

Designing risk-averse bidding strategies in sequential auctions for transportation orders

Valentin Robu, Han La Poutré

Abstract Designing efficient bidding strategies for agents participating in multiple, sequential auctions remains an important challenge for researchers in agent-mediated electronic markets. The problem is particularly hard if the bidding agents have complementary (i.e. super-additive) utilities for the items being auctioned, such as is often the case in distributed transportation logistics. This paper studies the effect that a bidding agent's attitude towards taking risks plays in her optimal, decision-theoretic bidding strategy. We model the sequential bidding decision process as an MDP and we analyze, for a category of expectations of future price distributions, the effect that a bidder's risk aversion profile has on her decision-theoretic optimal bidding policy. Next, we simulate the above strategies, and we study the effect that an agent's risk aversion has on the chances of winning the desired items, as well as on the market efficiency and expected seller revenue. The paper extends the results presented in our previous work (reported in [19]), not only by providing additional details regarding the analytical part, but also by considering a more complex and realistic market setting for the simulations. This simulation setting is based on a real transportation logistics scenario [20], in which bidders have to choose between several combinations (bundles) of orders that can be contracted for transportation.

1 Introduction

Design of electronic auctions is considered an important open area of research in electronic commerce, both from a theoretical and an application perspective. There are two main approaches to this problem. One concerns the design of the auction mechanism itself, such as it guarantees certain properties, such as efficiency, individual rationality or budget balance. However, for some auction designs, such as

CWI, Dutch National Research Institute for Mathematics and Computer Science
Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands
{robv, hlp}@cwi.nl

simultaneous ascending, sequential and repeated auctions, this is not possible and research has focused on designing the bidding strategies of the agents participating in such auctions.

As previously shown in [18, 4, 8, 21], the main problem that a bidder has to face in a sequential (or simultaneous ascending) auction is the exposure problem. Informally stated, exposure means that an agent has to commit to buying an item (and thus take a “sunk” cost [18]), before she can be sure that she will be able to secure other items in her useful set or bundle (i.e. the set of items that gives him positive utility).

In order to deal with this problem, several strategies have been proposed in existing literature. Boutilier et al. '99 examines the role of dynamic programming in computing bidding policies in sequential auctions, based on distributions over estimated prices. Reeves et al. '03 [18] study the problem of bidding in simultaneous ascending auctions (a problem closely related to the sequential settings) - in the context of market-based scheduling. Osepayshvili et al. '05 [1] continue this line of research, but use probabilistic prediction methods of final prices and introduce the concept of self-confirming price distribution predictions. Gerding et al. '07 [6] derive the optimal bidding strategy for a global bidding agent that participates in multiple, simultaneous second-price auctions with perfect substitutes. Unlike this work, however, they do not consider complementarities (i.e. agents requiring bundles of items), and the setting is slightly different, as all auctions are assumed to close exactly at the same time, not sequentially.

In a direction of work that considers a setting very related to this paper, Greenwald & Boyan '04 [8] study the bidding problem, both in the context of sequential and simultaneously ascending auctions. For the sequential auctions case, they consider a decision-theoretic model and show that marginal utility bidding represents an optimal policy. Their result applies, however, only to risk neutral agents. Hoen et al. '05 [21] look at the related problem of bidding in repeated auctions with complementarities and draw a parallel with the N-person iterated prisoner's dilemma.

The above approaches have been shown to be efficient in many situations, both in self play and against a wide variety of other strategies, in competitions such as the TAC. Although most do implicitly consider the aspect of risk, they do not explicitly model the risk-taking *attitude* of the bidding agents. By “explicitly model” we mean building a profile of the agent's risk preferences towards uncertain, future outcomes (such as the final allocation of a sequential auction). In standard economic theory, since the seminal work of K. Arrow and J. Pratt, preferences towards risk are considered essential in understanding and modeling decision making under uncertainty [2, 7, 17, 12]. In fact, a body of auction theory from economics [17, 13] identify risk preferences as an important, open research area. Existing economic approaches to risk modeling do not, however, consider sequential auctions over combinations of items or propose bidding heuristics.

From the point of view of multi-agent systems literature, only a limited number of papers discuss risk profiles. Babanov et al. '04 [3] use the concept of certainty equivalence, similar to our work, in the context of optimal construction of schedules for task execution. Liu et al. [11] do consider risk-aversion on the part of the agents (similar to the approach taken in this paper) - but their work is mostly con-

cerned with providing an analytical solution to the one-shot auction case. Finally, Vytelingum et al '04 [5] consider risk-based bidding strategies in a double-auction setting. However, both the auction setting (i.e. CDA) and the risk model used (which is not based the standard Arrow-Pratt model) make this work rather different in focus from ours. Finally, Vetsikas and Jennings [23] also consider a model that includes agent attitudes towards risk (among other factors, such as budget constraints and reserve prices), for the case on multi-unit, sealed-bid auctions. They provide a thorough theoretical analysis of this case, but they do not consider complementarities (i.e. agents desiring bundles of goods), nor sequential allocation.

1.1 Goals and organisation of this paper

The basic goal of this paper is to study the relationship between a bidder agent's attitude towards risk (measured by the standard Arrow-Pratt risk aversion model - more specifically the CARA model) and her perceived best available bidding policy in a sequential auction (modeled by a Markov Decision Process). First, we investigate analytically how an agent's perception of her optimal bidding policy, given her probabilistic expectation of future prices, is affected by her risk aversion profile. Similar to [4, 8, 18, 21, 9], we take a decision-theoretic approach to the design of bidding agents, meaning agents reason w.r.t. the probability of future price distributions, and do not explicitly deliberate over the preferences, risk profiles and strategies of other bidders. Next, we conduct an experimental study of how an agent's attitude towards risk affects not only her sequential bidding and chances of winning a desired bundle, but also the a locative efficiency of the market she participates in and the auctioneer's revenue.

The remainder of the paper is organized as follows. Section 2 presents the risk aversion model, which forms the foundation of the following sections. Section 3 describes the bidding model and discusses the optimal bidding policies for both first and second-price sequential auctions. Section 4 provides the experimental results, while Section 5 concludes the paper with a discussion.

