

Using Priced Options to Solve the Exposure Problem in Sequential Auctions¹

VALENTIN ROBU²

University of Southampton
Highfield Campus, Southampton, UK

LONNEKE MOUS³

CWI - Centrum Wiskunde & Informatica
Science Park 123, Amsterdam, The Netherlands

HAN LA POUTRÉ⁴

CWI - Centrum Wiskunde & Informatica
Science Park 123, Amsterdam, The Netherlands

vr2@ecs.soton.ac.uk, mous@cwi.nl, hlp@cwi.nl

We propose a priced options model for solving the exposure problem of bidders with valuation synergies participating in a sequence of online auctions. We consider a setting in which complementary-valued items are offered sequentially by different sellers, who have the choice of either selling their item directly or through a priced option. In our model, the seller fixes the exercise price for this option, and then sells it through a first-price auction. We analyze this model from a decision-theoretic perspective and we show, for a setting where the competition is formed by local bidders (which desire a single item), that using options can increase the expected profit for both sides. Furthermore, we derive the equations that provide minimum and maximum bounds between which the bids of the synergy buyer are expected to fall, in order for both sides of the market to have an incentive to use the options mechanism. Next, we perform an experimental analysis of a market in which multiple synergy buyers are active simultaneously. We show that, despite the extra competition, some synergy buyers may benefit, because sellers are forced to set their exercise prices for options at levels which encourage participation of all buyers.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems; K.4.4 [**Electronic Commerce**]: Distributed commercial transactions

General Terms: Algorithms, Economics, Theory, Experimentation

Additional Key Words and Phrases: Auction theory, priced options

¹This is a significantly extended and revised version of work initially presented at AMEC'08 [Mous et al. 2010]. An informal 4-page summary of the main ideas also appeared in ACM SIGEecom Exchanges [Mous et al. 2008].

²This work was performed during the author's employment at CWI, Amsterdam, supported by the NWO project DIACoDeM (Distributed Implementations of Adaptive Collective Decision Making).

³Also affiliated with the Econometrics Institute, Erasmus University Rotterdam, The Netherlands.

⁴Also affiliated with Utrecht University, Department of Information and Computing Sciences, Utrecht, The Netherlands.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2012 ACM 0000-0000/2012/0000-0001 \$5.00

1. INTRODUCTION

Online auctions play an important role in electronic commerce, as a method for allocating goods or services between self-interested agents. Single item auctions have been studied extensively in existing auction theory [Klemperer 1999], and several auction formats are known in which bidders can achieve their optimal utility by using simple, dominant bidding strategies. However, this property is generally true only for single-item, one-shot auction mechanisms, whereas in reality many of the auctions observed on the Internet today take place independently and sequentially, in the sense that they are run by different sellers and have different closing times. Furthermore, a buyer participating in a sequence of such auctions may desire a combination of items, rather than a single one. Whenever a buyer can obtain a synergy value⁵ between several goods sold sequentially, she faces an exposure problem.

The exposure problem has been studied before [Boutilier et al. 1999; Sandholm and Lesser 2002; Osepayshvili et al. 2005; Wellman et al. 2008; Greenwald and Boyan 2004] (among others). Informally, the problem occurs whenever an agent may buy a good at a higher price than what that good, by itself, is worth to her, in the hope of obtaining extra value through synergy with another good, which is sold later. However, if she then fails to buy this other good at a profitable price, she ends up with a loss. In this paper, we call such a global bidder a *synergy buyer*⁶.

The problem appears frequently on the Internet, under different forms. In retail electronic commerce, many goods sold on large online auction platforms (e.g. eBay) have complementary values to the bidders. For example, a buyer bidding on an expensive monitor may count on getting a corresponding configuration for the computer (and sound) system in a later auction. In the travel reservations domain, buyers need to reserve their flight, hotel and entertainment tickets as a package, and have little value for the different parts taken individually (this also being the setting of the Trading Agents (TAC) Travel Competition). The exposure problem also appears in business-business electronic commerce. For example, in transportation logistics, online freight exchange companies such as Teleroute (www.teleroute.com) list up to 150,000 transportation loads daily for different destinations across Europe, which are allocated on a competitive, auction-like basis⁷. However, the value of bidding for a transportation load for a carrier often depends on the probability of acquiring a return order, made available in a later auction.

Finally, another web domain where this problem appears is the dynamic allocation of web services, such as grid services, especially in domains where such services can be acquired from competing suppliers. A problem in this case is the co-location problem: two web services need to be acquired simultaneously, in order for an agent to extract value from them [Czajkowski et al. 1999; Stein et al. 2009]. For example, if a research lab secures a time slot to obtain observation data from an expensive telescope or reactor, it needs to

⁵The value of a combination of goods is super-additive with respect to the sum of the values of the goods, taken individually.

⁶Note that, since in the auction settings we consider in this paper, we always model a set of buyers bidding to acquire a good from a set of sellers (who conduct the auctions), we can use the terms “synergy buyer” or “synergy bidder” interchangeably, without loss of generality.

⁷In practice, allocation mechanisms used in multi-party logistics are not always strictly auctions, as the agent offering the order may decide which carrier bid to accept based on other criteria than just the lowest price offered (e.g. trust in that carrier, previous business relationship etc.).

ensure that the computing capacity required to process this data will be available at the required time.

In this paper, we use the generic term “goods” for the set of indivisible items to be allocated in a sequence of auctions. Without loss of generality, these can be thought of as either physical goods (such as computers or monitors) or virtual goods (such as web services, processing capacity, user attention space in online advertising etc.).

Some solutions for this problem have looked at designing the bidding strategies of individual agents participating in such a sequential auction market [Boutilier et al. 1999; Greenwald and Boyan 2004; Reeves et al. 2005; Vetsikas and Jennings 2008; Robu and La Poutré 2007; 2010]. Different classes of the TAC competition [Wellman et al. 2007] also require, among other capabilities, efficient sequential bidding from the participants. However, an automated bidding strategy participating in such a sequence of auctions faces a high degree of uncertainty, as its final utility depends on the outcome not only of the current, but also of future auctions. It is possible that bidding agents facing an exposure problem may choose not to participate in the market, because their optimal, decision-theoretic bidding policy does not give them a positive expected utility from the auction sequence. Furthermore, agents with an exposure problem may shade their bids, which reduces further both auctioneer revenues and market allocative efficiency.

For this reason, another important line of work takes the mechanism design point of view, and replaces sequential allocation with one-shot mechanisms, such as combinatorial auctions [Cramton et al. 2006; Sandholm 2002]. This approach, while it has been shown to be successful in theory and in practice for a range of settings, does have some important disadvantages. It typically requires a central point of authority, which receives the bids and computes the optimal allocation and payments, a process which can be computationally expensive. However, even assuming that the computational side of the combinatorial allocation problem can be addressed (and considerable work has focused in this direction, e.g. [Sandholm 2002]), many allocation problems occurring in practice are inherently decentralized and sequential, and cannot be mapped into one-shot, centralized mechanisms. Possible examples range from items sold on eBay by different sellers in auctions with different closing times, loads appearing over time from different shippers in distributed transportation logistics, to power allocation in dynamic electricity grids with competing suppliers.

In this paper, we consider a different approach, which preserves the sequential nature of the allocation problem, and propose a mechanism that involves auctioning *options* for the goods, instead of the goods themselves.

1.1 Options: basic definition

An option can be seen as a contract between the buyer and the seller of a good, subject to the following rules:

- The writer or seller of the option undertakes the *obligation* to sell the good for a pre-agreed *exercise price* on the demand of the buyer.
- The holder or buyer of the option gets the *right* to buy the good for the agreed *exercise price*, but not the obligation to do so.

Since the buyer gains the right to choose in the future whether or not she wants to buy the good, an option comes with an *option price*, which she has to pay regardless of whether she chooses to exercise the option or not.

Options can thus help a synergy buyer reduce the exposure problem she faces. She still has to pay the option price, but if she fails to complete her desired bundle, then she does not have to pay the exercise price as well and thus she limits her loss. So part of the uncertainty of not winning subsequent auctions is transferred to the seller, who may now miss out on the exercise price if the buyer fails to acquire the desired bundle. At the same time, the seller can also benefit indirectly, from the participation in the market by additional synergy buyers, who would have otherwise stayed out, because they faced a high risk of exposure to a potential loss.

1.2 Related work

In existing multi-agent literature, to our knowledge, there has been only limited work to study the use of options to address the exposure problem.

The first work to introduce an explicit option-based mechanism for sequential-auction allocation of goods to the multi-agent systems (MAS) community was by Juda & Parkes [Juda and Parkes 2009]. They create a market design in which synergy buyers are awarded free (i.e. zero-priced) options, in order to cover their exposure problem and, for this setting, they show that truth-telling is a dominant strategy. In this case model, the exposure problem is entirely solved for the synergy buyers, because they do not even have a possible loss consisting of the option price. Having a dominant bidding strategy for the buyers is a crucial property from a game-theoretic perspective, although in practice most real-life online markets do not exhibit this property.

