Multilateral Deferred-Acceptance Mechanisms

Liad Blumrosen[1] and Osnat Zohar[2]

[1] Department of Economics, The Hebrew University, Jerusalem, Israel
[2] Bank of Israel, Jerusalem, Israel

Abstract. We study the design of multilateral markets, where agents with several different roles engage in trade. We first observe that the modular approach proposed by Dütting et al. [5] for bilateral markets can also be applied in multilateral markets. This gives a general method to design Deferred Acceptance mechanisms in such settings; these mechanisms, defined by Milgrom and Segal [10], are known to satisfy some highly desired properties.
We then show applications of this framework in the context of *supply chains*. We show how existing mechanisms can be implemented as multilateral Deferred Acceptance mechanisms, and thus exhibit nice practical properties (as group strategy-proofness and equivalence to clock auctions). We use the general framework to design a novel mechanism that improves upon previous mechanisms in terms of social welfare. Our mechanism manages to avoid "trade reduction" in some scenarios, while maintaining the incentive and budget-balance properties.

1 Introduction

Markets are often characterized by multiplicity of participants with diverse characteristics. While much of the mechanism-design literature focuses on bilateral settings that distinguish between two types of agents - buyers and sellers, many real life applications require a more complex description of markets. For instance, dealing with markets in which buyers wish to purchase a bundle of items that are sold separately by different sellers might require a distinction between different types of sellers. Further and more complex distinctions will be needed as markets become more complex and involve trade between diversified agents.

In settings of bilateral trade, and nonetheless in more complex settings where agents engage in several bilateral transactions, it is generally impossible to achieve an efficient allocation while maintaining agents' participation constraints, incentive compatibility and budget balance [12]. Therefore, some goals needs to be sacrificed in order to fully achieve the others. For instance the VCG mechanism maintains agents' incentives and implements the efficient allocation but is generally not budget balanced. Other mechanisms relax incentive compatibility to achieve the efficient allocation and budget balance (see [7] for a survey).

In this paper we devise a family of mechanisms named multilateral deferred-acceptance (MDA) mechanisms. These mechanisms apply the methodology of deferred-acceptance (DA) auctions, introduced by Milgrom and Segal [10], to multilateral markets. DA auctions set allocations using an iterative process of rejecting the least attractive bid according to a carefully defined ranking function. Combining this sort of algorithm with threshold payments yields a mechanism with strong incentive properties - other than being truthful, the DA auction is also weakly group-strategy proof (WGSP). This means that no coalition of agents has a joint deviation from truthful bidding that is strictly profitable for all members of the coalition. Another desired feature of DA auctions is that they are equivalent to clock auctions, an auction format which is intuitive for bidders and is thus considered practical. WGSP and equivalence to clock auctions are desired properties of DA auctions that are not generally attained by other mechanisms. For example, the VCG mechanism and greedy mechanisms such as [6] do not possess these properties.

Dütting et al. [5] took a modular approach to adapt DA auctions to two-sided markets. We take the general concept introduced in [5] one step forward and observe that their modular approach can also be applied to multilateral markets with several types of agents. Similar to [5], the mechanism's operation is determined by two elements: separate rankings for each set of agents and a composition rule. In each period the composition rule selects few classes (or groups) of agents. The least desirable agent, according to the corresponding ranking, of each selected class will be rejected. When the mechanism terminates, all unrejected agents are declared winners and threshold payments are set. We generalize the result by [5] and show that any mechanism from this family (MDA mechanisms) is equivalent to a one-sided DA auction; thus, it is strategy-proof, individually rational, WGSP and equivalent to a clock auction.

After introducing the class of MDA mechanisms, we apply them in the context of *supply chains* (see, e.g.,[17, 16]). Supply chains are collections of markets where each agent engages in at least one bilateral trade, either as a seller of an item he produces, or as a buyer of items (consumption goods or inputs for production). Supply chains are thus composed of several two-sided markets and feature the same impossibilities that exist in bilateral-trade settings of maximizing social welfare while maintaining agents' incentives and budget balance.

We study a model for supply chains that was introduced by Babaioff and Nisan [1] and Babaioff and Walsh [2]. In this model, the supply chain can be viewed as a directed tree-graph with a node per each good. Ingoing edges to a node define the inputs for the production of the relevant good, and producers incur a manufacturing cost which is private information. [1] and [2] showed a dominant-strategy truthful, budget-balanced mechanism that waives only the least profitable trade (or *"procurement set"*, which is a minimal trade cycle in a supply chain and typically involves multiple agents).