2 Modeling Utility Functions Under Risk

The literature on risk aversion identifies several 3 main types of agents w.r.t. their risk profiles: risk averse, risk neutral (indifferent) and risk proclave ("risk loving") agents. In the following we will focus our attention mostly on the risk averse and risk neutral cases, since these are the cases that describe the behavior of economic agents in most practical situations (c.f. [2, 13, 17]). Denote the payoff achieved by an agent participating in an auction as z . The utility a risk-averse agent assigns to this payoff is described by the Arrow-Pratt utility function:

$$u(z) = 1 - e^{-rz} \text{ for } r > 0 \quad (1)$$

Our choice of defining Eq. 1 represents a standard form of defining utility functions under uncertainty [17] (the same choice is made in [17, 14], among others). This form ensures that the following relation holds:

$$r_u(z) = -\frac{u''(z)}{u'(z)} \quad (2)$$

As defined in Eq. 2, $r_u(z)$ corresponds to the Arrow-Pratt measure of absolute risk aversion [2, 17]. In this paper, we consider r constant for each agent, i.e. $r_u(z) = r, \forall z$, thus we use the constant absolute risk aversion (CARA) model.¹ Factor r represents a constant which differs for each agent, characterizing her own preference towards risk-taking.

We use a state-based representation, in which all possible future outcomes at time t is denoted by S_t . All $s \in S_t$ are assigned by the agent a monetary payoff z_s and an expected probability p_s (where $p_s > 0$ and $\sum_{s \in S_t} p_s = 1$). We define the lottery L_t (which the agent faces at time t) as the set of payoff-probability pairs, i.e. $L_t = \{(z_s, p_s)\}$ where $s \in S_t$. In this form, the definition is generic, but as we show in Sect. 3, there is a natural correspondence between lotteries and states in a sequential-auction game.

The *expected utility* of the agent at time t over the lottery L_t is described by a von Neumann-Morgenstern utility function:

$$E_u[L_t] = \sum_{(z_i, p_i) \in L_t} p_i u(z_i) \quad (3)$$

In case all the agents are risk-neutral (i.e. have $u(z) = z$), it is easy to compare expected utilities across agents. However, for risk averse agents this is not the case, and we need a measure that enables comparison of payoffs across agents with different attitudes to risk in uncertain domains. The utility functions of the agents are not directly comparable in this setting, since each agent has a different attitude towards future risk (different r factor).

The widely used concept in risk modeling is to identify a monetary value (i.e. amount of money), such that the agent is indifferent between receiving this value with certainty or entering the lottery. This amount is called the **certainty equivalent (CE)** of the lottery. It can be seen as the monetary payoff the agent would attach to the future, if all the uncertainty (and hence risk) were discounted.

Formally defined, the **certainty equivalent (CE)** of a lottery L_t is defined as the certain payoff value which is equivalent to the expected utility of the lottery L_t . That is:

$$u(CE(L_t)) = E_u(L_t)$$

Expanding both sides using Eqs. 1 and 3 above, we have:

¹ This is a widely used risk aversion model, which we deemed sufficient for the purpose of this work. We leave the study of Relative Risk Aversion (RRA) models to future research.

$$-e^{-rCE(L_t)} = \sum_{(z_i, p_i) \in L_t} -p_i e^{-rz_i}$$

Hence the following expression can be derived for the certainty equivalent of the lottery:

$$CE(L_t) = \begin{cases} -\frac{1}{r} \ln \sum_{(z_i, p_i) \in L_t} p_i e^{-rz_i} & \text{for } r > 0 \\ \sum_{(z_i, p_i) \in L_t} p_i z_i & \text{for } r = 0 \end{cases} \quad (4)$$

In other words, the certainty equivalent can be seen as the certain amount of money which has the same utility to the agent as the equivalent lottery, before the outcome of the lottery is known. In the following, we define and prove a recursive property of CE functions, which is relevant for their application to sequential games considered in this paper.

Property 1: Suppose we have a game that occurs in stages t ; at each time step t the game can transition into either one of 2 states: X_t^+ (having an associated reward z_t^+) with probability p_t^+ , or X_t^- (having an associated reward z_t^-), where $p_t^+ + p_t^- = 1$. In the sequential auction case considered here, X_t^+ , respectively X_t^- represent the states in which the agent wins / does not win an upcoming auction (the formal link is made in Sect. 2). The following relation holds:

$$CE[(CE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)], p_t^+), (z_t^-, p_t^-)] = CE[(z_{t+1}^+, p_t^+ p_{t+1}^+), (z_{t+1}^-, p_t^+ p_{t+1}^-), (z_t^-, p_t^-)]$$

Proof: The proof involves repeated application of Eq. (4) to the left-side term:

$$\begin{aligned} CE[(CE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)], p_t^+), (z_t^-, p_t^-)] &= \\ &= -\frac{1}{r} \ln[p_t^+ e^{-rCE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)]} + p_t^- e^{-rz_t^-}] \\ &= -\frac{1}{r} \ln[p_t^+ e^{-r[-\frac{1}{r} \ln[p_{t+1}^+ e^{-rz_{t+1}^+} + p_{t+1}^- e^{-rz_{t+1}^-}]]} + p_t^- e^{-rz_t^-}] \end{aligned}$$

After reducing $-r(-\frac{1}{r})$ and using that $e^{\ln X} = X$, we get:

$$\begin{aligned} &= -\frac{1}{r} \ln[p_t^+ p_{t+1}^+ e^{-rz_{t+1}^+} + p_t^+ p_{t+1}^- e^{-rz_{t+1}^-} + p_t^- e^{-rz_t^-}] \\ &= CE[(z_{t+1}^+, p_t^+ p_{t+1}^+), (z_{t+1}^-, p_t^+ p_{t+1}^-), (z_t^-, p_t^-)] \text{ (q.e.d.)} \end{aligned}$$

Note that the above property can be applied recursively to games with any number of stages. This property, while apparently straightforward, is important since it shows that performing local CE optimization at each time step gives the same result as CE optimization for the entire game (a property which is not obvious for non-linear functions). As such, it is used as an implicit assumption in our MDP model.

2.1 The importance of risk adversity in decision making: an example

In the following, we give an illustration why risk aversion can have an important effect on monetary values. Consider the case of two complementary-valued items: A and B, which are sold sequentially. Suppose the agent has to accept a sunk cost of \$5 (dollars or any monetary units) for item A. If she acquires both A and B, she makes a profit of \$10, but if she doesn't, she makes a loss of -\$5 (thus potential profit is double the size of potential loss). Supposing the agent estimates the probability of acquiring B at p_B , how large does p_B have to be in order for the agent to accept the gamble?




