

Multidimensional Fairness for Auction-based Resource Allocation

Albert Pla^a, Beatriz López^a, Javier Murillo^b

^a*Universitat de Girona, Spain*
{albert.pla,beatriz.lopez}@udg.edu

^b*Newronia, Spain*
javier.murillo@newronia.edu

Abstract

Multi-attribute resource allocation problems involves the allocation of resources on the basis of several attributes, therefore, the definition of a fairness method for this kind of auctions should be formulated from a multi-dimensional perspective. Under such point of view, fairness should take into account all the attributes involved in the allocation problem, since focusing in just a single attribute may compromise the allocations regarding the remainder attributes (e.g. incurring in delayed or bad quality tasks). In this paper, we present a multi-dimensional fairness approach based on priorities. For that purpose, a recurrent auction scenario is assumed, in which the auctioneer keeps track on winner and losers. From that information, the priority methods are defined based on the lost auctions number, the number of consecutive losing, and the fitness of their loser bids. Moreover, some methods contain a probabilistic parameter that enables handling wealth ranking disorders due to fairness. We test our approach in real-data based simulator which emulates an industrial production environment where several resource providers compete to perform different tasks. The results pointed that multi-dimensional fairness incentives agents to remain in the market whilst it improves the equity of the wealth distribution without compromising the quality of the allocation attributes.

Keywords: Multi-attribute resource allocation, Multi-agent systems, Fairness, Auctions

1. Introduction

Auction mechanisms offer the possibility to allocate resources and services in a market (e.g. a company which desires to externalize a production task) while optimizing the outcome of all of the participants (bidders and auctioneers). Thus, given a production task, resource providers bid for it, and the winner bid is the one that best fits the required resource specifications. The auction ends when the winner bid receive's the payment from the auctioener.

The auction designer's goals include optimizing the payoff or revenue of bidders and auctioneers, so that all the participating agents are gratified and keep in the market place. To evaluate how the bidders are satisfied with the auction outcome, as well as the revenue obtained by the auctioneer [3, 7] social welfare measures can be defined. The utilitarian view of social welfare has been the main followed approach, and consist on aggregating all of the agents' outcomes towards maximizing their payoff or revenue. In this utilitarian approach, the aggregation does not consider the fact that there could be big differences among agents' payoffs. When auctions are repeated over time (recurrent auctions), this situation may lead to the unsatisfaction of certain participants which eventually may decide to leave the market. When this occurs, only the most powerful bidders remain in the market, gaining the chance to create an oligopoly, control the market price and provoking a general fall of prices which can bankrupt the auctioneers. Literature often refers to these situations as the bidder drop problem [11] and the asymmetric balance of negotiation power [18].

To tackle these problems, fairness measures have been used in auction design, in what is known as egalitarian social welfare approaches. In this scenario, the behavior of bidders can be totally selfish as it is the auctioneer agent who uses fairness measures to distribute the revenues to keep bidders interested in participating and, therefore, reducing the bidder drop problem. At the end, more bidders mean a higher competition among them, leading the market to more competitive prices. Our work concerns recurrent auction mechanisms with a fairness functionality handled by the auctioneer in order to maintain its own revenue in the long run.

Particularly, we focus in the use of auctions for allocating resources to services where prices are one among other relevant aspects to take into account when clearing the auction. Attributes such as service time, distance among providers, ecological footprint, etc. can play an important part in the

process of the determining which suppliers best suits the production needs. Therefore, it is important to find a compromise between all the elements that condition the resources in order to obtain a satisfying allocation. Multi-attribute auctions offer the chance to consider different aspects besides the price, becoming an ideal option for the problem we are dealing with. Consistently fairness cannot be limited to the payoffs and revenues obtained by agents due to the fact that focusing the application of fairness in a single attribute (price) may involve undesirable consequences regarding the rest of attributes (e.g. unbalance workloads or produce delays). To prevent this issues, we propose to apply fairness mechanisms considering not only the economic aspects of the auction but also the remaining attributes involved in the resource allocation decision making process, in what we called the multidimensional fairness.

In this work, we explore a multidimensional fairness mechanism based on priorities in order to increase the social welfare resulting from a large sequence of auction allocations. Priorities are computed using information regarding all the attributes involved in the resource allocation process, avoiding unwanted behaviors, as for example, the situation in which the auctioneer achieves a cheap price but a large delay on performing some tasks, as other fairness mechanisms based exclusively on price could exhibit. We present a collection of different priority methods for including multi-dimensional fairness to a multi-attribute auction mechanism: two qualitative and two quantitative approaches with a deterministic version and an stochastic version of each one. To illustrate and test the auction mechanism, we simulate an industrial environment where different agents auction services which must be carried out by external service providers while trying to obtain a service in a reasonable price and time.

This paper is organized as follow: first we present a brief state of the art regarding multi-attribute auctions and fairness in auctions; second, in Section 3 we introduce some basic concepts regarding auctions and multi-attribute auctions; next, in section 4 we propose to endow the auction mechanism with multidimensional fairness by means of priority-based methods; in Section 5.2 we test the approach proposed in a task allocation simulator in order to analyze how it affects the social welfare of the allocations and, finally, in Section 6 we present the conclusions of our work and we point some possible future lines of research.

2. Related work

Despite it has been proven that preserving the number of participants in an auction increases its efficiency in the mid and long term [16] and that evaluating fairness from a multi-dimensional point of view can lead to a higher customer satisfaction, few research has been done in the field of fairness in recurrent auctions. From the best of our knowledge, there are no previous auction approaches dealing with fairness from a multidimensional point of view; however, there are some previous works involving a single factor, the price, conditioning fairness. In [11], the authors improve the welfare of the weakest agents by establishing a reservation price so that some goods cannot be sold if the prices offered by bidders are under this reservation price. The remaining goods are then distributed along the weakest agents. Regarding the implications of fairness in the resulting allocations, in [12] the authors make a comparison between fairness and efficiency concluding that there is a compromise between both. However, Murillo et al. [16] claims that this assertion may be true in the short-term but, if analyzing auction recurrence in the long-term, fairness-based methods also become efficient as auctioneers obtain as many benefits as with not fairness-based methods.

In [17], a first price single-dimension fair method based on priorities is proposed. In it, the auctioneers assign a priority to each agent according to its auction history. Bidders with a worst auction win record have a higher priority, meaning that they have higher chances to win the next auction. The authors propose to use the priority attribute to condition the auction clearance by aggregating the priority to the bid, however, in the payment rule, the priority attribute is omitted and does not conditions the payment the winners receive. Although the authors prove that their approach solves the bidder drop problem without incurring to a resource waste (leaving certain goods without a buyer), the mechanism suffers from an asymmetric balance of negotiation power [11, 16] caused by the lack of incentive compatibility (the strongest winners obtain a higher utility by underbidding). Our proposal, when using with a single dimension, can be compared to this previous approaches. In addition to presenting a solution to the bidder drop problem without incurring to a waste of resources, as the previous approaches, it reduces the asymmetric balance of negotiation power since using the priority attribute during the whole auction process (including the payment rule) allows us to follow a second price philosophy which encourages bidders to reveal their true values.

3. Background

This section provides some basic notions regarding auctions and multi-attribute auctions.

3.1. Auctions

An auction is a method for buying and selling goods using a bid system in which the winner bids obtain the auctioned goods [10]. When auctions are used to allocate resources or tasks, the mechanism commonly follows a reverse auction structure (the auctioneer is the seller and bidders are the buyers), which is the approach followed in this paper, and consists in the following steps:

1. Request for proposals: The auctioneer defines the item which is going to be auctioned (e.g. a task) and notifies the bidders.
2. Bidding: Bidders analyze the call for proposal and decide if they are interested in entering the auction. If so, they submit a bid B_i with their offer.
3. Winning determination problem (WDP): The auctioneer evaluates the bids. As this is a reverse auction, the cheapest one is considered the winner.
4. Payment mechanism: The auctioneer pays the winning bidder an amount that would depend on all of the received bids, depending on the payment rule (e.g. in a second price or Vickrey auction, it would pay the price offered by the second best bid [25]).

In recurrent auctions this process is repeated over time for each task or resource to be allocated. Be aware, that repetition does not mean that the same task is repeated over time, but a set of tasks, with different arrival times and resources demands. Repetition could be due, for example, to the just in time nature of manufacturing processes, when tasks arrives in an expected distribution time, but the actual arrival times are unknown.

A desirable property that an auction mechanism should provide is to ensure that bidders provide truthful bids (incentive compatible mechanism); which means that bidders obtain a better profit by revealing their real offer (true values) than by cheating. The most popular mechanism that guarantees

true bidding is the Vickrey auction [14], a second price sealed bid auction where the winner of the auction is the bid with the best offer and where it pays (receive in reverse auctions) the amount proposed by the second best bid.

On the other hand, a common assumption in auctions is the absence of externalities, which means agents do not take care of which are the other agents winning the auction and that agents are not influenced by aspects from outside the current auction [4]. We are making this assumption in our work, too.

Regarding the winner determination problem, it involves the use of a valuation function on bids, $V(B_i)$, so that the auctioneer can select the best bid according to it. In the case of a reverse auction, the auctioneer is faced then to the following problem:

$$\operatorname{argmin}_i(V(B_i)) \tag{1}$$

In the simplest auction mechanism, bids consist on the price b_i (i.e. $B_i = b_i$), and the valuation function is the value of the price $V(B_i) = b_i$. When the bid contains other information than prices, the auctions are known to be multi-attribute [20] and the winner determination problem becomes much more complex.

