

# Designing bidding strategies in sequential auctions for risk averse agents

Valentin Robu<sup>a,\*</sup> and Han La Poutré<sup>b</sup>

<sup>a</sup>*University of Southampton, School of Electronics and Computer Science, Southampton, SO17 1BJ, UK*

<sup>b</sup>*CWI, Center for Mathematics and Computer Science, Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands*

**Abstract.** Designing efficient bidding strategies for sequential auctions represents an important, open problem area in agent-mediated electronic markets. In existing literature, a variety of bidding strategies have been proposed and have been shown to perform with varying degrees of efficiency. However, most of strategies proposed so far do not explicitly model bidders' attitudes towards risk which, in mainstream economic literature, is considered an essential attribute in modeling agent preferences and decision making under uncertainty. This paper studies the effect that risk profiles (modeled through the standard Arrow-Pratt risk aversion measure), have on the bidders' strategies in sequential auctions.

First, the sequential decision process involved in bidding is modeled as a Markov Decision Process. Then, the effect that a bidder's risk aversion has on her decision theoretic optimal bidding policy is analyzed, for a category of expectations of future price distributions. This analysis is performed separately for the case of first price and second-price sequential auctions. Next, the bidding strategies developed above are simulated, in order to study the effect that an agent's risk aversion has on the chances of winning a set of complementary-valued items. The paper concludes with an experimental study of how the presence of risk-averse bidders affects both bidder profits and auctioneer revenue, for different market scenarios of increasing complexity.

## 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. Such approaches usually rely on combinatorial auction mechanism, where a trusted center collects a set of bids from all the agents and computes the final allocation and payments [25].

However, many real-life markets are much more decentralized and dynamic. Different items (or tasks) may be auctioned by different sellers, in auctions with different closing times, and sometimes even in different markets. Often, the items to be allocated may not even be known in advance. For example, in transportation logistics, orders may be given throughout a day, and carriers are expected to put in offers for such orders when they appear, often in a type of reverse second-price auction. In such *sequential* auction settings, research has mostly focused on designing the bidding strategies of the agents participating in such auctions, such as to guarantee maximal expected revenue for their owners.

As previous shown in [4,9,23,27], 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

---

\*Corresponding author. E-mail: vr2@ecs.soton.ac.uk.

agent has to commit to buying an item (and thus take a “sunk” cost [23]), 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 her a positive utility). If she does not manage to acquire the other items, she is exposed to the risk of a loss. For example, if we consider a decentralized transportation problem, a truck acquiring an order for a part-truck load in an auction may rely on acquiring other orders in the future to fill the remainder of the truck’s capacity in order to make a profit.

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 [23] 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 [7] 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 and Boyan ’04 [9] 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 [27] 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 have been considered essential in understanding and modeling decision making under uncertainty [2,8,13,21]. In fact, a body of auction theory from economics [17,21] identifies risk preferences as a very important, open research area. In recent econometrics and financial economics literature, this has led to considerable research interest in efficiently modeling and eliciting risk aversion from human users [5,21,22]<sup>1</sup>. Existing economic approaches to risk modeling do not, however, consider sequential auctions over combinations of items, nor propose bidding heuristics for this setting.

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. [15] do consider risk-aversion on the part of the agents (similar to the approach taken in this paper) – but their work is mostly concerned with providing an analytical solution to the one-shot auction case. Vytelingum et al. ’04 [6] 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

---

<sup>1</sup>A practical example of risk elicitation in finance are the questionnaires involving probabilistic choices between several scenarios that investment fund managers send to potential investors.

different in focus from ours. Finally, Vetsikas and Jennings [28] 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 organization 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). In this context, we consider the bidding decisions of an agent that desires a bundle of complementary-valued goods that are sold through a sequence of auctions.

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 [4,9,10,23,27], 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 with complementarities attitude towards risk affects her chances of winning a desired bundle, when bidding against a population of local bidders, desiring only one good. Furthermore, we also look at how this bidding policy affects the auctioneer's revenue. Our primary goal is to gain a qualitative understanding of how sequential auction markets are influenced when bidders with complementary valuations participating in them are risk averse.

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 proclive (“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 behaviour of economic agents in most practical situations (c.f. [2,17,21]). Denote the private payoff  $z$  achieved by an agent participating in an auction or lottery. The utility a risk-averse agent assigns to this payoff is described by the Arrow-Pratt utility function:

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

For the case of risk indifference ( $r = 0$ ), we take  $u(z) = z$ .

Note that the auction model we consider in this paper is a *private value* model. The payoff  $z$  of a bidder after participating in an auction is a difference between a private value  $v$  and the amount of money paid to acquire the item in the auction (or cost)  $c$ . Since the private value  $v$  is private to each agent, the payoff value is also private. Therefore, in a risk averse setting there are parameters describing the private preferences of a bidder: value  $v$  and risk aversion coefficient  $r$ . Preferences of agents cannot be directly compared by comparing their private values or payoffs, as the risk factors  $r$  must also be taken into account.

Our choice of defining Eq. (1) represents a standard form of defining utility functions under uncertainty [21] (the same choice is made in [18,21], 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 [2,21]. 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.<sup>2</sup> 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  $L_t$  over a set of payoffs  $z_s$  (corresponding to the state  $S_t$ ) 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 Section 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 and payoffs 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:

$$u(CE(L_t)) = E_u(L_t) \quad (4)$$

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

$$-e^{-rCE(L_t)} = \sum_{(z_i, p_i) \in L_t} -p_i e^{-rz_i} \quad (5)$$

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

$$CE(L_t) = \begin{cases} -\frac{1}{r} \ln \sum_{(z_i, p_i) \in L_t} p_i e^{-rz_i} & \text{for } r > 0 \\ \sum_{(z_i, p_i) \in L_t} p_i z_i & \text{for } r = 0 \end{cases} \quad (6)$$

<sup>2</sup>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.

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 Section 2). The following relation holds:

$$\begin{aligned} CE[(CE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)], p_t^+), (z_t^-, p_t^-)] = \\ = CE[(z_{t+1}^+, p_t^+ p_{t+1}^+), (z_{t+1}^-, p_t^+ p_{t+1}^-), (z_t^-, p_t^-)] \end{aligned} \quad (7)$$

Note that the notations  $z^+$  and  $z^-$ , for each time  $t$  only relate to whether the agent wins or does not win the auction. These numbers can actually be negative, in case the payoff for a state is negative. For example, an agent with a strictly complementary valuation for items sold at times  $t$  and  $t + 1$ , that wins the item at time  $t$  but does not win the item at time  $t + 1$ , gets the payoff  $z_{t+1}^-$ , which is negative.

