

A Market-Affected Sealed-Bid Auction Protocol^{*}

Claudia Lindner

Institut für Informatik, Universität Düsseldorf, 40225 Düsseldorf, Germany

Abstract. Multiagent resource allocation defines the issue of having to distribute a set of resources among a set of agents, aiming at a *fair* and *efficient* allocation. Resource allocation procedures can be evaluated with regard to properties such as budget balance and strategy-proofness. Designing a budget-balanced and strategy-proof allocation procedure that always provides a fair (namely, envy-free) and efficient (namely, Pareto-optimal) allocation poses a true challenge. To the best of our knowledge, none of the existing procedures combines all four properties. Moreover, in previous literature no attention is given to the allocation of unwanted resources (i.e., resources that seem to be of no use for all agents) in a way as to maximize social welfare. Yet, dealing inappropriately with unwanted resources may decrease each agent’s benefit. Therefore, we extend the scope of sealed-bid auctions by means of involving market prices so as to always provide an optimal solution under consideration of each agent’s preferences. We present a new market-affected sealed-bid auction protocol (MSAP) where agents submit sealed bids on indivisible resources, and we allow monetary side-payments. We show this protocol to be budget-balanced and weakly strategy-proof, and to *always* provide an allocation that maximizes both utilitarian and egalitarian social welfare, and is envy-free and Pareto-optimal.

Key words: multiagent systems, multiagent resource allocation, auctions

1 Introduction

In multiagent resource allocation, agents participate in an allocation procedure to obtain a *fair* and *efficient* allocation of a set of resources.¹ There are two types of procedures: centralized and decentralized. Auction protocols give a good example for centralized procedures: All agents are asked to state their preferences (utility values) for the resources given, and based on these the protocol makes a decision on the final allocation (see, e.g., [11, 6]). In contrast, in a decentralized environment the final allocation is the result of a sequence of conducted negotiations between single agents (see, e.g., [7]).

Many resource allocation approaches are about allocating each of the goods to any of the agents (see, e.g., [14, 1, 17]), and thus an important aspect is missed out: Among the goods to be allocated there may be “unwanted goods”, i.e., goods none of the agents is interested in. Though some procedures consider that a particular agent may not be interested in each good available (see, e.g., [17]), the issue of unwanted goods is no attention given to. The market-affected sealed-bid auction protocol (MSAP) to be presented

^{*} Supported in part by the DFG under grant RO 1202/12-1 (within the ESF’s EUROCORES program LogICCC: “Computational Foundations of Social Choice”).

¹ We will use “resource” and “good”, and “procedure” and “protocol” interchangeably.

in Section 4 fills this gap by involving market prices. This new auction protocol has highly relevant properties such as budget balance, weak strategy-proofness, and always providing an allocation that maximizes both utilitarian and egalitarian social welfare, and that is envy-free and Pareto-optimal (see Section 5). Utilitarian and egalitarian social welfare are the most common notions of social welfare (see, e.g., [5]). Informally speaking, for a given allocation, utilitarian social welfare states the sum of all agents' utilities, whereas egalitarian social welfare states the utility of the worst off agent. Both notions make meaningful statements about the quality of an allocation, utilitarian social welfare measures the overall benefit for society, and egalitarian social welfare measures the level of fairness as to satisfying minimum needs. Another substantial concept of fairness is envy-freeness: An allocation is said to be envy-free if none of the agents has an incentive to swap his or her share with any other agent's share. Regarding efficiency, the most fundamental concept is the notion of Pareto-optimality: An allocation is Pareto-optimal if any change to make an agent better off results in making another agent worse off. With regard to resource allocation protocols that involve monetary side-payments, budget balance states that all payments sum up to zero, i.e., the application of the protocol causes neither a profit nor a loss. These notions and the general framework are specified in Section 2. Moreover, the MSAP is proven to be "weakly strategy-proof", a notion to be motivated and introduced in Section 3. In short, weakly strategy-proof is a somewhat milder notion of strategy-proofness implying that a cheating attempt may be successful but is always at the risk of an overall loss.