3.2. Multi-attribute auctions

In multi-attribute auctions each bid is characterized by a set of attributes in addition to price (e.g. the auctioneer can take into account other aspects such as delivery times or CO₂ emissions). Therefore, the winner determination problem consists on finding the optimal bid regarding price but also the remainder attributes. Since the results depend on the auctioneer's goal, its objective optimization function, also known as scoring rule [1], is made available to the bidders. Whenever the scoring rule coincides or not with evaluation function $V(B_i)$ would depend on the auctioneers strategy [19].

From the optimization point of view, equation 3.1 require a more complex solution such as certain multicriteria methods [13, 21]. Moreover, multi-attribute auctions make the payment method more difficult; e.g. in a second-price auctions like the Vickrey one, the winner gets the good and pays the price of the second best. However, when several attributes are involved in the auction, the notion of second price is fuzzier [1]. A solution is to provide a set of attributes in such a way that its valuation is not lower than the

second best bid: e.g. keeping the bid attributes but modifying the economic value in order to equal the second best bid. An example of a Vickrey-based auction is VMA2 [22], as explained in Section 4.

3.3. Position auctions

Multi-attribute auctions handles attributes regarding the task or item auctioned, but other attributes can be considered in a valuation function, as attributes provided by the auctioneers. This is the case of Google's sponsored search position auctions. In position auctions an auctioneer (e.g. a search engine) auctions K slots to place an advertisement related to a given set of keywords. Then, bidders participate in order to achieve one of those slots, the auction is cleared by assigning the first slot to the best bid and the i^{th} slot to the i^{th} best bid until there are no more available slots. The payment is then done following a second price policy: the bidder which obtained the k slot pays the price offered by the $i^{th} + 1$ best bid.

In pay per click position auctions bidders offer the amount they aim to pay each time their ad is clicked, in other words, they pay per click (ppc). In this way, the bidder which will obtain the first slot will be the one with the highest ppc, and the i^{th} slot will go to the i^{th} highest ppc bid. However, this approach does not grant that the auctioneer utility will be maximized (e.g. a bidder which offers a high ppc may have a bad reputation, not receiving any click, thus the auctioneer would receive a payment of 0). To avoid that revenue loss, Google and Yahoo propose to add an attribute at describing the likelihood of an advertisement being clicked given the quality of its design and the destination web page [24, 6]. This attribute is incorporated into the auction by the auctioneer itself and it is aggregated with the ppc offered by bidders. Thus, the winner of the first slot is the bid which offers the best combination of ppc and at . The payment is then computed following a second price, the price per click of the i^{th} winner of the auction will correspond to ppc it should have offered to equal the $i^{th} + 1$ best bid.

It is important to note that the inclusion of this quality attribute does not affect the bidding strategy of bidders. The only influence of quality over bidders is that it encourages them to improve their ads and their web pages in order to make them more attractive to the auctioneer so they can get a higher quality valuation.

4. Multidimensional fairness

Our goal is to define a fairness mechanism that considers different attributes contributing in the decision of an auction-based resource allocation process. In particular, we suppose an environment where a set of resource agents compete for tasks that are repeated arbitrarily over time. Thus, the auctioneer can keep track of the historical outcomes of each bidder on past auctions, identify resources agents in risk of leaving the market, and making fair decisions to keep agents interested in the market.

To achieve our goal, we propose to assign priorities to resource agents according to their history in the resource allocation process (in a similar way that is done in other allocation problems where competitors are rated based in their past performances [2, 9]).

Priorities are defined in $w_i \in [0, 1]$. The higher is the ratio of lost auctions, the higher is the priority of the bidder and its probabilities of leaving the market. Consistently, priorities increase the chances of bidders of winning one of the succeeding auctions.

Priorities can be handled as a bidder attribute in the auction model, similar to the click likelihood in the sponsored search position auctions. However, position auctions only considers price to characterize tasks while we are dealing with several task attributes (price, time to deliver, etc.). Therefore, we use the VMA2 model that allows us to combine both, task attributes and bidders attributes.

In a nutshell, when a task should be fulfilled, an auctioneer calls for an auction, and all the interested bidders submit their bids B_i . Bids are vectors composed by all the attributes characterizing the tasks, in addition to price. Next, the auctioneer adds the priority attribute corresponding to each bidder into their bids ($B'_i = B_i \oplus \langle w_i \rangle$). As any other attribute, priority w_i will condition the winner of the auction and its revenue inasmuch as w_i is taken into account both in the winner determination problem and in the payment mechanism. Once the auction is finished, the auctioneer updates the priorities of the participants according to the results of the auction.

Fairness at the global level, then, is not a matter of a single attribute, as price or priority, but multidimensional as it comes from a set of all of the attributes involved in the auction: the ones characterizing the task (price, quality, delivery times, etc.) as well as the bidders (priority).

On defining the methods to compute the priority, one option is to directly relate the priority with the number of victories and defeats bidders

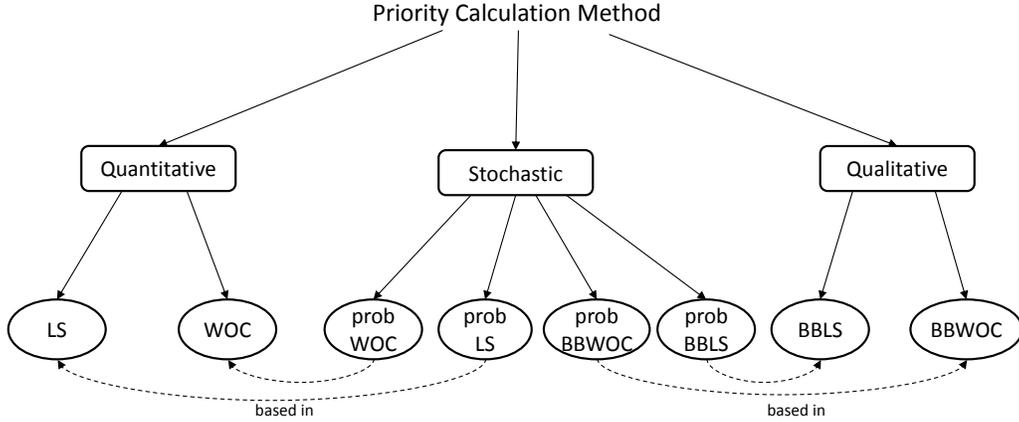


Figure 1: Classification of the priority calculation methods analyzed in this section.

have obtained in the participated auctions. This method offers a quantitative approach which just takes into account if a bidder has won or not. An alternative is to combine the information regarding victories with the bids offered by bidders and the valuation function used by the auctioneer. This approach allows to evaluate the priority in a qualitative way since the bid information can be used to determine how far from the victory a bid was, giving a higher priority to the bidders which do not win but offer high adjusted bids. Finally, both the quantitative and qualitative method can be also adapted under an stochastic model in order to deal with some possible side effects of fairness.

In the remaining of the section, we first provide the fundamentals where the priorities are defined and used, and then we describe the different methods are explained grouped according to the classification provided in Figure 1: quantitative, qualitative (or bid-based) and stochastic.

4.1. Underlying framework

In VMA2 [22] both, task attributes and bidder attributes are combined in a single mechanism. The protocol is composed by the following steps:

1. Request for proposals: The auctioneer defines the tasks requirements as time to deliver a tasks, minimum quality, etc.

2. Bidding: Interested bidders provides bids B_i , composed by the price b_i and values for the requested attributes, at_1, \dots, at_n , such that $B_i = \langle b_i, at_1, \dots, at_n \rangle$
3. The auctioneer receives the bids, and extends them with the bidder priorities, $B'_i = B_i \oplus \langle w_i \rangle$.
4. Winning determination problem (WDP): the auctioneer, selects the best bid according to

$$\operatorname{argmin}_i(V(B'_i)) \quad (2)$$

5. Payment mechanism: based on a second price schema, but with n attributes in addition to price (see below).
6. Attribute update: The auctioneer updates the information it has concerning the bidders. In this particular case, it computes the priority corresponding to each bidder.

The valuation function could be a multicriteria function (see required properties in [21]). For example, the product function could be used, so when bids B'_i contain information about task price b_i and duration t_i , and bidder priority w_i , we get

$$V(B'_i) = b_i * t_i * \frac{1}{1 + w_i} \quad (3)$$

Regarding the payment rule, it is a two case function. If the winner of the auction delivers the task as agreed (e.g. it respects the delivery time and the rest of bided attributes) it will receive the economical amount p it should have bid to obtain the same evaluation as the second bid. Otherwise, if the bidder cheats¹ and does not delivers the task as it was bid (e.g. it delivers the task to late or with different attributes than the agreed ones), it will receive the economical amount p which equals the valuation of the offered bid (with its original attributes and its economical bid) with the provided

¹Other issues regarding uncertainty in the estimation of the bid attributes on behalf of the bidder are out of the scope of this work, thus we consider that agents are able to accurately estimate their attributes and it is assumed that if agents provide bids that are not fulfilled it is because they cheat.

task (with the delivered attributes). In other words, given the winner bid B'_1 and the second best bid B'_2 (i.e. in a reverse auction as ours, it means that $V(B'_1) < V(B'_2)$, and $(V(B'_2) < V(B'_i)|\forall i > 2)$), the winner bidder receives and amount p equals to:

$$p = \begin{cases} V'(V(B'_2), B'_1 \setminus \{b_1\}) & \text{if receive attributes = bided attributes} \\ V'(V(B'_1), B'_1 \setminus \{b_1\}) & \text{otherwise} \end{cases} \quad (4)$$

where $V'(x, B'_x \setminus \{b_x\}) = b_x$ is the anti-function of $V(B'_x)$ regarding the price, given $V(B'_x) = x$ (see proves and details in [22]); and B'_1 is the bid vector composed by the real attributes provided by the bidder and its priority.