**Proof:** The proof involves repeated application of Eq. (6) to the left-side term:

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

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

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

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 aversion 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?

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

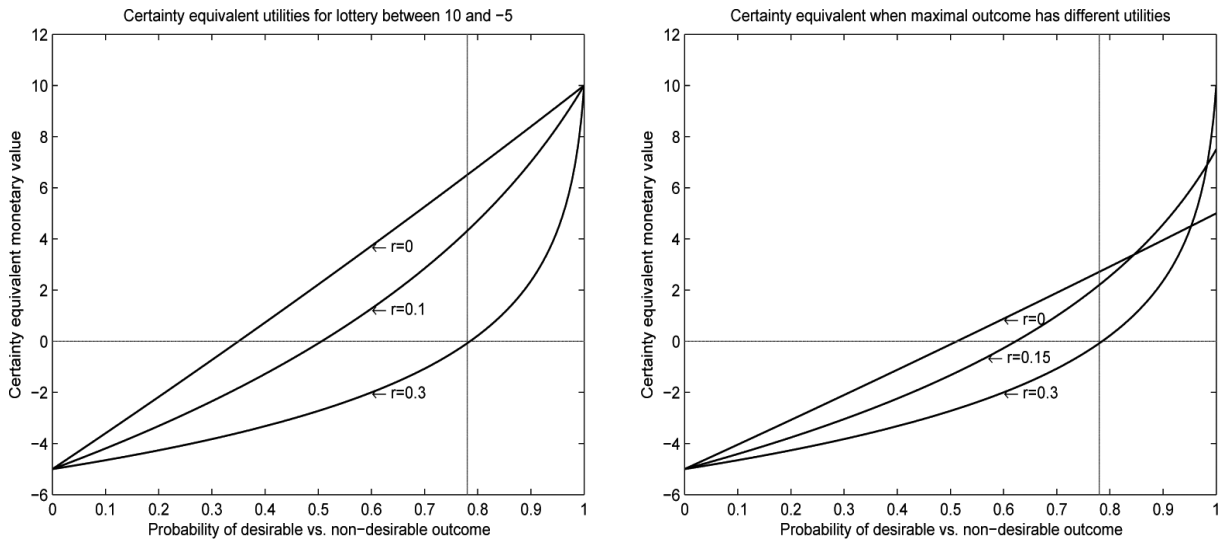


Fig. 1. Example of the certainty equivalents of 3 agents with 3 different risk profiles for a lottery with 2 possible outcomes:  $-\$5$  (non-desirable) and  $\$10$  (desirable). The figure illustrates 2 cases: A(left): The desirable outcome is assigned a monetary value of  $\$10$  by all agents. B(right):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

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

### 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 [9], 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  $\vec{B} = (b_1, \dots, b_n)$ , which assigns a bid  $b_t$  to each item sold at time point  $t$ . Fig. 2 illustrates this, for an auction with 2 items.

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

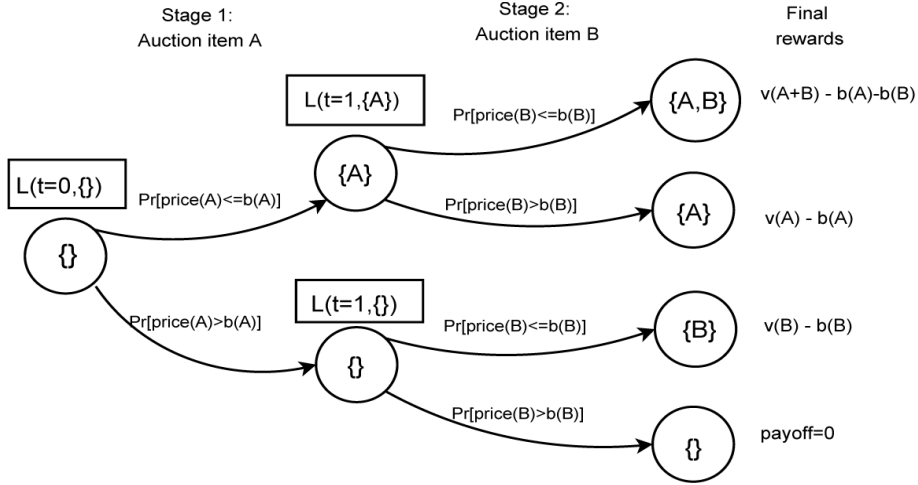


Fig. 2. The decision process faced by an agent in sequential auction, for a two stage example, with goods labeled A and B.

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 [1,4,9]) is more relevant to this setting than simply working with a vector of the average past prices (such as in [23,27]), 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  $b_t$ :

$$Prob(X_{t+1} = X_t \cup \{I_t\}) = Prob(ClosingPrice_t \leq b_t) = cdf_t(b_t) \tag{8}$$

where  $cdf_t(b_t)$  denotes the *cumulative density function* of the probability distribution over the closing prices, when bid  $b_t$  is placed.

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:

$$Q(X_t, b_t) = CE(L_t(X_t, b_t)) \tag{9}$$

The optimal bidding policy and the corresponding reward as:

$$b_t^* = \pi(X_t) = argmax_{b_t} Q(X_t, b_t) \tag{10}$$

$$V(X_t) = \max_{b_t} Q(X_t, b_t) \quad (11)$$

We can rewrite the above two equations, for the optimal bid at time  $t$   $b_t^*$  and the associated optimal certainty equivalent value  $CE^*$ , that can be obtained by taking the optimal bidding decision as:

$$b_t^* = \operatorname{argmax}_{b_t} CE(L_t(X_t, b_t)) \quad (12)$$

$$CE^*(L_t) = \max_{b_t} CE(L_t(X_t, b_t)) \quad (13)$$

Note that this optimization of the certainty equivalent value  $CE$  is performed for the current auction at time  $t$ , but assuming that the optimal bidding decisions are taken the whole sequence of future auctions, occurring at times  $t + 1, \dots, n$ . Therefore, the problem of determining the optimal bid  $b_t^*$  for time point  $t$  actually involves recursively determining the bids  $b_t^*, \dots, b_t^n$  that maximize the certainty equivalents of states at  $t_1, \dots, n$ . 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  $CE(L_0)$ . This means that standard MDP reasoning models can be applied to this problem, where the  $Q$  values of the standard MDP definition are the  $CE$  values of the lottery over future expectations at each step.