However, the mechanism proposed by Juda & Parkes relies on some assumptions that could limit its applicability in some real-life markets. In particular, market entry effects may not always be sufficient to motivate the sellers of the items to use options. Because the options are designed to be offered freely (zero-priced), there are cases in which sellers do not have a sufficient incentive to offer free options, because of the risk of remaining with their items unsold. The sellers could, however, demand a premium (in the form of the option price) to cover their risk. In such cases, only positively-priced options can provide sufficient incentive for both sides of the market (buyers and sellers) to prefer an options mechanism over direct auctions. Moreover, while their mechanism guarantees that truth telling is a dominant strategy for the buyers, this property may come at a loss of efficiency for some settings, and sellers are assumed to be willing to wait in the market (and get their payments marked downwards) until the buyers of their options leave.

Priced options have a long history of research in finance (see [Hull 2003] for an overview). However, the underlying assumption for all financial option pricing models is their dependence on an underlying asset, which has a current, public value that moves independently of the actions of individual agents (e.g. this motion is assumed to be Brownian for Black-Scholes models). This type of assumption does not hold for the online, sequential auctions setting we consider.

Another line of research in the business literature focuses on real options [Amram and Kulatilaka 1998; Smith and McCardle 1999], which do not rely on the price of an underlying, publicly traded asset. Most of the literature on real options we are aware of focuses on modeling long-term business investment decisions. A relevant work that studies the use of options in online auctions is [Gopal et al. 2005]. They discuss the benefits of using options to increase the expected revenue of a seller of multiple copies of the same good. In [Gopal et al. 2005], however, it is the seller that fixes both the option price and the exercise price when writing the option, which requires rather strong assumptions on the knowledge of the

seller and on the behaviour of the bidders.

There is also a connection between options and leveled commitment mechanisms [Sandholm and Lesser 2002; 2001; 't Hoen et al. 2005]. In the leveled commitment mechanism proposed by Sandholm and Lesser, both parties have the possibility to decommit (i.e. unilaterally break a contract), against paying a pre-agreed decommitment penalty. However, as [Sandholm and Lesser 2002] show, setting the level of the decommitment penalty can be hard, due to the complex game-theoretic reasoning required. There are situations in which both parties would find it beneficial to decommit but neither does, hoping the other party would do so first, in order to avoid paying the decommitment penalty. This differs from option contracts, where the right to exercise the option is paid by one party in advance. In our model, this right is sold through an auction, thus the option price is established through an open market.

An alternative direction of research that aims to tackle a similar challenge is online mechanism design [Friedman and Parkes 2003; Parkes 2007; Gerding et al. 2011; Robu et al. 2011]. However, the online mechanism design literature we are aware of is mainly concerned with the problem of declaring truthful entry and exit times in a market, and does not deal with complementary valuations or bidder exposure to risk.

Finally, recent work by Robu, Vetsikas, Gerding & Jennings [Robu et al. 2010a; 2010b] (which appeared after the publication of our initial paper [Mous et al. 2010]) starts from the priced options mechanism developed in this work, and proposes a more complex and flexible model for pricing options. The starting assumptions considered by the two lines of work are somewhat different, because this work considers a model with first priced options and hidden reservation values (following the transportation logistics business case that initially motivated the work), while Robu, Vetsikas, Gerding & Jennings consider a model with a sequence of complementary second-price auctions and no reservations. We refer interested readers to [Robu et al. 2010a; 2010b] for a detailed comparison of the two approaches.

1.3 Outline and contribution of our approach

The goal of this paper is to study the use of priced options to solve the exposure problem and to identify the settings in which using priced options benefits both the synergy buyer and the seller.

An option contract specifies two prices, so an adjustment needs to be made to the standard auction with bids of a single price. In this study, in order to make the analysis tractable, we have a fixed exercise price and a flexible option price. The basic way our mechanism works is that the seller determines the exercise price of an option for the good she has for sale and then sells this option through a first-price auction. Buyers bid for the right to buy this option, i.e. they bid on the option price. We note that this mechanism has the attractive property that direct auctioning of the items appears as a special case. If the seller fixes the future exercise price for the option at zero, then a buyer actually bids for the right to get the item for free. Since such an option is always exercised (assuming free disposal), this is basically equivalent to direct auctioning of the item itself.⁸

⁸An alternative would be to let the sellers fix the option prices, and the exercise prices be determined by the market. A potential downside of such a mechanism may be that, if the option price is set too low, bidders could hoard options without any intention of exercising them, just to block other bidders from competing in future auctions.

Based on the above description, we provide both an analytical and an experimental investigation of the setting. Our analysis of the problem can be characterized as decision-theoretic, meaning both buyer and seller reason with respect to expected future prices. There are both advantages and disadvantages to a decision theoretic approach. The disadvantage is that, unlike the existing game-theoretic approach to options [Juda and Parkes 2009] or related online mechanism design approaches [Parkes 2007], one cannot guarantee that bidders have a dominant bidding strategy. On the other hand, using a mechanism design approach often requires additional assumptions, such as the assumption in Juda & Parkes that sellers would be interested to provide options, in order to keep buyers truthful. Unlike such approaches, decision-theory tries to model directly the reasoning and bidding behaviour of agents acting in real life markets (in most real-life sequential auction markets, no dominant bidding strategy exists anyway). While forgoing some of the strong, game theoretic rationality concepts, this has the advantage that it makes the analysis computationally tractable for larger settings.

To summarize, our contribution to the literature can be characterized as twofold:

First, we consider a setting in which n goods (or options for them) are auctioned sequentially. In our setting, there is one synergy bidder with a complementary valuation over these goods, the rest of the competition being formed by local bidders desiring only one good. For this setting, we show analytically (under some assumptions) when using priced options can increase the expected profit for both the synergy buyer and the sellers, compared to the case when the goods are auctioned directly. In order to provide a rigorous formal characterization of these settings, we derive the equations that provide minimum and maximum bounds between which the bids of the synergy buyer are expected to fall, in order for both sides to have an incentive to use options.

In the second part of the paper, we consider market settings in which multiple synergy buyers (global bidders) are active simultaneously, and study it through experimental simulations. In such settings, we show that, while some synergy buyers loose because of the extra competition, other synergy buyers may actually benefit, because sellers are forced to fix exercise prices for options at levels which encourages participation of all buyers.

We note also that, while both parts of the paper study decision theoretic bidding behaviour, we consider different levels of information about the future available to the synergy bidder. In the analytical case, the exact order of the auctions is assumed to be known, and we consider a bidder that wants a bundle of all the items to be auctioned. In the experimental part, where the synergy bidder wants only a sub-bundle of the goods from a potentially large sequence, we assume that bidding agents know only the number of future buying opportunities for an item of each type, not their exact order. This is actually more realistic for the application scenarios we consider. For example, when bidding to acquire a part-truck order in transportation logistics, it is more realistic to assume that a carrier can approximate the number of future opportunities to buy a complementary load, but not the exact auction order in which future loads will be offered for auction.

The structure for the rest of this paper is as follows. Section 2 lays the foundation for further analysis by deriving the expected profits of synergy buyers and sellers for both the direct sale, respectively for a sale with options and clarifies some of the assumptions used in our model. Section 3 provides the analytical results and proofs of the paper, for a market of sequential auctions with one synergy buyer. Sections 4 and 5 present the results from our experimental study, while Section 6 concludes with a discussion.

2. EXPECTED PROFIT FOR A SEQUENCE OF N AUCTIONS AND 1 SYNERGY BUYER

In Section 3 of this paper, we show analytically that options can be profitable to both synergy buyer and seller. This section provides a basis for these proofs, by first deriving the expected profit functions (which depend on the bids of the synergy buyer) for the synergy buyer and the seller. Throughout this study it is assumed that both sellers and buyers are risk neutral and that they want to maximize their expected utility or, in this case, their expected profit.

2.1 The market setting

We consider a market set-up in which n unique, complementary goods, are sold individually in auctions with sequential closing times.

Formally, let G be the set of n goods for sale in a temporal sequence of auctions and $v_{syn}(G_{sub})$ be the valuation the synergy buyer has for $G_{sub} \subseteq G$. In this section, we further assume that $v_{syn}(G) > 0$ and $\forall G_{sub} \subsetneq G, v_{syn}(G_{sub}) = 0$. In other words, to somewhat simplify the theoretical analysis, we consider a synergy buyer that desires the bundle of all the goods considered in the model ($G_{sub} = G$).

The goods $G_1..G_n \in G$ are sold individually through sequential, first-price, sealed-bid auctions. The main reason for this choice is that, in many settings where sequential auctions occur in practice, such as request-for-quotes (RFQ) auctions in logistics or supply chains, a model close to first-price auctioning is often used.

Moreover, in a setting with sequentially closing auctions (unlike in single-shot auctions), the usual reason for preferring second-price auctions to first-price ones (i.e. that bidding one's value is a dominant strategy) does not apply. In sequential setting with valuation complementarities of the agents, second-price auctions do not have the dominant strategies properties described by Vickrey for a single auction (see also [Boutilier et al. 1999; Greenwald and Boyan 2004] for a discussion of this issue).

