# Market-Based Allocation with Indivisible Bids*

**L. Julian Schvartzman** and **Michael P. Wellman**

University of Michigan
Artificial Intelligence Laboratory
Ann Arbor, MI 48109-2110 USA
{lschvart, wellman}@umich.edu

**Abstract.** We study multi-unit double auctions accepting bids with indivisibility constraints. We propose different price-quote policies and study their influence on the efficiency of market-based allocation. Using a reconfigurable manufacturing scenario where agents trade large quantities of multiple goods, we demonstrate potential benefits of supporting indivisibility constraints in bidding. These benefits are highly sensitive to the form of price quote provided, indicating interesting tradeoffs in communication and allocation efficiency.

## 1 Introduction

Consider a scenario with $N$ manufacturing facilities with capabilities to produce various industrial parts. The facilities are controlled by different agents (e.g, firms, or profit-center divisions within the same large firm), and may vary in capacity, fixed and variable costs for producing the different part types, time for reconfiguring to switch between parts, transportation costs, and perhaps other factors. Each facility also has a set of customer orders, each representing a promise to pay a fixed amount contingent on delivery of a specified quantity of a particular type of part in the current period.

Since the facilities face heterogeneous cost structures, they stand to achieve significant gains in efficiency by exchanging orders among themselves. We can formulate the order allocation problem as a global optimization, but of course the agents may not have the appropriate incentives to reveal their private information about costs and orders, or comply with the resulting order exchanges. Economic mechanisms such as combinatorial auctions [1] can address these incentives problems, and provide an elegant solution when in fact they can be instituted. However, there are several organizational and computational impediments to holding large-scale (measured in numbers of goods and agents, and units per good) two-sided combinatorial auctions, and these are as yet uncommon in practice. It is substantially simpler to deploy individual two-sided multi-unit auctions for each of several goods, and these more ad hoc markets can address the allocation problem to a useful degree. For example, idealized models of such configurations as general-equilibrium systems demonstrate the potential of computational markets to achieve efficient allocations in convex, competitive environments [2, 3].

* This work was supported by the NSF Engineering Research Center for Reconfigurable Manufacturing. An extended report is available from the authors.

However, realistic versions of this scenario vary from the idealization in several ways.[1] One particularly important characteristic of this application domain is nonconvexity in preferences and production technology, as manifest (for example) in fixed costs, reconfiguration switching costs, and preset order sizes. The most straightforward multi-unit auction mechanisms assume divisibility of offers: an agent willing to buy $q$ units at some price would also be willing to accept $q' \leq q$ units at that price. This assumption will not generally hold given nonconvex preferences and costs, and therefore agents with these characteristics may be hesitant to bid at all absent some condition that its offer be accepted in whole or not at all.

We investigate the design of multi-unit auctions accommodating such indivisibility constraints. Our focus is on how such auctions can be operated in a computationally efficient manner, and on the auctions' *price quote* policies for revealing information to agents to guide their bidding. We evaluate our designs experimentally, employing a version of the manufacturing scenario sketched above. Our main finding is that supporting indivisibility constraints can in fact improve the quality of global allocations achieved through trading, but that this improvement does depend pivotally on the form of the price quote. In the full version of this paper we also show how the computational costs of optimizing bid matching and producing meaningful quotes can be amortized over the auction's operation, calculated incrementally throughout the dynamic bidding process.

## 2 Auction Mechanisms

We consider two-sided auctions for multiple units of a single good. The auctions clear periodically at predefined intervals, and thus implement a *call market*. We distinguish two major versions of this auction, differing in their treatment of offer quantities. In the first (called "standard" for purposes of this paper), quantities appearing in bids are assumed divisible, and so the bidder effectively expresses a willingness to trade any amount up to the specified quantity at the associated unit price. In the second, offers are considered "all-or-none" (AON), and so agents explicitly specify the payment at which they would be willing to trade any acceptable discrete quantity. We refer to this version as the "AON" auction henceforth.

In both auctions, agents may submit *bid schedules*, indicating the prices offered to trade various quantities (with negative quantities indicating amounts offered to sell). The points on the schedule are exclusive (i.e., treated as "XOR" [4]), in that the resulting allocation will employ at most one of them. For divisible (standard) bids, the prices are per unit, and consistency requires that unit prices be nonincreasing in quantity. For indivisible (AON) bids, the prices represent total payments for the associated quantity, and these totals (not the per-unit payments) must be nondecreasing in quantity. Assuming only free disposal, with AON bids, agents can express arbitrary valuations for the good [5, 4]. Standard divisible bids can express only convex valuations.