A naive alternative to this method would be the 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 states.<sup>3</sup> This is done in the following Sections, which 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 and Boyan [9] 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:

$$b_t^* = CE(L_{t+1}(X_{t+1}^+)) - CE(L_{t+1}(X_{t+1}^-)) \quad (14)$$

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

*Proof.* The proof resembles the proof in the textbook of Krishna [14], which refers, however, only to risk-neutral bidders. First, we simplify the notation by denoting the certainty equivalent of the state when the item is acquired by  $CE_{t+1}^+ = CE(L_{t+1}(X_{t+1}^+))$  and the certainty equivalent of the state when

---

<sup>3</sup>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 concept as dominant bidding strategy from standard auction theory (i.e. independent of the behaviour of other bidders). As discussed in the introduction, dominant strategies are not known to exist for the sequential settings considered in this paper.



the item is not acquired by  $CE_{t+1}^- = CE(L_{t+1}(X_{t+1}^-))$ . There is a set of  $n_a$  independent bidders in each auction, that only desire the item sold in that auction. All auctions being second price, they always have a dominant policy of bidding their true value. Let the valuation function over these  $n_a$  bidders be denoted by  $G(x) = G_i(x)^{n_a}$ , where  $G(x)$  is the cumulative distribution that bids of all  $n_a$  agents are smaller than  $x$  ( $G_i(x)$  here refers to a single independent bidder, but we can consider them in aggregate, without loss of precision). Then  $g(x)$  is the density function of this distribution, i.e. it denotes the probability that the highest bid of the independent bidders is exactly  $x$ .

Note that the state of winning the auction and having to pay  $x$  brings a monetary gain of  $CE_{t+1}^+ - x$  for the agent, while loosing brings a monetary gain of  $CE_{t+1}^-$ . In this case, however, the amount to be paid depends on the highest bid of independent bidders, so the standard certainty equivalence definition needs adjusting. Basically, the  $CE$  of bidding  $b_t$  in a state at time  $t$  can be expressed as:

$$CE(b_t) = -\frac{1}{r} \ln \left\{ \int_0^{b_t} g(x) e^{-r(CE_{t+1}^+ - x)} dx + (1 - G(b_t)) e^{-rCE_{t+1}^-} \right\}$$

The optimal bid  $b_t^*$  can be obtained by taking the derivative of the above expression, i.e. when  $\frac{dCE(b_t)}{db_t} = 0$ . This gives:

$$\frac{-g(b_t)}{r \int_0^{b_t} g(x) e^{-r(CE_{t+1}^+ - x)} dx + (1 - G(b_t)) e^{-rCE_{t+1}^-}} (e^{-r(CE_{t+1}^+ - b_t)} - e^{-rCE_{t+1}^-}) = 0$$

Of this expression, the first fraction is never zero and can be reduced, which basically gives:

$$e^{-r(CE_{t+1}^+ - b_t)} = e^{-rCE_{t+1}^-}$$

Which finally, after applying the logarithm and dividing by  $(-r)$  gives:

$$CE_{t+1}^+ - b_t = CE_{t+1}^-$$

Resulting in the final expression for  $b_t$  as:

$$b_t^* = CE_{t+1}^+ - CE_{t+1}^-$$

So basically, the marginal optimal bid  $b_t^*$  in a sequence of second price auctions is always the marginal difference between the certainty equivalents of the next two states. Note that this was known from standard auction theory for the case of risk-neutral bidders [14]. Basically, since a rational agent views all previous payments as sunk costs, they can be discounted and do not have to be accounted for in future bids. The intuitive reason why we find this result in the case of risk-averse bidders as well is that certainty equivalent functions, although not linear, are basically monotonically increasing in the monetary payoffs of future states, so it is rational for the agent to increase her bid until the difference  $CE_{t+1}^+ - CE_{t+1}^-$  is covered.

### 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 [15] 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




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  $r \rightarrow 0$  (left) and  $r = 0.3$  (right).

bid level  $b_t^*$  given in Eq. (14) above for the second price auction represents an upper bound on the bid level a rational agent would place in a first-price auction.

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 [16]). 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  $v_{\{A,B,C\}} = 15$  and 0 for all other subsets (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 show 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:

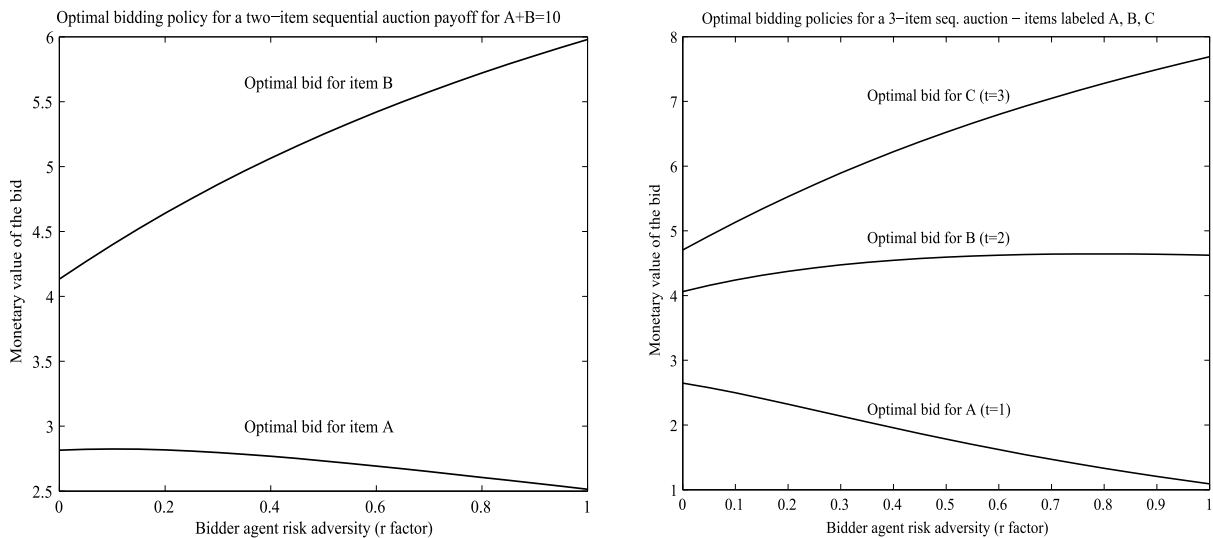


Fig. 4. The optimal bidding policy available to an agent having risk aversion  $r$ , in a 2, respectively 3-stage sequential auction. The items have a complementarity value of \$10 (resp. \$15) if acquired together, but no value if acquired separately. The costs for all items are drawn from a normal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

- The more averse a risk agent is, the higher she will bid for the second item in a 2-stage auction 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. We also performed tests with other mean expectation costs, and found that, if these costs become 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. Bidding strategy for multiple copy auction sequences

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 Section 4 below, we make 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.<sup>4</sup> Suppose items are

<sup>4</sup>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.