Pioneering work in the field of auction procedures for indivisible goods was done by Knaster. His auction protocol of sealed bids² is about agents that submit sealed bids for single goods, and the agent whose bid is the highest is assigned the good; monetary side-payments are used for compensation. Knaster's procedure always provides an efficient allocation but does not guarantee envy-freeness (see, e.g. [3]). Just as in Knaster's procedure of sealed bids, the MSAP asks agents to submit sealed bids on single goods reflecting their individual welfare from receiving these goods. This has the advantage of making winner determination easy (as opposed to the hardness of winner-determination problems in combinatorial auctions, see [11, 6]). In this regard, in order to create a mutual bid-basis and to account for unwanted goods, the option to sell goods on the open market is included, i.e., the MSAP combines the actual value of a good with each agent's preferences. Furthermore, we allow monetary side-payments, and a central authority manages the allocation procedure. Put simply, the MSAP extends the scope of sealed-bid auctions by means of involving market prices in order to always provide a *fair* and *efficient* allocation—even when taking unwanted goods into account.

The New Economy, which is rooted in continuous information-technological progress, provides the basis for the MSAP to act on the global market. The internet and related technologies overcome geographical borders and increase market transparency, and hence create some kind of perfect information environment [4]. Being generated and used by internet users worldwide, information (e.g., with reference to demand and supply) turns out to be the crucial factor when performing global market activities. The global market provides a shared platform for market research, giving each agent the same chances to sell unwanted goods, and thus to make bids on a comparable basis.

² This protocol has been proposed by Knaster and was first presented in Steinhaus [13].

Application areas could be the allocation of inheritance items or collective raffle prizes, or in the case of two agents even the allocation of household goods within a divorce settlement. In general, the MSAP can be applied to a mixture of *insignificant replaceable goods* (e.g., a car that is of no (considerable) personal significance to a particular agent) and *significant replaceable goods*, i.e., goods that are sort of irreplaceable to an agent, but a certain monetary compensation would be accepted (e.g., a car that is of considerable personal significance to a particular agent for some reason—such as this agent has been born in this car—but the agent would be willing to accept some money to set the car aside). The MSAP can also be applied to a set of either solely significant or solely insignificant replaceable goods.

2 Preliminaries and Notation

Let $A = \{a_1, a_2, \dots, a_n\}$ be a set of n agents, and let $G = \{g_1, g_2, \dots, g_m\}$ be a set of m indivisible and nonshareable goods (i.e., each good is to be allocated in its entirety and to one agent only). If some amount of money is among the goods to be allocated, this money is excluded from G and its value is split equally among all agents in A . Moreover, the number of agents as well as the number of goods are not restricted, and there is no limitation on how many goods are to be allocated per agent. While in previous literature the focus is mostly on scenarios where only one single good is to be assigned per agent (see e.g., [14, 15, 9]), we do not need this restriction (for related work, see, e.g., [2, 17]).

Let $U = \{u_1, u_2, \dots, u_n\}$ be a set of n utility functions representing each agent's preferences (i.e., bids), where $u_j : G \rightarrow \mathbb{R}$ specifies agent a_j 's utility of each single good in G . An *allocation of G* is a mapping $X : A \rightarrow G$ with $X(a_j) \cap X(a_k) = \emptyset$ for any two agents a_j and a_k , $j \neq k$. At this, $u_j(X(a_j))$ gives agent a_j 's *additive utility* of the subset (bundle) of goods allocated to him or her by allocation X ; to simplify notation, we will write $u_j(X)$ instead of $u_j(X(a_j))$. Note that agents do not have any knowledge about the utility values of other agents. Let $C = \{c_1, c_2, \dots, c_n\}$ be a set of n side-payments that agents a_j in A either make (i.e., $c_j \in \mathbb{R}^-$) or receive (i.e., $c_j \in \mathbb{R}^+$) in conjunction with an allocation. At this, the value of money is supposed to be the same for all agents.