Operation of the standard auction is relatively simple, as described, for example, by Wurman et al. [6]. Mechanisms resembling the AON auction have been described in the

[1] Our work does not address all important variations from the ideal. In particular, we maintain the assumption of competitiveness, modeling our agents essentially as price takers.

literature, and employed in practice.[2] For example, van Hoesel and Müller [7] consider the special case of combinatorial auctions where all goods are the same, and point out that optimal allocations can be found by dynamic programming. This corresponds to a one-sided, one-shot version of the AON auction. Kothari et al. [8] present a one-sided, one-shot auction that supports AON bidding in the form of a minimum trade quantity, but then assumes divisibility for quantities beyond this minimum. Several other authors have considered indivisibility constraints in multi-unit auctions [9, 10, 5], and have also identified the connection to knapsack methods for matching bids. We describe details of the allocation algorithm, as well as other AON auction policies, in the sections below.

### 2.1 Bids Matching Algorithm

As pointed out most explicitly by Kelly [5], optimal winner determination for single-good, two-sided, multi-unit auctions with indivisible XOR bids (i.e., our AON auctions) reduces to the *Multiple Choice Knapsack Problem* (MCKP) [10]. Here we present a formulation of MCKP specialized (slightly) to the auction setting.

Consider a set of $N$ agents, with agent $i$ submitting bid $B_i$. Each $B_i$ is comprised of $m_i$ *bid points*, $(p_{ij}, q_{ij})$, specifying a payment $p_{ij}$ offered to exchange quantity $q_{ij}$. Each bid includes a dummy point $(0,0)$. Offers to buy are expressed as positive payment-quantity pairs, and offers to sell as negative payment-quantity pairs. Because the standard MCKP requires positive coefficients, we define transformed bid points $(p'_{ij}, q'_{ij}) = (p_{ij} + \bar{p_i}, q_{ij} + \bar{q_i})$, where $\bar{p_i} \equiv -\min_{j \in B_i} p_{ij}$, and $\bar{q_i} \equiv -\min_{j \in B_i} q_{ij}$. (Note that this transformation affects only bids with sell points; for buy-only bids, $\bar{q_i} = \bar{p_i} = 0$.) We then define the knapsack *capacity* $c \equiv \sum_i \bar{q_i}$. Conceptually, the capacity $c$ is equivalent to the total number of units that are offered for sale. We denote by $C$ the maximum number of units that could be traded. In order to ensure that $C$ is bounded, we assume agents have a limited ability to take short positions in the goods traded. The MCKP is formulated as:

$$\begin{aligned} \text{maximize: } & \sum_{i=1}^{N} \sum_{j=1}^{m_i} p'_{ij} x_{ij} \\ \text{subject to: } & \sum_{i=1}^{N} \sum_{j=1}^{m_i} q'_{ij} x_{ij} \leq c, \\ & \sum_{j=1}^{m_i} x_{ij} = 1, i \in A, x_{ij} \in \{0,1\}. \end{aligned}$$

We assume free disposal of units, reflected in allowing the auction to match bids with more sales than purchases. Excess units are allocated arbitrarily among sellers. Note that formulating and implementing the same problem without the assumption of free disposal is straightforward.

[2] Our understanding is that practical trading mechanisms admitting AON bids typically handle them in an ad hoc manner. For example, such bids might be matched in a greedy manner, as in common electronic stock trading systems, which just pass over AON bids if the entire quantity cannot be fulfilled.

Solving MCKP is NP-hard, which is shown by reduction to a basic knapsack problem [10] (p. 318). Using dynamic programming, however, the problem can be solved in pseudopolynomial time [11]. Let $Z_l(d)$ be the value of the optimal solution to the MCKP defined to include only the first $l$ agents, $1 \leq l \leq N$, and with restricted capacity $0 \leq d \leq c$. We further define $Z_0(d) = 0$ for $0 \leq d \leq c$, and $Z_l(d) = -\infty$ for $d < 0$.

We can characterize $Z_l(d)$, $1 \leq l \leq N$, $0 \leq d \leq c$, using the following recursion:

$$Z_l(d) = \max_{1 \leq j \leq m_l} Z_{l-1}(d - q'_{lj}) + p'_{lj}. \tag{1}$$

The optimal solution is obtained when $l = N$ and $d = c$. Given $N$ bids of maximum size $m = \max_i m_i$, the running time to solve MCKP using dynamic programming is $O(mNc)$ [10]. If agents choose to submit full demand curves (i.e., with up to $C$ bundles each), the running time becomes $O(NC^2)$.