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 bundle the agent wishes to buy.

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:

$$\text{Prob}(\text{ClosingPrice}_A \leq b_A) = 1 - [1 - \text{cdf}_A(b_A)]^{n_A} \quad (15)$$

The above formula can be used to determine the probabilities of the getting an item of type A in the final state (i.e. after all auctions for a good of type A have closed). One still needs to apply the MDP to determine the best policy based on these probabilities, but this is straightforward, as it does not require computing the whole tree.

Note that this policy only uses as input the number of future auctions of each type remaining before the auction has to leave the market, not their exact order. In fact, if one knows the exact order that future auctions take place in, then it might be better to compute the whole tree (although that's exponentially more expensive). However, not knowing the exact order that future auctions will take place in, only the number of auctions of each type, is more realistic in many real-world settings. So, for example, in the simplified transportation case shown in Section 4.5 below, in practice, planners may know that a number of opportunities (i.e. transportation orders) to fill a truck may appear before the truck needs to start driving, but they don't know exactly the order in which these will be offered.

For the experiments reported in this paper, because goods are all of the same type (even if a bidder may desire only a bundle consisting of several such goods), this heuristic approximates very well the optimal bidding policy. In this case (i.e. same-type goods), there is basically only one possible sequence of future auctions, and the length of this sequence basically represents the full information needed to describe it. If there are several possible types of bundles, then the difference in performance may depend on the exact auction sequence. However, even this problem can be mitigated by randomizing over all possible auction sequences when performing experimental evaluation.

As we discussed in the numerical example, having multiple future opportunities to buy a good may determine risk-neutral agents to reduce their bids (since there is a higher chance of winning one of them), but it may also encourage risk-averse bidders to join the bidding, bidders which would otherwise find a short sequence of auctions to be too risky to participate.

#### 4. Experimental analysis

The goal of the experimental results presented in this paper is to test how different sequential auction market settings are influenced by the presence of a complementary valuation bidders, with different risk aversion levels. We look at how risk aversion influences the expected profit that the synergy bidder makes over a sequence of auctions, as well as the probability of completing the desired bundle and ending up with an incomplete bundle (which can result in a loss). Furthermore, we also study how the expected revenue of the seller(s) is influenced by the presence of a synergy bidder in a market, as well as how the

number of buying opportunities (i.e. length of the auction sequence) influences the expected profits of seller and buyer.

The first part of this paper studies these questions for a market with a single synergy bidder participates in a sequence of auctions for items of the same type. More concretely, we assume a market consisting of a sequence of auctions, each populated by a set of local (single-item) bidders and one synergy (global) bidder that desires exactly one bundle of two items. The number of auctions that the synergy bidder can stay in the market is fixed for each simulation round (although this parameter will be varied between different experiments). In the second part of our experimental study, we introduce bundle differentiation, i.e. the auction sequence consisting of auctions for two types of items and a synergy buyer that can choose between the two possible bundles. This setting was motivated by a transportation logistics setting described in Section 4.5.

#### 4.1. *Experimental hypotheses*

In order to better structure the presentation, we first formulate three hypotheses, that should be confirmed or disproved in the experimental tests. These hypotheses are intuitively formulated based on the properties observed in the theoretical part of this paper and should help the reader understand better the focus and choices made in the simulation model.

**Hypothesis 1:** A more risk averse agent will have a lower chance of ending the sequence of auctions with an incomplete bundle (i.e. a bundle in which the first item is acquired, but not the subsequent ones, hence resulting in a loss).

Note that the statement in Hypothesis 2 appears obvious: it is more a control hypothesis. If we do not find this, then there may be reasons to believe something is wrong in our experimental set-up. The most important side effect is stated as:

**Hypothesis 2:** A synergy bidder with a higher risk aversion will obtain a lower average profit from bidding in a sequence of auctions than a synergy bidder which is less risk averse.

The final hypothesis refers to the case of different auction lengths.

**Hypothesis 3:** For all risk aversion levels, the expected profit of a synergy buyer desiring a bundle of items will be higher if there are more auctions in the sequence (i.e. more opportunities to buy), while the chance of ending up with an incomplete bundle will be lower.

Besides these hypotheses, referring to the buyer, in our experiments we also look at the average revenue that a seller of a set of items sold in sequence will be lower if the synergy buyer present in the market is more risk averse. There are a further two further hypotheses, related to markets with different item types, but they will be introduced them later.

#### 4.2. *Experimental setup*

The experimental set-up used is as follows. We consider a sequence of  $n$  closed, first-price auctions, in all of which exactly one item of the same type  $A$  is sold. In each of these auctions, there are an (unspecified) number of local bidders, assumed myopic, that desire exactly one item of type  $A$ . Since these agents are assumed myopic (i.e. they only consider the current auction they participate in), we can

model their bids in each auction through some random distribution. Note that the myopicity assumption of local bidders is important here: if the bidders are able to strategize over the sequence of auctions, or over the presence of a synergy buyer in this sequence, then the model we use for their bidding behaviour may not hold. In this model, because in all auctions an identical good is sold, we can model the maximum bid received from the competition in each of the  $n$  auctions in the sequence through identical, independent probability distributions (i.i.d.) – a choice that is also made in other decision-theoretic bidding models, e.g. [7]. Since we do assume any prior information about the way independent bidders place their bids, we take the most general case and assume they follow normal distributions  $N(\mu, \sigma)$ .

In each sequence of auctions there is exactly one synergy (or global) bidder participating. This bidder desires exactly one bundle of two items of the same type  $A$  (and is assumed to have no value for an individual item). The synergy bidder must acquire this bundle in exactly  $n$  auctions (here, the number  $n$  of auctions that a bidder can stay in the game can also be thought of as a way to model a shorter or longer deadline that a bidder has). The value  $v(2 * A)$  that the synergy bidder assigns to the useful bundle, the number of auctions  $n$ , as well as the parameters  $\mu$  and  $\sigma$  that model the behaviour of single-item bidders are all parameters of the simulation. For each market configuration, average results are reported over 1000 runs.