Fig. 1 Example of the certainty equivalents of 3 agents with 3 different risk profiles for a lottery with 2 possible outcomes: -\$5 (non-desirable) and \$10 (desirable). The figure illustrates 2 cases: A. The desirable outcome is assigned a monetary value of \$10 by all agents. B. The optimal outcome is assigned a monetary value of \$5 (for the risk indifferent agent ($r \rightarrow 0$), \$7.5 by the slightly risk averse one ($r=0.15$) and \$10 by the strongly risk-averse agent ($r=0.3$)).

We plot the CE payoffs in this lottery for 3 risk attitudes of the agents, from $r \rightarrow 0$, $r = 0.15$ and $r = 0.3$. The left-hand side of Fig. 1 shows the case when all agents have the same evaluation for both the desirable (i.e. +\$10) and the non-desirable (-\$5) outcome. From this figure, one can already see that a risk neutral agent ($r = 0$) would “join in” this lottery or sequence of auctions, if the probability of winning (getting the desirable outcome) exceeds 33.3%. However, a relatively risk-averse agent ($r = 0.3$) would need to have at least 78% probability of winning in order for it to assign a positive CE value to this lottery (and thus have an incentive to participate in the game). In the right-hand side of Fig. 1, we keep the payoff of the non-desirable outcome constant at -\$5, but we vary the maximal payoff from \$5 (for the risk indifferent agent), to \$7.5 (for $r = 0.15$) and \$10 (for $r = 0.3$). Even if the estimated probability of acquiring the bundle $\{A, B\}$ is exactly the same for all

agents, the probability of winning has to be above 97% in order for the agent with the highest valuation among the 3 agents to also assign the sequence of auctions the highest CE value.

3 Bidding in sequential auctions with complementarities

As shown in the introduction, the main problem that a bidder has to face in a sequential auction with complementarities is the exposure problem. Following Boutilier et. al. [4] and Greenwald & Boyan [8], we model the decision problem that the bidder agent has to face in sequential auctions as a Markov Decision Process.

Assume there is a set of items I_t , sold in sequential auctions held at time points $t = 1..n$. A state in this game is specified by a set of goods X_t acquired up to time t (where $X_t \subseteq I$ for $t = 1..n$). The *bidding policy* of an agent in this game is described by a vector of bids $\mathbf{B} = (b_1, \dots, b_n)$, which assigns a bid b_t to each item sold at time point t . Fig. 2 illustrates this, for an auction with 2 items.

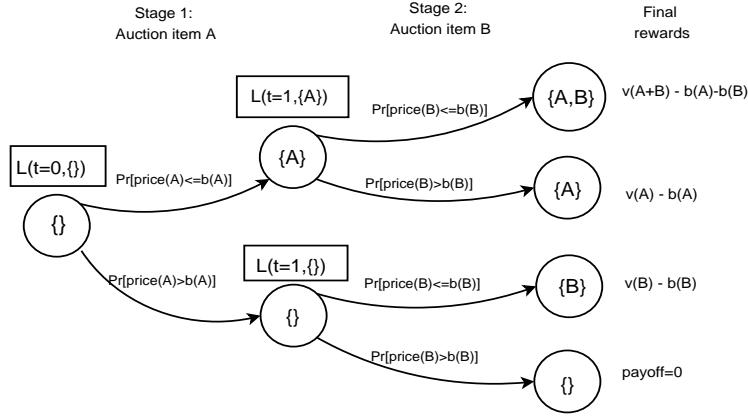


Fig. 2 The decision process faced by an agent in sequential auction, for a two stage example, with goods labeled A and B

The bidding agent maintains a probabilistic expectation of the closing prices for items I_1, \dots, I_n , in the form of n distributions. In the current model, these distributions are assumed independent of each other and stationary during one bidding round of n auctions (n could also be seen as the number of auctions the agent can stay in the game before its deadline). This definition of stationarity does not exclude the agent being able to learn, or refine its distributions of closing prices between episodes but, in this paper, we assume they are stationary for the duration of n auctions (i.e. one episode).

Considering the probabilistic distribution of future prices (a similar choice as in [4, 8, 1]) is more relevant to this setting than simply working with an vector of

the average past prices (such as in [18, 21]), since the thickness of the tails of the distribution may be of particular importance if the agents are risk averse. Note that in this form, we do not make any assumption on the type or shape of the expected future distributions: they can be normal, log-normal (usually used to model future prices in financial markets), uniform, binomial etc. For the results reported in this paper, we employed the normal distribution, but the generic approach can be applied to other distributions as well. The transition probabilities between different states are the cumulative distribution probabilities that the agent wins the lottery with its current bid b_t :

$$Prob(X_{t+1} = X_t \cup \{I_t\}) = Prob(ClosingPrice_t \leq b_t) = cdf_t(b_t)$$

We model the utility of a future outcome at each time step t (except the final one when all the goods have been allocated) as equivalent to a lottery $L_t(X_t, b_t)$. The payoffs of this lottery are determined by the agent's utility function, the set of items acquired so far X_t and bid b_t . The probabilities over outcomes depend on the bid b_t and expectation of future price distributions. The decision problem the agent faces, at each time point is to choose a bid b_t that provides the right balance between expected payoff and probability of winning, given her risk aversion r . This means choosing b_t which maximizes the certainty equivalent of lottery $CE(L_t(X_t, b_t))$. Using formal MDP notation, the value at each state is:

$$Q(X_t, b_t) = CE(L_t(X_t, b_t))$$

The optimal bidding policy and its reward being:

$$b_t^* = \pi(X_t) = \operatorname{argmax}_{b_t} Q(X_t, b_t)$$

$$V(X_t) = \max_{b_t} Q(X_t, b_t)$$

We can rewrite the above two equations as:

$$b_t^* = \operatorname{argmax}_{b_t} CE(L_t(X_t, b_t))$$

$$CE^*(L_t) = \max_{b_t} CE(L_t(X_t, b_t))$$

Due to the recursive property of the CE function (captured by Lemma 1 above), maximizing $CE(L_t)$ at each state leads to maximizing the initial certainty equivalent expectation for the entire sequence of auctions, i.e. maximizing $CE(L_0)$. An alternative to this method would be the straightforward application MDP optimization directly to the utility function of the agent (as done in [4] for risk neutral agents). For risk-averse agents, however, due to the non-linear nature of the utility functions, definitions of bidding policies in sequential auctions can only be defined in terms of the CE values of future states². This is done in the following Sections, which

² We stress that the term “optimal” used in this paper, should be interpreted as optimal w.r.t. the bidder's aversion to risk and estimation of future price distributions. This is not the same as dom-

also include a numerical example and an illustration that provides insight into the dynamics of the problem.