Many different methods exist to solve MCKP, including a number of branch-and-bound techniques and hybrid algorithms (see Kellerer et al. [10] for an extensive review). Our implementation is customized for the dynamic auction context, which may call for repeated solution of the MCKP for small changes in the set of bids. We therefore developed an incremental version of the clearing algorithm, designed to minimize the average solution time over a sequence of auction operations. Space constraints preclude a description of the algorithm here; see the full paper for details.

### 2.2 Clearing and Pricing

*Clearing* the auction is the process of identifying the subset of bids that match and produce the highest possible surplus. The outcome of a clear operation is to determine the deals resulting from this matching, and removing the matched bids from the order book. Given our incremental algorithm, most of the work is performed when bids are inserted into the order book. Once this is done, identifying the match takes constant time. Extracting the deals takes time linear in the number $N'$ of bids matched. Modifying the order book to include only unmatched bids requires $N'$ deletions or $(N - N')$ insertions.

Matching bids with indivisibility constraints and three or more agents requires non-uniform pricing [12]. It is easy to understand the impossibility of having uniform pricing through an example. Consider the bids in Table 1. In such example, the auction would match all four bids. However, it is impossible to make all buyers pay at most what they offered while making all sellers receive at least what they requested with a unique price.

**Table 1.** Matching these bids requires non-uniform prices.

| | |
|---|---|
| (13 @ 26) | buy 13 units at $2 each |
| (2 @ 1) | buy 2 units at $0.5 each |
| (-11 @ -11) | sell 11 units at $1 each |
| (-4 @ -10) | sell 4 units at $2.5 each |

There are many ways to determine non-uniform prices consistent with a given set of bids and an allocation. For a fixed allocation of goods, monetary transfers do not affect overall efficiency. Therefore, since we are not addressing strategic issues in this work, the choice is not pivotal for our experimental analysis. Nevertheless, to fully specify the mechanism one must identify a pricing rule. Ours starts with Vickrey prices and adjusts them proportionally to ensure budget balance.[3] The Vickrey calculation requires that we compute the total surplus with each agent's bid excluded, for which $O(N')$ deletion and insertion operations need to be performed.

### 2.3 Quoting

After each bid, the auction issues a *price quote*, providing to the agents some information regarding the state of the order book, intended as a guide to future bidding. In the standard auction, the quote comprises a BID-ASK pair, representing the prices at which an agent could successfully trade at least one unit. The BID quote defines the price at which an agent could sell one unit, and the ASK quote the corresponding price to buy. For standard (divisible-bid) auctions, we can incrementally maintain the order book so that price quotes can be provided in constant time once the bids are inserted [6].

For the AON auction, it is not immediately apparent how the auction should define its price quotes. We identified four candidate quoting policies, described here and compared experimentally in Section 4 below.

**Standard Quote** One possibility is for the AON auction to provide a "standard" quote, defined as the BID-ASK pair reflecting the order book interpreted *as if the bids were divisible*. Constructing this interpretation requires some care, since simply treating each bid point as a divisible offer may violate the standard auction's consistency condition requiring that quantity be nonincreasing in unit price. To ensure this monotonicity, we transform each bid $B_i$ by first sorting the bid points $(p_{ij}, q_{ij})$ (not including the dummy point with $p_{ij} = q_{ij} = 0$) in decreasing order of unit price. We then traverse the list, translating each to a unit-price bid point, skipping any that would violate the monotonicity condition with respect to those already seen. These translated bids can then be handled by the order book and quoting algorithm of the standard auction.

**Marginal Unit Quote** A second quote candidate attempts to maintain the interpretation of the standard quote as a price threshold sufficient to trade one unit, but respecting the indivisibility constraints of AON bids. Calculating this quote requires solving the MCKP for the bids in the order book. Under this interpretation, the ASK quote is always defined as long as there is any sell offer in the order book. The same is not true for the BID quote, however, because it could be the case that no existing offer or combination of offers can be satisfied by contributing a single additional unit. The marginal unit quote takes the same form as the standard quote, but provides more conservative values.

[3] Intuitively, this should tend toward reducing the incentive to behave strategically, though we know this is not strictly true. Of course, no budget-balanced mechanism can achieve this exactly. A serious treatment of strategic concerns in this context is a subject of future work.

Indeed, it is even possible (and consistent) that the ASK price quoted be lower than the BID price, something that cannot happen in the divisible case. It would also be possible to define this quote with any particular quantity defined as "marginal" (e.g., ten units instead of one). Given our incremental computation (detailed in the full version of the paper), these quotes can be extracted from the order book in constant time.