The time these auctions take place in is $t = 1 \dots n$, such that at time t good $G_t \in G$ is auctioned. The above assumptions mean that if the synergy buyer has failed to obtain G_t , then she cannot achieve a bundle, for which she has a positive valuation. So if G_{t+1} is auctioned with a positive reserve price, then obtaining G_{t+1} would cost the synergy buyer money. If the synergy buyer fails to obtain G_t , then it is rational for her to not place bids in subsequent auctions.

Therefore, in this paper, we consider a model in which the number of future opportunities to buy the good (i.e. auctions) is known, but there is uncertainty over the outcome of these auctions. This models well decentralized settings, in which sellers are independent and/or the items are auctioned off as they arrive. One such practical example [Robu et al. 2008; Robu et al. 2011] is decentralized transportation logistics, where transportation orders are auctioned off by different sellers (called shippers) at different points of arrival in the market, as they become available⁹.

The bids of the synergy buyer are $\vec{B} = (b_1, \dots, b_n)$, where b_t is the bid the synergy buyer will place for good G_t , conditional on having won the previous auctions. Because

⁹In future work, we plan to look at extending this model to deal with uncertainty about the *number* of future auctions, as well as their outcomes. However, this would require a further approximation in the way that future uncertainty in completing the desired bundle is computed, which would make getting clear analytical results difficult.

of the first-price auction format, b_t is also the price the synergy buyer has to pay if she wins the auction. Throughout this analysis, we assume the competition the synergy buyer faces for each good G_t (sold at time t) is formed by local bidders that desire only the good G_t . We further assume that these local bidders do not consider the bids placed by the synergy buyer in their bidding. Therefore, from the perspective of the synergy bidder, the competition can be modeled as a distribution over the expected closing prices at each time point t , more precisely as a distribution over a value $b_{t,max}$, which is the maximal bid placed by the competition not counting b_t .

An important part of the reasoning of the synergy buyer's strategy, in our model, is the availability, for each auction held at time t , of a probability distribution $F_t(b_t)$, which gives the buyer her probability of winning the item sold at time t by placing bid b_t in that auction. There are several ways in which, in a realistic scenario, the synergy buyer could acquire this information. First, it may be that the synergy buyer knows, for each auction, the number of local bidders she is competing against, and has a distribution over their valuations. In such a case, it would be easy to aggregate this local competition in a single probability distribution function, that returns the probability of winning, given a bid. More generally, however, this distribution could be learned from repeated interactions/participation in the market, and may not necessarily require knowledge about the number of competitors in each auction.

We can exemplify this type of probabilistic reasoning in a realistic application scenario, which initially motivated this theoretical work - distributed transportation logistics [Robu et al. 2008; Robu et al. 2011]. In such a market, carriers (i.e. companies owning the actual trucks) have to bid in request for quotes auctions¹⁰. A logistic planner (representing a carrier), knows what an order from Amsterdam to London costs to execute, on average, given the market conditions on a given day. If she bids an amount b_t , she can estimate the probability of being awarded that order. Note that, in this case, she may not know exactly which other carrier companies are present in the market, but from her experience she can estimate her chances of winning the order by placing a certain bid.

2.2 Hidden reservation values

For each good G_t , there exists a strictly positive reservation value of $b_{t,res}$, which is the seller's own valuation for that good, or, alternatively, it can be seen as a resale value if she fails to sell the good in the current auction. To explain, in many real sequential auction markets where options can be applied, sellers have the option of trying to resell their goods later, even if the expected revenue of selling later is less than the expected revenue from selling now. For instance, someone who can't sell his/her computer monitor or bike frame on Ebay today will try again in the future, although there is a cost involved in waiting. While we do not model resale explicitly in our model, it is realistic to allow the goods to have a residual resale for sellers, because a seller would not want to sell her good now if the maximum offer received would be less than what she could get by waiting to sell in a future auction.

In order to model this formally, we allow sellers of each good G_t to set a hidden reserve value $b_{t,res}$, not visible to the bidders before the auction starts. The way such a model with hidden reserves works is that, after all the bids have been received, a seller can keep the

¹⁰Note that, while in this paper, for simplicity, we consider direct, not reverse auctions in which the lowest bid wins, the exposure problem over bundles of orders is identical.

goods if the maximal bid received falls under her hidden reserve value. Having a secret (as opposed to a public) reservation value in a *first price* auction motivates sellers to be truthful in setting their reserves - see [Bajari and Hortacsu 2003; Elyakime et al. 1994] for a discussion of this point¹¹.

Note that, while hidden reservation values ensure sellers are not forced to sell their items unless they want to, they also have the advantage of preventing the “hold up” effects, that would appear with publicly posted reservation values. For example, the seller in the last auction in the sequence could post such a high reservation price as to extract the entire valuation from the synergy bidder. However, with a hidden reservation value, the game is two-stage: the seller decides on whether to accept or reject offers *after* the bids are received (i.e. she cannot pre-commit to a reservation price), which precludes this undesired effect¹².

In order to model the reasoning of the synergy bidder in the presence of the hidden reservation value, we introduce an additional joint variable bm_t defined as:

$$bm_t = \max\{b_{t,maxl}, b_{t,res}\} \quad (1)$$

where $b_{t,maxl}$ denotes the maximum bid by one of the local bidders in the auction at time t , while $b_{t,res}$ is the reservation price of the seller. Thus, bm_t can be seen as the *maximum alternative bid* in the auction at time t , which can come either from one of local bidders or the seller (representing its hidden reserve value, below which the item won't get sold). Using a decision theoretic approach, we can model the strategy of the synergy bidder with respect to only a single probability over variable bm_t , which is essentially a probability over the maximum of variables $b_{t,maxl}$ and $b_{t,res}$. This can be easily computed if the probabilities over $b_{t,res}$ and bm_t are available separately, or it could be learnt directly over time, from repeated participation in the market.

2.3 Synergy buyer's profit with n unique goods, without options

Formally, we denote by $F_t(b_t)$ the probability that the synergy buyer wins good G_t with bid b_t - where $F_t(b_t)$ depends on whether b_t can outbid the maximal bid bm_t of by the competition, excluding b_t (as defined in Equation 1 above). To deal with ties, we assume the synergy buyer wins only G_t if $b_t > bm_t$ and not if the bids are equal. Then $F_t(b_t)$ can be defined as follows:

$$F_t(b_t) = Prob(b_t > bm_t) \quad (2)$$

The synergy buyer has only a strictly positive valuation for the bundle of goods G , which includes all the goods G_t , sold at times $t = 1..n$. Therefore, in a market without options, the a-priori expected profit π_{syn}^{dir} of the synergy buyer is:

$$E(\pi_{syn}^{dir}) = \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] + \left[\sum_{j=1}^n (-b_j) \prod_{k=1}^j F_k(b_k) \right] \quad (3)$$

The synergy buyer wants to maximize her expected profit. So her optimal bids $\vec{B}^* =$

¹¹Note, however, that, as shown in [Elyakime et al. 1994], having a publicly posted reserve value may actually bring sellers more revenue, but in our model we don't allow this. The reason is that this would not keep sellers truthful, and place an additional computation burden on the bidders, due to the presence of the public reserve price parameter.

¹²In addition, in practical settings it would be hard for the seller to know its exact place in the auction sequence

(b_1^*, \dots, b_n^*) maximize equation 3:

$$\vec{B}^* = \operatorname{argmax}_{\vec{B}^*} E(\pi_{sym}^{dir}) \quad (4)$$

Note that, with a decision theoretic model, the synergy bidder takes into account the reservation values $b_{t,res}$ as part of the expectation probability $F(b_t)$ (recall that $F(b_t)$ is defined as the probability of winning by placing bid b_t).

2.4 Seller's expected profit and strategy assumptions

Next, the profit of the sellers are examined. We assume that all sellers have their own valuation for the good that they sell and that they set their reserve price of $b_{t,res}$ equal to this private valuation. So when the good is sold for b_t , the seller of G_t has a profit π_t^{dir} of $b_t - b_{t,res}$. As previously shown, the synergy buyer participates only when she has won the previous auctions; otherwise bm_t is the maximal placed bid (or, if $bm_t = b_{t,res}$, the seller keeps the item).

Additionally, we also need an assumption on the seller's patience, because when an option is sold to a synergy buyer, that buyer must be able to decide whether to exercise it or not after all the other auctions of interest finish. To avoid such timing issues, in our model we explicitly assume that the n auctions that a synergy buyer can participate in are conducted by sellers with longer deadlines than the buyers.

Given the above modeling assumptions, the expected profit of the seller of the good G_t sold at time t can be written as:

$$\begin{aligned} E(\pi_t^{dir}) = & (E(bm_t) - b_{t,res}) \left(1 - \prod_{i=1}^{t-1} F_i(b_i) \right) + \left(F_t(b_t)(b_t - b_{t,res}) \right. \\ & \left. + (1 - F_t(b_t))(E(bm_t | bm_t \geq b_t) - b_{t,res}) \right) \prod_{i=1}^{t-1} F_i(b_i) \end{aligned} \quad (5)$$

Intuitively explained, the equation defines the expected utility over 3 disjoint cases: one in which the optimal bids b_i of the synergy bidder were not sufficient to win all auctions up to time t , in which case the expected profit of the seller is the highest expected bid of the local bidders, captured by $E(bm_t)$, minus its own reservation value $b_{t,res}$ (or, in the case $bm_t = b_{t,res}$, possible according to Equation 1, this term becomes zero); the second case in which the synergy bidder wins all previous auctions, including the current one (i.e. the one at time t), in which case the expected profit is this bid minus reservation $b_t - b_{t,res}$, and the third in which the synergy buyer won all previous auctions but fails to win the current one, in which case still the highest bid by the local bidders is taken.