3.1 *Optimal bidding policy for sequential 2nd price (Vickrey) auctions*

Greenwald & Boyan [8] show that the optimal bidding strategy for a risk-neutral agent in a second-price sequential auction is to bid the difference between the expected value of the state when the auction is won and the expected value of the state when the auction is not won. Here we can extend these results to the risk-averse case as follows.

Suppose at time t (after a set of t previous auctions) the agent is in a state in which he has the set of items X_t . At the next step (i.e. after the auction occurring at t), he can transition in either one of two possible states: one in which he obtains the set of items $X_{t+1}^+ = X_t \cup \{I_{t+1}\}$ (if the auction is won) or $X_{t+1}^- = X_t$ (if the auction is not won). If the auction at time t is a second-price one, the optimal bidding policy available to the agent is:

$$b_t^* = CE(L_{t+1}(X_{t+1}^+)) - CE(L_{t+1}(X_{t+1}^-)) \quad (5)$$

assuming that at all subsequent steps $t+1, \dots, n$ the locally optimal bids are chosen.

This can be shown similarly as in Greenwald et al. [8] and results straightforwardly from the definition of the CE function. We do note, however, that although Formula 5 does give a closed-form optimal bidding policy, applying it in practice remains a computationally challenging problem: it requires the agent to compute the certainty equivalents of 2 whole subtrees, representing the remaining game starting from the two possible resulting states.

3.2 *Optimal bidding policy for sequential 1st price auctions. Numerical solutions*

For first-price auctions no closed form optimal bidding policy can be formulated, because agents have, as in the case of risk-neutral agents, an incentive to shade their bid. Liu et. al. '03 [11] show, for the case of single-shot first-price auctions that, on average, risk averse agents shade their bid less than risk neutral agents, since they want to minimize the chance of losing the auction. In this case, the optimal bid level b_t^* given in Eq. 5 above for the second price auction represents an upper bound on the bid level a rational agent would place in a first-price auction.

inant strategy from standard auction theory (i.e. independent of the behavior of other bidders), as dominant strategies are not known exist for this setting.




Fig. 3 Example of the certainty equivalent payoff in a two-stage sequential auction for 2 items: A (at time $t=1$) and B (at $t=2$). The graph shows the CE value of the corresponding 2-stage game, if the costs for both items are drawn from $N(\mu = 2.5, \sigma = 1.5)$, for an agent with $r=0.3$




Fig. 4 The optimal bidding policy available to an agent having risk aversion r , in a 2, respectively 3-stage sequential auction. The items have a complementarity value of \$10 (resp. \$15) if acquired together, but no value if acquired separately. The costs for all items are drawn from a normal distribution $N(\mu = 2.5, \sigma = 1.5)$.

For the sequential case, in order to get insight into the case, we computed the numerical solutions of the optimal bidding policy as perceived by the agents at time $t = 0$ (before entering the sequence of auctions). This is done for a sequence of 2, respectively 3 upcoming auctions (items are numbered alphabetically, by the order they are being auctioned). The analysis can be extended to any number of auctions, and the results are largely similar).

We take the expected distributions for future prices for individual items are drawn from identical, independent *normal* distributions (i.i.d.s are a choice widely used

in economic modeling [12]). In this case, we chose normal distributions with mean $\mu = 2.5$ and dispersion $\sigma = 1.5$. The chosen valuations levels are: $v_{\{A\}} = 0$, $v_{\{B\}} = 0$ and $v_{\{A,B\}} = 10$ (for the 2-stage auction), respectively $v_{\{A,B,C\}} = 15$ (for the 3-stage auction). This choice of values is such that the sum of the mean expectation of the costs is exactly half the bundle payoff.

A bidding policy is defined as a combination of bids for item b_A, b_B , with the note that the bid for B is only placed if the agent wins A in the preceding auction (otherwise, it has a dominant policy to bid 0 and earns a reward of 0). Using a mathematical optimization package (in our case Matlab), we computed the optimal bid levels of this game (b_A^*, b_B^*), for each level of risk aversion from 0 to 1, as well as the expected CE level of this optimal bidding policy, i.e. $\max_{b_A, b_B} CE(L_{t=0})$.

In Fig. 3, we illustrate the CE value of the initial choice to enter the set of auctions (i.e. $CE(L_{t=0})$) for one level of risk aversion r and all possible combinations of bids for the first, respectively second good in the sequence. As can be seen in Fig. 3, the surface of possible bids has a single optimum point, for each level of risk aversion. In Fig. 4 we plot the optimal bid levels for a sequence of 2, respectively 3 auctions. Basically, each point on the left (i.e. two-item) side of Fig. 4 corresponds to the coordinates of the optimum point in exactly one bidding surface, such as shown in Fig. 3. The same can be said about the right side (i.e. the 3 item case), although in this case the bidding surface cannot be actually visualized (being 4-dimensional).

From the analysis of Fig. 4, we can already highlight some important effects:

- The more averse a risk agent is, the higher she will bid for the second item in the sequence. Intuitively, a risk averse agent is more concerned with reducing as much as possible the probability she will lose the auction for B and not cover her sunk cost for item A. By contrast, a more risk-neutral agent is willing to accept a slightly higher probability she will have a sunk cost, if the potential gain is greater. Otherwise stated, agents with different risk profiles have different levels of awareness of costs already incurred.
- By contrast, the optimal bid level for item A slightly decreases as the agent becomes more risk averse. Risk averse agents are not willing to accept a high sunk cost - thus their optimal policy is to avoid bidding aggressively in the first round. They may prefer not to participate at all in the sequence of auctions, than to win the first auction with a high sunk cost, which would be difficult to cover. Furthermore, note that in this example, the average mean expectation of cost of the first item is only a quarter (\$2.5) of the maximal possible cost, but if this becomes higher, the effect is considerably more pronounced - and risk-averse agents' optimal bid policy may simply be not to participate at all in the auction sequence.