The MSAP asks agents to bid on single goods only. Hence, direct synergetic effects caused by the allocation of bundles of goods are disregarded. However, in Section 5 it is shown that the final allocation involves bundles when accounting for side-payments, and that statements regarding social welfare, fairness, and efficiency can be made.

Various criteria have been introduced to measure the quality of an allocation such as the concepts of social welfare, envy-freeness, and Pareto-optimality. Concerning a given society of agents, concepts of social welfare measure the benefit an allocation yields. Transferring this measure to individual agents gives the notion of individual welfare. Definitions 1, 2, 3, and 4 each are given with reference to a resource allocation setting where monetary side-payments are allowed.

Definition 1. Consider an allocation X of a set G of goods to a set A of agents, where the agents' preferences are represented by utility functions U . Let C be the set of side-payments agents in A either make or receive, as appropriate. The *individual welfare* of agent a_j obtained through allocation X and side-payment c_j is defined as $iw_j(X(a_j)) = u_j(X(a_j)) + c_j$. We will write $iw_j(X)$ instead of $iw_j(X(a_j))$.

Note that for an allocation X and any two agents a_j and a_k in A , $a_j \neq a_k$, the individual welfare agent a_j would obtain through the assignment of agent a_k 's share is defined as $iw_j(X(a_k)) = u_j(X(a_k)) + c_k$. In terms of social welfare, in this paper, we focus on the following two types (see, e.g., [5]).

Definition 2. Consider an allocation X of a set G of goods to a set A of agents, where $iw_j(X)$ is the individual welfare agent a_j in A obtains through allocation X . (1) The *utilitarian social welfare* is defined as $sw_u(X) = \sum_{a_j \in A} iw_j(X)$. (2) The *egalitarian social welfare* is defined as $sw_e(X) = \min\{iw_j(X) \mid a_j \in A\}$.

We now define envy-freeness and Pareto-optimality. Informally speaking, an allocation is envy-free if every agent is at least as happy with his or her share as he or she would be with any other agent's share. An allocation is Pareto-optimal if no agent can be made better off without making another agent worse off.

Definition 3. Let a set G of goods and a set A of agents be given, and let X and Y be two allocations of G to A . Let $iw_j(X)$ and $iw_j(Y)$ be the individual welfares agent a_j in A obtains through allocations X and Y . (1) X is said to be *envy-free* if for any two agents a_j and a_k in A , we have $iw_j(X(a_j)) \geq iw_j(X(a_k))$. (2) Y is said to be *Pareto-dominated* by X if for each agent a_j in A , we have $iw_j(X) \geq iw_j(Y)$, and there exists some agent a_k in A such that $iw_k(X) > iw_k(Y)$. An allocation is said to be *Pareto-optimal* (or, *Pareto-efficient*) if it is not Pareto-dominated by any other allocation.

The notion of budget balance makes a statement on the quality of a resource allocation protocol that involves monetary side-payments.

Definition 4. A resource allocation protocol is said to be *budget-balanced* if for every allocation obtained it holds that all side-payments sum up to zero, i.e., $\sum_{c_j \in C} c_j = 0$.

3 Motivation

Considering the design of a multiagent resource allocation procedure, the common goal is to guarantee a fair (namely, envy-free) and efficient (namely, Pareto-optimal) allocation. In this context, two more desirable properties for a resource allocation protocol to possess are *budget balance*, and *strategy-proofness* (i.e., none of the agents has an incentive to bid dishonestly). While the well-known Groves mechanisms [8] satisfy both Pareto-optimality and strategy-proofness, they in general do not guarantee to provide an envy-free allocation and are not budget-balanced (see, e.g., [10]). In fact, Tadenuma and Thomson [16] showed that envy-freeness and strategy-proofness are mutually exclusive. Thus, in this paper, we aim at guaranteeing envy-freeness, Pareto-optimality and budget balance while weakening the requirement of strategy-proofness. In literature it is common to define strategy-proofness by incentive compatibility which states that truthfulness is the dominant strategy (see, e.g., [12]). In order to deal with the impossibility result given, we focus on a somewhat weaker notion of strategy-proofness.