Our first result for supply chains shows that the trade reduction mechanisms of [1] and [2] can be implemented as MDA mechanisms. Thus, other than being IR, strategy-proof and budget balanced (as proven in [1] and [2]), these trade reduction mechanisms are also WGSP and equivalent to clock auctions. This is shown under the assumption of homogeneous demand (i.e., all end consumers demand the same bundle).[3]

Our second and main result shows how to use the machinery of MDA mechanisms to construct a novel mechanism that provides an improved outcome in terms of social welfare compared to the above trade reduction mechanisms. It operates by iteratively rejecting procurement sets, but unlike the trade reduction mechanism that always waives one procurement set that engages in trade in an efficient outcome, our mechanism will sometimes result in the efficient allocation (with no reduction of valuable trades). Using the values of agents that have already been declared losers, we bound the payments of active agents and identify situations where the efficient outcome can be implemented while maintaining a balanced budget. In markets where the efficient allocation consists of a small number of procurement sets, this improvement may be substantial. We ran computer simulations of simple supply chain networks. In the simulations, our new mechanism improves upon the trade reduction mechanisms in around 17% of the instances and saves up to 100% of the overall efficiency. This provides a good indication that the improvement in efficiency is not a rare phenomenon and can be significant.

The seminal paper by McAfee [8] introduced the trade reduction technique.[4] McAfee's mechanism was given for two-sided markets with unit-demand buyers and unit-supply sellers. This mechanism either implements the efficient allocation or reduces the least valuable profitable trade. [5] proved that the trade reduction mechanism for two-sided markets (a simplified version of [8] that always eliminates one valuable trade) can be implemented via a DA mechanism and it is therefore WGSP and equivalent to a clock auction.[5]

The paper is organized as follows: Section 2 defines deferred acceptance mechanisms. Section 3 defines MDA mechanisms and presents their main properties. Section 4 introduces the applications of MDA mechanisms to supply chains.

2 General Deferred-Acceptance Auctions

Consider a set N of single-parameter agents and let $\mathcal{F} \subseteq 2^N$ be the set of feasible sets of agents. An allocation in this setting is represented by a set of winning

[3] Our model is a generalization of the linear model in [1]; [2] did not require homogeneous demand.

[4] There is a vast literature on the efficiency of two-sided auctions that followed [8], see, e.g., [13, 14, 3]. The efficiency of DA auctions was studied in [4].

[5] Our work is inspired by [8] in several ways. First, we sacrifice efficiency in order to satisfy incentive constraints and budget balance and our mechanism loses at most the least valuable procurement set. In addition, McAfee's mechanism computed a price as a function of the "best" losing bids, and if this price cleared the market, no trade reduction would take place. Our mechanism acts in the same spirit and sometimes implements the efficient allocation, but it is not a generalization of McAfee's mechanism. In fact, our mechanism always omits one trade when applied to the degenerate supply chain of a single two-sided market with unit-demand buyers; the benefits of our mechanism stem in more complex markets.

agents $A \in \mathcal{F}$. Every agent $i \in N$ is characterized by a type t_i such that given an allocation A and payments $\{p_i\}_{i \in N} \subseteq \mathbb{R}$, agent i's utility is $t_i + p_i$ if $i \in A$ and p_i if $i \notin A$. An agent's type is assumed to be private information.

In this setting, [10] define DA auctions. Each agent $i \in N$ is required to submit a single bid from a finite set of possible bids $B_i \subseteq \mathbb{R}$. According to submitted bids, an iterative process of rejecting agents is preformed and all agents that are not rejected in the process are declared winners. We now describe this process in detail.

An agent is considered **active** in iteration t if he has not been rejected in any iteration prior to t. Let $A_t \subseteq N$ denote the set of active agents at the beginning of iteration t. Each active agent is assigned a score which is a function of his bid and the bids of all previously rejected agents:

Definition 1. *[10] A **DA scoring function** is a function of the form* $\sigma_i^t : B_i \times B_{N \setminus A_t} \to \mathbb{R}_+$ *that is non-decreasing in the first argument.*

The scoring functions form a ranking over the set of active agents in which higher ranked agents are considered less attractive. Following this logic, the DA algorithm iteratively rejects the agents with the highest score, until all agents have a score of zero. Formally:

Definition 2. *[10] Given DA scoring functions, a **DA algorithm** is defined as follows: All bidders are initially active. If all active bidders have a score of zero, the algorithm terminates and the remaining active bidders are declared winners. Otherwise, the algorithm rejects the active bidders with the highest score, removing them from the active set, and iterates.*

A DA auction can now be formally defined as follows. We then define two desired properties of DA auctions.

Definition 3. *[10] A **DA auction** is a sealed-bid auction that computes an allocation using a DA algorithm and makes the corresponding threshold payments to winners.*[6] *Losing agents are paid zero.*

Definition 4. *A mechanism is **weakly group strategy-proof (WGSP)** if for every profile of truthful reports* b*, every set of agents* $S \subseteq N$ *and every strategy profile* b'_S *of these agents, at least one agent in* S *has a weakly higher payoff from the profile of truthful reports* b *than from the strategy profile* $(b_{N \setminus S}, b'_S)$*.*

Namely, a mechanism is WGSP if no coalition of agents can do strictly better by misreporting their values, given that all other agents report their true values.