3.3 Heuristic addressing multiple copies

The MDP-based bidding strategy outlined above can lead to an optimal bidding policy, but only if all CE values of the states for the entire game are computed. This can become computationally expensive, especially if the sequence contains many

stages (auctions). In the simulations presented in Sect. 4 below, we have made an approximation that enable us to significantly prune the state tree in solving the multiple copies problem. This problem appears when the bidding agent is interested in only a limited number of items to form a useful bundle, but these are offered for sale repeatedly.³ Suppose items are divided into several types. The agent’s expectation of closing price distributions for all items of a given particular type is the same (thus she does not model the future expectation probability per auction or per item, but per type of item). If this expectation remains the same during the number of bidding rounds the agent stays in the game, then it is possible to reduce the state tree representation from a representation dependent on the number of future auctions to a representation which depends only on the size of the desired bundle.

Formally, if there are several items of type A and the agent knows that there are n_A more auctions of items of type A to take place. Then the probability of transition from any state X to a state $X \cup \{A\}$ (i.e. winning at least one item of type A at some point in the next sequence of n_A opportunities), given that the agents bids b_A in each of the auctions in that sequence is:

$$Prob(ClosingPrice_A \leq b_A) = 1 - [1 - cdf_A(b_A)]^{n_A}$$

The effect of having multiple future opportunities to buy a good may determine risk-neutral agents to reduce her bids (since there is a higher chance of winning one of them), but it may also encourage risk-averse bidders to join the game, bidders which would otherwise find a short sequence of auctions to be too risky. Next, we show the simulation results of these bidding policies, for populations of agents with different attitudes to risk.

4 Experimental analysis

There are several goals (or “research questions”) that could be followed in such a sequential auction model. Our goal is to highlight some of the complex interaction effects between the risk aversion of the agent, his/her valuations, his/her expectation of future closing prices, the number of expected future buying opportunities. We identify three main questions to be explored:

- The first of these is how does the risk profile of an agent affect her bidding and hence her chances of winning the desired (or target) bundles, against a population of agents with other risk profiles. We have studied this for different expectations the agents have about costs in future auctions and for different lengths of auction sequences (or number of auctions can stay in the game, before their deadline).
- We also investigated what happens in a larger market, where several different types of goods are sold. Especially, we look into how the order in which

³ Multiple copies can be seen as an instance of the substitutability problem - though substitutability is wider, if we allow for partial substitutes. These are not considered in the current work.

complementary-valued goods are sold affects the final allocation and the agents that manage to get their desired bundles.

- Finally, we look into how the presence of risk-averse bidders can affect the overall allocation efficiency of the market (i.e. “social welfare” of the allocation) and, as a related question, how it affects the seller’s (or sellers’) expected revenues.

4.1 Experimental analysis of distribution of risk profiles of winning bidders

To illustrate the aspect of winner risk profiles, consider a population of 10 bidding agents, each having a risk profile coefficient r ranging from 0 to 0.9 (in 0.1 increments). The agents bid in a sequence of n_A first-price auctions, all of which involve the sale of one type of item, A. Each agent needs exactly one bundle containing two items of A. The valuation the agents assign for each bundle of 2 A-s is the same for all agents - equal to \$10 and this also holds for their initial expectations of future prices, which are drawn from independent, identical normal distributions $N(\mu_A, \sigma_A)$. The rationale for having these parameters uniform across agents is that we wish to study the effect of risk aversion on bidding, by isolating it from the effect of higher valuations or different models of the future. Two sets of tests are illustrated as a bar plot in Fig. 5.⁴

In the first setting (left-hand side of Fig. 5 the mean expectation of closing prices per item μ_A is varied from \$1 to \$10, but the number of future auctions (or the number of auction the agents can stay in the game before their deadline) is kept constant: $n_A = 7$ steps. The spread $\sigma_A = \$2.5$, in all cases. Each vertical line represents the risk adversities of the winners in an auction sequence: the solid circle represents the first agent who manages to complete the desired bundle, triangle the second, a diamond shape the third, etc.

In the second test (the right hand side of Fig. 5, the mean expected cost is kept constant at $\mu_A = \$4$ (hence the agents have the expectation of closing prices for each auction: $N(\mu_A = 4, \sigma = 2.5)$). What is varied in this case is the number of future auctions in which item A is offered - from 1 (in which no agent can make her desired bundle, hence nobody bids) to 10 (in which case 5 agents can get their desired bundle). Two main effects are observed:

- If there is a high chance of getting the desired bundle in the sequence of auctions, then agents which are more risk averse tend to bid more aggressively, and hence win the bundle first. This is because, as shown with exact solutions in Sect. 3-4, agents that are more risk averse try to minimize their risk of not covering the sunk costs by more aggressive bidding.
- However, if the sequence of auctions is perceived as “too risky” (chances of acquiring both items are too low), risk averse agents simply drop out of bidding

⁴ Actually, many more tests have been conducted, but the results of this test configuration has been chosen to illustrate the results.

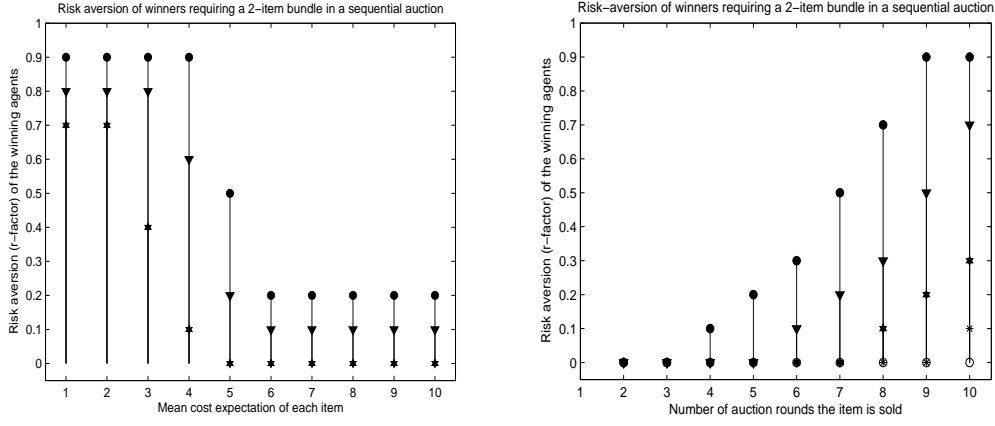


Fig. 5 Risk aversion of winners in a sequential auction. The risk profile coefficient r in the bidder population varies from 0 to 0.9, but each agent requires a bundle of exactly 2 items (thus requires to win 2 auctions). This bundle has a value of 10 for each agent. On the left-side graph the X-axis shows the mean cost expectation the agents have per item. On the right side, the X axis stands for number of auction rounds. In both cases, the Y-axis gives the r -coefficients of the agents who actually acquire the desired bundle.

