many types of side constraints that have been studied in the literature including assignment constraints [12], budget constraints, volume/capacity constraints (see Example 3), and number of winners constraints (see Example 4) [13]. In this section we present some examples of side constraints that can lead to arbitrage opportunities even in variants where arbitrage is impossible without side constraints.

Example 3. Consider a single-unit surplus-maximizing CE with free disposal. Recall from Theorem 1 that arbitrage is impossible in this type of exchange. Now suppose that the exchange administrator limits the volume of any one consumer to at most 40% of the traded demand volume. (The exchange administrator may, for example, desire to place this constraint on the allocation so as to limit the effects of a single consumer going bankrupt.) Let $\mathcal{B} = \{B_1, B_2, B_3, B_4, B_5\}$ be the current bids. Let $B_1 = \langle(-1), -2\rangle$, $B_2 = \langle(-1), -2\rangle$, and $B_3 = \langle(-1), -3\rangle$ be single-unit supply bids asking for \$2, \$2, and \$3, respectively. Let $B_4 = \langle(1), 4\rangle$ and $B_5 = \langle(1), 5\rangle$ be single-unit demand bids willing to pay \$4 and \$5, respectively. Without any side constraints, this market will clear, accepting B_1 , B_2 , B_4 , and B_5 in the optimal allocation with a surplus of 4. However, in this allocation, the demand bids B_1 and B_2 are each demanding 50% of the total traded demand volume, violating the constraint. Taking the side constraint, into consideration, the market does not clear and no trade takes place. Now suppose an agent submits an arbitraging bid $B_6 = \langle(1), -1\rangle$ indicating that she is willing to accept one unit of the item and also receive \$1. Now the exchange will clear by accepting all items with a surplus of \$1.

Example 4. Consider a single-unit surplus-maximizing CE with free disposal. Recall from Theorem 1 that arbitrage is impossible in this type of exchange. Now suppose that the exchange administrator requires there to be at least 3 winning bidders in the exchange. If there are only two bidders that have submitted bids, then there is potentially an arbitrage situation. If ignoring the minimum number of winners constraint results in a positive surplus, then there is an arbitrage opportunity. A third bidder could place an arbitraging bid which, as long as the price on the bid is such that its value is more than the negative of the surplus in the unconstrained exchange, will allow the market to clear, thus resulting in an arbitrage opportunity for the third bidder.

7 Conclusions and Future Research

Combinatorial exchanges are trading mechanisms that allow agents to specify preferences over bundles of goods. When the agents' preferences exhibit complementarity and/or substitutability, this can lead to more efficient allocations than is possible using traditional exchanges. Before CEs can be widely implemented, many aspects need to be further examined. This paper has studied arbitrage, a risk-free profit opportunity, in the context of CEs. Arbitrage is a particular aspect of CEs that has not been studied until now.

In this paper, we have shown that some CEs allow agents to perform arbitrage and thus extract a positive payment from the market while contributing nothing, something that is not possible in traditional exchanges. We analyzed the extent to which arbitrage is possible and computationally feasible in CEs. In particular we showed that while arbitrage is impossible in surplus-maximizing CEs with free disposal, it is possible in surplus-maximizing CEs without free disposal and in volume-maximizing and liquidity-maximizing CEs with and without free disposal. We also showed that even in the surplus-maximizing CE with free disposal, the use of side constraints can lead to arbitrage opportunities. For volume-maximizing and liquidity-maximizing CEs we have proposed an improved CE that achieves the same economic objective as the original exchange, but it curtails a particularly undesirable form of arbitrage in which an arbitraging bid is included in the optimal allocation while a non-arbitraging bid on the same bundle is not included in the optimal allocation. Despite the existence of arbitraging bids in most types of CEs, we showed that in order for an agent to take advantage of an arbitrage opportunity, she needs to be equipped with all of the information about an exchange in order to generate arbitraging bids that are guaranteed to be in an optimal allocation.

In this paper we have mentioned the possibility of an exchange administrator computing arbitrage opportunities as a method for providing useful feedback. Agents participating in a CE face a difficult decision problem. The space of possible bundles to bid on is large. Further, agents often are constrained by time and computational limits. Any guidance they receive from the exchange administrator would be quite useful. Future research includes exploring this possibility

of using arbitrage to provide more effective feedback and guidance to agents in a CE, and to study this in the context of other preference elicitation mechanisms.

References

1. Bossaerts, P., Fine, L., Ledyard, J.: Inducing liquidity in thin financial markets through combined-value trading mechanisms. Social science working paper 1095R, California Institute of Technology, Pasadena, California (2000)
2. Sandholm, T., Suri, S., Gilpin, A., Levine, D.: Winner determination in combinatorial auction generalizations. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Bologna, Italy (2002) 69–76 Early version appeared at the AGENTS-01 Workshop on Agent-Based Approaches to B2B, pp. 35–41, Montreal, Canada, May 2001.
3. Sandholm, T., Suri, S.: BOB: Improved winner determination in combinatorial auctions and generalizations. *Artificial Intelligence* **145** (2003) 33–58 Early version: Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations. National Conference on Artificial Intelligence (AAAI-00), pp. 90–97, Austin, TX, July 31 – August 2.
4. Smith, T., Sandholm, T., Simmons, R.: Constructing and clearing combinatorial exchanges using preference elicitation. In: AAI-02 workshop on Preferences in AI and CP: Symbolic Approaches. (2002) 87–93
5. Parkes, D.C., Kalagnanam, J., Eso, M.: Achieving budget-balance with Vickrey-based payment schemes in exchanges. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI), Seattle, WA (2001)
6. Parkes, D.C., Kalagnanam, J., Eso, M.: Achieving budget-balance with VCG-based payment schemes in combinatorial exchanges. Technical report, IBM Research RC22218 W0110-065 (2002)
7. Deng, X., Li, Z., Wang, S.: On computation of arbitrage for markets with friction. In Du, D.Z., Eades, P., Estivill-Castro, V., Lin, X., Sharma, A., eds.: COCOON. Volume 1858 of Lecture Notes in Computer Science., Sydney, Australia, Springer-Verlag (2000) 310–319
8. Cai, M., Deng, X.: Arbitrage in frictional foreign exchange market. In Harland, J., ed.: *Electronic Notes in Theoretical Computer Science*. Volume 78., Elsevier (2003)
9. Deng, X., Li, Z., Wang, S.: Computational complexity of arbitrage in frictional security market. *International Journal of Foundations of Computer Science* **13** (2002) 681–684
10. Hull, J.C.: *Options, Futures, and Other Derivatives*. Fifth edn. Prentice Hall (2003)
11. Rothkopf, M.H., Pekeč, A., Harstad, R.M.: Computationally manageable combinatorial auctions. *Management Science* **44** (1998) 1131–1147
12. Kalagnanam, J., Davenport, A., Lee, H.: Computational aspects of clearing continuous call double auctions with assignment constraints and indivisible demand. *Electronic Commerce Research Journal* **1** (2001)
13. Sandholm, T., Suri, S.: Side constraints and non-price attributes in markets. In: IJCAI-2001 Workshop on Distributed Constraint Reasoning, Seattle, WA (2001) 55–61