4.2. Quantitative priority methods

Quantitative priority methods take into account the number of victories and defeats achieved by each agent. Using this information the auctioneer gives more priority to bidders with a high defeat record, reducing the risk of these agents leaving the market. Two base methods have been defined: won auction coefficient (WOC) and losing streak (LS).

	Timestamp	0	1	2	3	4	5	6	7	8
a_1	Bid Sequence	-	×	×	✓	×	×	✓	×	×
	WOC w_1	0.00	0.50	0.66	0.50	0.60	0.66	0.57	0.63	0.66
	LS(8) w_1	0.00	0.13	0.25	0.00	0.13	0.25	0.00	0.13	0.25
a_2	Bid Sequence	-	✓	✓	×	×	×	×	×	×
	WOC w_2	0.00	0.00	0.00	0.25	0.40	0.50	0.57	0.63	0.66
	LS(8) w_2	0.00	0.00	0.00	0.13	0.25	0.38	0.50	0.63	0.75
a_3	Bid Sequence	-	×	×	×	✓	✓	×	✓	×
	WOC w_3	0.00	0.50	0.66	0.75	0.60	0.50	0.57	0.50	0.55
	LS(8) w_3	0.00	0.13	0.25	0.375	0.00	0.00	0.13	0.00	0.13
a_4	Bid Sequence	-	×	×	×	×	×	×	×	✓
	WOC w_4	0.00	0.50	0.66	0.75	0.80	0.83	0.86	0.88	0.77
	LS(8) w_4	0.00	0.13	0.25	0.375	0.50	0.63	0.75	0.88	0.00

Table 1: Illustrative example of the auction record of three different agents and the priority they would obtain when using the won auction coefficient and the losing streak (with a ml of 8). ✓ means that the bidder has won the auction whilst × means it has not.

4.2.1. Won auction coefficient (WOC).

The won auction coefficient is the most intuitive way of computing the priority of an agent, it establishes the ratio between the number of auctions in which a bidder has participated ($par(a_i)$) and the number of auctions in which it has won ($won(a_i)$), as shown in the following expression:

$$w_i = 1 - \frac{1 + won(a_i)}{1 + par(a_i)} \quad (5)$$

This measure gives the auctioneer the notion of which has been the proportion of won auctions for each bidder since the auctioneer entered the auction market. Using this priority, a bidder which has won all the auctions will have a priority of $w = 0$ while a bidder which has never won an auction will have a priority close to 1.

To illustrate the behavior of WOC consider, for example, the sequence of wins (\checkmark) and defeats (\times) obtained by a set of agents shown in Table 1.

The WOC method takes into account information regarding the whole period where an auctioneer has been active, thus, in the long run the auctioneer will tend to compensate the bidders which lost many times during the first auctions and it will ballast which had won the most. On the one hand, this means that if a bidder wins an auction after a large period of defeats, its priority will remain high as it still has a low ratio of victories. In the Table 1 example, a_4 won the last auction but it still has the highest WOC priority because it is the agent which have won less auctions. On the other hand, a bidder experimenting a long losing streak will have a low priority if he obtained a high number of victories in the past. For instance, in Table 1, a_2 has a long losing streak but, at the end, it has the same priority as a_1 due to the fact that they both have won 2 auctions (a_2 at the very beginning, a_1 in a more spaced way). To alleviate somehow such a situation, a time window can be adopted (e.g. computing only the auctions summoned during the last week); similarly, limiting the auction record which is taken into account (e.g. computing only the last 100 auctions summoned by the auctioneer) could smooth such compensation and ballast side effects of WOC caused by its long term memory.

4.2.2. Loosing streak (LS).

Under certain circumstances, especially when some of the bidders are directly controlled by human entities, certain bidders may experiment cognitive distortions which can affect their reasoning [15]. E.g. a human-controlled

bidder might be more susceptible to leave the market after a succession of two won auctions and 6 defeats (as agent a_2 of Table 1 with a WOC of 0.66) than after a succession with the same amount of won auctions but where the victories are more spread (e.g. the case of agent a_1 with 2 defeats, 1 victory, 2d, 1v, 2d).

For this kind of scenarios the priority can be defined in the basis of the bidders losing streak $ls(a_i)$ (the number of consecutive auctions which agent a_i has lost). The longer is the losing streak of an agent, the higher his priority. For that purpose, we define a tolerance threshold ml such that after losing a certain number of consecutive auctions as long as ml , a bidder is as susceptible to leave the market as one which has lost $ml + 1$ auctions. Each auctioneer fixes ml according to its believes about bidders behaviors. The ideal strategy would be to assign an specific ml threshold for each bidder. However since the auctioneer does not knows the different tolerances to the losing streaks of the different bidders, it must estimate a generic ml which will be the same for every bidder. Joreover, using a different ml for each bidder leads to an unjust starting point as not all the bidders would be treated under the same rules (e.g. if two agents exchange their bid sequences they may obtain different results since they are participating under different conditions). There fore, the same ml is assigned to each bidder.

We use this threshold to define the highest priority of the bidder as follows:

$$w_i = 1 - \frac{\max(0, ml - ls(a_i))}{ml} \quad (6)$$

where ls_i is the losing streak of bidder a_i . In this way, a bidder which has just won an auction has a priority of $w = 0$ whilst a bidder which has been ml auctions without winning has a priority of $w = 1$. Conversely to the previous presented method, we can say that the LS method has no long term memory as a bidder will see its priority reduced to 1 after winning an auction independently to which has been his record of results in the past.

Following the example of Table 1, we can see how despite agent a_1 and a_2 have the same number of victories and same priority when using WOC, a_2 has a higher priority when it is computed using LS due to the fact that it has the highest losing streak.

4.3. Qualitative bid-based priorities

The above presented methods have an absolute behavior, since they treat all non-winning bidders in the same way without taking into account the

attribute values provided on their bids. In other words, the priority method treats in the same way the bidder which almost won the auction (the second best bid) and the bidder which offered the worst bid. On the one hand, this situation may encourage some bidders to provide dummy bids with no chances of winning in order to increase their priorities, obtaining higher chances to win one of the following auctions. On the other hand, some bidders offering almost-winning bids could feel frustrated when being aware that attribute values of bids have no effect on the priority outcome. In order to avoid that, we propose to take into account the attribute values of the offered bid in the priority definition.

For that purpose, we define the fitness q_i of a bid B'_i in relation to the winner bid of the auction B'_1 depending on whether the auction is reverse (left) or not (right):

$$q_i = \frac{V(B'_1)}{V(B'_i)} \qquad q_i = \frac{V(B'_i)}{V(B'_1)} \qquad (7)$$

Following this definition we can say that the winner bid of the auction has the maximum fitness $q_1 = 1$ and that the higher q_i is for a bid, the closer it has been to win the auction. Taking bid fitness into account, we propose to redefine the *Won Auction Coefficient* and the *Loosing Streak* priority calculation methods presented above.

4.3.1. Bid-based won auction coefficient (BBWOC).

The priority is calculated in the same way that with the WOC method, however, instead of taking into account the number of auctions in which an agent has participated, it gathers the fitness of the bids it offered in the j auctions where it has participated.

$$w_i = 1 - \frac{1 + \text{won}(a_i)}{1 + \sum_{j=0}^{c-1} q_i^j} \qquad (8)$$

where q_i^j is the fitness of the bid B'_i in the j auction and c indicates the current auction.

In this way, if there are two bidders which have won the same number of auctions, the bidder which has been providing higher quality bidders will be the one with the highest priority. This can be observed in the example provided in Table 2 where bidders a_2 and a_3 have won the same number of auctions, however, at the end, a_2 has a higher priority (0.623) than a_3 (0.549) as a_2 have been providing better bids.

	Timestamp	0	1	2	3	4	5	6	7	8
a_1	Bids valuation	-	7	10	8	6	5	10	9	9
	Bid fitness q_1	-	0.875	1.000	0.888	0.856	0.555	1.000	1.000	1.000
	bbWOC w_1	0.000	0.467	0.304	0.469	0.567	0.614	0.514	0.443	0.388
	bbLS w_1	0.000	0.109	0.000	0.111	0.218	0.288	0.000	0.000	0.000
a_2	Bids valuation	-	7	8	9	5	9	9	8	7
	Bid fitness q_2 A2	-	0.875	0.800	1.000	0.714	1.000	0.900	0.889	0.778
	bbWOC w_2	0.000	0.467	0.626	0.456	0.544	0.443	0.523	0.582	0.623
	bbLS w_2	0.000	0.109	0.209	0.000	0.089	0.000	0.113	0.224	0.321
a_3	Bids valuation	-	8	5	7	7	7	6	5	4
	Bid fitness q_3	-	1.000	0.500	0.778	1.000	0.778	0.600	0.556	0.444
	bbWOC w_3	0.000	0.000	0.200	0.390	0.299	0.407	0.470	0.517	0.549
	bbLS w_3	0.000	0.000	0.063	0.160	0.000	0.097	0.172	0.242	0.297

Table 2: Illustrative example of the auction record of three different agents, the valuation received by their bids and the priority they would obtain when using the bid based won auction coefficient and the bid based losing streak (with a ml of 8). Winning bids are marked in bold.