Definition 5. *A descending **clock auction** is a dynamic mechanism that presents a decreasing sequence of prices to each bidder. Each presentation is followed by a decision period in which each bidder decides whether to exit or continue. When the auction ends, the bidders that have never exited are declared winners and are paid their last (lowest) accepted prices.*

[6] Threshold payments will be formally defined in the next section. Informally, these are the highest bids for a winning agent such that he remains a winner.

Theorem 6. *[10] Any DA auction is individually rational (IR), strategy-proof, WGSP and equivalent to a clock auction.*

Milgrom and Segal [10] show that several previously known mechanisms (e.g., [11, 9, 15]) can be implemented as DA auctions and thus inherit all the properties specified in Theorem 6. In section 4.1 we take a similar approach by showing that the mechanisms in [1] and [2] can be implemented as MDA mechanisms and thus inherit their properties.

3 Multilateral Deferred-Acceptance Mechanisms

We now turn to settings of multilateral markets in which agents might differ in several aspects other than their types. Consider a setting in which agents have some distinct and known characteristics which allow sorting them into different classes. Let K be the number of agents' classes and denote by N_k the set of agents of class $k \in \{1, ..., K\}$. Thus, the set of all agents N can be described as a union of K disjoint sets $N = N_1 \cup N_2 \cup \cdots \cup N_K$.[7]

3.1 Multilateral Deferred-Acceptance Algorithms

For all $i \in N$ let $B_i \subseteq \mathbb{R}$ be the finite set of possible bids for bidder i, so the input for the MDA algorithm is a vector $b \in \Pi_{i \in N} B_i$.[8] In the spirit of [5], the MDA algorithm is composed of two elements: **scoring functions** and **composition functions**. We now define these two elements and describe how they construct an MDA algorithm. Scoring functions are defined similarly to [10] (Definition 1):

Definition 7. *For each $k = 1, ..., K$ and $i \in N_k \cap A_t$, agent i's* ***scoring function*** *$s^t_{k,i} : B_i \times B_{N_k \setminus A_t} \to \mathbb{R}_+$ is non-decreasing in the first argument and assures no ties between agents of the same class.*[9]

The scoring of agent $i \in N_k \cap A_t$ is compared to the scores of all other active agents of class k in period t to form a ranking on the set $N_k \cap A_t$. The "no ties"

[7] In a two-sided market with producers and consumers of a homogeneous good, N_1 might be the set of producers and N_2 might be the set of consumers. In that case producers' types will be thought of as production costs, so a producer with a cost c_i will have utility of $-c_i + p_i$ if $i \in A$, and p_i otherwise. Consumers' types will be thought of as their value from possessing one item of the traded good.

[8] As mentioned in Section 2, [10] define DA auctions with finite bid spaces. In order to use their results, we do the same. This also requires a more delicate definition of truthfulness and we follow the definition of strategy-proofness in [10] which uses the standard dominant-strategy truthfulness, only with taking care of the finite bid space. We refer the readers to [10] for the exact definition.

[9] In order to keep notation simple, the scoring functions are denoted with superscript t yet they are allowed to depend on the entire history of active agents $(A_1, ..., A_t)$ and not just on the t-period information. In the remainder of the paper, all objects denoted with superscript t are allowed to be history dependent.

requirement does not appear in [10] and it is made here as we sometimes wish to carefully control the number of agents of each class that are rejected. In order to simplify the presentation of our mechanisms, from now on, whenever possible ties occur, we assume the existence of a tie-breaking rule instead of formally defining scoring functions with no ties.

Composition Functions: In each period a different composition function is defined. Its inputs are the bids of all previously rejected agents and its output is a subset of $\{1, ..., K\}$.

Definition 8. *In each period t, a* ***composition function*** *C^t is a function of the form $C^t : B_{N \setminus A_t} \to 2^K$.*

Multilateral Deferred-Acceptance Algorithms: Given a set of scoring functions and composition functions, an MDA algorithm is defined as follows: In each period t, the composition function $C^t(b_{N \setminus A_t})$ outputs a subset of classes. For each class $k \in C^t(b_{N \setminus A_t})$, the algorithm queries the active agents of class k and rejects the highest scoring one, according to the scoring functions $\{s^t_{k,i}\}_{i \in N_k \cap A_t}$. This means that in period t the number of agents rejected is $|C^t(b_{N \setminus A_t})|$. If $C^t(b_{N \setminus A_t}) = \emptyset$, the algorithm terminates and all active agents are declared winners. The algorithm's operation can be described in the following manner:

1. Initialize the algorithm with $A_1 = N$.
2. For each $t \geq 1$, if $C^t(b_{N \setminus A_t}) = \emptyset$, stop. Accept all currently active agents A_t.
3. If $C^t(b_{N \setminus A_t}) = C^t \neq \emptyset$, define: $A_{t+1} = A_t \setminus \bigcup_{k \in C^t} \operatorname*{argmax}_{i \in N_k \cap A_t} s^t_{k,i}(b_i, b_{N_k \setminus A_t})$ and return to 2.