#### 4.3. Experimental results for one-type item auctions

In this section we give the results for a market setting where only one type of item  $A$  is sold, and the synergy bidder desires exactly one bundle of two such items. Initially, we do this for sequences of  $n = 7$  auctions, where the values of independent bidders are drawn from a distribution  $N(\mu = 4, \sigma = 2)$ . The synergy bidder assigns a value of  $v(A, A) = 10$  (notice that this means the synergy buyer values a bundle of two item, on average, with 25% more than the independent bidders). The bids of the synergy buyer are computed according to the heuristic in Section 3.3, based on the risk aversion coefficient shown on the abscissa. Results (with averages over 1000 runs) are shown in Fig. 5. Returning to the hypotheses stated in Section 4.1, we see that indeed, Hypothesis 2 is confirmed for this setting: the higher the risk aversion of a synergy bidder, the lower his/her average expected profit (the drop is quite considerable – from 4 to around 2.5). However, one can also notice the variance of the results decreases slightly for the risk-averse bidder. These results are consistent with expectations.

The right side of Fig. 5 shows the average revenues of the seller for this setting. These are somewhat surprising, given that one would intuitively expect seller revenues to drop if bidders in the market are more risk averse. In fact, it seems there is some type of revenue equivalence, in the sense that the revenue that a seller can expect in such a sequence of sequential auctions with a risk averse bidder does not depend on his/her risk aversion. From further examination, this can be explained as follows: since these are first price auctions, a more risk averse bidder will bid less often in this sequence. However when he does bid, he will bid considerably more which, on average, has a compensating effect for the reduced participation. Nevertheless, a more in-depth investigation is needed (with a larger market and more synergy bidders), in order to formulate a hypothesis regarding this point.

#### 4.4. Results for one item and different auction lengths

In this Section, we extend the above analysis to a setting where we vary not only the risk aversion  $r$  of the synergy buyer, but also the number of auctions he/she can participate in to acquire the desired two-item bundle. Furthermore, we made this setting more competitive: while the valuation of each synergy buyer for a bundle of 2 A-s remains  $v(A, A) = 10$ , the competition is slightly more aggressive

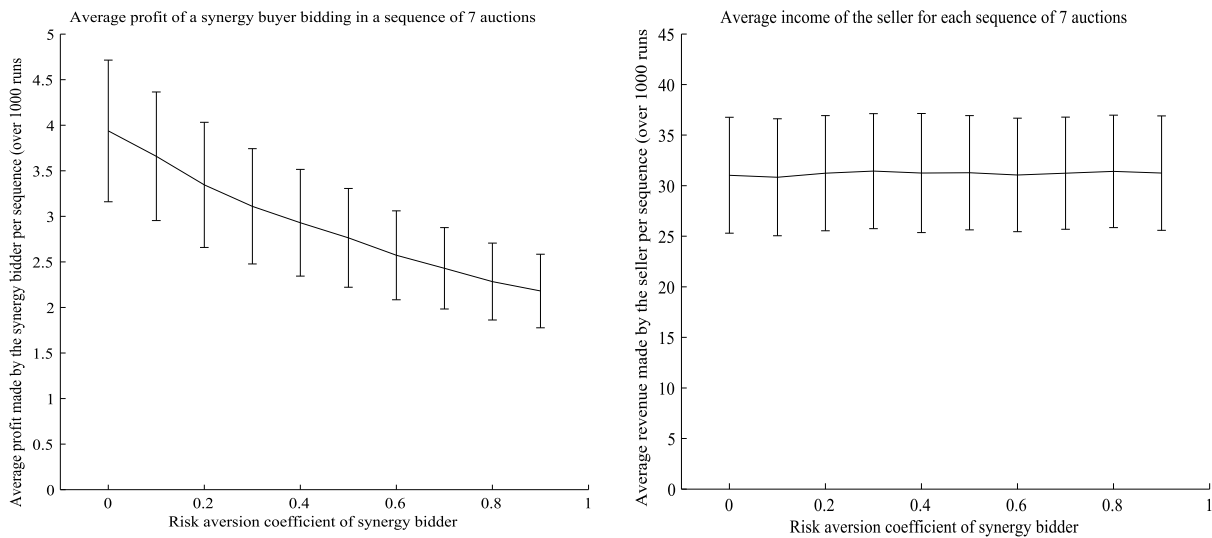


Fig. 5. Setting with 7 sequential auctions and one synergy buyer with  $v(A, A) = 10$ , and independent agents for each auction, bidding according to  $N(4, 2)$ . The left side gives the average profit of the synergy buyer, while the right the average revenue of the seller (both averaged over 1000 runs).

and will be according to  $N(4.5, 2)$ . This basically means that we reduce from 20% to 10% the advantage in valuation that a synergy buyer has, on average, over the independent bidders. The motivation for this is that it seems more relevant to study how the success rate is influenced by multiple buying opportunities (i.e. varying number of auctions), for a more competitive setting.

Results from these tests are shown in Fig. 6, respectively 7. Returning to the above stated hypotheses, we can conclude that, indeed Hypothesis 2 can be confirmed. Even for this modified setting, there is a marked decrease in average synergy bidder profits, as his/her risk aversion increases. This effect is more noticeable for the higher auction lengths (5, 7 or 10 items). The reason for this is that, for this competitive setting, the participation rates, even for the more risk-neutral agents are on the low side (see left side of Fig. 7) and relatively constant over  $r$ . Furthermore, it seems that the revenue the seller can expect is also relatively constant over the risk aversion of the synergy buyer, even for these more general tests.

From looking at Figs 6 and 7 one can clearly see the very large effect that the number of available auctions has, both on the expected profit and success rate of the synergy buyer (thus confirming Hypothesis 3 above). This effect clearly holds for all risk aversion coefficients of the synergy buyer and all configurations. From Fig. 7, one can also see that a more risk averse bidder, while making, on average, less profit, does have some advantages in this type of auctions: his/her chances of ending up with an incomplete bundle before the deadline (hence making a loss in that particular auction run) decrease considerably. Thus, these results support Hypothesis 1, as expected. Since in many real life bidding situations (one will be discussed in the following Section), agent consider only the possibility of profit/loss in an immediate run (not the long-term statistical average), minimizing chance of a loss, even if it has only 5%–10% probability, can be an important consideration.