4.3.2. Bid-based losing streak (BBLs).

The same can be done with the LS method. This method follows the priority function in Equation 6; however, instead of using the losing streak accumulated by bidders to define the priority, we propose to use the summation of the bid fitness obtained during the losing streak. In this way a long streak of almost winning bids will increase the priority faster than a long losing streak of low quality bids.

$$w_i = 1 - \frac{\max(0, ml - \sum_{j=c-ls(a_i)}^{c-1} q_i^j)}{ml} \quad (9)$$

In Table 2 it can be seen how, despite a_2 (3) has a shorter losing streak than a_3 (4), a_2 ends with a higher priority (0.321 against 0.297) due the fact that its bids have been closer to the victory than the ones bid by a_3 .

4.4. Stochastic priorities.

A problem that may arise due to priorities is the alteration of the agents wealth ranking (see Section 5). E.g. the richest agent when not using the fairness mechanism could become the second richest agent when priorities are used. This problem can be minimized by updating the priority of the bidder following stochastic or probabilistic methods, so sometimes the wealth rank is

altered and some times it is not. In addition, a propbabilistic approach may prevents agents from learning the priority mechanism. Thus, the step 6 of the protocol is modified so the bidders priority is updated or not according to an update priority probability parameter $up_j \in [0, 1]$ defined by the auctioneer a_j . The priority of the bidder's can be computed using any of the previously presented methods, the difference is that with this stochastic approach the priority is not always updated.

In Algorithm 1 we present a stochastic algorithm which can use equations 5, 6, 8 and 9 as functions to calculate the priority of bidders. Depending on the equation choosen to determine the probability we can refer tot he method as probWOC, probLS, probBBWOC or probBBLs.

Algorithm 1 Stochastic priority algorithm (auction step 6)

```

1: for all Bidder  $a_i \in$  auction_participants do
2:   if  $Random(0, 1) < up_j$  then
3:     update  $a_j.priority[a_i]$ 
4:   else
5:      $a_j.priority[a_i] = a_j.priority[a_i]$ 
6:   end if
7: end for

```

Using a stochastic method bidders will not have the certainty of which is their priority and they will not be able to take advantage of it. Moreover, the use of an update probability can smooth the alteration of the wealth ranking as it is reasonable to expect that a low update priority probability will reduce the fairness of the final allocation but it will also reduce the wealth ranking alteration. Therefore this approach rises the issue of finding a good compromise between fairness and wealth ranking preservation.

5. Experimentation

To test and illustrate the behavior of the presented mechanism, we simulated the processes and the services required by a company during a certain time period.

5.1. Experimental set up

The problem domain used to analyze the properties of the presented mechanism consists on the operation of an industry which needs to solve its

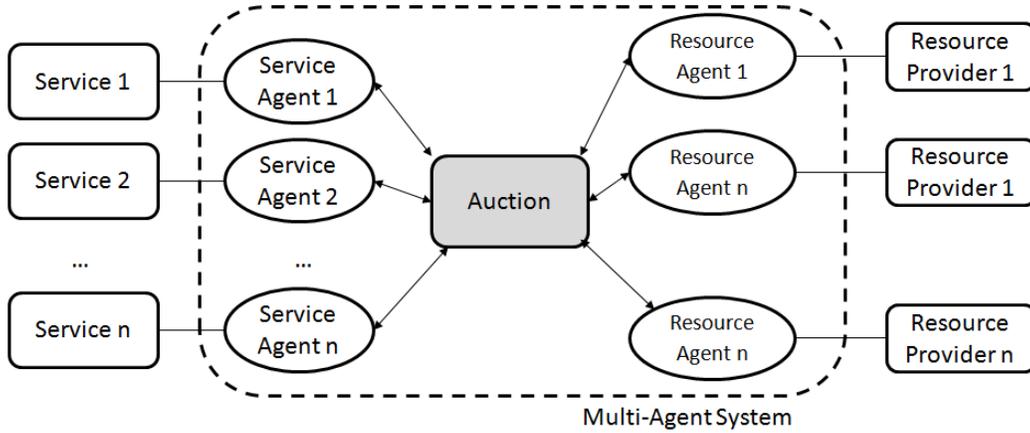


Figure 2: Multi-agent System architecture for the simulation

	S1	S2	S3	S4	S5	S6	S7		S1	S2	S3	S4	S5	S6	S7
RP1			✓		✓		✓	RP5		✓	✓		✓		
RP2	✓	✓	✓	✓	✓	✓	✓	RP6					✓		
RP3		✓	✓		✓		✓	RP7	✓	✓	✓	✓	✓	✓	✓
RP4	✓	✓	✓	✓	✓	✓	✓	RP8	✓		✓	✓	✓	✓	✓

Table 3: Capacity of a resource provider (RP) to develop a service (S).

customers' issues (e.g. unforeseen faults, local system crashes, etc.). When a customer informs about an incidence, a preliminary diagnosis process determines its urgency and its type, then a task is generated to solve the incidence, which should be assigned to an external resource provider. The urgency parameter determines the maximum time in which each incidence must be solved whilst the typology determines the kind of service which must attend the issue, in consequence urgency and type determine the task configuration.

The data used to define the models for the simulation contains real information regarding the incidences notified in a manufacturing company during two years.² In this data there are 7 different types of service (S1 to S7), 8 resource providers (RP1 to RP8) with differentiated skills as each resource provider can perform only certain type of services (see Table 3). For each type of service the occurrence probability distribution among a day has been estimated from the data provided, such as to generate the task arrival times

²Due to confidentiality reasons, we cannot mention the name of the company.

in the simulation environment following the same distribution pattern. The same has been performed for the urgency parameter (to generate the tasks' time window requirement) and for the task service durations of each resource provider (each resource provider has different time distributions for each kind of task).

To ensure the reproducibility of the experiments and to allow other researchers to test their approaches, an anonymized version of the data used and the exact specifications of the distributions can be found in [23].

The simulation is built on the basis of a multi-agent system. Each kind of service and each resource provider is managed by its own agent (See Figure 2). Thus, there are two kinds of agents: service agents SA (which want to solve an incident by performing some tasks and need a resource in order to do so) and resource providers RA (which can supply resources to carry out tasks). Therefore, service agents play the role of auctioneers, while resource providers the bidders ones. Overall, there are 15 agents (7 SA and 8 RA) which compete to buy and to provide resources. All SA have the same configuration regarding the priority methods (namely ml and up_j). When an SA needs to allocate a task to a resource, it summons an auction (defined by the service type and the task execution time window) and all the resource providers which can carry the required type of service offer their bids. Bids are composed by two attributes: the bid amount b_i corresponding to the task's cost and the task delivery time t_i . The valuation function used during the winner determination problem and the payment is the product, as shown in Equation 3. The antifuction used to compute the payment is:

$$p = \begin{cases} \frac{(b_2 \cdot t_2) \cdot (1 + w_1)}{t_1 \cdot (1 + w_2)} & \text{if } t' \leq t_1 \\ \frac{(b_1 \cdot t_1)}{t'} & \text{otherwise} \end{cases} \quad (10)$$

where p is the payment the winner will receive, b_1 the economic bid of the winner bid, t_1 the delivery time offered in the winner bid, b_2 and t_2 the economic and delivery time attributes of the second best bid, t' the final delivery time the winner provides and w_1 and w_2 the priorities of the best and the second best bids.

According to the goal of this simulation we propose the following three studio scenarios:

- *Scenario 1:* The goal of this scenario is to show the differences between multi-dimensional fairness and uni-dimensional fairness measures. The

uni-dimensional fairness is based only on the price, as the previous mechanisms in the literature. Particularly, we define the priority in the uni-dimensional approach based on the price alone [16] (traditional concept of 'revenue') as follows:

$$w_i = 1 - \frac{\sum_{au=0}^{par(a_i)} r_i^{au}}{\sum_{l=0}^{l=n} \sum_{au=0}^{au=par(a_l)} r_l^{au}} \quad (11)$$

where r_l^{au} is the revenue that the bidder l obtained in the auction au and n the number of bidders in the market.

This priority relates the economic amount a bidder has won compared to the total economic amount which has been paid by the auctioneer. To complement the experiment, a mechanism without fairness (no priority) is also taken into account. Regarding the methods presented in this paper, we test WOC, LS, BBWOC and BBLs, leaving the stochastic approaches for another scenario. About the parameters required by the priority methods, the values of ml in the LS and BBLs methods have been set to 70 (as in the simulation, in average, a resource provider participates in 10 auctions per day and we considered that a week without working is enough time for an agent to leave the market).

- *Scenario 2:* The goal of this experiment is to analyze the role which probability plays in the stochastic priority methods. Particularly, we focus on the influence of the priority in the wealth ranking achieved by agents. In this scenario the experiment is repeated several times with different priority update probabilities up_j , starting from $up_j = 0$ and increasing gradually by 0.1 until $up_j = 1$. All SA have the same up_j value.
- *Scenario 3:* The aim of this scenario is to evaluate how the use of priorities, defined in a multidimensional way, minimize the bidder drop problem and other related problems. To this end, bidders are designed in such a way that if after a certain number of auctions they have not obtained benefits, they leave the auction market. All the bidders have the same skills (RP6) and true values (55€ per task, and the same execution times) but follow different bidding strategies: half of them are honest and adaptive, while the other half are cheaters.