3.2 Multilateral Deferred-Acceptance Mechanisms

To complete describing the MDA mechanism we now define the payments.

Definition 9. *Given an MDA algorithm and a vector of bids $b \in \Pi_{i \in N} B_i$, let $A(b)$ denote the set of winning agents. The* ***threshold payment*** *of a winning agent $i \in A(b)$ is defined as $\sup\{b'_i \in B_i | i \in A(b'_i, b_{-i})\}$*

Definition 10. *An* ***MDA mechanism*** *computes allocation using an MDA algorithm and makes the corresponding threshold payments to winning agents. Losing agents pay zero.*

In the full version of the paper we show that every MDA mechanism can be implemented as a DA auction. The method we use is similar to the one used in [5] for the generalization of DA auctions to two-sided markets.

Proposition 11. *For every MDA mechanism there is an equivalent DA auction.*

Proposition 11 establishes that MDA mechanisms inherit all the properties of DA auctions. Together with Theorem 6, we conclude that:

Corollary 12. *Any MDA mechanism is IR, strategy-proof, WGSP and equivalent to a clock auction.*

4 Applications to Supply Chains

We now present a model of supply chains which follows [2] and generalizes the linear supply-chain model in [1]. Consider an economy with K types of items, denoted $1,..,K$. We begin by describing the production of these items, assuming that all items of a specific type are manufactured in the same manner, using the same inputs.[10]

Assumption 1 *Each product is manufactured with a* ***unique manufacturing technology****.*

Production in this economy can be described as a directed a-cyclical graph with K nodes representing the K different types of items. In this graph an edge (j,k) indicates that the production of item k uses item j as an input and the weight of the edge is the number of items of type j needed for production. Since any directed a-cyclical graph has a topological ordering,[11] assume WLOG that this ordering is given by the numbering of items' types. This means that the manufacturing of a type-k item makes use only of items of types $1,...,k-1$.

This production structure allows us to characterize the production of a type-k item with a **production vector** $q^k = (q_{1,k},...,q_{K,k})' \in \mathbb{Z}_{\leq 0}^{k-1} \times \{1\} \times \{0\}^{K-k}$. Arguments $1,...,k-1$ of the production vector are non-positive integers representing the quantities of inputs required for production (i.e., $-q_{j,k}$ for $j = 1,...,k-1$ is the weight of the edge (j,k)).[12] $q_{k,k} = 1$ indicates that one unit of item k is being produced in the process. Items $k+1,...,K$ are not involved in the production of item k, so $q_{k+1,k} = \cdots = q_{K,k} = 0$.

We further assume that each producer in the economy can manufacture a single item. Thus, all the producers that manufacture an item of type $k \in \{1,...,K\}$ are substitutes and will be regarded as agents of class k. Let $c_i \in [0, \bar{c}_k]$ denote the production cost of producer i of class k. A producer's class and production vector are assumed to be common knowledge but his cost is private information.

We now turn to describe end consumers in the economy. These agents lack any production abilities but they benefit from consumption. All consumers are single minded, meaning that each consumer values only one particular bundle of items and will gain zero utility from consuming any smaller bundle.

Up until now all our assumptions follow the ones made in [2]; We now add an additional assumption, which was also assumed in [1]:

[10] In the full paper, we relax the assumption of unique manufacturing technology and design an MDA mechanism for scenarios where a certain good can be produced from different types of inputs.

[11] A topological ordering of a directed a-cyclical graph is an ordering of the nodes such that for every edge (j,k), the node j comes before k in that ordering.

[12] The assumption that $q_{j,k}$ for $j < k$ is an integer, rather than a real number, is without loss of generality since any amount of items can be regarded as one unit. For example, if item j is flour and all items $k > j$ are produced using amounts of flour in multiples of 0.5 kg, set one unit of item j to be 0.5 kg of flour.

Assumption 2 ***Homogeneous demand*** *- All end consumers demand the same bundle* $d = (d_1, ..., d_K)' \in \mathbb{N}^K$ *where* d_k *is the demand from item* k*.*

The demanded bundle d is commonly known but consumer i's valuation for this bundle, $v_i \in [0, \bar{v}]$, is private information. We refer to consumers as agents of class $K+1$.

Define the **production matrix** as $Q = (q^1, ..., q^K)$. If there are μ_k producers of class k and $\mu = (\mu_1, ..., \mu_K)'$ then the supply of items is given by the vector $Q \cdot \mu$ where the k'th argument is the supply of item k. Denote by $\tilde{\mu}_k$ the number of class-k producers needed to meet the demand d of one consumer, i.e., the vector $\tilde{\mu} = (\tilde{\mu}_1, ..., \tilde{\mu}_K)$ solves $Q \cdot \tilde{\mu} = d$.[13]

Assume WLOG that initially there is no excess demand of any item, i.e., that $Q \cdot \mu \geq |N_{K+1}| \cdot d$, where $\mu = (|N_1|, ..., |N_K|)'$. An equivalent requirement is that $\mu \geq |N_{K+1}| \cdot \tilde{\mu}$. If this condition is not met, reject the highest bidding consumers until there is no excess demand.