(since their certainty equivalence value for the sequence is not high enough to justify taking the risk), leaving the more risk-neutral agents to acquire the goods. In the left-hand side of Fig. 5, there is a natural transition point, when the expected mean cost of good μ_A is \$5, which is exactly half the value of the desired bundle. For the right-hand side plot, the transition can be seen in almost each step: as the number of possible future opportunities to buy the item (i.e. no of future auctions) diminishes, risk-averse agents increasingly drop out of the game. When only two or three auctions are available ($n_A = 2, 3$), only risk-neutral agents ($r = 0$) are interested in bidding.

Basically, these results can be partially explained by looking into the analysis in Fig. 4. As shown there, the bid for the second item in a sequence is always higher than any of the bids for the first item, regardless of the risk aversion of the agents. This is because, when bidding on a second item to complete the bundle, agents have to cover an already made cost for the first item - otherwise, they would incur a loss.

The difference between which agent, i.e. with which risk aversion, actually gets the bundle is made by how much they would be willing to bid for the first item. If the sequence of auctions, taken as a whole, given the number of possible future buying opportunities and items, is promising enough, then the risk-averse agents will bid more. If there are several future opportunities to buy the second item, this helps reduce their risk of an expected loss. It is possible that the sequence of auctions (given the number of future opportunities) is simply not “attractive enough” for a risk-averse buyer, so his best option is to bid nothing, earning a sure zero profit (because the whole sequence of auctions has a certainty equivalence value below zero).

In the following Section, we see that this decision problem is further complicated if the agents actually have several alternative bundles they can choose from.

4.2 *Complex bundle preferences. Effect of auction orders*

While the above Section has already highlighted the complexity of bidding in sequential auction to get a bundle of two items of just one possible type of item. However, in most real-life scenarios, on top of the question of how to divide their bids between complementary items in a sequence, agents are confronted with several alternatives that they must choose from in bidding. In fact, the potential complexity of the space possible preferences is very large. In this Section, while we do not completely model the full potential complexity of preferences, we show that simply having a second type of good to choose from introduces a whole different dimension to the dynamics of decision-theoretic bidding in sequential auctions.

We should mention, that our choice for the valuation structure of the bundles, while simple, is motivated by a transportation logistics application setting and does capture much of the dynamics of that use case. Therefore, before we describe the experimental set-up and results, we motivate it by briefly describing how the choices made fit our application setting.

4.2.1 Bidding in sequential auctions for transportation orders

The problem setting we considered in our auction model is that of distributed transportation logistics with partial truck loads [22]. A real-life auction platform for this case, developed in collaboration with a large logistic company, is described in [20].

In the logistic setting we consider, transportation orders (either from one, but usually from different sellers/shippers) are usually sold at different points in time through spot market type mechanisms (usually auctions). The bidders for these loads are small transportation companies who try to acquire suitable set (bundle) of orders that would fit their trucks. In this model, we assume all orders are ready for pick-up or return delivery at one central transportation depot⁵. Fig. 6 shows just such a topology, with delivery point group into 2 main delivery regions).

Acquiring suitable combinations (bundles) of orders to fit the same trip with one truck is crucial for profitability in this setting. A truck acquiring, for example, an order for 1/2 truckload to be delivered to a certain region usually counts on acquiring another 1/2 truckload order from the same region, in order to make a profit. In this case, item types represent different delivery regions - each trucking company expecting different costs/profit structure per region, depending on its transportation network. Another possibility for bundling can be when orders represent symmetrical outgoing/return orders which originate in the same region.

⁵ This is a realistic assumption, as in many cases in many domains, especially if there is just one shipper, or several small shippers who aggregate their demand to one central distribution point.




Fig. 6 Example transportation scenario with one central depot D and two disjoint transportation regions: A and B.

In the utility model used in our auction simulations, we abstract the main characteristics of this setting. In this way, bidders can be considered as truck owners (i.e. carriers), the items are transportation orders, item types correspond to different delivery or pick-up regions. In practice, auctions for transportation orders are reverse auctions: the bidders that offer the lowest cost get the order. However, we choose to have a model with direct, first-price auctions, which is basically equivalent to it and is easier to compare with other models and in existing literature.

4.2.2 Effect of different auction orders

The sequential model we consider is as follows. The number of auction rounds is still fixed at 7, but there are two types of goods: A and B. In this setting, we introduce a differentiation between the items: items of type B are relatively “rarer” (they are sold only in 2 auctions out of the 7), while items of type A are sold in 5 auctions. Hence there are $C_7^2 = 21$ possible auction orders in each test set.

The risk-aversion coefficients of the bidding agents are again varied between 0 and 0.9. All agents require either a bundle containing exactly two A-type items or a bundle containing two B-type items. There are two valuation schemes for the bundles: a lower valuation scheme and a higher one. The averages of the values chosen are, for the lower scheme $v_L(A,A) = \$9 \text{ XOR } v_L(B,B) = \12 for and for the higher scheme: $v_H(A,A) = \$15 \text{ XOR } v_H(B,B) = \20 . Note that in both schemes, a bundle of the rarer item type B was assigned a higher value, on average.

In Fig. 7, all agents are assigned the higher valuation scheme, i.e. $v_H(A,A) = \$15 \text{ XOR } v_H(B,B) = \20 , and the same initial expectations about closing prices, described by probability distribution $N(\mu_A = 4, \sigma = 2.5)$, for both items A and B). There is, however, one important difference to the model where only a single type of item is sold: the order that items of different types are auctioned in plays a crucial role. Basically, for 7 auctions and 2 types of items there are $\binom{7}{2} = 21$ possible orders in which items can be sold. The table below Fig. 7 give the index of the auction orders. The dots on the vertical bars give the risk aversion of the agent that actually acquired a bundle of a certain type: from the 7 items sold, there are two bundles of type A (denoted by a circle) and just one possible bundle of type B




Fig. 7 Risk aversions of winners in a sequence of 7 auctions for 2 types of goods: A (more common - 5 auctions) and B (rarer - 2 auctions). Agents have different risk aversions, but all assign the same valuation for useful bundles: $V_H(2 * A) = \$15$, $v_H(2 * B) = \$20$. The Y-axis gives and result bars give the risk aversion coefficient r of the agents that manage to acquire one bundle of 2 items of type A (filled circle) or the bundle of type B (triangle). As a convention, $r = -0.1$ (i.e. dot below the axis) means that a bundle that type was left unsold. On the X axis, the order of the 7 auctions for 2 types of items is varied, among the 21 possible combinations. The index of auction orders is shown in the table below the figure.