#### 4.5. Setting with different item types and more complex preferences

The previous Section has already highlighted the complexity of bidding in sequential auction to get a bundle of two items, even for the simplified setting with one possible item type. However, in most

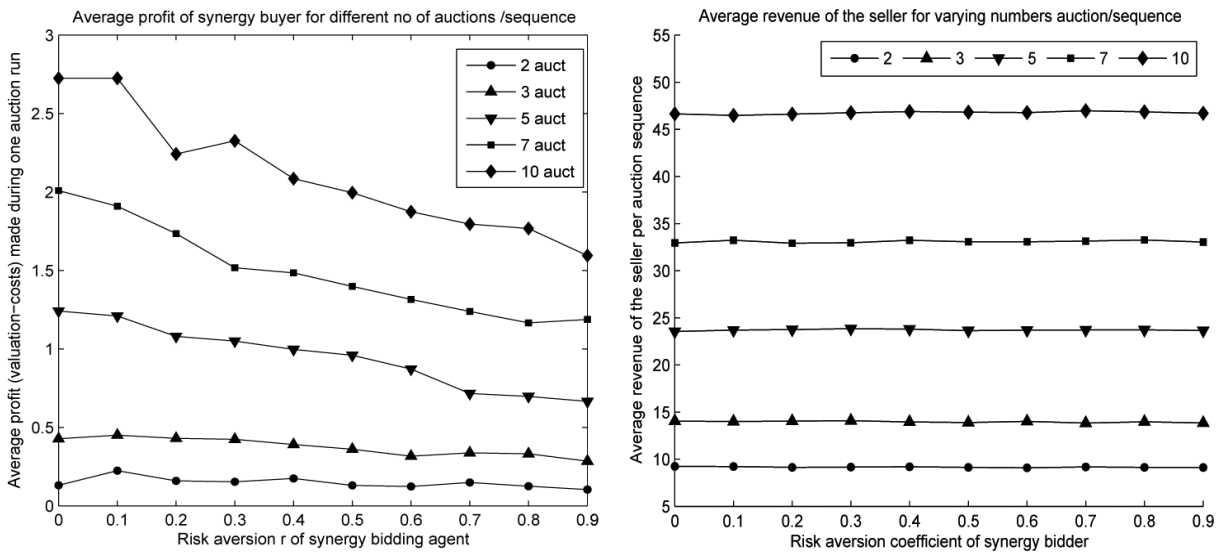


Fig. 6. Results for the average profit of the synergy buyer (left) and seller revenue (right), for a setting with one synergy agent with  $v(2A) = 10$  and a set of independent agents bidding according to  $N(4.5, 2)$ . The number of auctions the synergy agent can stay in the game to acquired the desired bundle, as well as his/her risk aversion coefficient are varied for the different settings. All results are averages over 1000 runs, but to avoid overloading the picture, error bars were not included.

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 during 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 possible preferences, we show that 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 experimental choices made could plausibly fit a real-life application setting.

#### 4.5.1. 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 (a real-life, business-oriented platform for this case, developed in collaboration with a large logistic company, is described in [24]).

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 a suitable set (bundle) of orders that would fit the capacity of their trucks. In this model, we assume all orders are ready for pick-up or return delivery at one central transportation depot.<sup>5</sup> Figure 8 shows just such a topology, with delivery point group into 2 main delivery regions).

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



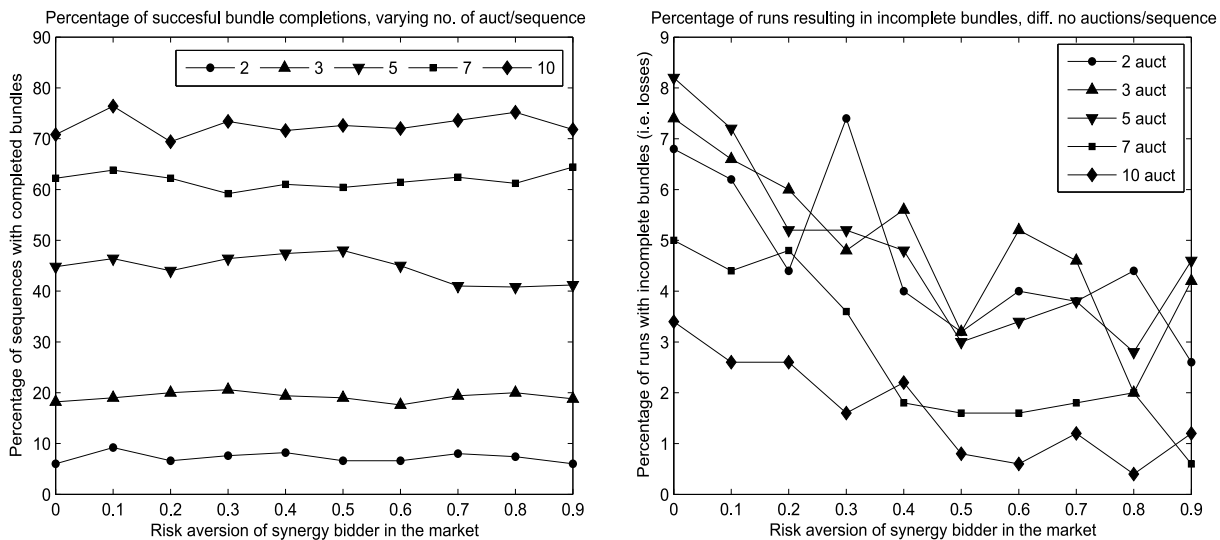


Fig. 7. Success and failure rates for a setting with one synergy agent with  $v(2A) = 10$  and a set of independent agents bidding according to  $N(4.5, 2)$ . The left side graph shows the percentage (among the 1000 runs) in which the agent acquired his/her target bundle of two A items. The right side graph shows the percentage of runs with incomplete bundles (i.e. runs in which the synergy buyer obtained the first item, but not the second, hence resulting in a loss).




Fig. 8. Example transportation scenario with one central depot D and two disjoint transportation regions: A and B. A truck starting from the central depot (D) can only drive to one of the delivery regions, assuming it has to return to the base the same day. Therefore, a bidding agent representing this truck has to acquire part-load orders from exactly one of the two 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 concern symmetrical outgoing/return orders which originate in the same region.

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, the corresponding model with sequential ascending auctions studied in this paper is basically equivalent to this, and it's easier to compare with other models and in existing literature.