2.5 Synergy buyer and seller profits in a model with options

Previous sections derived the expected profit functions for the synergy buyer and the sellers in a market without options. The next step is to do the same for a market with options. This section has the same setting as the general model with n goods being sold, only now an option on G_t is auctioned at time t . Therefore, all the sellers in the market will sell options for their goods, instead of directly the goods themselves. After the n auctions have taken place, the buyers need to determine whether or not they will exercise their option. It is assumed that an option is exercised only if a buyer has obtained her entire, desired bundle. The local bidders are only interested in G_t , so they will always exercise an option on G_t

should they have one. The synergy buyer is only interested in a bundle of all goods, so she will only exercise an option (and pay the corresponding exercise price) if she has options on all the goods required.

The option consists of a fixed exercise price K_t and the synergy buyer's bids on the option price are $\vec{OP} = (op_1, \dots, op_n)$. The maximal bid without the synergy buyer was bm_t , but now opm_t is the maximal placed option price.

Recall that we assume that the competition is formed by local bidders, who cannot reason about the presence in the market or the bids placed by the synergy buyer. Moreover, all local bidders in an auction only want the one good sold in that auction, hence they do not benefit from having an option and they will always exercise any option they acquire. Because of these assumptions, it follows that the competition will keep bidding the same total price, which is the bid without options minus the exercise price. Thus the distribution of the competition is only shifted horizontally to the left, by the reduction of the exercise price: $opm_t = bm_t - K_t$ (since the seller can set $K_t \leq b_{t,res}$, this ensures that always $bm_t \geq K_t$, c.f. Equation 1). Thus, if the synergy buyer bids the same total price (option + exercise), then she has the same probability of winning the auction in both models.

Let $F_t^o(op_t)$ be the probability that op_t wins the auction for the option on G_t . So if $op_t + K_t = b_t$, then $F_t^o(op_t) = F_t^o(b_t - K_t) = F_t(b_t)$.

The synergy buyer's expected profit with options then is:

$$E(\pi_{syn}^{op}) = \left[v_{syn}(G) - \sum_{h=1}^n K_h \right] \prod_{i=1}^n F_i^o(op_i) + \sum_{j=1}^n \left[(-op_j) \prod_{k=1}^j F_k^o(op_k) \right] \quad (6)$$

So her optimal bids $\vec{OP}^* = (op_1^*, \dots, op_n^*)$ maximize the profit equation 6:

$$\vec{OP}^* = \underset{\vec{OP}^*}{\operatorname{argmax}} E(\pi_{syn}^{op}) \quad (7)$$

The main difference for the seller of G_t , is that if the synergy buyer wins, then she earns op_t immediately when the options is sold, and an additional $K_t - b_{t,res}$ when (and if) the option is exercised. The probability of exercise is the probability that the synergy buyer wins all the subsequent auctions. As in the model without options, the seller of good G_t can set a hidden reservation value for her good $b_{t,res}$. As before, this represents a potential resale value for her, in case the item remains unsold, or in case the option for the item is sold, but it is not exercised. Basically, the way the mechanism works is that each seller announces the exercise price level of K_t , and receives a number of option price bids. After all these option price bids are received, the seller has the option to cancel the auction and keep the good if the maximal bid received falls under $b_{t,res} - K_t$.

Given this model, the total expected profit of the seller of good G_t sold at time t is:

$$\begin{aligned} E(\pi_t^{op}) &= (E(opm_t) + K_t - b_{t,res}) \left(1 - \prod_{i=1}^{t-1} F_i^o(op_i) \right) \\ &+ \left(F_t^o(op_t) \left[op_t + (K_t - b_{t,res}) \prod_{h=t+1}^n F_h^o(op_h) \right] \right) \\ &+ \left((1 - F_t^o(op_t)) (E(opm_t | opm_t \geq op_t) + K_t - b_{t,res}) \right) \prod_{i=1}^{t-1} F_i^o(op_i) \quad (8) \end{aligned}$$

Briefly explained, this equation has the same 3-case structure as Eq. 5 above. In two cases: when the synergy buyer loses an auction for one the earlier items in the sequence (before the items sold at time t), or when she wins all the earlier auctions, but not the auction at time t , the expected payoffs are equivalent to the direct auctioning case, although this time expressed slightly differently, based on both the exercise and option price. However in one case, when the synergy buyer acquires all the previous items and the current one (middle line in Eq. 8), the payoff is composed of two amounts. The option price op_t will be gained for sure, in this case. However, the difference between the exercise and reserve price $K_t - b_{t,res}$ (which signifies the item actually changes hands) is acquired only if the synergy bidder also wins all the subsequent auctions at times $h = t + 1..n$.

This is an important difference, since in one important case, part of the amount she is about to receive depends on the outcome of future auctions. The key, however, rests in the key observation that the synergy buyer should be willing to bid more in total (i.e. $K_t + op_t$) than in the direct auctions case. This will be analyzed in the next section.

Note that the order in the auction sequence is important, and sellers placed towards the end of the auction sequence are likely to benefit more from the fact that a synergy bidder is present in the auction. In practice, it would be desirable to establish the agenda such that the most valuable items are sold first - see Fatima [Fatima 2006] for a discussion. The theoretical analysis provided in the next section, however, starts from very general framework, and would allow us to model any auction order.

Before presenting our analytical and experimental study, we summarize for clarity the assumptions used in the model in the form of Table I.

3. ANALYTICAL STUDY OF THE CASES IN WHICH OPTIONS CAN BENEFIT BOTH SYNERGY BUYER AND SELLER

In Section 2, we derive the a-priori, expected profit for the synergy buyer and the sellers as a function of the synergy buyer's bids for a market with and without options. In this section, we use these functions to determine the difference in profit between the two markets, which is $\pi_{\delta t}$ and $\pi_{\delta syn}$ for the seller of good G_t and the synergy buyer respectively, where:

Definition 3.1.

$$\begin{aligned}\pi_{\delta t} &= \pi_t^{op} - \pi_t^{dir}, \\ \pi_{\delta syn} &= \pi_{syn}^{op} - \pi_{syn}^{dir}\end{aligned}$$

So if $\pi_{\delta t}$ and $\pi_{\delta syn}$ are positive, then both agents are better off with options.

3.1 Bidding strategies which ensure that both parties benefit from using options

Let \vec{B}^* denote the synergy buyer's optimal bidding policy in a market where goods are sold directly (without options). We assume for the rest of Section 3 that for $1 \leq t \leq n$, $F_t(b_t^*) > 0$ and $F_t(b_t^*) < 1$. So she may complete her bundle, but may also end up paying for a worthless subset of goods. Thus she faces an exposure problem. For the market with options, we define a benchmark strategy \vec{OP}' for the synergy buyer, so that the two markets can easily be compared.

Definition 3.2. Let b_t^* be the optimal bid that the synergy buyer would place in the auction at time t if no options are offered, and K_t the exercise price of the option sold at time t (pre-set by the seller of good G_t). Under the assumptions that $b_t^* > K_t$ and that local

Synergy buyer	Requires all the goods G_t , sold at times $t = 1..n$ (full complementarity) Decision-theoretic reasoning w.r.t. two distributions: $F_t(b_t)$ in the direct auctions model $F_t^o(op_t) = F_t(b_t - K_t)$ in the model with options
Local bidders	Only want good G_t auctioned at time t Do not reason about bids placed by the synergy bidder Maximal bid placed by local bidders modeled as $Eb_{t,maxl}$ Behaviour can be captured by joint stochastic variable $bm_t = \max\{b_{t,maxl}, b_{t,res}\}$
Sellers	Only sell one good G_t sold at time t through a closed, first price auction Are patient (stay in the market longer) than synergy bidders Have a residual (resale) value $b_{t,res}$ in case the good is unsold
Reservation values	<i>Hidden:</i> Seller cannot pre-commit and announce reservation value Seller may keep the good if maximal bid received under its reservation In first price auctions, seller will use its residual value $b_{t,res}$ truthfully
Option model	Each seller sets and announces exercise prices K_t In the analysis, all bids $b_t^* \geq K_t$, otherwise bidder drops out.
Sellers' prior knowledge	<i>Analytical part:</i> Sellers know their position in auction sequence <i>Experiments:</i> Sellers may not know their exact position in advance
Type of distribution considered	Any type of distributions can be handled by the bound formulas shown. For some distributions, the bids can only be determined numerically. To give a closed form expression for optimal synergy bids, uniform distribution are used (but in Section 3.2 only).