Definition 5. Given that all agents have no knowledge about the utility functions of other agents, a resource allocation protocol is said to be *weakly strategy-proof* if a cheating agent is always risking a loss and is never guaranteed to cheat successfully.

In the context of additive valuation functions, Knaster’s procedure of sealed bids satisfies Pareto-optimality and budget balance, but lacks of envy-freeness and strategy-proofness—though it is weakly strategy-proof according to Definition 5 (see, e.g. [3]). Willson [17] presented a procedure that indeed is envy-free, Pareto-optimal, budget-balanced and weakly strategy-proof (according to Definition 5). However, this procedure does not give consideration to welfare maximization when having unwanted goods.

According to Willson’s procedure, agents are expected to state negative values for those goods they do not want (i.e., to specify some monetary compensation to be paid to the agents in order to persuade them to accept those goods nonetheless), but there are no restrictions in terms of some sort of value limit. Thus, regarding an unwanted good (i.e., a good that is a burden on each of the agents), it is possible that the absolute equivalent of each single negative value is higher than the overall value of all other goods to be allocated. This allows agents to receive compensations that, in the end, may cause a money-losing allocation. As an example, consider the setting that we have two agents a_1 and a_2 , three goods g_1, g_2 and g_3 , and the following utility values $u_1(g_1) = 100$, $u_1(g_2) = 50$, $u_1(g_3) = -300$, $u_2(g_1) = 80$, $u_2(g_2) = 60$ and $u_2(g_3) = -250$. According to the procedure given in [17], each good is assigned to the highest bidder. Thus, good g_1 is assigned to agent a_1 , and goods g_2 and g_3 are assigned to agent a_2 . The overall benefit sums up to $100 + 60 + (-250) = -90$, and hence agent a_1 has to pay side-payments worth -145 to agent a_2 , resulting in a *negative* share of -45 for each of the agents. Just one unwanted good, here g_3 , can smash a whole allocation. But, if an agent does not want a good for personal use, this does not necessarily mean that this good is worth nothing to the agent as he or she may have good selling opportunities. For example, let us assume good g_3 is a car and both agents have no use for it, hence they only see the cost involved such as the cost for scrapping or insurance. By missing the option to sell unwanted goods and to distribute the related profit, agents may end up paying rather than benefiting.

Moreover, there may be goods that, though being wanted by some agents, do not make up a high personal significance. Having no common basis for the specification of utility values, agents with a similar preference for one particular good (e.g., considering the good to be of no personal significance) may state significantly different utilities. Without a common basis the values stated may neither be related to the actual value of the good nor to one another, and thus a lower overall benefit may be caused.

To address the issues mentioned above, we present and analyze a new market-affected sealed-bid auction protocol that is proven to be budget-balanced and weakly strategy-proof, and to always provide an allocation that is envy-free, Pareto-optimal, and that *always* maximizes both utilitarian and egalitarian social welfare.

4 A Market-Affected Sealed-Bid Auction Protocol

The MSAP is about allocating goods that are to give away in as fair a way as possible. However, agents may have diverse preferences for the goods in G , and thus some allocations may result in an advantage for one agent and in a disadvantage for another—which is unfair as every agent is to be treated equally. For the purpose of achieving not only an efficient but also a fair allocation, the aim is to assign all goods in G in such a way that the individual welfares of all agents are equalized according to how valuable the goods

are to them. Regarding a good that is significant to an agent, the utility value reflects the level of personal significance of this good, whereas, regarding an insignificant good, the utility value states the profit the agent could make by selling this good.³

Note that the central authority (CA) managing the allocation procedure is not one of the agents. If the CA needs to be paid for its job, this is done proportionally by all agents once the protocol is finished. It is assumed though that the CA generally does not have to be paid for organizing the allocation. Moreover, the MSAP may involve side-payments, but, as opposed to other approaches (see, e.g., [16, 17]), none of the goods in G needs to be infinitely divisible. Furthermore, neither any agent nor the CA will lose any value by the application of the MSAP, because all side-payments are included in the overall value of the goods in G . Let X^Σ denote the final allocation obtained by the MSAP. We write sw^Σ instead of $sw(X^\Sigma)$, and iw_j^Σ instead of $iw_j(X^\Sigma)$.