Furthermore, in reporting the results, we also make the assumption that there is a single seller for all the goods (or orders) in the sequence. While in practice transportation orders may originate from multiple shippers (or customers), the aggregate revenue of the single seller can be seen as indicative of a

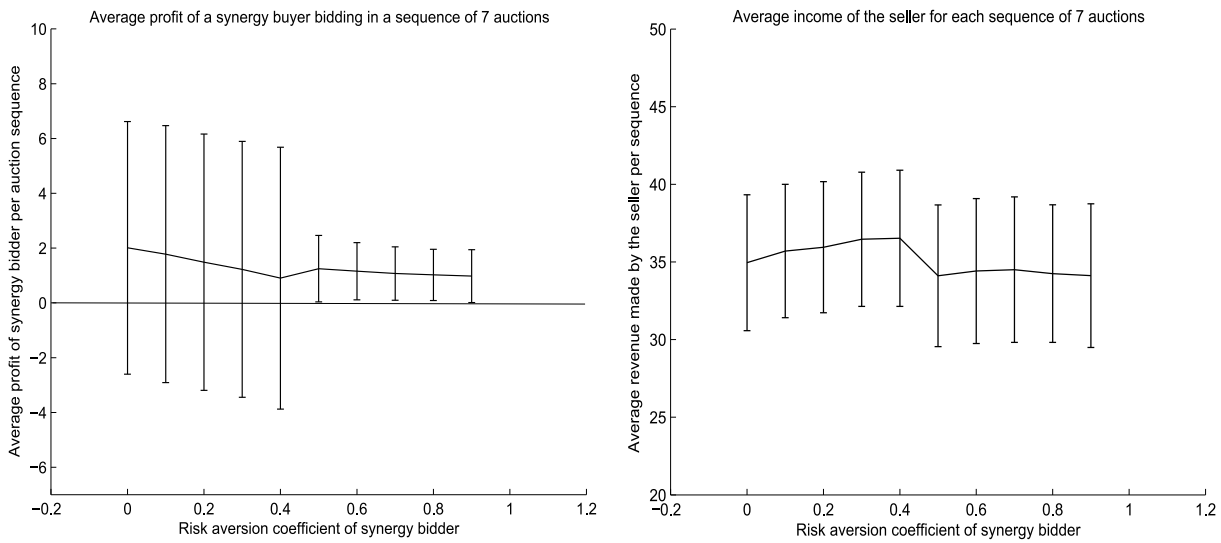


Fig. 9. Results for the average profit of the synergy buyer (left) and seller revenue (right), for a setting with two items A (of value  $v_{AA} = 10$ ) sold in 5 auctions, and B (of value  $v_{BB} = 20$ ) sold in 2 auctions. Notice there is a transition, because agents with risk aversion  $r \geq 0.5$  do not try to get the higher value bundle (of item B).

global average, that a seller, without knowing his/her precise place in the sequence of auctions, has from selling items in this sequence.

#### 4.5.2. Multiple item simulation set-up

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 more common, and sold in 5 out of 7 auctions.

However, the value the synergy buyer assigns to a bundle of such items is also asymmetrical. A bundle of 2 items of type B has a valuation of  $v(B, B) = 20$ , while a bundle of 2 items of type A:  $v(A, A) = 10$ . The competition coming from single-item bidders for those goods is also different. For goods of type A, the bids from this competition are modeled through a normal distribution  $N(\mu_A = 4, \sigma_A = 2)$ , while for items of type B through a distribution:  $N(\mu_B = 6, \sigma_B = 2)$ . Therefore, the additional valuation of the synergy agent for a bundle of two items, compared to the *average* paid by independent bidders is only  $2/10 = 20\%$  for a bundle of type A, but  $8/20 = 40\%$  for a bundle of type B. at the same time, a bundle of type B is twice as rare.

#### 4.6. Multiple item setting: hypotheses

Before we present the result graphs for this setting, we follow the format of the previous Sections, and formulate two additional hypotheses:

**Hypothesis 4:** In a market with two types of items, one of which is rarer, but also more valuable than the other, the synergy bidders with a risk aversion coefficient above a certain level may select to bid for a bundle of the more common item, in order to maximize their chances of completing the bundle. This can reduce the synergy bidder’s average expected profit from the auction sequence.

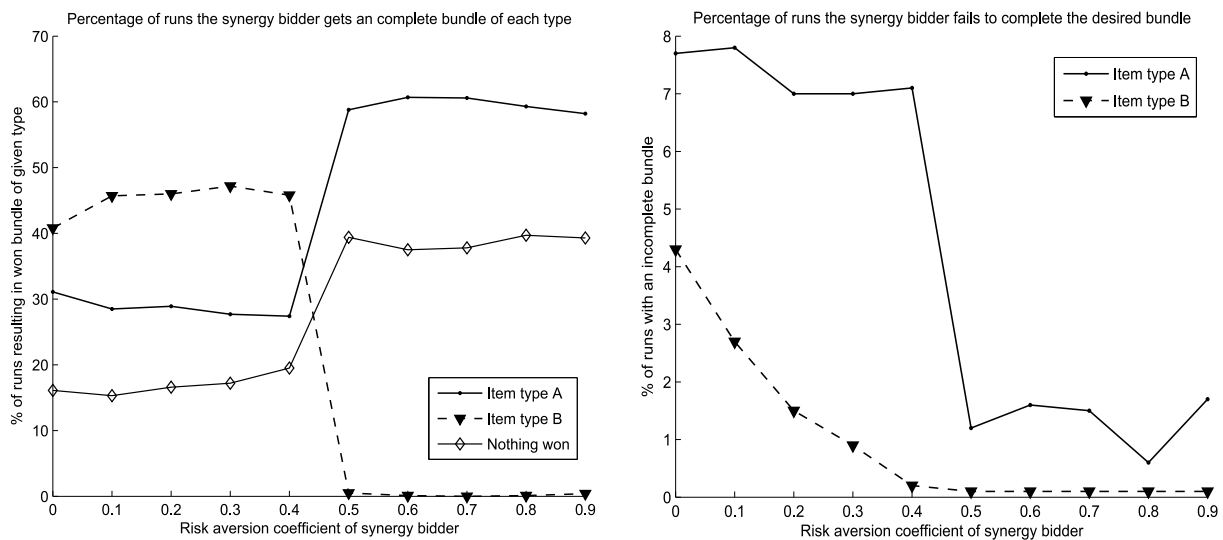


Fig. 10. Percentages of success and failure per 1000 simulation runs, for a setting with two different types of items described above. Notice the transition at  $r \geq 0.5$ , showing that risk averse bidders do not try to get the bundle with the rarer item B, but only one of A.

**Hypothesis 5:** In the above setting, if risk averse agents prefer to bid for more common, but less valuable goods, this also reduces the revenues of the auctioneer.

#### 4.7. Results for two-item case

Experimental results for the above setting are presented in Figs 9 and 10. Figure 9 gives the average profit of the synergy agent and the seller, while Fig. 10 gives the percentages of auctions that agents complete (or fail to complete) bundles of items or either type. All results reported are averages over 1000 runs.

Basically, Hypotheses 4 and 5 are, on the whole, confirmed by these tests: there is a decrease in the expected profit of the synergy buyer and seller. However, there is an important caveat: there seems, for this parameter settings, to be an important threshold effect as the risk aversion factor of the agent becomes  $r = 0.5$ . The reason for this threshold effect is clear from Fig. 10: at this level, the more risk averse agents stop trying to bid for the more valuable, but also “riskier” bundle of item B (for which there are only 2 available auctions), and go for a bundle of item A, from which there is less absolute profit to be made, but for which there are 5 available auctions.