(denoted by a triangle). The value of $r = -0.1$ (below the horizontal axis) is not a real value, it is a convention that a possible bundle of that item type remained unsold.

From examining Fig. 7, several conclusions can be drawn. First, note that the more risk averse agents do acquire a bundle (thus bid more aggressively for the first item). However, this is always an item of type A, a bundle of which has a lower valuation than a bundle of type B, but with 5 possible opportunities to choose from. This can be explained by observing that, in this case, unlike the experiments reported in Sect. 4.1 above, the decision facing each bidder is not simply between going into the bidding for the sequence of auctions, or deciding (based on the estimated certainty equivalent value), that the potential loss makes not participating at all the best strategy. The decision, in this case, is also between going for a bundle of type A, which has a lower expected profit, but there are 4 more auction opportunities to buy the second item after the first one, or a bundle of type B, but for which there is

only one other opportunity to complete the bundle. In this case, the more risk averse agents simply prefer the “safer”, but lower expected profit option.

Another interesting effect of the auction order in Fig. 7 is that, if both type B items are sold later in the sequence, then there is a higher chance that more risk averse agents will bid for them. This seems at first counter-intuitive, but it can be explained by seeing that, after a few auctions for an item of type A have passed, the relative attractiveness of an item of type B increases, by comparison to an item of type A. Most interesting, in the last 3 sequence types (where both items of type B are sold in the last 3 auctions), 2 items of type A remain unsold. We will explain this effect in the next subsection, after introducing other valuation schemes.

4.2.3 Effect of different valuation schemes. Unsold bundles

In experiments from this Section, we keep to the same settings as above, but first lower the valuations for the 2 bundles of items to $V_L(2 * A) = \$9$, $v_L(2 * B) = \$12$. Results are presented in Fig. 8.




Fig. 8 Risk aversions of winners in a sequence of 7 auctions for 2 types of goods: A (more common - 5 auctions) and B (rarer - 2 auctions). The settings are the same in Fig. 7, but all agents have lower valuations for the item bundles: $V_L(2 * A) = \$9$, $v_L(2 * B) = \$12$.

From this Figure, first we can see that all agents with risk aversions $r \geq 0.5$ are not interested in bidding in either of the two bundles (the risk/reward ratio being too great). However, there is another perhaps more surprising effect. For the first 10 (out of 21) auction orders - all the ones in which an item of type B is sold in either the first or the second auction in the sequence, the bundle of the higher-valued item B actually remains unsold. However, if the item of type B appear in the last 5 auctions, then this bundle of the rarer item does get sold, albeit to the risk neutral agent. The reason for this can be deduced by examining the position of the more risk-averse agents. In a sequence of 7 items, 2 of type B and 5 of type A, their best policy is to wait to try to bid for a bundle of type A (which is relatively “safe”). However, if only 5 auctions are left, 3 of type A and 2 of type B, the items of type B suddenly appear, by comparison, more attractive (as the bundle of type B has a higher value).

Therefore, even if it would be highly beneficial for a bundle of items to be allocated and even if the decision-theoretic policy of (some of) the agents is to bid for these items, the availability of less risky alternatives may distract the agents from bidding. The sequence for selling the items plays a crucial role in the dynamics of the bidding.




Fig. 9 Risk aversions of winners in a sequence of 7 auctions for 2 types of goods: A (more common - 5 auctions) and B (rarer - 2 auctions). Half of all agents (the more risk-averse half, those with $r \geq 0.5$) have valuations scheme $V_H(2*A) = \$15, v_H(2*B) = \20 , and the other half (with $r \leq 0.4$) have valuation scheme $V_L(2*A) = \$9, v_L(2*B) = \12 . The meaning of all other factors are the same as in Fig. 7

A similar, though not identical, effect is observed for the last 3 possible orders of auction sequences, for all settings in Figures 7, 8 and 9. There is again, a dependency of the choice between the two types of items, but this time in reverse. In case all the items of type B are among the last 3 in the sequence, agents prefer to wait to bid for a bundle of this type, ignoring the possibility of acquiring the last bundle of type A.

In Fig. 9, we examine what happens if half of the agents (the more risk-averse half) have the higher valuation scheme $V_H(2 * A) = \$15, v_H(2 * B) = \20 and half of the agents have valuation scheme $V_L(2 * A) = \$9, v_L(2 * B) = \12 (the less risk averse half). Results are in-between the two cases presented above. For the first 6 auction orders, i.e. when an item of type B is sold exactly in the first auctions, the bundle of 2 items of type B remains unsold, while for the last 3 auction orders (i.e. both items of type B sold in the last 3 auctions), one bundle of type A remains unsold. In the next Section we examine at the aggregate level (i.e. averaged over all auction sequences) what happens to market efficiency and auctioneer revenues for this setting.

4.3 Experimental analysis of market efficiency and auctioneer revenue

In order to study the effects of risk aversion on market efficiency, we looked at the case when only a (varying) proportion of the risk averse agents have the higher valuation scheme. Thus, in the results presented in Fig. 10, what is varied along the X-axis is the minimum risk aversion coefficient of the agents having the higher valuation scheme for the two bundles, i.e. the “border” between the v_L and v_H valuation schemes defined above. This means that at 0.3, for instance, agents with risk aversion at or above 0.3 value bundles using v_H values, while agents with risk aversion at 0.2 or less value bundles using v_L values. The simulation results are plotted in Fig. 10. For each point in the plot, all 21 possible auction orders were considered, for 5 different combinations of v_L and v_H values (randomly generated as discussed above). Thus, each point represents an average (with the corresponding dispersion) for over 100 simulated auction threads. The simulation results, averaged over all possible auction orders, for several value configurations are plotted in Fig. 10.

A first effect that can be observed from Fig. 10 is that, if only risk averse agents have a higher value for the bundles, overall market efficiency decreases.