A *procurement set* ([2]) is a set of agents that contains one consumer and the minimal amount of producers needed to meet his demand. Formally:

Definition 13. *A* ***procurement set*** *is a set of agents containing one consumer and* $\tilde{\mu}_k$ *producers of class* k *for every* $1 \leq k \leq K$*.*

Example 14. Figure 1 depicts a simple supply chain; fabric will be referred to as item 1, hats as item 2 and shirts as item 3. The production structure is such that producing either one hat or one shirt requires one roll of fabric. This implies that the production vectors are $q^1 = (1,0,0)'$, $q^2 = (-1,1,0)'$, $q^3 = (-1,0,1)'$ and the production matrix is $Q = (q^1, q^2, q^3)$.

Each consumer in this example demands a bundle of one hat and two shirts, i.e., $d = (0,1,2)'$. This implies that $\tilde{\mu} = Q^{-1} \cdot d = (3,1,2)'$ which means that a procurement set in this example contains one consumer, three producers of fabric, one producer of hats and two producers of shirts.

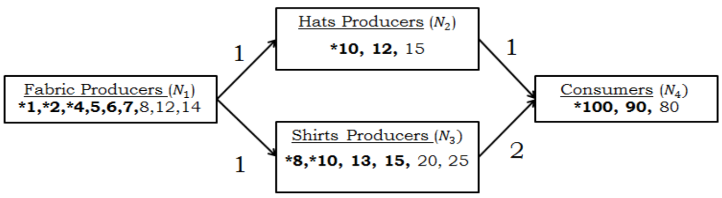

Fig. 1. A simple supply chain. The optimal allocation is marked in bold and the Trade Reduction allocation is marked with asterisks.

[13] Q is a unitriangular matrix with negative integers on the entries above the main diagonal. Since $d \in \mathbb{N}^K$, it can be shown that $\tilde{\mu} = Q^{-1} \cdot d$ is a vector of non-negative integers and thus appropriately represents numbers of agents.

Definition 15. *[2] Given a set of bids, the* ***Trade Reduction allocation*** *is obtained from the optimal allocation by reducing one procurement set. The procurement set reduced is the one with the lowest value consumer and the highest cost producers out of all the procurement sets in the optimal allocation. Together with threshold payments, determined by submitted bids, this allocation rule establishes the* ***Trade Reduction mechanism****.*

Example 16. In Figure 1 the optimal allocation is marked in bold and the Trade Reduction allocation is marked with asterisks. The threshold payments for winning agents in the examples: fabric producers are paid 5 each, the hat producer is paid 12, shirts producers are paid 13 and the consumer pays 90.

Proposition 17. *[2] The Trade Reduction mechanism is IR, strategy-proof, budget balanced and incurs the loss of the least valuable procurement set in the optimal allocation.*

4.1 The MDA Trade Reduction Mechanism

The MDA mechanism we present in this section implements the Trade Reduction allocation in a process that is equivalent to iterative rejection of procurement sets. In each iteration, the MDA algorithm examines the most valuable procurement set out of the ones that were rejected so far. The algorithm calculates the net cost of this procurement set, i.e., the sum of costs of all the producers in the set minus the value of the consumer in the set. While this net cost is strictly positive, the algorithm keeps on rejecting more procurement sets. Immediately after the mechanism rejects one procurement set with a non-negative net cost, it terminates and accepts all currently active agents. This way only one efficient procurement set is rejected, but all the rest are accepted.

We now turn to the formal definition of the mechanism. For each $k = 1, ..., K$ let B_k be a finite set of possible bids for all producers of class k, such that $B_k \subseteq [0, \bar{c}_k]$ and $\max B_k > \bar{c}_k$. Let $B_{K+1} \subseteq [-\bar{v}, 0]$ be a finite set possible bids for all consumers, such that $\max B_{K+1} > 0$.[14]

Scoring Functions: Producers of each class are ranked in an ascending order of costs and consumers are ranked in a descending order of values, i.e., all agents are ranked in a descending order of attractiveness. Formally:

$$\begin{array}{ll} \forall\, t,\, k \in \{1, ..., K\},\, i \in A_t \cap N_k, & s^t_{k,i} = c_i \\ \forall\, t,\, i \in A_t \cap N_{K+1}, & s^t_{K+1,i} = \bar{v} - v_i \end{array} \tag{1}$$

The composition rule is based on two auxiliary functions:

[14] Consumers' bid spaces are defined as subsets of $[-\bar{v}, 0]$ so we can treat all agents, producers and consumers, in a similar manner such that higher bidding agents are less attractive. Thus, the mechanism will determine negative monetary transfers for consumers and positive transfers for producers. The maximal (minimal) possible bid of a producer (consumer) is set to be higher (lower) than his highest possible cost (lowest possible value) to insure that participation is strictly preferable to non-participation (see [10]).