The left-hand side of Fig. 9 shows that, while going for the bundle of item B brings, on average, slightly more profit, this result also is subject to a much higher variance, i.e. the bidding agents are more likely to loose money by failing to complete their desired bundle. By contrast, bidding for a bundle of type A (as the more risk-averse agents do), can slightly decrease the average expected profit, but the bidder is less likely to loose money. In fact, the lower interval of the variance bars, in this case, seem to be all above the zero axis.

The seller revenue (right side of Fig. 9) is also influenced by the risk aversion (and, hence, the bidding behaviour) of the synergy bidder, but the decrease in seller revenue that occurs at the threshold level is relatively slight (of only a few percentage points). Nevertheless, one should note this is a setting with only one synergy bidder present in the market, therefore the average effect may be understated.

## 5. Conclusions and further work

To summarize, the main contributions of the work presented in this paper are as follows.

First, we establish a formal link between bidding strategies in sequential auctions and standard (Arrow-Pratt) risk aversion models from economics. Next, we derive a useful property of certainty equivalence functions and it shows how such functions can be naturally applied to sequential auction games. 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.

Then, we study experimentally the effect that this decision-theoretic bidding behaviour of risk averse bidders has on his/her expected profit, for markets in which the competition is formed of "myopic", local bidders (i.e. bidders that require only one particular good). We show that, as expected, more risk-averse bidders have less of a chance to end up with an incomplete bundle, and hence make a loss. But, on average (i.e. assuming a market with repeated interactions), they make less expected profit. When bundles of two possible items are available, we show that more risk averse bidders may prefer to bid for the more common one (even if it has less absolute value), rather than risk the chance of making a loss. This is rational for them, as it reduces their risk, although it also reduces their average expected profit.

The paper, while providing some important results regarding the complexity of the sequential bidding problem for risk averse agents, leaves several issues to be answered in further work. An important one is deriving optimal bidding strategies in markets in which several synergy agents (i.e. bidders with complementary valuations) bid against each other, not only against myopic, single-value bidders, such as in this work. 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 decommitment [26] and options [11,12, 20,19] in reducing or eliminating the exposure problem that risk-averse agents face is another promising direction for further work.

## References

- [1] A. Osepayshvili, 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 IJCAI'99*, pages 527–523, San Francisco, 1999. Morgan Kaufmann.
- [5] K.-Y. Chen, L.R. Fine and B. Huberman, Predicting the future, *Information Systems Frontiers* 5(1) (2003), 47–61.
- [6] R.K. Dash, N.R. Jennings and D.C. Parkes, Computational mechanism design: A call to arms, *IEEE Intelligent Systems*, 2003, pages 40–47.
- [7] E.H. Gerding, R.K. Dash, D.C.K. Yuen and N.R. Jennings, Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *Proc 6th Int Joint Conf on Autonomous Agents and Multi-agent Systems (AAMAS'07), Honolulu, Hawaii, 2007*, pages 267–274.
- [8] W.H. Green, *Econometric Analysis*, Prentice Hall, New Jersey, 1993.
- [9] A. Greenwald and J. Boyan, Bidding under uncertainty: Theory and experiments. In *Proceedings of the 20th Conf. on Uncertainty in AI (UAI-04)*, 2004, pages 209–216.

- [10] A.X. Jiang and K. Leyton-Brown, Bidding agents for online auctions with hidden bids, *Machine Learning* **67**(1–2) (2007), 117–143.
- [11] A.I. Juda and D.C. Parkes, An options-based method to solve the composability problem in sequential auctions. In *Agent-Mediated Electronic Commerce VI*, Springer 2006, pages 44–58.
- [12] A.I. Juda and D.C. Parkes, The sequential auction problem on ebay: An empirical analysis and a solution. In *Proc 7th ACM Conf on Electr Commerce*, ACM Press, June 2006, pages 180–189.
- [13] D. Kahneman and A. Tversky, Prospect theory: An analysis of decision under risk, *Econometrica* **47**(2) (1979), 263–292.
- [14] V. Krishna, *Auction Theory*, Academic Press, 2002.
- [15] Y. Liu, R. Goodwin and S. Koenig, Risk-averse auction agents. In *Proc of AAMAS'03*, 2003, pages 353–360.
- [16] A. Mas-Collel, M.D. Whinston and J.R. Green, *Microeconomic Theory*, Oxford University Press, 1995.
- [17] P. Milgrom, *Putting Auction Theory to Work*, Cambridge University Press, 2004.
- [18] P. Milgrom and R.J. Weber, A theory of auctions and competitive bidding – part 2, *Economic Theory of Auctions*, 2000.
- [19] L. Mous, V. Robu and H. La Poutré, Can priced options solve the exposure problem in sequential auctions? In *ACM SIGEcom Exchanges* **7**(2), ACM Press, 2008.
- [20] L. Mous, V. Robu and H. La Poutré, Using priced options to solve the exposure problem in sequential auctions. In *Proc of Agent-Mediated Electronic Commerce (AMEC'08)*. Springer LNCS (to appear), 2008.
- [21] H.J. Paarsch and H. Hong, *An Introduction to the Structural Econometrics of Auction Data*, MIT Press, 2006.
- [22] M. Rabin and R.H. Thaler, Risk aversion, *Journal of Economic Perspectives* **15**(1) (2001), 219–232.
- [23] D.M. Reeves, M.P. Wellman, J.K. MacKie-Mason and A. Osepashvili, Exploring bidding strategies for market based scheduling, *Decision Support Systems* **39** (2005), 67–85.
- [24] V. Robu, H. Noot, H.L. Poutré and W.-J. van Schijndel, An agent platform for auction-based allocation of loads in transportation logistics. In *Proc 7th Int Joint Conf on Autonomous Agents and Multi Agent Systems (AAMAS'08), Industry Track*, FAAMAS Press, 2008, pages 3–10.
- [25] T. Sandholm, Algorithm for optimal winner determination in combinatorial auctions, *Artificial Intelligence* **135**(1–2) (2002), 1–54.
- [26] T. Sandholm and V. Lesser, Leveled-commitment contracting: a backtracking instrument for multiagent systems, *AI Magazine* **23**(3) (Fall 2002), 89–100.
- [27] P.J. 't Hoen and J.A. La Poutré, Repeated auctions with complementarities. In *Proc of 7th Agent-Mediated Electronic Commerce Workshop (AMEC'05), Utrecht*, 2005.
- [28] I.A. Vetsikas and N.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*, 2008, pages 1621–1624.