Table I. Summary of assumptions underlying the model.

bidders in the auction at time t do not reason about the bids of the synergy bidder, we define the benchmark strategy for the synergy buyer's bids with options $\vec{OP}' = (op'_1, \dots, op'_n)$ for $1 \leq t \leq n$ as:

$$op'_t = b_t^* - K_t$$

The benchmark strategy implies that the synergy buyer will bid the same total amount for the good, as if she used her optimal bidding policy in a direct sale market. Clearly this does not have to be her profit-maximizing bid in a market where priced options are used. In fact, it is almost always the case that the synergy buyer will bid a different value in a market with priced options. This deviation from the benchmark is denoted by λ_t :

Definition 3.3. Let λ_t denote the deviation in the bid of the synergy buyer on the item G_t sold at time t , in a model with options, with respect to her profit-maximizing bid b_t^* in a model without options. So her bid on an option for G_t will be $op'_t + \lambda_t$.

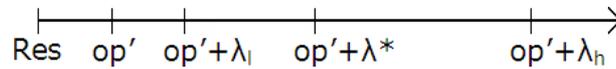


Fig. 1. A possible situation in which options are desirable.

These definitions enable us to define the bounds within which the use of options (with a given exercise price) are desirable for both the synergy buyer and the seller, for each good

in the auction sequence (except the last one, for which there is no uncertainty, so the use of options is indifferent). Fig. 1 gives the visual description of a generic setting in which options are beneficial for both sides. It shows the possible bids a synergy buyer can place for an option. First, bids have to be bigger than the reserve price Res , for each good in the sequence. The point op' is where the synergy buyer keeps bidding the same total price as in a market without options, c.f. Def. 3.2.

The deviations, in an option model, from the benchmark bid op' is measured by three levels, all denoted by λ :

- λ_l : The minimal premium the seller requires to benefit from using options over op' (due to the risk of remaining with the item unsold)
- λ_h : The maximal additional amount the synergy buyer is willing to pay for an option, over his bid in an auction without options op' , such that her expected profit is at least as high as in the no-options case.
- $\lambda^* = op^* - op'$, where op^* is the synergy buyer's profit-maximizing bid in the market with options.

Given these definitions, if it is rational for the synergy buyer to bid an additional quantity between λ_l and λ_h (as shown in Fig. 1), then both she and the seller are better off with options.

In the rest of Sect. 3, we derive the analytical expressions which can be used to determine the values for λ_l , λ_h and λ^* and compare them. Before this, however, we describe an important assumption behind the proofs in the remainder of this section.

3.1.1 Overview of our proof technique. In order to derive the λ bounds defined above, we use a recursive argument structure. First, we look at what happens when we introduce an option for just the first good, leaving the remaining goods to be allocated using the benchmark strategy, which mirrors the allocation of a direct auctions. Given the assumptions defined above regarding the bidding behaviour of the local bidders, the use of a benchmark strategy by the synergy buyer would provide the same outcomes as that of a direct auction, without options. The availability of options in the remaining auctions at times $t = 2..n$ would only increase her chances of winning the rest of the items needed to complete her bundle which, in turn, will only increase what the synergy bidder is willing to bid in the first auction.

Formally, we only consider one of the λ parameters: the one corresponding to the first good. Recall that, for this good, the buyer's probability of not completing her desired bundle, hence her exposure problem, is the greatest. Our proof structure could be generalized as a recursive procedure: if one shows that options are beneficial to use for the first item in a sequence, given a remaining [non-empty] sequence of auctions, this can be generalized to all remaining sub-sequences (except for the very last item, for which the analysis is trivial, as options cannot bring a benefit by comparison to direct auctions).

In order to analytically examine the benefits of deviating from the benchmark strategy op'_1 in the first auction, the proofs will use the supposition that the synergy buyer will use the benchmark strategy from Def. 3.2 for the remaining goods in the sequence. The use of the benchmark bidding strategy for the remaining items can be seen as giving an "upper bound" for the lower lambda value expected by the seller (i.e. λ_l) and a lower bound for the highest value that can be offered by the buyer (i.e. λ_h). We can see this by examining the effect of this assumption on each of the parties:

- For the *synergy buyer*: Being offered the opportunity to use options also in future auctions can only increase her expected profit from future auctions (since $\lambda^* \geq 0$ and $op^* \geq op'$). Otherwise, the synergy buyer will revert to using her benchmark strategy op' , which brings the same expected profit as the direct sale case. Her expected profit is at least as large in the options case as in the direct sale case i.e. $E(\pi_{syn,t \geq 2}^{op}) \geq E(\pi_{syn,t \geq 2}^{dir})$.
- For the *seller of the first item*: Because for each of the following items $op^* \geq op'$, the probability that the agent will get all the future items can only increase, for each of the items in the sequence. Formally: $F_h^o(op_h^*) \geq F_h^o(op_h') = F_h(b_h^*), \forall h = 2..n$. This implies that $\prod_{h=2}^n F_h^o(op_h^*) \geq \prod_{h=2}^n F_h(b_h^*)$, therefore the probability that the option for the first item is exercised can only increase. Therefore, this benchmark case acts as a lower bound for the expected profit of the seller, and as an upper bound on the λ_l .

In future auctions the synergy seller and buyer can use options, but this will not negatively affect the initial decisions, i.e. at the beginning of the auction sequence. Therefore, the lambda values referred to in the equations in the following sections could be formally denoted as λ_l^{as} and λ_h^{as} , where in the general case it holds that $\exists \lambda_l, \lambda_h$ such that $\lambda_l \leq \lambda_l^{as}$ and $\lambda_h \geq \lambda_h^{as}$. To avoid overloading the notation, we still use λ_l and λ_h , but the reader should be aware these refer to the tightest bounds on these lambda values, under the assumption that the benchmark bidding strategy is used in all auctions subsequent to the current one.

3.1.2 *When synergy buyer is better off with options.* This part of Section 3.1 examines for which bids the synergy buyer is better off with options. This is done by determining the maximal amount she is willing to pay for options.

LEMMA 3.4. *Let $\vec{B}^* = \langle b_t^* \rangle$ for $1 \leq t \leq n$ be the vector of optimal bids of the synergy buyer in the model without options, and $op'_t + \lambda_t$ be the bids in a model with options. Then the expected gain (i.e. difference in expected profit) from using options $E(\pi_{\delta syn})$ can be written as:*

$$\begin{aligned}
E(\pi_{\delta syn}) &= \left[v_{syn}(G) \left(\prod_{i=1}^n F_i(b_i^* + \lambda_i) - \prod_{i=1}^n F_i(b_i^*) \right) \right] \\
&+ \left[\sum_{j=1}^n K_j \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right) \right] \\
&+ \sum_{j=1}^n (-\lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \\
&+ \left[\sum_{j=1}^n (-b_j^*) \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{k=1}^j F_k(b_k^*) \right) \right]
\end{aligned}$$

PROOF. We compute the difference in profit between a model with options and a model without options, using expected profit equations (6) and (3), as defined in the previous section. In a model without options, the optimal bids of the synergy buyer at each time step t are given by b_t^* . In a model with options, we express the bidding policy as a deviation

with respect to the benchmark strategy with options, i.e. $op'_t + \lambda_t$. This gives the difference:

$$\begin{aligned} E(\pi_{\delta_{syn}}) &= \left[v_{syn}(G) - \left[\sum_{h=1}^n K_h \right] \prod_{i=1}^n F_i^o(op'_i + \lambda_i) \right] \\ &+ \left[\sum_{j=1}^n (-op'_j + \lambda_j) \prod_{k=1}^j F_k^o(op'_k + \lambda_k) \right] \\ &- \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] - \left[\sum_{j=1}^n (-b_j^*) \prod_{k=1}^j F_k(b_k^*) \right] \end{aligned}$$

We can now replace op'_t with the definition of the benchmark strategy (i.e. same total bid amount, as in the case without options), using the properties: $op'_t = b_t^* - K_t$ and $F_t^o(op'_t + \lambda_t) = F_t(b_t^* + \lambda_t)$. This gives:

$$\begin{aligned} E(\pi_{\delta_{syn}}) &= \left[v_{syn}(G) - \left[\sum_{h=1}^n K_h \right] \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right] \\ &+ \left[\sum_{j=1}^n (-b_j^* + K_j - \lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \right] \\ &- \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] - \left[\sum_{j=1}^n (-b_j^*) \prod_{k=1}^j F_k(b_k^*) \right] \end{aligned}$$

This formula is now re-grouped, separating the terms $v_{syn}(G)$, $\sum_{j=1}^n K_j$, $\sum_{j=1}^n (-\lambda_j)$ and $\sum_{j=1}^n (-b_j^*)$, each with its corresponding probabilities to complete the proof the proof:

$$\begin{aligned} E(\pi_{\delta_{syn}}) &= \left[v_{syn}(G) \left(\prod_{i=1}^n F_i(b_i^* + \lambda_i) - \prod_{i=1}^n F_i(b_i^*) \right) \right] \\ &+ \left[\sum_{j=1}^n K_j \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right) \right] \\ &+ \sum_{j=1}^n (-\lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \\ &+ \left[\sum_{j=1}^n (-b_j^*) \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{k=1}^j F_k(b_k^*) \right) \right] \end{aligned}$$

□

To explain intuitively Lemma 3.4, the difference in expected profits between the two models is formed of 4 parts (corresponding to the 4 lines). First, in an options model, the synergy bidder has a higher probability of getting the desired bundle and extract its value, since she bids more in total (line 1). Furthermore, in an options model, the bidder does not have to pay exercise prices unless she acquires all n items in the desired bundle (line 2). On the minus side, she does have to pay a set of additional amounts λ (line 3) for all

items she bids on until one is lost (line 3) and, for these items, the chance of acquiring them increases slightly, which also increases the chance of lost bids (line 4).