**Anonymous Full Schedule Quote** The third quote we consider provides to all agents a full schedule of payments that would be required to exchange any feasible quantity given the current state of the order book. This can be viewed as a collection of marginal unit quotes, one for each feasible quantity. The quote is anonymous because the same values are provided to every agent. Note that only relevant payment-quantity pairs need to be communicated to an agent: for a given payment, a quote for the minimum number of units the agent needs to sell to get such payment, and the maximum number of units the agent can buy with such payment. As for the marginal unit quote, the schedule may not be monotone: the unit price to exchange various quantities may be increasing or decreasing or mixed along the schedule.

Also like the marginal unit quote, the full schedule quote can be extracted directly from the order book given our incremental computation scheme, though of course extracting and communicating it will take time proportional to its size, $O(C)$.

**Non-Anonymous Full Schedule Quote** The final quote we consider is similar to the previous one, but each agent is provided with personalized values based on its existing bid. More specifically, this quote provides agent $i$ the schedule of payments calculated by excluding from the order book the bid sent by $i$.

**Quote Discussion** The four candidate quotes present distinct tradeoffs. The standard and marginal-unit quotes are compact, but the more accurate one may be excessively conservative when used as a guide for quantities greater than a single unit. The full schedule quotes provide high-fidelity information, but may be too large to be reasonably communicated in some applications.

We explore the implications of the various quote policies in our experiments below. (Results for non-anonymous quotes are pending from ongoing experiments, not reported here.) Of course, the worth of a quote is intimately tied to how the agents use this information in their bidding. We discuss our assumptions about agent behavior in Section 4.2 below.

### 2.4 Implementation

We implemented the AON auction as an extension of AB3D, a configurable auction and market-game server developed at the University of Michigan [13]. AB3D provides a flexible bid-processing architecture, with a rule-based scripting language to specify particular auction policies and temporal control structure. The standard call market was already supported by AB3D. To handle indivisible bidding, we added a new bid language specifying quantity-payment schedules, and new matching, pricing, and quoting modules implementing the algorithms and policies described above. In particular, we implemented all four quoting candidates as selectable options. Parameters in the auction script determine whether to allow AON bids, and if so, which of the available quoting and pricing policies to employ.

## 3 Manufacturing Domain

We evaluate the AON auction in a market-based allocation problem based on the manufacturing scenario sketched in the Introduction. The setting comprises a set of $N$ manufacturing *modules*, defined as arrangements of manufacturing machines, with capabilities for producing a variety of parts. Each module is controlled by an agent, whose objective is to maximize profit by operating the module to produce parts fulfilling customer orders over an $L$-day production period. In our market-based model, agents may increase their individual and collective profit by exchanging orders among themselves, thus exploiting their comparative advantages and configuration decisions.

We provide a full specification of the model below, describing the goods traded, utility and cost functions of the manufacturing modules, and the market configuration. Specific parameter settings for the model, and trading policies implemented by agents in our simulations, are described in Section 4.

### 3.1 Goods traded

The core allocation problem in this domain is deciding which manufacturer will produce what quantity of each of $M$ types of parts in the current period. The total quantity demanded of part type $r$ is $D_r$, and initially each agent is given orders for some share of that demand. Producing part $r$ entitles the manufacturer to a fixed income of $I_r$ per unit, up to the number of units for which it holds orders.

The purpose of the market is to enable trading of orders among manufacturing modules. The goods traded are the rights to produce parts for orders. A unit of good $r$, therefore, entitles the holder to produce a unit of the corresponding part and receive the corresponding payment $I_r$ from the customer.

Note that the parameter $D_r$ bounds the maximum quantity of good $r$ that can be exchanged at one time, and thus plays the role of $C$ in the definition of the AON auction.

### 3.2 Agent objectives

Agents aim to maximize profit, defined as

$$\text{income} - \text{production costs} + \text{trading cash flow}.$$

Income is simply the total payment for producing parts. Trading cash flow represents the balance of payments from trading orders with other agents. Production costs include several components, depending on the quantity and types of parts produced. These are defined by a set of agent-specific parameters:

– $FC_i$: Fixed cost, a one-time payment if module $i$ produces one or more parts.