The MSAP is a multi-stage resource allocation protocol and consists of three phases: the bidding phase, the assignment phase and the compensation phase. In the course of the bidding phase agents are asked to specify a utility value for each of the goods. During the assignment phase each good is assigned to the agent whose benefit from receiving this good is the highest, where ties can be broken arbitrarily. Finally, the compensation phase is about equalizing the individual welfares of all agents by means of monetary side-payments. The MSAP is presented in detail in Figure 1.

Remark 1. Some remarks on the steps of the protocol in Figure 1 are in order:

1. **B5.** Agent a_j states $u_j(g_i)$ (i.e., the individual welfare agent a_j would obtain from receiving good g_i) according to the following rules.
 - (a) If a_j is not interested in g_i , a_j would not keep g_i but would sell it and thus states a utility value fulfilling $u_j(g_i) \leq M(g_i)$. Agent a_j would receive revenue $M(g_i)$ for selling g_i but he or she may also have some expenses $S_j(g_i)$ caused by selling g_i , i.e., $u_j(g_i) = M(g_i) - S_j(g_i)$. If agent a_j wants to make sure to receive good g_i by no means, he or she states a utility value of zero.⁴
 - (b) If a_j would like to have g_i for him- or herself, a_j states a utility value fulfilling $u_j(g_i) \geq M(g_i)$. At this, the value difference between $u_j(g_i)$ and $M(g_i)$ expresses the degree of significance of g_i to a_j , i.e., the higher the difference the more significant g_i is to a_j .
2. **A3.** If $u_k(g_i) < P_{CA}(g_i)$ holds true, the CA has the lowest selling cost for g_i and none of the agents in A is interested in keeping g_i for him- or herself, i.e., g_i is an *unwanted good*.
3. **C1.** If none of the goods in G needs to be sold, it holds that $sw^\Sigma \geq M^\Sigma = \sum_{i=1}^m M(g_i)$. After the assignment phase is finished, all agents concerned and the CA go into the matter of selling those goods that are to be sold, since in the compensation phase all related profits are involved. However, if agents concerned have sufficient cash at hand, the selling could be done later on, though at the risk of financial loss and the chance of additional profit as market prices may change over time.

³ There is no need to consider selling opportunities for significant goods, since agents are interested in keeping those goods due to their significance.

⁴ Note that stating “0” would simplify the process for agent a_j (i.e., no selling activity would be required), but this may cause a lower overall benefit compared to when a_j would sell g_i .