1. **The Net Cost Function** NC^t: For each period t, denote by $NC^t(b_{N\setminus A_t})$ the net cost of the most valuable procurement set rejected so far. Formally:

$$NC^t(b_{N\setminus A_t}) = \sum_{k=1}^{K}\sum_{l=1}^{\tilde{\mu}_k} c^t_{k,(l)} - v^t_{max} \tag{2}$$

where $c^t_{k,(l)}$ is the l'th lowest cost reported in $b_{N_k\setminus A_t}$ and v^t_{max} is the highest value reported by a rejected consumer (i.e., $-v^t_{max}$ is the minimal bid in $b_{N_{K+1}\setminus A_t}$).[15]
2. **The Excess Supply Function** ES^t: Let μ^t_k denote the number of active agents of class k in period t, i.e., $\mu^t_k = |N_k \cap A_t|$. The aggregate demand in period t is equal to $\mu^t_{K+1}\cdot d$ and the number of producers of class $k \in \{1,...,K\}$ needed to meet it is $\mu^t_{K+1}\cdot\tilde{\mu}_k$. If there are more producers of class k than that, regard the class-k producers as being in excess. According to this logic, the excess supply function indicates the classes of producers that are in excess in period t. Formally, $ES^t(A_t) = \left\{k | 1 \le k \le K,\ \mu^t_k > \mu^t_{K+1}\cdot\tilde{\mu}_k\right\}$.

Composition Functions: Now we can define the composition functions, using the auxiliary functions NC^t and ES^t. For every period t, define:

$$C^t(b_{N\setminus A_t}) = \begin{cases} ES^t(A_t) & \text{if } ES^t(A_t) \neq \emptyset \\ \{1,...,K+1\} & \text{if } ES^t(A_t) = \emptyset \text{ and } NC^t(b_{N\setminus A_t}) > 0 \\ \emptyset & \text{if } ES^t(A_t) = \emptyset \text{ and } NC^t(b_{N\setminus A_t}) \le 0 \end{cases} \tag{3}$$

In words, the algorithm first rejects excess producers, as determined by the first line in (3). This is repeated until there are no excess producers, i.e., until $ES^t = \emptyset$, which means that supply equals demand (recall that we assumed that initially there is no excess demand). From this point, the algorithm's operation can be described as an iteration of three steps:

1. If there is no excess supply ($ES^t = \emptyset$), examine the net cost of the most valuable procurement set rejected so far, NC^t. If NC^t is non-positive - terminate (third line in (3)). Otherwise, continue to step 2.
2. Reject one agent of each class $1,...,K+1$ (second line in (3)).
3. As long as there is excess supply ($ES^t \neq \emptyset$), reject one agent of each class $k \in ES^t$ (first line in (3)). Once $ES^t = \emptyset$, return to step 1.

It is worth noting that each time steps 1-3 are completed, one procurement set is rejected. The rejection begins with eliminating the highest bidding agent of each class and continues with eliminating excess supply. Since each procurement set includes only one consumer, this process is equivalent to rejecting one procurement set, and that is the highest costing active procurement set.

As we show in the full version of the paper, the MDA trade reduction mechanism is in fact equivalent to the Trade Reduction mechanism:

[15] For any $k = 1,..,K$ such that $|N_k \setminus A_t| < \tilde{\mu}_k$, set $c^t_{k,(|N_k\setminus A_t|+1)} = c^t_{k,(|N_k\setminus A_t|+2)} = ... = c^t_{k,(\tilde{\mu}_k)} = \max B_k$ and if no consumer was rejected prior to period t, set $v^t_{max} = 0$. Specifically, for $t = 1$ set $NC^1(\emptyset) = \sum_{k=1}^{K}\tilde{\mu}_k \max B_k$.

Proposition 18. *Consider an MDA mechanism that is defined by the scoring functions (1) and the composition functions (3). This mechanism is equivalent to the Trade Reduction mechanism (Definition 15).*

We can now use Proposition 18 (together with Proposition 12) to infer additional properties of the Trade Reduction mechanism of [1] and [2].

Corollary 19. *The Trade Reduction mechanism is WGSP and equivalent to a clock auction.*

4.2 The Modified Trade Reduction Mechanism

The class of MDA mechanisms allows considerable freedom in design while maintaining the incentive properties common to all MDA mechanisms. We use this feature to modify the Trade Reduction mechanism and improve its social welfare. The improvement is possible since the MDA trade reduction mechanism (Section 4.1) uses an inaccurate measure of the deficit - the net cost of the last rejected procurement set. This causes the mechanism to reject more trades than is actually needs in order to keep the budget balanced. The new mechanism uses a different measure of deficit and thus waives efficient trades less frequently. For this, we need the following definition:

Definition 20. *[10] For each agent $i \in N_k \cap A_t$, the **t-period threshold payment** $p^t_{k,i}(b_{N \setminus A_t})$, is the maximal bid that would have kept i active until iteration t, holding all other bids fixed.*

First note that the mechanism's final threshold payment for a winning agent is equal to his T-period threshold payment, where T is the final period. This follows directly from the definition of threshold payments (Definition 9).