– $LC_i$: Labor cost, paid for every day in which the module is in production.
– $VC_{i,r}$: Variable cost, be paid for each unit of part $r$ that gets produced.
– $CF_i$: Set of possible *configurations*. Each manufacturing configuration provides distinct production capabilities. Only one configuration can be used in any given day. For each configuration $f \in CF_i$, each module has:
  - $PC_{f,r}$: Production capacity per part type, the quantity of parts of type $r$ that the module produces per day.
  - $RC_f$: Reconfiguration cost to be paid if the configuration is used.
  - $RT_f$: Reconfiguration time (in days) that it takes to set up configuration $f$, during which no part can be produced.

The configuration capacities and times, along with the period length $L$, define the production possibilities for module $i$. The various cost parameters define the total cost for any feasible production plan.

Although complicated, the foregoing determines well-defined optimization problems for the agent:

– Determining an optimal production plan given holdings of goods $r$.
– Determining optimal demand for goods $r$ given current holdings and market prices.

### 3.3 Market Configuration

The overall market system comprises the agents representing manufacturing modules, plus one auction for each part type. We simulate an instance of this setup by generating parameter values from prespecified probability distributions, and communicating these values to the respective agents. Each agent is initially allocated customer orders corresponding to equal shares, $D_r/N$, of the overall demand for each part $r$.

The simulations are implemented using our configurable market game server, AB3D [13]. Each game instance lasts twenty minutes, with each auction clearing periodically every 48 seconds. The auctions are staggered, so that the initial clears occur at multiples of $48/M seconds$.

The agents operate asynchronously, submitting bids to the auctions according to the policy described in Section 4.2. Agents can request price quotes reflecting the latest auction state, and retrieve notices of any transactions from prior bids.

At the end of a game instance, the server calculates final holdings based on cumulative transactions, and determines a score for each agent. The score depends on an agent's production plan given its total available orders, which entails solving an optimization problem for each agent. AB3D solves these using a commercial optimization package (AMPL/CPLEX), given an integer linear programming (ILP) formulation specified as part of the game description.

The overall value of the resulting allocation is simply the sum of the scores over the $N$ agents. For comparison, we can also calculate (offline if necessary) the global optimum of the system without trading, assuming a central planner that can allocate orders across manufacturing modules.

## 4 Experiments

We ran a set of 58 paired trials with both standard and AON auctions. For AON auctions, we tested *standard*, *marginal*, and *full schedule* quotes. The following sections describe the specific problem instance we chose for our manufacturing scenario, the behavior of the agents, and the results obtained.

### 4.1 Manufacturing Problem Setting

For each of the 58 trials run, we obtained a new set of randomly chosen parameter values, as specified in Table 2. Each paired trial used the same set of parameter values, and compared standard auctions and AON auctions with the quoting alternatives discussed.

**Table 2.** Settings of the manufacturing scenario used for our experiments. Parameters specifying a range are drawn from a uniform distribution. (*) parameter specifies total for all parts in a configuration, each part getting a random proportion.

| Parameter | | | Values |
|---|---|---|---|
| General | # of agents ($N$) | | 4 |
| | # of parts ($M$) | | 4 |
| Public information | $I_r$ | | [1000, 2000] |
| | $D_r$ | | [2000, 6000] |
| | L | | [250, 300] |
| Private information for agent $i$ | $FC_i$ | | [300000, 400000] |
| | $LC_i$ | | [15000, 20000] |
| | $VC_i$ | | [250,350] |
| | $\|CF_i\|$ | | 2 |
| | For each $f \in CF_i$ | $PC_f$ | [20,60] (*) |
| | | $RC_f$ | [400000, 800000] |
| | | $RT_f$ | [5,15] |

### 4.2 Agent Bidding

Agents bid in a set of auctions $G$, each corresponding to a different good $r$. Each agent follows an incremental bidding approach similar to the one described by Cheng and Wellman [2]. The main loop that controls an agent's behavior is as follows.

**repeat**
    Get price quotes.
    Get transactions (i.e., matching bids).
    **for** each auction $g \in G$ **do**
        Select a new point to be added to the bid in $g$.
        Fix inconsistencies in bid.[4]

[4] Make smallest possible changes to the old points in the bid in order to maintain divisible prices nonincreasing in quantity and indivisible payments nondecreasing in quantity.

```
        Submit updated bid to g.
    end for
until Timeout {allocation process is over}
```

The results described in Section 4.4 were obtained by using the same agent structure, with some variations in terms of selection of new bidding points which are explained below.

### 4.3 Selection of New Bidding Points

In each iteration $a$ of the main loop, an agent updates its bid for good in auction $g$ with one new point $(p_{g,a}, q_{g,a})$, taking into account current holdings and assuming other goods (not in auction $g$) could be freely bought or sold at the most recent quote. We used two different methods for picking incremental points, one for dealing with divisible bids and another for indivisible ones.