In the following, we turn our attention to providing equations that allow us to deduce the λ parameters that give the synergy buyer an incentive to use options. As previously explained in Sect. 3.1.1 above, we simplify the proof structure by only focusing on the most important option for the synergy buyer: the one on the first good (when bidding for this good, the probability of not completing her entire bundle is the greatest). This is done under the assumption that for the goods in the sequence, we assume the benchmark strategy is used (i.e. $\lambda_t = 0$ for $t > 1$). For the rest of the items in the sequence, the same proof technique can be applied recursively.

THEOREM 3.5. *Let λ_1 be the deviation in the bidding strategy, compared to the benchmark strategy op_1^b , as defined in Def. 3.2. If $\lambda_t = 0$ for $1 < t \leq n$, then by definition, $E(\pi_{\delta_{syn}}) \geq 0$ if $0 \leq \lambda_1 < \lambda_h$. The value of λ_h (corresponding to $E(\pi_{\delta_{syn}}) = 0$) can be solved as the numerical solution to the following equation:*

$$\begin{aligned} F_1(b_1^* + \lambda_h)\lambda_h &= F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\ &+ (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{syn}(G) \prod_{i=2}^n F_i(b_i^*) - \sum_{j=1}^n (b_j^*) \prod_{k=2}^j F_k(b_k^*) \right] \end{aligned}$$

PROOF. The proof is based on the difference in profit function derived in Lemma 3.4, using the assumption that $\lambda_t = 0$ for $1 < t \leq n$. As the expectation function of the synergy bidder is descending in the value of λ , we determine when $E(\pi_{\delta_{syn}}) = 0$.

$$\begin{aligned} &\left[v_{syn}(G)(F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \prod_{i=2}^n F_i(b_i^*) \right] \\ &+ \left[\sum_{j=1}^n K_j (F_1(b_1^* + \lambda_h) \prod_{k=2}^j F_k(b_k^*)) - (F_1(b_1^* + \lambda_h) \prod_{i=2}^n F_i(b_i^*)) \right] \\ &+ (-\lambda_h) F_1(b_1^* + \lambda_h) \\ &+ \left[\sum_{j=1}^n (-b_j^*) (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \prod_{k=2}^j F_k(b_k^*) \right] = 0 \end{aligned}$$

Isolating the values of λ_h yields the formula in Th. 3.5.

$$\begin{aligned} F_1(b_1^* + \lambda_h)\lambda_h &= (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{syn}(G) \prod_{i=2}^n F_i(b_i^*) \right] \\ &+ F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\ &+ (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[\sum_{j=1}^n (-b_j^*) \prod_{k=2}^j F_k(b_k^*) \right] \end{aligned}$$

Which give the following equation for determining λ_h :

$$\begin{aligned} F_1(b_1^* + \lambda_h)\lambda_h &= F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\ &+ (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{syn}(G) \prod_{i=2}^n F_i(b_i^*) - \sum_{j=1}^n (b_j^*) \prod_{k=2}^j F_k(b_k^*) \right] \end{aligned}$$

□

3.1.3 When the first seller is better off with options. We now determine the minimum or lower bound λ_l (the level of λ that, according to Def. 3.3, keeps the seller of G_1 indifferent about options). In order to compare this bid with the λ_h from the previous section, it is again assumed that $\lambda_t = 0$ for $1 < t \leq n$.

THEOREM 3.6. *If without options the synergy buyer bids \vec{B}^* and with options $op'_1 + \lambda_1$ for G_1 and op'_t for $1 < t \leq n$, then $E(\pi_{\delta_1})$ for the seller of G_1 is:*

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^*)(\lambda_1 + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right]) \\ &+ (F_1(b_1^* + \lambda_1) - F_1(b_1^*)) (b_1^* + \lambda_1 - E(bm_1 | b_1^* + \lambda_1 \geq bm_1 > b_1^*)) \\ &+ (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] \end{aligned}$$

By definition, λ_l is the lower bound for λ_l that guarantees that the expected profit of the seller $E(\pi_{\delta_1}) > 0$. The value of λ_l can be obtained as the solution to the equation $E(\pi_{\delta_1}) = 0$, which using the equation above gives:

$$\begin{aligned} F_1(b_1^* + \lambda_l)(-\lambda_l) &= F_1(b_1^* + \lambda_l) \left((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] \right) \\ &+ (F_1(b_1^* + \lambda_l) - F_1(b_1^*)) (b_1^* - E(bm_1 | b_1^* + \lambda_l \geq bm_1 > b_1^*)) \end{aligned}$$

PROOF. The difference in profit is equation (8) minus equation (5):

$$\begin{aligned} E(\pi_1^{op}) - E(\pi_1^{dir}) &= \left(F_1^o(op_1) \left[op_1 + (K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op_h) \right] \right) \\ &+ (1 - F_1^o(op_1)) \left(E(opm_1 | opm_1 \geq op_1) + K_1 - b_{1,res} \right) \\ &- \left(F_1(b_1^*) (b_1^* - b_{1,res}) + (1 - F_1(b_1^*)) (E(bm_1 | bm_1 \geq b_1^*) - b_{1,res}) \right) \end{aligned}$$

Recall that the the price op_1 bid in an options model can be expressed in terms of the benchmark strategy op'_1 and the deviation λ_1 .

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1^o(op'_1 + \lambda_1) (op'_1 + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h) \right]) \\ &+ (1 - F_1^o(op'_1 + \lambda_1)) \left(E(opm_1 | opm_1 \geq op'_1 + \lambda_1) + K_1 - b_{1,res} \right) \\ &- F_1(b_1^*) (b_1^* - b_{1,res}) - (1 - F_1(b_1^*)) \left(E(bm_1 | bm_1 \geq b_1^*) - b_{1,res} \right) \end{aligned}$$

Furthermore, we can make the substitution to replace op'_1 with its definition, as follows:
 $op_1 = op'_1 + \lambda_1 = b_1^* - K_1 + \lambda_1$ and $F_1^o(op_1) = F_1^o(op'_1 + \lambda_1) = F_1(b_1^* + \lambda_1)$:

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^* + \lambda_1)(b_1^* - K_1 + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_{oh}(op'_h) \right]) \\ &+ (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(-E(bm_1|b_1^* + \lambda_1 \geq bm_1 > b_1^*) + b_{1,res}) \\ &- F_1(b_1^*)(b_1^* - b_{1,res}) \end{aligned}$$

Split $F_1(b_1^* + \lambda_1)$ into $F_1(b_1^*)$ and $F_1(b_1^* + \lambda_1) - F_1(b_1^*)$ and combine some K_1 and $b_{1,res}$.

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^*)(-K_1 + b_{1,res} + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h) \right]) \\ &+ (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(b_1^* - K_1 + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h) \right]) \\ &- E(bm_1|b_1^* + \lambda_1 \geq bm_1 > b_1^*) + b_{1,res} \end{aligned}$$

Thus:

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^*)(\lambda_1 + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right]) \\ &+ (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(b_1^* + \lambda_1 - E(bm_1|b_1^* + \lambda_1 \geq bm_1 > b_1^*)) \\ &+ (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] \end{aligned}$$

Since, by definition, $E(\pi_{\delta_1}) = 0$ gives the value of λ_l , this value can be solved via the equation in Th. 3.6.

$$\begin{aligned} F_1(b_1^* + \lambda_l)(-\lambda_l) &= F_1(b_1^* + \lambda_l)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right]) \\ &+ (F_1(b_1^* + \lambda_l) - F_1(b_1^*))(b_1^* - E(bm_1|b_1^* + \lambda_l \geq bm_1 > b_1^*)) \end{aligned}$$

□

Intuitively, the difference in profit has two parts: the cases where the synergy buyer wins the auction in both markets and the ones where she only wins with options. With the first, the synergy buyer pays more than she used to and with the second, the synergy buyer pays more than the local bidders, who used to win if $\lambda_1 < \lambda_l$. But both cases have the downside for the seller that the synergy buyer may now not exercise her option.