Second, note that for an active agent $i \in N_k \cap A_t$, the t-period threshold payment is determined only by bids of rejected agents of class k. This is true since agent i's bid can not affect $C^{t'}(b_{N \setminus A_t})$ for $t' < t$ and thus can not affect which classes of agents are chosen for rejection prior to period t. The only effect agent i's bid has is on his ranking relative to other agents of class k. Furthermore, in cases where agents are ranked solely by their bids (as will be the case here), the t-period threshold payment of a class-k agent is equal to the bid of the last rejected agent of his class. This means that all active agents of class k have the same t-period threshold payments.

Definition 21. *Consider an MDA mechanism in which agents are ranked only by their bids. For all t and k, the **t-period threshold payments for active agents of class k** is:*

$$p^t_k = \begin{cases} \min\limits_{j \in N_k \setminus A_{t-1}} b_j & \text{if } N_k \setminus A_{t-1} \neq \emptyset \\ \max B_k & \text{if } N_k \setminus A_{t-1} = \emptyset \end{cases}$$

We now turn to the formal definition of the Modified Trade Reduction mechanism by defining the scoring and composition functions.

Scoring Functions: Similar to Section 4.1, producers are ranked in an ascending order of costs and consumers are ranked in a descending order of values (see (1) for the formal definition).

The composition rule is similar to the one presented in Section 4.1 with the slight difference that it rejects procurement sets according to a *lower bound* on net costs, instead of the net costs themselves. The lower bound, or **minimal net cost**, is a function of the t-period threshold payments:

$$MNC^t(b_{N\setminus A_t}) = \sum_{k=1}^{K} \tilde{\mu}_k p_k^t + p_{K+1}^t \tag{4}$$

Let the excess supply function ES^t be defined as in Section 4.1. Now define the **composition function** for each period t as follows:

$$C^t(b_{N\setminus A_t}) = \begin{cases} ES^t(A_t) & \text{if } ES^t(A_t) \neq \emptyset \\ \{1, ..., K+1\} & \text{if } ES^t(A_t) = \emptyset \text{ and } MNC^t(b_{N\setminus A_t}) > 0 \\ \emptyset & \text{if } ES^t(A_t) = \emptyset \text{ and } MNC^t(b_{N\setminus A_t}) \leq 0 \end{cases} \tag{5}$$

Definition 22. *The **Modified Trade Reduction mechanism** is an MDA mechanism with the scoring functions (1) and the composition functions (5).*

The Modified Trade Reduction mechanism operates as follows. First it rejects the least valuable procurement set. Fixing this allocation, the t-period threshold payments are calculated together with the implied deficit, which is proportional to MNC^t. If the deficit is non-positive, the mechanism terminates. Otherwise, the least valuable active procurement set is rejected, and so on. The critical stage of the mechanism comes after it removes enough procurement sets and reaches an efficient allocation. Then, it computes a "within-class" threshold payments for each class of agents: the value of the most valuable agent that does not win in the efficient allocation. It then checks what would happen if all active agents paid their within-group threshold: if there is no deficit, then the mechanism outputs the efficient allocation. Otherwise, a trade reduction is performed.[16]

Theorem 23. *The Modified Trade Reduction mechanism satisfies:*

1. *It is IR, strategy-proof, WGSP and equivalent to a clock auction.*
2. *It is weakly budget balanced.*

[16] Note that this procedure is not equivalent to the following mechanism: run VCG if it is budget balanced, otherwise run a trade reduction. This mechanism is not truthful, as the VCG payment of an agent can be determined by agents of other classes (who therefore can manipulate the outcome). The Modified Trade Reduction mechanism uses bounds on the payments that are determined only by the agents of each class, and therefore it is strategy-proof.

3. *For every realization of values and costs, the mechanism either sets the optimal allocation or incurs the loss of the least valuable procurement set.*

Item 1 holds since the modified mechanism is an MDA mechanism. For the proof of item 2 see the full version of the paper. Next, we prove item 3.

Note that in each period, $p^t_{K+1} = -v^t_{max}$ and $p^t_k = c^t_{k,(1)}$ for all $k = 1, ..., K$. Now use the definitions of $NC^t(b_{N \setminus A_t})$ and $MNC^t(b_{N \setminus A_t})$ ((2) and (4) respectively) to get that for all t:

$$NC^t(b_{N\setminus A_t}) = \sum_{k=1}^{K} \sum_{l=1}^{\tilde{\mu}_k} c^t_{k,(l)} - v^t_{max} \geq \sum_{k=1}^{K} \tilde{\mu}_k c^t_{k,(1)} - v^t_{max} = \sum_{k=1}^{K} \tilde{\mu}_k p^t_k + p^t_{K+1} = MNC^t(b_{N\setminus A_t}) \quad (6)$$

Since the modified mechanism terminates once $MNC^t \leq 0$ and the MDA trade reduction mechanism terminates once $NC^t \leq 0$, the former mechanism terminates (weakly) prior to the latter. Since the procurement sets are ordered in the same manner in both mechanisms, the allocation determined by the modified mechanism contains the Trade Reduction allocation but is possibly larger.