To have enough kind of agents of both kinds, we have increased the amount of bidder agents up to 50 with respect to the previous scenarios.

	Service Agents (auctioneers)	Resource Providers (bidders)	Bidders Strategy	Priorities Used
Scenario 1	7 (S1 to S7)	8 (RP1 to RP8)	Adaptative	No priority, WOC, Uni-dimensional, BBWOC, LS (ml=25), BBLs (ml=25)
Scenario 2	7 (S1 to S7)	8 (RP1 to RP8)	Adaptative	probWOC probBBWOC probLS (ml=25), probBBLs (ml=25)
Scenario 3	10 (type S5)	50 (R6)	25 Adaptative (mcal = 50) 25 Cheater (mcal = 100)	No Priority, WOC, BBWOC, probWOC, LS (ml = 25), BBLs (ml=25)

Table 4: Simulation parameters. Each scenario has a duration of 400 cycles (days) and it is repeated 200 times (mcal = maximum consecutive auctions lost before leaving the market)

On the one hand, 25 of them are follow an adaptive strategy to improve their bids in order to win auctions [11]; they can stay up to 50 auctions without obtaining benefits. On the other hand the cheater agents start following an underbidding strategy (working below their true value cost and affording economic losses for a while) in order to force the adaptive agents to leave the market for creating an oligopoly where they can fix the market price. Once this happens, they try to rise prices in order to increase their benefits. In order to carry this strategy, cheating agents can stay up to 100 auctions without obtaining benefits before leaving the market. To feed the 50 bidders, there are 10 auctioneers which only auction tasks of type S5 (Table 3). The experiment is performed with no priority, WOC, BBWOC, probBBWOC (up=0.7), LS and the BBLs.

Table 4 summarizes the parameter used in the different simulated scenarios. The simulations last 400 cycles which correspond, approximately, to a year of work in the industry and each experiment is repeated 200 times in order to obtain significant data.

5.2. Experimental results

This subsection presents and discusses the results of the proposed scenarios. Results are measured in terms of:

- Price: To evaluate how the priorities affect the price, we compared the revenue among bidders (the higher, the best), the auctioneer's expenditure (the lower the best) and the task allocation mean cost on behalf of the auctioneer (the lower, the best).
- Delivery time: Particularly, since tasks are defined to be performed inside a time window, we are interested in delays caused by tasks with the aim to avoid or minimize them. A one-dimensional fair mechanism that does not take into account this attribute, but only the price, should fail in reducing the delays; while a multi-dimensional fair mechanism as the one presented here, should achieve better results (minimize delays).
- Fairness:
 - The Gini's index [8] is used to evaluate the fairness of the resulting allocations at the end of the simulation (the lower, the best). Gini's index is thought to analyze the social welfare differences among populations, thus, we have adapted it to our domain. The percentage of the population which corresponds to each resource provider is determined using the number of types of task which can perform each service agent. E.g. RP2 has a population of 18.42% as it can perform 7 tasks out of 38, RP5 a population of 7.89% (3/38), etc.
 - The agent's wealth rank modification (WRM). We wanted to evaluate the influence of fairness by comparing the position of an agent in the wealth rank when no fairness is used and when it is. For that purpose, we use Spearman's footrule [5] distance (SFD). In this way, we define the WRM as the SFD between the original wealth rank obtained when no priorities were used (R^o) and the wealth rank when fairness is used (R^f) divided by the SFD between the original wealth rank (R^o) and its reverse rank (R^r);

$$WRM = \frac{sfd(R^o, R^f)}{sfd(R^o, R^r)} \quad (12)$$

where $sfd(X, Y)$ is the spearman’s footrule between the ranks X and Y and $WRM \in [0, 1]$.

With this metric, a WRM of 0 means that the priority has not altered the wealth ranking at all whilst a WRM of 1 means that the wealth rank has been completely altered.

For instance, if we have three agents a, b, c which obtained a wealth of 10,9 and 3 respectively when none priority was used and which obtained a wealth of 8,9 and 5 when fairness is used we can see that there has been 2 displacements in the wealth rank (b goes from the second to the third whilst c varies from the third to the second). Given that inverting the original wealth rank produces 4 displacements (a varies 2 positions by going from the first position to the last and c 2 more positions from going from the last to the first), the WRM is 0.5 ($2/4$) since there has been 2 displacements out of 4 possibles.

- The number of agents which remain in the system is used to evaluate both fairness the bidder drop problem. In a fairer allocation bidders will prefer to stay on the market, a higher number of bidders reflects a fair allocation.

5.2.1. Results on scenario 1:

Figures 3 and 4 show the mean revenue and expenditure obtained by agents during the simulations with the different mechanism considered in this scenario: no priorities, uni-dimensional fairness, and multi-dimensional fairness methods (WOC, LS, BWOC and BBLs). On the one hand, Figure 3 points that the use of priorities changes the costs of the service allocation. It can be seen that, generally, when using priorities, service agents have not needed to increase their expenditure to allocate resources. The only case when an agent had to spent more money to externalize a task when using a priority has been with SA2: when using the LS method SA2 had to pay an average of 34,365.75 euros per simulation whilst when no priority was used it paid 31,652.89. In the rest of cases the agents expenditure when using priority methods has been lower than when no priority has been used. These results are the expected ones in a fairness mechanism which pursues to maintain agents interested in the auction, and then, more agents means more competitive prices.

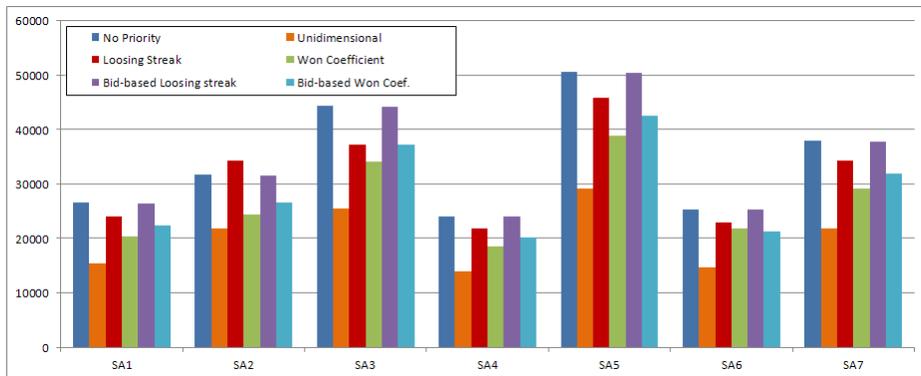


Figure 3: Service agents (auctioneers) mean expenditure at the end of simulations. No priority (dark blue), uni-dimensional priority (orange), LS (red), WOC (green), BBL (purple), BBWOC (clear blue).

On the other hand, regarding the benefits of the resource providers (Figure 3), it can be seen that the use of priorities considerably changes the distribution of the revenues: the differences between agents revenue has been reduced when including fairness into the mechanism as expected. Another remarkable fact is that bidders which obtained poor revenues (e.g. SP1, 5 and 6) increase their benefits. However it is also important to notice that the inclusion of fairness into the mechanism has modified the ranking of the richest service providers (e.g. RP4 is the second richest agent when no fairness mechanism is applied whilst using a bid-based fairness mechanism makes him the fourth richest agent). This result is further analyzed in the next experimental scenario.

To deepen on the analysis of fairness in the allocations, Figure 5 presents a box plot of the Gini's index [8] corresponding to the 200 executions of the scenario. It is important to notice that the differences between agents revenue are conditioned by the skills. In this sense, we can see that the four poorest agents (RP1, RP3, RP5 and RP6) are the ones with less skills and it is reasonable to assume that other agents with better skills will have the chance to perform more tasks. This fact is taken into account in the fairness analysis performed: in the computation of the Gini's index each resource provider is weighted according to the number of skills (population) it is able to perform.

Results obtained clearly show that the use of priorities lowers the inequality coefficient: when no priority is used the index has an average of

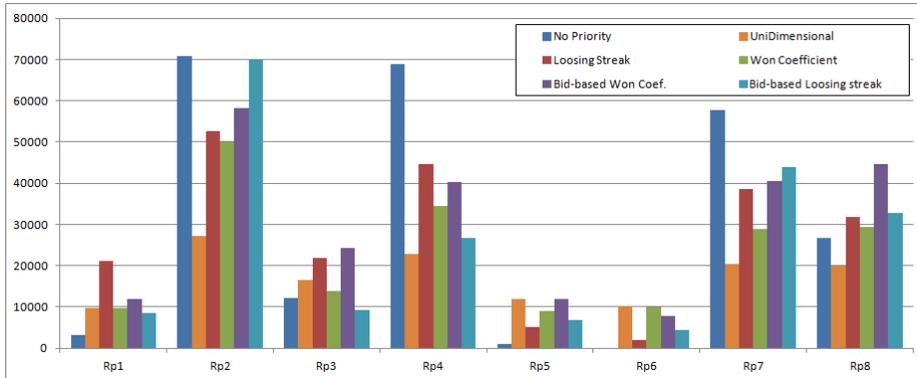


Figure 4: Resource providers (bidders) mean revenue at the end of simulations. From left to right: No priority (dark blue), uni-dimensional priority (orange), LS (red), WOC (green), BBLS (purple), BBWOC (clear blue).