This decrease in market efficiency can be explained by the dynamics explained in detail in Sect. 4.2.2 above. The first is that, in many of the possible orders, it is possible that some bundles, even for the more valuable good remain unsold. It is possible that agents bid too cautiously, preferring to bid for a bundle of the lower-valued good A (which is more common, being sold in 5/7 auctions), instead of a higher-valued bundle of good B (sold in 2/7 auctions). Their perceived optimal bidding policy is to try to minimize the probability they will not cover their sunk costs. This usually leaves the more risk-neutral agents to take the chances to acquire




Fig. 10 Allocation efficiency and seller revenue in a sequential auction with 7 rounds and 2 items types, each agent requiring (exclusively) two items of the same type. There are two valuation schemes for the bundles. The minimal risk aversion coefficient of the agents having the higher valuation scheme is varied on the X axis, while the Y axis shows the efficiency, respectively revenue.

the bundle of B-type items, but it also leads to a decrease in market efficiency. A similar effect is observed for auctioneer revenues.

5 Conclusions and further work

To summarize, the main contributions of our work are the following. First, we establish a formal link between bidding strategies in sequential auctions and standard (Arrow-Pratt) risk aversion models from economics. We derive a useful property of certainty equivalence functions and it shows how such functions can be naturally applied to sequential auction games. Next, we study the way in which the perceived optimal bidding strategy computed by a risk averse agent, given her probabilistic model of the future, differs from the optimal strategy of a risk neutral agent. Risk averse agents tend to bid more aggressively throughout the sequence of auctions, in order to cover their sunk costs for the initial items in the sequence. However, if the future sequence of auctions is initially perceived as too risky (given the agent's initial estimation of future closing prices), the best strategy available to a risk averse agent is simply not to participate at all. Finally, we show this behavior can have consequences not only in determining the winners in such sequences of auctions, but also on the allocation efficiency of the market and the revenue of the auctioneer(s). The reluctance of risk averse agents to join in risky auctions may reduce their own exposure, but has a distorting effect on the allocation mechanism.

The paper leaves several issues to be answered in further work. Some pertain to more theoretical analysis of the bidding strategies. Our current efforts are focused on

deriving closed form expressions for the best available bidding policies agents have, based on their attitude towards risk, in restricted categories of settings. From a more experimental perspective, new bidding heuristics could be developed for software agents that do not only target raw efficiency, but also allow their owners to select a balance between expected profit and risk, based on their personal preferences. Finally, the role of mechanisms such as options [10] in reducing or eliminating the exposure problem that risk-averse agents face is another promising direction for further work. We have begun investigating this possibility in another direction of work [16, 15].

References

1. Osepayshvili A., M.P. Wellman, D.M. Reeves, and J.K. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Proc. of the 21st Conf. on Uncertainty in AI (UAI-05)*, 2005.
2. K.J. Arrow. *Aspects of the Theory of Risk-Bearing*. Y. Hahnsson Foundation, Helsinki, 1965.
3. A. Babanov, J. Collins, and M. Gini. Harnessing the search for rational bid schedules with stochastic search. In *Proc. of AAMAS'04, New York, USA*, pages 355–368, 2004.
4. C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proc. of the 16th Int. Joint Conf. on AI (IJCAI-99), Stockholm*, pages 527–534, 1999.
5. Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, pages 40–47, 2003.
6. Enrico H. Gerding, Rajdeep K. Dash, David C.K. Yuen, and Nicholas R. Jennings. Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *Proc. of AAMAS'06, Honolulu, Hawaii*, pages 267–274, 2007.
7. W. H. Green. *Econometric Analysis*. Prentice Hall, 1993.
8. A. Greenwald and J. Boyan. Bidding under uncertainty: Theory and experiments. In *Proceedings of the 20th Conf. on Uncertainty in AI (UAI-04)*, pages 209–216, 2004.
9. A.X. Jiang and K. Leyton-Brown. Bidding agents for online auctions with hidden bids. *Machine Learning* (to appear), 2006.
10. Adam I. Juda and David C. Parkes. An options-based method to solve the composability problem in sequential auctions. In *Agent-Mediated Electronic Commerce VI*, pages 44–58. Springer, 2006.
11. Y. Liu, R. Goodwin, and S. Koenig. Risk-averse auction agents. In *Proc. of AAMAS'03, Melbourne, Australia*, pages 353–360, 2003.
12. A. Mas-Collel, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
13. P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
14. P. Milgrom and R.J. Weber. A theory of auctions and competitive bidding - part 2. *Economic Theory of Auctions*, 2000.
15. Lonneke Mous, Valentin Robu, and Han La Poutré. Can priced options solve the exposure problem in sequential auctions? In *ACM SIGEcom Exchanges 7(2)*. ACM Press, 2008.
16. Lonneke Mous, Valentin Robu, and Han La Poutré. Using priced options to solve the exposure problem in sequential auctions. In *Proc. of the 10th Int. Workshop on Agent-Mediated Electronic Commerce (AMEC'08), Estoril, Portugal*. Springer LNAI (to appear), 2008.
17. H.J. Paarsch and H. Hong. *An Introduction to the Structural Econometrics of Auction Data*. MIT Press, 2006.
18. D.M. Reeves, M.P. Wellman, J.K. MacKie-Mason, and A. Osepayshvili. Exploring bidding strategies for market based scheduling. *Decision Support Syst.*, 39:67–85, 2005.

19. Valentin Robu and Han La Poutré. Designing bidding strategies in sequential auctions for risk averse agents. In *Proc. of the 9th Int. Workshop on Agent-Mediated Electronic Commerce (AMEC'07), Honolulu, Hawai'i*. Springer LNBI (to appear), 2007.
20. Valentin Robu, Han Noot, Han La Poutré, and Willem-Jan van Schijndel. An agent platform for auction-based allocation of loads in transportation logistics. In *Proc. of AAMAS'08 (Industry Track), Estoril, Portugal*, pages 3–10. IFAAMAS Press, 2008.
21. P.J. 't Hoen and J.A. La Poutré. Repeated auctions with complementarities. In *Proc. of 7th Workshop on Agent-Mediated Electronic Commerce (AMEC'05), Utrecht*, 2005.
22. S. van der Putten, V. Robu, J.A. La Poutré, A. Jorritsma, and M. Gal. Automating supply chain negotiations using autonomous agents: a case study in transportation logistics. In *Proc. 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'06), Industry Track*, pages 1506–1513. ACM Press, 2006.
23. Ioannis A. Vetsikas and Nicholas R. Jennings. Towards agents participating in realistic multi-unit sealed-bid auctions. In *Proc. 7th Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS'08), Estoril, Portugal*, pages 1621–1624, 2008.