According to equation (6), MNC^t is a lower bound on the net costs of all previously rejected procurement sets. While MNC^t is positive, all rejected procurement sets have positive net costs and the mechanism rejects more procurement sets. The mechanism terminates once a procurement set with a non-positive lower bound on its net cost is rejected. By that time all the procurement sets with positive net costs were rejected. The proof of 3 follows.

The Modified mechanism improves upon the Trade Reduction mechanism in scenarios where there is variance in the values of agents, such that the "within-class" threshold is sufficiently far from the values of the next losing agent. It follows that we need consumers to demand more than one unit of some item for having a different outcome than the Trade Reduction mechanism. (Indeed, when applying the modified mechanism to the two-sided market in [8], it identifies with the Trade Reduction Mechanism and eliminates one efficient trade.)

Simulations: We ran some computer simulations of simple supply chains where the advantages of the modified mechanism can kick in. Our simulations considered economies with $n = 2, ..., 10$ buyers and $2n$ sellers of a homogeneous good, where each buyer is interested in a bundle of two units and values were drawn from the uniform distribution on $[0, 1]$. The modified mechanism improves upon the Trade Reduction mechanism in around 17% of the instances, even when the market grows substantially. In small economies the improvement can be up to a 100% of the overall efficiency. More details appear in the full version.

Diverse Manufacturing Technologies: The MDA framework allows us to explore more general settings of supply chains. Specifically, we are able to relax the assumption of unique manufacturing technologies (Assumption 1) and examine a setting in which part of the production process may be conducted with different types of inputs. The details appear in the full version of the paper.

Acknowledgments: We thank Moshe Babaioff for helpful discussions. Liad Blumrosen was supported by the Israel Science Foundation (grant No. 230/10).

References

1. Babaioff, M., Nisan, N.: Concurrent auctions across the supply chain. Journal of Artificial Intelligence Research 21, 595–629 (2004)
2. Babaioff, M., Walsh, W.E.: Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. Decision Sup. Sys. 39(1), 123–149 (2005)
3. Cripps, M.W., Swinkels, J.M.: Efficiency of large double auctions. Econometrica 74(1), 47–92 (2006)
4. Dütting, P., Gkatzelis, V., Roughgarden, T.: The performance of deferred-acceptance auctions. In: Proceedings of the fifteenth ACM conference on Economics and computation. pp. 187–204. ACM (2014)
5. Dütting, P., Roughgarden, T., Talgam-Cohen, I.: Modularity and greed in double auctions. In: Proceedings of the fifteenth ACM conference on Economics and computation. pp. 241–258. ACM (2014)
6. Lehmann, D., Oćallaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. Journal of the ACM (JACM) 49(5), 577–602 (2002)
7. Lubin, B., Parkes, D.C.: Approximate strategyproofness. Current Science 103(9), 1021–1032 (2012)
8. McAfee, R.P.: A dominant strategy double auction. Journal of economic Theory 56(2), 434–450 (1992)
9. Mehta, A., Roughgarden, T., Sundararajan, M.: Beyond moulin mechanisms. In: 8th ACM Conference on Electronic Commerce. pp. 1–10. EC '07 (2007)
10. Milgrom, P., Segal, I.: Deferred-acceptance auctions and radio spectrum reallocation. In: Proceedings of the fifteenth ACM conference on Economics and computation. pp. 185–186. ACM (2014)
11. Moulin, H.: Incremental cost sharing: Characterization by coalition strategy-proofness. Social Choice and Welfare 16(2), 279–320 (1999)
12. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. Journal of economic theory 29(2), 265–281 (1983)
13. Rustichini, A., Satterthwaite, M.A., Williams, S.R.: Convergence to efficiency in a simple market with incomplete information. Econometrica 62(5), 1041–63 (1994)
14. Satterthwaite, M.A., Williams, S.R.: The optimality of a simple market mechanism. Econometrica 70(5), 1841–1863 (2002)
15. Singer, Y.: Budget feasible mechanisms. In: The 51st Annual Symposium on Foundations of Computer Science. pp. 765–774. FOCS '10 (2010)
16. Vorobeychik, Y., Kiekintveld, C., Wellman, M.P.: Empirical mechanism design: methods, with application to a supply-chain scenario. In: ACM Conference on Electronic Commerce. pp. 306–315 (2006)
17. Walsh, W., Wellman, M., Ygge, F.: Combinatorial auctions for supply chain formation. In: In Proc. ACM Conference on Electronic Commerce. pp. 260–269 (2000)