0.3623 whilst the use of priorities results in indexes lower than 0.3078 (with the BBLS). The lowest inequality index is obtained with the uni-dimensional priority (0.2134), followed by the BBWOC (0.2262) and the WOC (0.2371). The variance of the results does not allow us to determine which is the best priority method between those three (2-sample T tests do not point significant differences between the results), however, it can be clearly seen that including a priority into the auction improves the equity of the resulting allocation. Since revenue only takes into account the first attribute of the bids, it is reasonable that the uni-dimensional priority obtained the best Gini's index. However, if we also take into account another attribute (time), results show a different face.

Figure 6 shows the mean number of delays produced during the simulations when using different types of priorities. In this Figure, it is possible to observe that the uni-dimensional priority produces the highest mean number of delays, due to the fact that the uni-dimensional priority is only concerned by the cost attribute, pointing the need of applying fairness in a multidimensional way. When using multidimensional fairness, delays are significantly reduced, as shown in the box-plots. Using Student's t-test with a 95% of confidence we can say that, inside the multidimensional priority methods, the LS is the worst of them and the WOC the second worst. Regarding the WOC and the BB-WOC, the test does not points significant differences between them. It can be seen that when using qualitative methods the number of delays produced (7.242 and 8.003 time units) is similar to when no priority

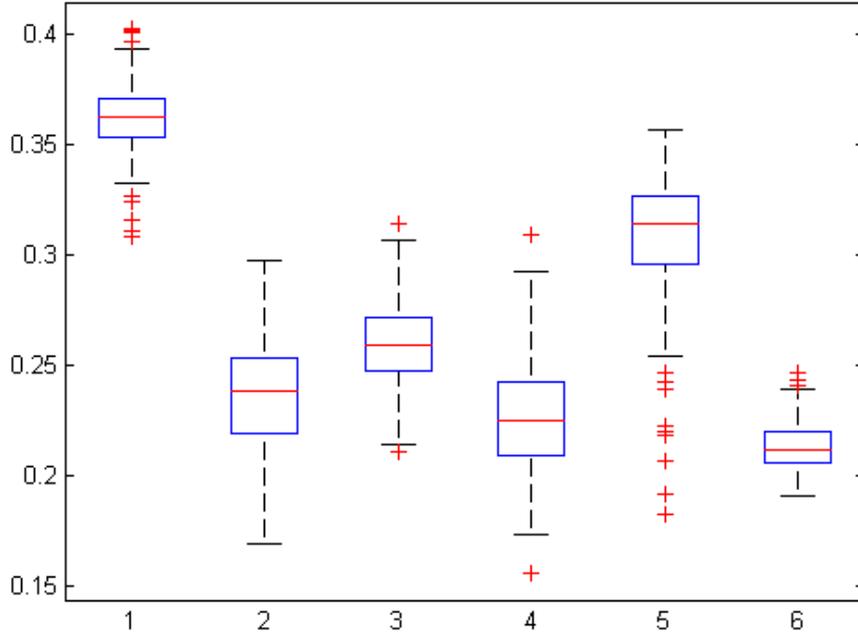


Figure 5: Resource agents' Gini's index for the resulting allocations in terms of revenue: 1. No priority 2. Losing Streak 3. Won coefficient 4. Bid-based WOC 5. Bid-based LS 6. Unidimensional.

is used. This is caused by the fact that bidders which offer almost-winning bids obtain high priority faster than low-quality bidders, in this way, weaker bidders still obtain advantage after a long period losing without jeopardizing the quality of the allocations.

In the boxplot of Figure 6, it is possible to observe that when no priority is used the number of delays produced is also low (7.911). This is due the fact that the auction mechanism chooses the bids which minimize cost and execution time without taking into account priorities. However, the results obtained are worse than the BBWOC method (7.424), and very close to the BBLs method (8.003). If the Gini's index (Figure 5) is put together with the delay results (Figure 6), we can conclude that multidimensional fairness is providing at the long run the better allocation outcomes than no priorities and uni-dimensional approaches; and particular, the BBWOC qualitative

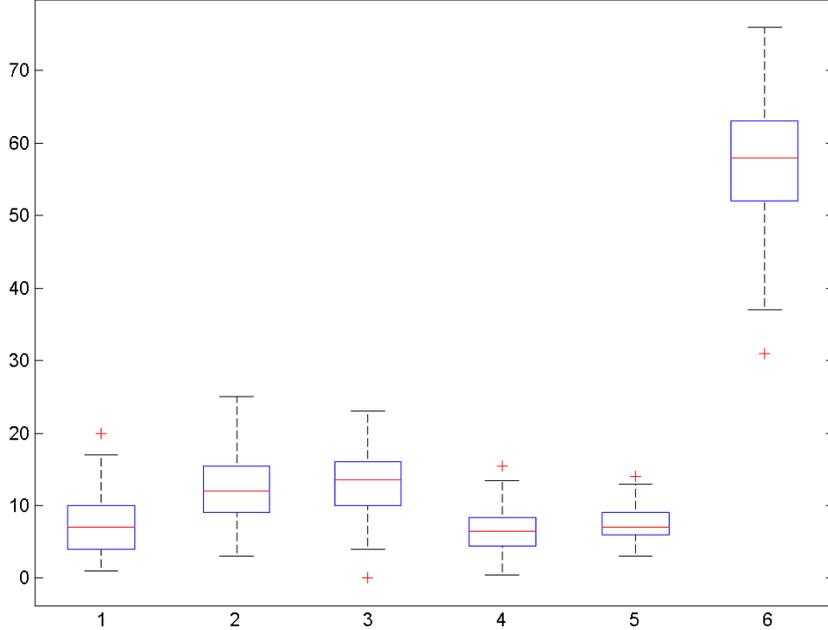


Figure 6: Delays produced: 1. No priority 2. Loosing Streak 3. Won coefficient 4. Bid-based WOC 5. Bid-based LS 6. Unidimensional.

methods is the best.

Summarizing the results of scenario 1, we can conclude that the use of multidimensional priorities can improve the fairness of the task allocation by assigning some services to the weaker agents without increasing the global cost of the allocation nor incurring in significant delays; however by using this methodologies, some of the strongest bidders may experience envy as their incomes could be reduced for the sake of equity.

5.2.2. Results on scenario 2:

This scenario focus the discussion on the wealth rank disorder stated in the previous scenario. For this end, we use the WRM measure defined above (Equation 12) which is based on the the number of positions in the wealth rank that each agents wins or loses due to the fairness mechanism.

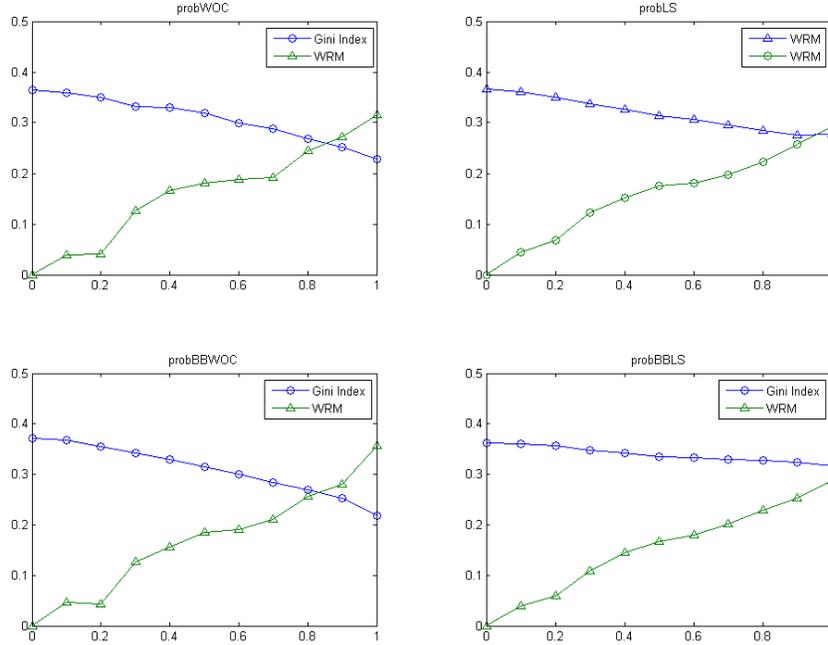


Figure 7: Mean Gini Index (blue circles) and mean WRM (green triangles) for the stochastic priority methods when using WOC (top left), LS (top right), BBWOC (bottom right) and BLS (bottom left).

As stated in the scenario description, simulation consists on varying the up_j parameter from 0.0 to 1.0 (from changing the priority always according to the priority method, to never changing the priority). It is reasonable to expect that results when up_j is 0.0 will be similar to the results of the previous scenario when no priority is used and that when $up_j = 1.0$ results will be similar to the ones obtained in the previous scenario with WOC, BBWOC, LS and BLS. Respect the WRM, ranking disorders are expected to increase as priorities are more often used (the higher the up_j , the higher WRM).