<p>Bidding Phase: For each good g_i in G perform steps B1 to B5.</p> <p>B1. Based on market research, the CA determines market price $M(g_i)$, i.e., the revenue the CA or an agent would receive when selling good g_i on the market.</p> <p>B2. The CA determines selling cost $S_{CA}(g_i)$, i.e., the cost caused by the CA selling g_i on the market (e.g., the cost for meeting a potential buyer, or the cost for shipping the good).</p> <p>B3. The CA calculates profit $P_{CA}(g_i) := M(g_i) - S_{CA}(g_i)$.</p> <p>B4. The CA discloses market price $M(g_i)$, but conceals selling cost $S_{CA}(g_i)$ and profit $P_{CA}(g_i)$.</p> <p>B5. Each agent a_j in A specifies utility value $u_j(g_i) \geq 0$ and submits this one to the CA.</p> <p>Assignment Phase: For each good g_i in G perform steps A1 to A3.</p> <p>A1. Find agent a_k in A such that there is no agent a_j in A with $u_j(g_i) > u_k(g_i)$, i.e., find a highest bidder for g_i. (Ties can be broken arbitrarily.)</p> <p>A2. If $u_k(g_i) \geq P_{CA}(g_i)$ and $u_k(g_i) > 0$, good g_i is allocated to agent a_k and the highest bid for good g_i is recorded by setting $u^*(g_i) := u_k(g_i)$.</p> <p>A3. If $u_k(g_i) < P_{CA}(g_i)$ or $u_k(g_i) = P_{CA}(g_i) = 0$, the CA is going to keep g_i for the time being and the highest bid for g_i is recorded by setting $u^*(g_i) := P_{CA}(g_i)$.</p> <p>Compensation Phase:</p> <p>C1. For final allocation X^Σ calculate the overall social welfare by $sw^\Sigma := \sum_{i=1}^m u^*(g_i)$.</p> <p>C2. In compliance with values $u_j(X^\Sigma)$, divide set A into three disjoint sets, R, S, and T with $A = R \cup S \cup T$, such that: (1) $u_r(X^\Sigma) > (1/n) \cdot sw^\Sigma$ for all a_r in R; (2) $u_s(X^\Sigma) < (1/n) \cdot sw^\Sigma$ for all a_s in S; (3) $u_t(X^\Sigma) = (1/n) \cdot sw^\Sigma$ for all a_t in T.</p> <p>C3. All agents a_r in R (i.e., all advantaged agents) have to make side-payments $c_r \in \mathbb{R}^-$ such that $iw_r^\Sigma = u_r(X^\Sigma) + c_r = (1/n) \cdot sw^\Sigma$.</p> <p>C4. For goods g_ℓ in G with $u^*(g_\ell) = P_{CA}(g_\ell)$ that had to be sold by the CA, the CA has to make side-payments $c_{CA} \in \mathbb{R}^-$ such that $c_{CA} = -\sum_{\lambda \in \{\ell\}} P_{CA}(g_\lambda)$.</p> <p>C5. All agents a_s in S (i.e., all disadvantaged agents) receive side-payments $c_s \in \mathbb{R}^+$ such that $iw_s^\Sigma = u_s(X^\Sigma) + c_s = (1/n) \cdot sw^\Sigma$. Note that $-\sum_{\sigma \in \{s\}} c_\sigma = c_{CA} + \sum_{\rho \in \{r\}} c_\rho$.</p> <p>The CA discloses social welfare sw^Σ, the assignment of goods g_i according to X^Σ and side-payments c_j for each agent a_j in A.</p>
--

Fig. 1. A market-affected sealed-bid auction protocol for any number of goods and agents.

4. **C5.** Side-payments c_r and c_s can be made to or received from several agents, and agents either make side-payments or receive side-payments. Agents a_t in T have been in possession of a proportional share of social welfare sw^Σ after the assignment phase already, i.e., $c_t = 0$ and $iw_t^\Sigma = u_t(X^\Sigma) = (1/n) \cdot sw^\Sigma$ for all a_t in T .

5 Results and Discussion

The easiest way of equitably allocating all goods in G would be if the CA itself would sell all goods on the market, and distribute the profit made in a proportional manner among all agents. In this case, the overall social welfare sw^Σ would equal $P^\Sigma = \sum_{i=1}^m P_{CA}(g_i)$. Thus, $(1/n) \cdot P^\Sigma$ specifies the minimal individual welfare each agent in A is guaranteed to obtain through the allocation of all goods in G by the MSAP. However, taking each agent's preferences into consideration may increase all individual welfares, and the overall social welfare accordingly, up to any amount. Concerning

unwanted goods, the MSAP includes the option to sell those goods with the best possible profit, which gives an opportunity to increase the overall social welfare, and which guarantees that the overall social welfare is not devaluated by some “out-of-favor” good. Furthermore, the MSAP takes into account that agents may have exceptionally low selling costs for one or the other good, and by this keeps all selling costs as low as possible.