DIVISIBLE: For divisible bids, an agent selects a new bidding point for the good in auction $g$ by picking a price $p_{g,a}$ and calculating the quantity $q_{g,a}$ the agent would be willing to buy or sell at such price in order to maximize its profit. Calculation is done using an ILP model that encodes the agent's utility function as explained in Section 3.2. Prices $p_{g,a}$ are selected in the following arbitrary order:

1. $p_{g,a} = \text{BID}$
2. $p_{g,a} = \text{ASK}$
3. If the bid in $g$ already contains prices for 1 and 2 above, $p_{g,a}$ is selected from a normal distribution $N(\mu, 1)$,

$$\mu = \begin{cases} (hb + \text{ASK})/2 & \text{if } q_{a-1} > 0 \vee (q_{a-1} = 0 \wedge pr < .25) \\ (ls + \text{BID})/2 & \text{if } q_{a-1} < 0 \vee (q_{a-1} = 0 \wedge pr < .5) \\ lb & \text{if } (q_{a-1} = 0 \wedge pr < .75) \\ hs & \text{otherwise} \end{cases}$$

   where $hb$ ($hs$) and $lb$ ($ls$) are the highest buy (sell) and lowest buy (sell) offers already in the bid and $pr$ is a random value uniformly distributed between 0 and 1.

(Note that BID and ASK refer to the most recent quote obtained by the agent.)

The basic idea behind the approach described above is to help agents find feasible trades by gradually making them place their highest buy and their lowest sell offer. We empirically tested other alternatives to ensure that our comparison of divisible versus indivisible bidding was not biased by an unreasonable point-selection approach. Specifically, we compared the procedure described with a random selection of points, and also with another in which prices are picked by finding the maximum possible gap between any two consecutive pairs of (sorted) prices already in the bid and selecting their average. Our results indicated that the approach chosen provided the best average performance among the alternatives we evaluated.

INDIVISIBLE: For indivisible bidding, the agent selects a new bidding point for the good in auction $g$ by picking a quantity $q_{g,a}$. The payment $p_{g,a}$ is given by the maximum (minimum) value at which the agent is willing to buy (sell) $q_{g,a}$ units, which is calculated using an ILP model that encodes the agent's utility function as explained in Section 3.2. Quantities $q_{g,a}$ are selected in the following order:

1. $q_{g,a} = -H_{g,a}$ (sell all holdings available in iteration $a$)
2. $q_{g,a} = D_g - H_{g,a}$ (buy all available items, i.e., demand minus holdings)
3. $q_{g,a} =$ random value uniformly distributed in the range $[-H_{g,a}, D_g - H_{g,a}]$ (excluding 0)
4. If the bid already contains quantities for 1, 2, and 3 above:

$$q_{g,a} = \begin{cases} -H_{g,a} & \text{with probability .1} \\ D_g - H_{g,a} & \text{with probability .1} \\ \text{average between any two} & \\ \text{consecutive (sorted) quantities} & \\ \text{that are further apart} & \text{with probability .8} \end{cases}$$

The method described gradually fills the largest gaps in the bid being constructed, and "refreshes" each extreme occasionally with a 0.1 probability.

### 4.4 Results

The average performance relative to a global optimal allocation (i.e., assuming a central planner) as calculated from our 58 trials is given in Table 3. Results show that AON auctions quoting an anonymous full schedule provided the best performance, and that AON auctions using either a standard or marginal unit quote performed worse than standard auctions with divisible bids. Using AON auctions with standard quotes provided the worst average performance, although the differences with AON auctions quoting marginal units are not statistically significant at reasonable levels.

**Table 3.** Results of 58 paired trials calculated as average performance in terms of global optimal. Differences between 1 and 2 are significant at the $10^{-7}$ level, 2 and 3 at the .03 level, but 3 and 4 only at the 0.07 level.

| # | Auction | Quote | Average performance |
|---|---|---|---|
| 1 | AON | Full schedule | 91.3 % |
| 2 | Standard | Standard | 79.2% |
| 3 | AON | Marginal unit | 70.7 % |
| 4 | AON | Standard | 61.7% |

We are not suggesting based on this particular experiment that the differences shown in Table 3 are an indicator of the differences to be found under any possible parameter configuration of our scenario or other settings. Before we ran the systematic paired tests described above, we informally experimented with other parameter settings. Even though we observed that AON auctions quoting full schedules always provided the best average performance, in several settings the differences detected were not as stark.