Figure 7 shows the results obtained in terms of the Gini's index and the WRM with the probabilistic versions of WOC, BBWOC, LS and BLS. As expected, in all cases, the resulting allocation is fairer as the higher up_j is because, as commented before, when $up_j = 1$ the method behaves in a deterministic way. When analyzing revenue of the agents, we can see how Gini's

index progressively varies from an initial value (common to all the methods) situated around 0.36 when $up = 0.0$ to different values conditioned by the priority calculation when $up = 1.0$. In all the cases the mean Gini's index obtained when $up = 1$ (probWOC=0.2379, probBBWOC=0.2217, probLS=0.2688 and probBBL=0.3092) are very similar to the values obtained in the previous scenario, confirming that the stochastic methods with $up = 1$ are equivalent to the deterministic approaches.

On the other hand, focusing in the wealth rank disorder which the presented fairness mechanisms may produce, Figure 7 clearly shows that fairness and rank preservation are clearly related: the higher the fairness, the higher the disorder. In all the cases, it can be clearly seen that being fairer implies changing the wealth ranking, and this can cause some complaints or dissatisfaction among bidders who lose some positions in the ranking. Consequently to that, it can also be seen that probBBWOC (with the best fairness) it is also the method with the highest WRM (0.3489) followed by probWOC (WRM=0.3092), probLS (0.2973) and probBBL (0.2901).

Regarding the compromise between fairness and rank preservation, it can be seen that the probWOC and probBBWOC methods with $up_j = 0.7$ and 0.8 may present the best solution since they obtain best fairness than probLS and probBBL but with a lower rank modification. Concordantly to this fact, probWOC and probBBWOC are also the priority methods where the WRM curve and Gini's index curves cross at a lowest up (around $up = 0.85$) whilst in probLS it crosses around 0.95. Finally, in the probBBL the two curves never cross.

To sum up, the stochastic priority methods proposed offer the chance to reduce the alteration of the wealth rank provoked by the priority, especially in the case of probWOC and probBBWOC. However, not silver bullet can be provided and the analyzed methods must be configured according to the mechanism designer needs, by tuning the up_j parameter that best suits to each problem.

5.2.3. Results on scenario 3:

This last scenario analyzes how the use of priorities affects the bidder drop problem, the problem which auction mechanisms must face. Another related problem we analyze is the asymmetric balance of negotiation power between auctioneers and bidders caused by the progressive desertion of agents, decreasing the competitiveness of the market. For that purpose the number of

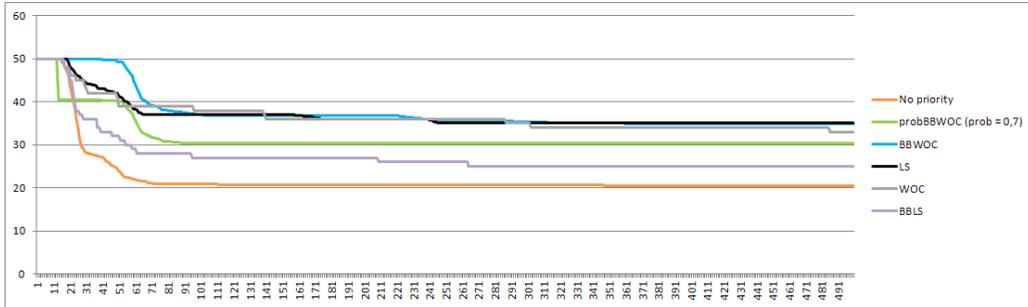


Figure 8: Active bidders in the market during a simulation where half of the agents (25) try to manipulate the auction working under their production cost and where the other half are adaptative.

participants which take part in the auction market among time and the task mean cost are used as metrics.

Figure 8 shows the number of active bidders which are inside the auction market along time when using different priority strategies and Figure 9 show a boxplot of the mean active agents after the simulations. The first interesting result is that priority methods maintain a higher number of agents interested in the market than a mechanism without priorities. In the non-priority case, the chart shows how the number of participants drastically drops from 50 to 20.09, with all the adaptative agents left out (in this case all the adaptative bidders and even some cheater bidders are forced to leave the auction due to the greedy behavior of the most powerful cheater agents). This descent is minimized when priorities are included in the auction mechanism. If we compare this result with an ideal case where all the bidders act honestly and no bidder leaves the market, it means that only the 40% of the bidders are maintained within the auction.

In Figure 8 it can be also observed that between the twentieth time cycle and the 70th, using any of the proposed priority methods, there is a reduction on the participation. This fact is produced by the 25 agents which start underbidding in order to collapse the market and create an oligopoly. However, the influence of these bidders is lower depending on the kind of priority method used.

The chart shows that the BBWOC (clear blue), the LS (black) and the WOC (grey) are the priority methods which best preserve the number of active bidders inside the auction (between 34 and 35 active bidders, which corresponds to the 68% and 70% of bidders of the ideal scenario). In these

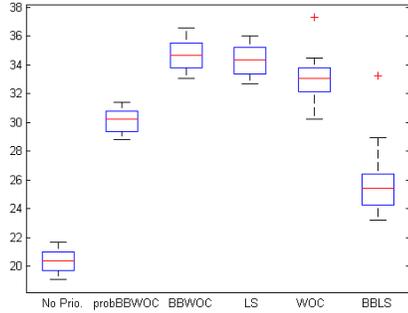


Figure 9: Box plot showing the number of active bidders after the simulations when no priority is used and when using different methods of calculating the priority.

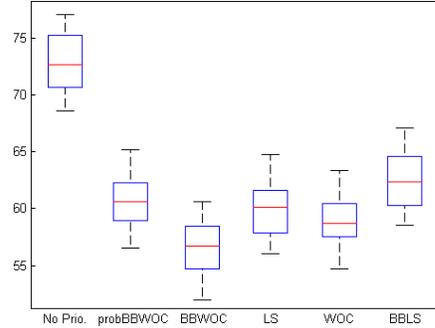


Figure 10: Box plot showing the mean cost of a workflow allocation after the simulations when no priority is used and when using different methods of calculating the priority.

cases, priority has increased the chances of winning of the adaptive agents, keeping their interest in the auctions and avoiding the creation of an oligopoly by the side of the underbidding agents.

The probBBWOC priority calculation method and the BBLs priority do not appear to be as effective to preserve the number of active bidders: 30(60%) and 25 (50%) respectively. In the first case this fact is probably produced by the lower influence of the priority. In the second, the lower number of active bidders is caused by the way priorities are assigned, based on the bid fitness. As underbidding agents start offering *better* bids, in spite that they do not win, their bids have high fitness and, therefore, they obtain higher priorities than adaptative agents. In consequence, the chances of many adaptive agents to win an auction during the first part of the simulations decrease, and they leave the auction after the given losing streak.

Regarding the asymmetric balance of negotiation power, it can be measured through the value of the mean cost of a task allocation. Figure 10 clearly points that the allocation costs are reduced when priorities are used. When no priority is used the mean allocation costs rises to an average of 72.65 per task; however, when a priority method is used the mean cost is significantly reduced up to 56.44 (a 22.65% cheaper). Taking into account this boxplot and the one in Figure 9 we can say that, under certain circumstances, the use of our priority methods can prevent and difficult the bidder agents to fix the market price.

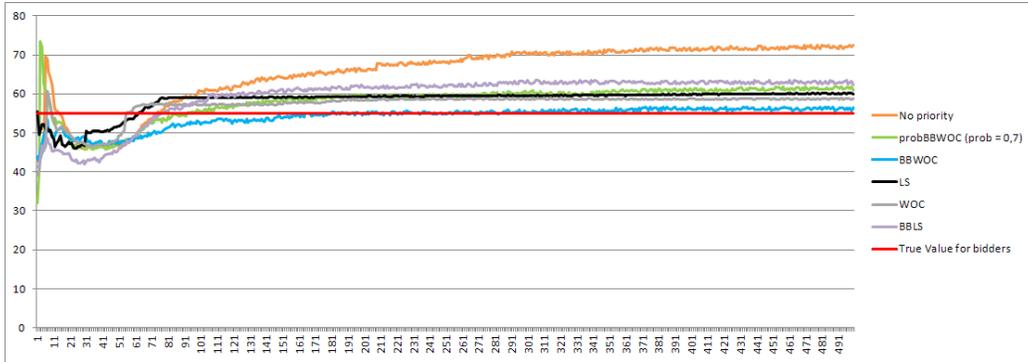


Figure 11: Mean task cost after a simulation where half of the agents try to manipulate the auction working under their production cost and where the other half adopt an adaptative strategy.

To deepen on the results regarding the asymmetric balance of negotiation power, Figure 11 shows the evolution of the mean task costs during the simulation. The red line indicates the mean task cost which a task would have had in the ideal case where all the bidders of the mechanism had bid truthfully (55 per task). The first remarkable fact which can be observed in the chart is that, independently from the priority method used, there is an initial drop of prices. This descent is caused by the underbidding agents which pretend to exclude the rest of bidders. After this initial gap it can be seen how the mean cost of the allocations increases progressively. However, the roof of this increase is different for each priority used. On the one hand, when no priority is used, as there are no adaptative bidders remaining in the market, the underbidding bidders can arise the price as much as they want (72.65, a 33.09% higher than the true values) because they have more power to fix the market price. On the other hand, this increment is lower when priorities are used. The BBWOC is the type of priority which obtains the mean cost per task most similar to the one fixed by bidder's true value (56.93 against 55, only a 3.51% more expensive), very close to the ideal situation. The rest of priority methods obtain higher mean price per task (between 57.97 and 62.48) being the WOC the second lowest mean price and the BBLs method the worst. It is important to notice that the two methods which obtained the best performance are both based on the won auction coefficient. The supremacy of WOC versus LS methods is due to the fact that losing streak methods favours agents which have not won in

the last auctions, in consequence, when the underbidding bidders have a big losing streak they may benefit of the priority. Conversely, the auction won coefficient takes into account all the victories agents have obtained along time. Thus, underbidding agents are ballasted since they have won almost all the auctions at the beginning of the simulation, having fewer chances to win the rest of the auctions.