3.1.4 Condition for both synergy buyer and seller to be better off with options. The previous parts of Section 3.1 give the equations for the cases when the individual agents are better off with options. These results will now be combined to give the formal condition for when they are both better off. Intuitively, this condition is equivalent to stating that the minimum bid the seller of G_1 requires should be below the maximal value the synergy buyer is willing to pay. As shown the beginning of Section 3.1.1, the equations for λ_l and

λ_h that are derived in Theorems 3.5 and 3.6 above are the narrowest possible interval values, under the assumption that all remaining auctions are direct auctions. Let the solutions to the equations in Theorems 3.5 and 3.6 be denoted by λ_h^{as} and λ_l^{as} . We show that $\exists \lambda_l, \lambda_h$ such that $\lambda_l \leq \lambda_l^{as}$ and $\lambda_h \geq \lambda_h^{as}$. Next, we summarize the results in a final theorem:

COROLLARY 3.7. *Under the condition that the optimal decision of the synergy buyer is to bid λ_x additionally for an option on G_1 (where $\lambda_l^{as} < \lambda_x < \lambda_h^{as}$), then both the seller of G_1 and the synergy buyer have a higher expected profit in a market with only options compared to one without options.*

PROOF. This corollary follows from the results of previous theorems. Say that the synergy buyer bids $op'_1 + \lambda_x$ for the first good in the sequence, where $\lambda_l^{as} < \lambda_x < \lambda_h^{as}$ and op'_t for the other goods. Then the synergy buyer bids more than $op' + \lambda_l^{as} \geq op' + \lambda_l$ (because $\lambda_l \leq \lambda_l^{as}$), so according to Theorem 3.6 the seller of G_1 has a higher expected profit with options. Also, the synergy buyer bids between $0 < \lambda_x \leq \lambda_h^{as} \leq \lambda_h$ extra (as $\lambda_h \geq \lambda_h^{as}$), so according to Theorem 3.5 she too has a higher expected profit with options with these bids. Therefore \exists a non-empty interval $[\lambda_l, \lambda_h]$ for which both parties prefer using options, rather than a direct sale. \square

3.2 Synergy buyer's profit-maximizing bid with uniform distributions

In the previous sections, we focused our attention on deriving equations for the bounds λ_l and λ_h between which the additional bids of the synergy buyer have to fall in order for both parties to be incentivised to use options. Note that those previous results are quite general and hold for any type of distribution that the maximal bid from the other agents in the first auction $F(b_1^*)$ might follow.

While these bounds were defined in relation to the expected-profit maximizing bid b^* in a model *without* options, the optimal (i.e. expected profit maximizing) bid op^* in a model with options have yet to be defined. The reason for this is that deriving this is much more involved than the optimal policy in a model without options. In this section, we look at the synergy buyer's profit-maximizing bids op^* , but with the additional assumption that $F_1(b_1)$ follows a uniform distribution in the range of the possible bids. Note that, while the analytical result provided here is for a uniform distribution, the same effects hold for Gaussian distributions. In fact, the optimal bids can be derived for Gaussian distributions, but just not in a closed analytical form, as is done in this section for uniform distributions.

In order to derive the optimal λ^* , we do this by use the same framework introduced in Def. 3.3 and Fig. 1 above. That means, we compute the deviation λ^* between the optimal bid in a model with options and the optimal bid in a model without options, i.e. the difference $\lambda^* = (K_1 + op_1^*) - b_1^*$ (the reason to do this will become apparent in the proof, but, basically, by taking the difference, several terms drop out). Note that in this section, we still apply the above results and assumption regarding bidding the benchmark strategy in future auctions, but to simplify the notation, we still use λ_l and λ_h , instead of λ_l^{as} and λ_h^{as} .

If the profit-maximizing bid $op_1^* > op'_1 + \lambda_l$, then according to Theorem 3.6 the seller of G_1 is better off with options. Therefore, it is in the rational interest of the seller to set the exercise price for selling her good such that the expected optimal bid of her buyers, in a model with options, will provide sufficient incentive for the seller to also use options, and thus the following condition holds: $op_1^* > op'_1 + \lambda_l$. Note that in order to use Theorem 3.6, the bids for the other goods are fixed at op'_t . First op_1^* and λ_l are derived.

LEMMA 3.8. *If $F_1(b_1)$ follows a uniform distribution between ua and ub , then $op_1^* + K_1 - b_1^* = \lambda^*$, where:*

$$\lambda^* = \begin{cases} 0.5(K_1(1 - \prod_{i=2}^n F_i(b_i^*)) + \sum_{j=2}^n K_j(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*))), \\ \text{if } ua \leq E(\pi_{syn,k \geq 2}^{dir}) \leq ub + (ub - ua) \\ 0, \text{ otherwise} \end{cases}$$

PROOF. With a uniform bid distribution between ua and ub , the probability of winning with bid b_1 has the following shape:

$$F_1(b_1) = \begin{cases} 0, & \text{if } b_1 < ua \\ (b_1 - ua)/(ub - ua) = \alpha(b_1 - ua), & \text{if } ua \leq b_1 \leq ub \\ 1, & \text{if } b_1 > ub \end{cases} \quad (9)$$

$$f_1(b_1) = \begin{cases} 1/(ub - ua) = \alpha, & \text{if } ua \leq b_1 \leq ub \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

For F_1^o the variables α_o , ua_o and ub_o are used, where $ua_o = ua - K_1$ and $ub_o = ub - K_1$, so that $F_1(b_1) = F_1^o(op_1)$ when $b_1 - K_1 = op_1$.

First, we determine, for this type of distribution, the equation for the optimal bid b_1^* in a model without options. To do this, we start from the expected profit equation (3):

$$E(\pi_{syn}^{dir}) = F_1(b_1)[v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + F_1(b_1)(-b_1) + F_1(b_1)[\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)]$$

$$E(\pi_{syn}^{dir}) = F_1(b_1) \left[-b_1 + [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \right]$$

So the derivative with respect to b_1 :

$$\begin{aligned} \frac{\partial E(\pi_{syn}^{dir})}{\partial b_1} &= f_1(b_1) \left[-b_1 + [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] \right. \\ &\quad \left. + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \right] + F_1(b_1)(-1) = 0 \end{aligned}$$

Filling in the equations for f_1 and F_1 leads to:

$$[v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] + ua = 2b_1^*$$

Nevertheless, the b_1^* obtained through this formula still has to satisfy the interval constraints $ua \leq b_1^* \leq ub$. This means:

$$ua \leq \frac{[v_{syn}(G) \prod_{i=2}^n F_i(b_i)]}{2} + \frac{[\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)]}{2} + \frac{ua}{2} \leq ub$$

Which yields:

$$ua \leq [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \leq 2ub - ua$$

Note that the middle expression is, in fact, the expression for the expected profit of a direct synergy bidder, from the second auction onwards (i.e. for $k \geq 2$), discounting the bid to be paid for the first item. Therefore, we can rewrite this condition as:

$$ua \leq E(\pi_{syn,k \geq 2}^{dir}) \leq ub + (ub - ua)$$

From this form, it is easier to explain why outside this interval, $\lambda^* = 0$. If the expected profit of the future sequence $E(\pi_{syn,k \geq 2}^{dir}) < ua$, there is no point in the buyer to continue bidding (either direct or with options), as she cannot afford her desired bundle anyway. Therefore, both b^* and λ^* should be zero. If the expected profit of the future sequence exceeds the value of ub with a whole interval $ub - ua$ (i.e. $E(\pi_{syn,k \geq 2}^{dir}) > ub + (ub - ua)$), then the direct bid assures the bidder of winning the item (as uniform distributions are bounded). But this means that options are also not useful, so again $\lambda^* = 0$ (there is no point of bidding more than in a direct model).

To get the value of λ^* outside these trivial cases is more involved. First, we compute the optimal bid op_1^* in a model *with* options:

$$E(\pi_{syn}^{op}) = \left[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=1}^n F_i^o(op_i) \right] + \left[\sum_{j=1}^n (-op_j) \prod_{k=1}^j F_k^o(op_k) \right]$$

First, we isolate op_1 in the above equation:

$$\begin{aligned} E(\pi_{syn}^{op}) &= F_1^o(op_1) \left[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i^o(op_i) \right] \\ &+ F_1^o(op_1)(-op_1) + \left[\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k) \right] \end{aligned}$$

$$\begin{aligned} E(\pi_{syn}^{op}) &= F_{o1}(op_1) \left[-op_1 + [(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_{oi}(op_i)] \right] \\ &+ \left[\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_{ok}(op_k) \right] \end{aligned}$$

We take the derivative with respect to op_1 :

$$\begin{aligned} \frac{\partial E(\pi_{syn}^{op})}{\partial op_1} &= f_1^o(op_1) \left[-op_1 + [(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i^o(op_i)] \right] \\ &+ \left[\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k) \right] + F_1^o(op_1)(-1) = 0 \end{aligned}$$

In order to determine the optimal value op_1^* , we add the condition $\frac{\partial E(\pi_{syn}^{op})}{\partial op_1} = 0$:

$$\begin{aligned} & \alpha_o \left[-op_1^* + [(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i^o(op_i)] \right. \\ & \left. + [\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k)] \right] + \alpha_o (op_1^* - ua_o)(-1) = 0 \end{aligned}$$

Which finally yields the following equation for determining op_1^* :

$$[(v_{syn}(G) - \sum_{h=1}^n K_h) \prod_{i=2}^n F_i^o(op_i)] + [\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k)] + ua_o = 2op_1^*$$