### 4.5 Influence of Quoting

Quoting marginal prices with standard auctions makes sense from two perspectives. First, it provides an accurate value for marginal units in order for agents to construct their bids. Second, it provides a lower (upper) bound on both the unit price and total payment to be paid (received) when bidding to buy (sell) an arbitrary number of units. On the contrary, marginal values for AON auctions do not contain similarly valuable information. By assuming divisibility with AON auctions, the marginal value provided is neither accurate for the marginal nor a bound on the price for additional units. In this case, the quote provides a very loose approximation of value. Similarly, if we took into consideration indivisibility constraints when quoting marginal values with AON auctions, the result would be very conservative. It is often undefined for the BID (we need an agent or combination of them intending to buy a single unit), and the ASK can often be excessively high (when the auction matches bids with sell quantities much larger than the marginal). Moreover, this quote provides no information about the unit price beyond the marginal, showing only that the total payment for more units will be at least the price for a single unit.

The effects described above were confirmed in part by measurements applied to our simulation results. We define a trade as *desired* (A) with respect to agent $i$, if, once executed, it increases or maintains the profit (i.e., income minus costs plus cashflow) of $i$ assuming that no further trades occur. We identified three possible reasons for agent $i$ to engage in undesired trading: *outdated information* (B1), *misleading non-anonymous quotes* (B2), or *misleading anonymous quotes* (B3). A bid can contain outdated information because its points were calculated incrementally or due to the asynchronous nature of the bidding process. Outdated information (B1) thus refers to the case in which $i$ engaged in a trade that it would have rejected had it reevaluated its bid using the most up-to-date information (i.e., current holdings and quotes). Misleading non-anonymous (B2) or anonymous (B3) quotes are those that made $i$ believe that it could buy or sell goods at the quote, when that was actually not possible. Non-anonymous quotes are those calculated by excluding from the order book the bid sent by agent $i$. Finally, every transaction that decreased utility and cannot be otherwise explained must have occurred because of a *dependency on other auctions* (C). Such dependencies exist because agents construct their bids in an auction assuming they could trade in other auctions at the quote. Since communication is asynchronous and auctions clear at different times, some intermediate decreases in utility are normal and expected.

Suppose we had $T$ transactions, and transaction $t$ occurred in auction $g$ for quantity $q_t$ and payment $p_t$. We perform two different optimizations for each agent $i$:

- $R^*(H)$ is the profit achieved by $i$ when calculating its optimal production plan based on holdings $H$, assuming that $i$ cannot trade further.
- $P^*(H,Q)$ is the highest payment that $i$ is willing to offer to trade quantity $q_t$ in auction $g$, assuming it holds goods $H$ and that it could freely trade goods in auctions other than $g$ at the prices given by quotes $Q$.

We further define $Q_i$ and $Q$ as the most up-to-date non-anonymous and anonymous quotes, respectively. Holdings $H^t$ are the goods held by $i$ in all auctions right after $t$ occurred; holdings $H^0$ are initial endowments; holdings $H^{t'_i}$ and $H^{t'}$ are the goods held right after $t$ and an hypothetical clear of all auctions other than $g$ occurred, assuming that $i$ bid to achieve optimal holdings as calculated for $P^*(H^{t-1}, Q_i)$ and $P^*(H^{t-1}, Q)$, respectively; and $p'_t$ is the lowest hypothetical payment that $i$ could have bid in order to trade $q_t$ in $g$. Negative coefficients for payments and quantities are used for sell offers. Given these definitions, we can classify transaction $t$ for agent $i$ as follows.

| | |
|---|---|
| • (A) Desired | If $R(H^{t-1}) \leq R(H^t)$ |
| • (B1) Occurred due to outdated information | If $P^*(H^{t-1}, Q_i) < p'_t$ |
| • (B2) Occurred due to a misleading non-anonymous quote | If $R(H^{t'_i}) < R(H^{t-1})$ |
| • (B3) Occurred due to a misleading anonymous quote | If $R(H^{t'}) < R(H^{t-1})$ |
| • (C) Necessary due to auction dependencies | If $t \notin A, B1, B2, B3$ |

Using the same data obtained for the experiments reported in Table 3, we measured the percentage of transactions in (A), (B1), (B2), (B3), or (C) for the different quoting mechanisms we tested with AON auctions. The results are shown in Table 4. Note that the Full Schedule quote was the least misleading. Had we used a non-anonymous full schedule quote for decision making, we would have entirely avoided (by definition) misleading quotes, potentially increasing the overall performance even further. The marginal unit and standard quotes were similarly misleading, providing comparable results regardless of quote anonymity. Thus, personalizing these two quotes for decision making does not seem likely to help. Marginal unit quotes appeared to provide a relatively high percentage of desired trades (but poor overall results), which is somewhat expected given such a highly conservative quote.