After the application of the MSAP, every agent a_j in A possesses a bundle of goods (which may be empty) and some side-payments (which may be positive, negative, or zero). Note that each agent’s individual welfare iw_j^Σ is at least as high as a proportional share of the overall social welfare that could have been achieved if all goods in G would have been allocated to this agent, i.e., $iw_j^\Sigma \geq 1/n \cdot \sum_{i=1}^m u_j(g_i)$. Consequently, by including all agents’ preferences each agent experiences an increase of what he or she actually receives over what he or she anticipated to receive according to his or her measure. Combining utility values and side-payments, individual welfare iw_j^Σ can be interpreted as the bundle (consisting of goods and/or money) agent a_j received by final allocation X^Σ . Analogously, social welfare sw^Σ can be linked to the concept of utilitarian social welfare, i.e., $sw_u(X^\Sigma) = sw^\Sigma$, and to the concept of egalitarian social welfare, i.e., $sw_e(X^\Sigma) = 1/n \cdot sw^\Sigma$, again in consideration of monetary side-payments.

Theorem 1. *Every allocation obtained by the MSAP maximizes both utilitarian social welfare and egalitarian social welfare according to the agents’ valuations.*

Proof. The allocation of all goods g_i in G is conducted in a way such that each good is assigned to a highest bidder, or to the CA, respectively. By this, the overall social welfare to be distributed, which turns out to correspond to the utilitarian social welfare of final allocation X^Σ (including side-payments), is maximized on the basis of all agents’ valuations of the goods. Given each agent’s utility that would result from receiving good g_i , every other allocation that assigns at least one good to another agent than one of the highest bidders would result in a lower utilitarian social welfare.

Maximization of egalitarian social welfare follows immediately from maximization of utilitarian social welfare, since the MSAP makes each agent to receive a proportional share of utilitarian social welfare $sw_u(X^\Sigma)$ according to his or her measure. \square

From an inter-agent perspective, each agent values every other agent’s bundle at most as much as his or her own, and thus an envy-free allocation is guaranteed. Moreover, no agent can be made better off without making any other agent worse off.

Theorem 2. *Every allocation obtained by the MSAP is Pareto-optimal and envy-free.*

Proof. The notion of efficiency implies that there is no better overall outcome for the set of agents involved [3]. Taking this statement into account, Pareto-optimality follows immediately from Theorem 1. Envy-freeness is easy to see when considering that each good g_i in G is allocated to one of the agents that bid the highest value $u^*(g_i)$, and that the very same agent has to make side-payments (each valued $(1/n) \cdot u^*(g_i)$) to the $n - 1$ other agents.⁵ Since each of the $n - 1$ other agents values good g_i at most $u^*(g_i)$, this guarantees that none of those $n - 1$ agents envies this agent for having received good g_i . Concerning goods that are sold by the CA, each agent is receiving the same proportional monetary share of the profit made, and hence, in this case too, no envy is created. \square

⁵ For the sake of convenience, all side-payments of agent a_j sum up to one final side-payment c_j .

Budget balance of the MSAP (in the sense of Definition 4) follows immediately from steps C3, C4, and C5 as given in Figure 1.

Corollary 1. *The MSAP is budget-balanced.*

Our last result shows that the market-affected sealed-bid auction protocol presented in Figure 1 is weakly strategy-proof (in the sense of Definition 5).

Theorem 3. *The MSAP is weakly strategy-proof.*

Proof. With reference to any good g_i in G , a “cheater” (i.e., an agent in A not telling the truth) could cheat by either stating a higher or a lower than his or her true utility value. In the former case, if the cheater wins the bid for g_i , he or she, just as all other agents, will obtain a higher individual welfare, but at the expense of the cheater as he or she has to compensate for the difference—a difference which in fact is only fictitious. That is, the cheater has to pay compensations out of a fund that does not exist and which is based on untruthful values only. Consequently, all agents but the cheater would benefit from this type of cheating. A cheating attempt would be reasonable only if a utility value lower than the true one is stated. Referring to this, if an agent’s true utility value is higher than the market price $M(g_i)$ (i.e., if this agent wants good g_i for him- or herself due to its high significance), he or she is motivated not to cheat by stating a lower utility value, since in this case he or she may end up not getting g_i at all. On the other hand, if an agent that is not interested in keeping good g_i would cheat by stating a lower than the true utility value, this cheating attempt would succeed if, firstly, the cheater has the highest bid, and secondly, the cheater’s bid is at least as high as profit $P_{CA}(g_i)$. In contrast, if the second condition is not fulfilled, good g_i would be sold by the CA causing a decreased individual welfare for all agents, including the cheater. This is the reason why for all g_i in G selling cost $S_{CA}(g_i)$ and profit $P_{CA}(g_i)$ are not disclosed, aiming to motivate all agents to state true utility values as otherwise they would risk their share. To sum up, trying to cheat by stating a higher than the true utility value is of no advantage to the cheater, and trying to cheat by stating a lower than the true utility value always bears the risk of ending up with even less. Thus, the MSAP is weakly strategy-proof. \square