In summary, this experiment has shown that the use of priorities can reduce the effects of the bidder drop problem, minimizing the effects of the unbalance power of negotiation between auctioneers and bidders. The fact the number of bidders leaving the market is reduced also minimizes the asymmetric balance of negotiation bower between bidders and auctioneers as having a high number of active bidders prevents the formation of an oligopoly which fixes the price of the market (as happens when no priority is used).

5.2.4. Summary

Our experiments pointed that in multi-attribute auctions priority must be computed taking into account all the attributes involved in the auction, otherwise, the quality of the allocation may be compromised. In our approach, being the winner or the loser of an auction means to take into account all of the attributes involved in the auction. This was clearly observed when a uni-attribute priority was used: unlike the multidimensional priority methods which obtained fair allocations without increasing the number of delays, the unidimensional priority obtained slightly fairer allocations but jeopardizing the quality of the allocations since the number of delays has increased 6 times.

On the other hand results show that improving the fairness of an allocation can alter the wealth distribution among agents and may annoy the richest agents. The stochastic methods proposed in Section 4.4 can soften the this alteration without compromising to much the fairness of the allocation. The experiments suggest that the best compromise between a fair allocation and wealth rank preservation is achieved when using the probBBWOC with a *up* of 0.7.

Finally, experiments showed that the use of a priority within the auction process can reduce the bidder drop problem and prevent the creation of oligopolies which can fix the market prices. This fact combined with the results of the first experiments make the use of multi-dimensional priorities an interesting tool for those mechanism designers which want to take into

account the equity of allocations in the long run in multi-attribute auction domains.

Note that, if we assume that agents cannot change their identity, using this kind of priorities an agent can enter the market and leave it whenever it wants since the measures are calculated using all the auctions which each auctioneer have summoned. In other words, if an agent leaves the market and, after a while, reenters to it, the agent will maintain the priority it had when it left since it will have the same track of won and participated auctions. Additionally, an agent which just entered the market is treated in the same way that the rest of the participants, the only particularity of this case is that its participated and won auction record will be empty (this fact will have more or less relevance according to the method used to calculate the priority). Nevertheless, future research should take into account dynamic scenarios.

6. Conclusions

Fairness mechanisms for multi-attribute resource allocation problems, from a multi-dimensional perspective, need to take into account all the attributes which affect the allocation since focusing fairness only on the revenues of the participants may compromise the quality of the rest of the attributes involved in the decision process (e.g. delays in deliveries, alterations in energy consumptions, etc.).

In this paper we propose a multidimensional fairness approach for auction-based multi attribute resource allocation which faces the problem in a more holonic way, taking into account priorities along all the auction process (in the winner determination problem but also in the payment mechanism steps). The proposed approach assigns higher priorities to the bidders which are susceptible to leave the auction market due to their bad results in previous auctions, in this way, agents which have a high probability of leaving the market have higher chances of winning an auction, thereby obtaining an incentive to remain in the market.

We presented a collection of eight methods for computing the priority of the agents based on the ratio of auctions they win and lose or their losing streaks. Two of the methods are designed under a quantitative schema (absolute values regarding auctions won or lost) whilst two others are designed under a qualitative schema (giving higher priorities to those bidders which offered good bids but lost the auction); furthermore, an stochastic version

of these four methods is also provided. The stochastic approaches proposed allow to tune the consequences of the priority methods regarding the wealth ranking disorder.

The performed experiments, based on real data extracted from a manufacturing company, point that the inclusion of multi-dimensional priorities improves the fairness of the resulting allocations (measured using the Gini's index) whilst not compromising the rest of attributes involved in the allocation. The best results have been obtained by the methods which compute the priorities using a ratio of won and participated auction both in the quantitative (WOC) and qualitative (BBWOC) versions. The allocations of WOC and BBWOC were slightly less fair in terms of revenue than when using a uni-dimensional fairness method, however the multi-dimensional fairness priorities clearly outperformed the results of uni-dimensional priorities when analyzing the rest of attributes. Experiments also suggested that wealthiest agents may feel uncomfortable with the obtained allocations due to an hegemony loss. This problem is minimized when using the stochastic priority calculation methods as it allows to adjust the influence of priorities, achieving a fair distribution without altering so much the hegemony of richest agents. Finally, the results also showed that the use of priorities can reduce the bidder drop problem and showed that preserving a high number of bidders prevents the strongest agents to fix the market price.

Further research should concern problems that may arise due to use of false identities by bidders, a common problem within fairness literature [16]. Moreover, despite in this work we have studied some issues related to underbidding agents, further research should be conducted regarding incentive compatibility in the long run.

Acknowledgements

This research project has been partially funded through the BR10/18 Scholarship granted to the first author of the paper.

- [1] Che, Y.-K. (1993). Design competition through multidimensional auctions. *The RAND Journal of Economics*, 24(4):668–680.
- [2] Cheng, K., Zhang, H., and Zhang, R. (2013). A task-resource allocation method based on effectiveness. *Knowledge-Based Systems*, 37(0):196 – 202.
- [3] Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lematre, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-aguilar, J., and Sousa, P. (2006). Issues in multiagent resource allocation. *Informatica*, 30:2006.

- [4] Conitzer, V. (2010). Algorithms and theory of computation handbook. chapter Auction protocols, pages 16–16. Chapman & Hall/CRC.
- [5] Diaconis, P. and Graham, R. L. (1977). Spearman’s footrule as a measure of disarray. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 262–268.
- [6] Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259.
- [7] Endriss, U., Maudet, N., Sadri, F., and Toni, F. (2003). Resource allocation in egalitarian agent societies. In Herzig, A., Chaib-draa, B., and Mathieu, P., editors, *Secondes Journées Francophones sur les Modèles Formels d’Interaction (MFI-2003)*, pages 101–110. Cépaduès-Éditions.
- [8] Gastwirth, J. L. (1972). The estimation of the lorenz curve and gini index. *The Review of Economics and Statistics*, 54(3):306–16.
- [9] Hermoso, R., Billhardt, H., and Ossowski, S. (2013). Trust-based role coordination in task-oriented multiagent systems. *Knowledge-Based Systems*, 52(0):78 – 90.
- [10] Krishna, V. (2002). *Auction Theory*. Academic Press.
- [11] Lee, J. and Szymanski, B. (2005). A novel auction mechanism for selling time-sensitive e-services. In *IEEE Conference on ECommerce Technology (CEC’05)*, pages 75–82. Press.
- [12] Lematre, M., Verfaillie, G., Fargier, H., Lang, J., Bataille, N., and m. Lachiver, J. (2003). Equitable allocation of earth observing satellites resources. In *In Proc of 5th ONERA-DLR Aerospace Symposium (ODAS03)*.
- [13] Luo, X. and Jennings, N. R. (2007). A spectrum of compromise aggregation operators for multi-attribute decision making. *Artificial Intelligence*, 171(23):161 – 184.
- [14] MacKie-Mason, J. K. and Varian, H. R. (1994). Generalized vickrey auctions.
- [15] Mañoso, V., Labrador, F. J., and Fernández-Alba, A. (2004). Tipo de distorsiones cognitivas durante el juego en jugadores patológicos y no jugadores (in spanish, cognitive distortion types during gambling in pathological players and unusual players). *Psicothema*, 16(4):576–581.

- [16] Murillo, J., Lpez, B., Muoz, V., and Busquets, D. (2012). Fairness in recurrent auctions with competing markets and supply fluctuations. *Computational Intelligence*, 28(1):24–50.
- [17] Murillo, J., Munoz, V., Lopez, B., and Busquets, D. (2008). A fair mechanism for recurrent multi-unit auctions. In Bergmann, R., Lindemann, G., Kirn, S., and Pchouek, M., editors, *Multiagent System Technologies*, volume 5244 of *Lecture Notes in Computer Science*, pages 147–158. Springer Berlin Heidelberg.
- [18] Murillo Espinar, J. (2010). *Egalitarian behaviour in multi unit combinatorial auctions*. PhD thesis, Universitat de Girona.
- [19] Parkes, D. C. and Kalagnanam, J. (2005). Iterative multiattribute vickrey auctions. *Management Science*, 51:435451.
- [20] Parsons, S., Rodriguez-Aguilar, J. A., and Klein, M. (2011). Auctions and bidding: A guide for computer scientists. *ACM Comput. Surv.*, 43(2):10:1–10:59.
- [21] Pla, A., Lopez, B., and Murillo, J. (2012). Multi criteria operators for multi-attribute auctions. In Torra, V., Narukawa, Y., Lpez, B., and Villaret, M., editors, *Modeling Decisions for Artificial Intelligence*, Lecture Notes in Computer Science, pages 318–328. Springer Berlin Heidelberg.
- [22] Pla, A., Lopez, B., Murillo, J., and Maudet, N. (2013). Multi-attribute auctions with different types of attributes: Enacting properties in multi-attribute auctions. *Expert Systems with Applications*, page submitted.
- [23] University of Girona (2013). Production factory data set.
- [24] Varian, H. R. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6):1163 – 1178.
- [25] Vickrey, W. (1961). Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37.