**Table 4.** Analysis of transactions in AON auctions. (*) Offline simulation only.

| # | Quoting | (A) | (B1) | (B2)* | (B3) | (C) |
|---|---|---|---|---|---|---|
| 1 | Full schedule | 62.1% | 20.7% | 0% | 17.8% | 16.1% |
| 2 | Marginal unit | 64.8% | 17.6% | 26.2% | 25.2% | 3.5% |
| 3 | Standard | 57.4% | 22.8% | 24.8% | 23.1% | 7.1% |

## 5 Discussion and Future Work

Our study of the AON auction provides evidence for the viability and potential benefits of accounting for indivisibility constraints in market-based allocation, without resorting to fully combinatorial auction designs. Whether one should adopt an AON auction or standard divisible auction depends on the specific setting. Relevant factors include:

1. Expressivity. Divisible bids allow expressing only convex valuations, whereas indivisible ones do not have such limitation. Would agents with nonconvex valuations refrain from participating in auctions with mandatory divisibility?
2. Undesired trades. If agents with nonconvex valuations do participate in a divisible auction, they risk loss-producing transactions. How much they suffer as a result depends on the degree of nonconvexity.

3. Computation. Our incremental algorithms provide a relatively efficient way to operate AON auctions, which should be fast enough for several practical applications. Standard auctions, however, are still faster to operate, and more predictable since performance is less dependent on the number of units offered for sale.
4. Quote communication. Our experiments showed that the level of detail in price-quote information can play an important role in overall efficiency, and in particular that simple marginal quotes were not enough for AON auctions to improve on the performance of a standard auction. Even though much work remains to be done in this area, it is obvious that the communication burden of the quote used should be evaluated when choosing an auction mechanism over the other.

Further work will refine our comparisons and evaluate additional quoting policies. For example, it would be interesting to measure the potential benefit of providing full schedule quotes in standard auctions, as we have for AON auctions. It would be particularly beneficial to identify intermediate quoting policies for AON auctions that provide much of the benefit of full schedule quotes without the full expense. Understanding this tradeoff remains an important goal. Finally, we are interested in exploring the strategic bidding issues posed by indivisibility constraints as well as alternative quoting policies.

## References


1. Cramton, P., Shoham, Y., Steinberg, R., eds.: Combinatorial Auctions. MIT Press (2005)
2. Cheng, J.Q., Wellman, M.P.: The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. Computational Economics **12** (1998) 1–24
3. Ygge, F., Akkermans, H.: Decentralized markets versus central control: A comparative study. Journal of Artificial Intelligence Research **11** (1999) 301–333
4. Nisan, N.: Bidding and allocation in combinatorial auctions. In: Second ACM Conference on Electronic Commerce, Minneapolis (2000) 1–12
5. Kelly, T.: Generalized knapsack solvers for multi-unit combinatorial auctions: Analysis and application to computational resource allocation. In: AAMAS-04 Workshop on Agent-Mediated Electronic Commerce (AMEC-VI), New York (2004)
6. Wurman, P.R., Walsh, W.E., Wellman, M.P.: Flexible double auctions for electronic commerce: Theory and implementation. Decision Support Systems **24** (1998) 17–27
7. van Hoesel, S., Müller, R.: Optimization in electronic markets: Examples in combinatorial auctions. Netnomics **3** (2001) 23–33
8. Kothari, A., Parkes, D.C., Suri, S.: Approximately-strategyproof and tractable multi-unit auctions. In: Fourth ACM Conference on Electronic Commerce, San Diego, CA (2003) 166–175
9. Kalagnanam, J.R., Davenport, A.J., Lee, H.S.: Computational aspects of clearing continuous call double auctions with assignment constraints and indivisible demand. Electronic Commerce Research **1** (2001) 221–238
10. Kellerer, H., Pferschy, U., Pisinger, D.: Knapsack Problems. Springer Verlag (2004)
11. Dudzinski, K., Walukiewicz, S.: Exact method for the knapsack problem and its generalizations. European Journal of Operations Research (1987) 28:3–21
12. Wurman, P.R., Wellman, M.P., Walsh, W.E.: A parametrization of the auction design space. Games and Economic Behavior **35** (2001) 304–338
13. Lochner, K.M., Wellman, M.P.: Rule-based specification of auction mechanisms. Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-04), New York (2004) 818–825