6 Conclusions

We have proposed a new market-affected sealed-bid auction protocol that can be applied to a set of any number of agents and a set of any number of goods. We have shown this budget-balanced and weakly strategy-proof protocol to possess nice properties such as always providing an envy-free and Pareto-optimal allocation with maximal utilitarian and egalitarian social welfare. In addition, this protocol guarantees each agent to receive a bundle that is worth at least as much as a proportional share of all goods, according to his or her measure. These advantages notwithstanding, we mention the following limitations of this protocol. Depending on the scenario, agents not selling those goods received may need to have some cash at hand or to hold sufficient liquid assets in order to make side-payments. In this regard, poorer agents not holding sufficient liquid assets could, to play safe, accept each good for selling only, and by this avoid any trouble.

However, the total of side-payments to be made by an agent never exceeds the value of all to this agent assigned goods, and thus each agent is guaranteed to gain in individual welfare by the application of the MSAP. Moreover, having huge amounts of the same good to be sold may have an impact on the market price of this good. Note also that weak strategy-proofness (in the sense of Definition 5) is a quite softer concept than the common notion of strategy-proofness.

In terms of future work, one direction to go could be the involvement of each agent's wealth by using weighted utility values for all significant goods. In this way, poorer agents could be motivated to make bids that reflect each good's true significance; rather than, fearing side-payments, to accept each good for selling only.

References

1. E. Aragones. A solution to the envy-free selection problem in economies with indivisible goods. Technical Report 984, Northwestern University, Center for Mathematical Studies in Economics and Management Science, Apr. 1992.
2. C. Beviá. Fair allocation in a general model with indivisible goods. *Review of Economic Design*, 3(3):195–213, 1998.
3. S. Brams and A. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
4. B. Cassiman and S. Sieber. The impact of the internet on market structure. Technical Report D/467, IESE Business School, July 2002.
5. Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
6. V. Conitzer, T. Sandholm, and P. Santi. Combinatorial auctions with k -wise dependent valuations. In *Proceedings of the 20th National Conference on Artificial Intelligence*, pages 248–254. AAAI Press, 2005.
7. P. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
8. T. Groves. Incentives in teams. *Econometrica*, 41(4):617–31, 1973.
9. S. Ohseto. Implementing egalitarian-equivalent allocation of indivisible goods on restricted domains. *Journal of Economic Theory*, 23(3):659–670, 2004.
10. S. Pápai. Groves sealed bid auctions of heterogeneous objects with fair prices. *Social Choice and Welfare*, 20(3):371–385, 2003.
11. T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Winner determination in combinatorial auction generalizations. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 69–76. ACM Press, July 2002.
12. Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, 2009.
13. H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
14. L. Svensson. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4):939–54, 1983.
15. K. Tadenuma and W. Thomson. No-envy and consistency in economies with indivisible goods. *Econometrica*, 59(6):1755–67, 1991.
16. K. Tadenuma and W. Thomson. Games of fair division. *Games and Economic Behavior*, 9(2):191–204, 1995.
17. S. Willson. Money-egalitarian-equivalent and gain-maximin allocations of indivisible items with monetary compensation. *Social Choice and Welfare*, 20(2):247–259, 2003.