We now focus our attention at computing the difference λ^* between the optima decision-theoretic bid in a model with options vs. a model without options. By definition, we have that: $\lambda^* = (K_1 + op_1^*) - b_1^*$, so $2\lambda^* = 2op_1^* + 2K_1 - 2b_1^*$. When taking this difference, $ua_o = ua - K_1$ and op_k are replaced according to $op_k = op_k' = b_k^* - K_t$ (because for the other auctions, the benchmark strategy is used) and $F_k^o(op_k') = F_1(b_1^*)$. Then all variables cancel each other out, except for the K_t :

$$\begin{aligned} 2(b_1^* + \lambda^* - K_1) &= \left[[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i(b_i^*)] \right. \\ & \left. + [\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*)] \right] + ua - K_1 \end{aligned}$$

hence

$$\begin{aligned} 2\lambda^* &= \left[[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i(b_i^*)] \right. \\ & \left. + [\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*)] \right] + ua + K_1 - 2b_1^* \end{aligned}$$

thus

$$\begin{aligned} \lambda^* &= 0.5 \left(\left[[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i(b_i^*)] \right. \right. \\ & \left. \left. + [\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*)] \right] + ua + K_1 \right. \\ & \left. - \left(\left[[v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \right] + ua \right) \right) \end{aligned}$$

After some re-writing:

$$\lambda^* = 0.5 \left(\left(-\sum_{h=1}^n K_h \right) \prod_{i=2}^n F_i(b_i^*) + \sum_{j=2}^n K_j \prod_{k=2}^j F_k(b_k^*) + K_1 \right)$$

Re-arranging the parantheses:

$$\lambda^* = 0.5(K_1 - K_1 \prod_{i=2}^n F_i(b_i^*) - \sum_{h=2}^n K_h \prod_{i=2}^n F_i(b_i^*) + \sum_{j=2}^n K_j \prod_{k=2}^j F_k(b_k^*))$$

Which finally leads to the equation in Lemma 3.8:

$$\lambda^* = 0.5(K_1(1 - \prod_{i=2}^n F_i(b_i^*)) + \sum_{j=2}^n K_j(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*))) \quad (11)$$

□

The main intuition behind this formula is that, in an options model, the synergy buyer saves the exercise price when she fails to complete her bundle. Therefore, it is her profit-optimizing strategy, in a model with options, to increase her bid with a part of the potential savings on the exercise prices of subsequent auctions.

LEMMA 3.9. *If $F_1(b_1)$ follows a uniform distribution, then the lower bound is:*

$$\begin{aligned} \lambda_l = & -(b_1^* - ua + [1 - \prod_{h=2}^n F_h(b_h^*)](b_{1,res} - K_1)) + \\ & + \sqrt{(b_1^* - ua + [1 - \prod_{h=2}^n F_h(b_h^*)](b_{1,res} - K_1))^2} \\ & - 2(b_1^* - ua) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right] (b_{1,res} - K_1) \end{aligned}$$

PROOF. Take the λ_l equation from Theorem 3.6. With a uniform distribution, $F_1(b_1) = \alpha(b_1^* - ua)$ and $E(bm_1 | b_1^* + \lambda_l \geq bm_1 > b_1^*) = b_1^* + 0.5\lambda_l$. So the equation becomes:

$$\begin{aligned} \alpha(b_1^* + \lambda_l - ua)(-\lambda_l) = & \alpha(b_1^* + \lambda_l - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ & + \alpha\lambda_l(b_1^* - b_1^* - 0.5\lambda_l) \end{aligned}$$

Dividing both sides by α and reducing b_1^* in the last parenthesis gives:

$$(b_1^* + \lambda_l - ua)(-\lambda_l) = (b_1^* + \lambda_l - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) + \lambda_l(-0.5\lambda_l)$$

After re-arranging the terms and moving the left -hand side to the right, this yields:

$$(b_1^* + \lambda_l - ua)(\lambda_l + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) - 0.5\lambda_l^2 = 0$$

The above equation can be brought to standard, 2nd order polynomial form in the unknown

λ_l :

$$0 = 0.5\lambda_l^2 + \lambda_l(b_1^* - ua + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ + (b_1^* - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right])$$

This polynomial equation can then be solved via the quadratic formula:

$$\lambda_l = -(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1)) \\ \pm \sqrt{(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1))^2 - 2(b_1^* - ua) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1)}$$

Note that, formally, the condition $ua \leq b^* \leq ub$ should also be imposed in the above equation. However, if b^* for the direct sale case falls outside this interval (i.e. if $\leq E(\pi_{syn,k \geq 2}^{dir}) < ua$ or $\leq E(\pi_{syn,k \geq 2}^{dir}) > ub + (ub - ua)$), we know that the the lambda of the seller $\lambda^* = 0$, so there is no point in the seller even considering offering options. Outside this interval, it makes no sense to compute an expression for λ_l . \square

The next and final step involves comparing the equations for λ^* (from Lemma 3.8) and λ_l (from Lemma 3.9), such as to derive a condition for when $\lambda_l < \lambda^*$. We found that getting a closed form expression for this condition is not possible for these two equations. However, the framework developed above is sufficient to enable the seller to solve this condition numerically using a standard solver and, thus, choose the optimal level for the exercise price K_1 .

Note that all the analysis performed in this section (and, overall, in this paper) refers to using options when [at least] one of the buyers participating in the sequential auction market is a synergy buyer (and, thus, she has an exposure problem, as defined in the introduction). All the optimal price bounds for options given here refer to the case when options serve to relieve this exposure problem of a buyer with complementarities.

It is conceivable, however, that options might also prove useful in cases when agents do not have synergy valuations. For example, options could also be used in the case of substitutabilities (i.e. when agents have to choose between a set of items sold in sequence). However, the complementarity problem is arguably the hardest to address, and this is why we focus on it here, leaving the study of the usefulness of options in other cases to future work.

3.3 Numerical illustration of option pricing

In this section, before we provide the full experimental analysis of the model, we provide some details of the optimal pricing window (i.e. the interval for which $op' + \lambda_l \leq op' + \lambda^* \leq op' + \lambda_h$). To this end, we use a configuration similar to the settings used in the experiments reported in Section 4.

We consider a basic setting with $n = 2$ auctions, and a synergy bidder wanting both items. Her valuation for getting both of these items is v_{syn} . Now, in each of the 2 auctions the bidder faces a number of local bidders only interested in acquiring the item in that

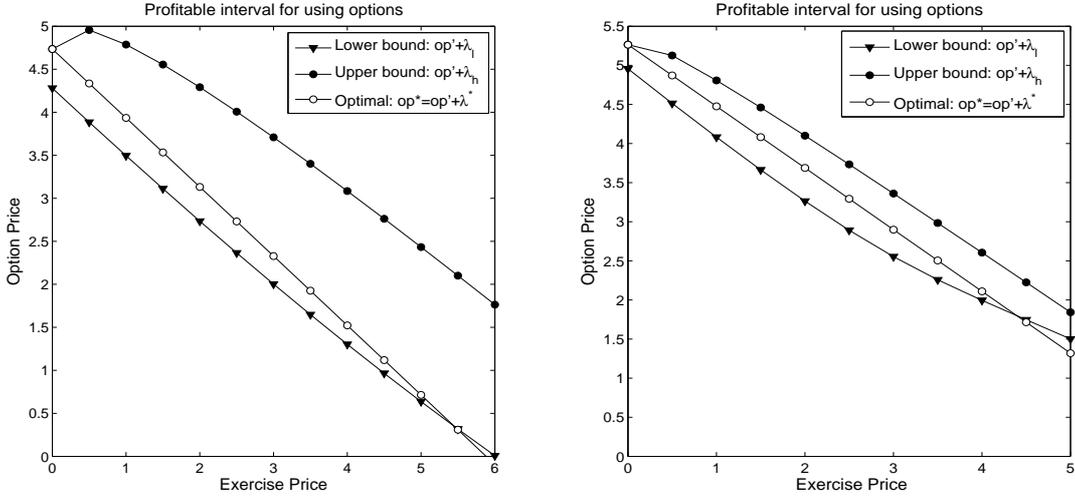


Fig. 2. Illustration of the size of the window for which $op' + \lambda_l \leq op' + \lambda^*$ (for which using priced options is beneficial for both seller and synergy bidder), in a sequence of 2 auctions and 2 cases: A (left). A synergy bidder with a valuation $v_{syn} = 24$ faces one local bidder in each auction, with valuations drawn at random from $N(\mu = 10, \sigma = 4)$. B (right). A synergy bidder with a valuation $v_{syn} = 28$ faces 5 local bidders in each auction, with valuations drawn at random from $N(\mu = 10, \sigma = 4)$.

auction, whose valuations are drawn from a normal distribution $N(10, 4)$. We consider two settings: in the first one, a synergy bidder with a valuation for 2 items of $v_{syn} = 24$ faces exactly one local bidder with valuation drawn from $N(10, 4)$ per auction. In the second setting, the valuation of the synergy bidder is $v_{syn} = 28$, but she faces 5 local bidders with valuations drawn from $N(10, 4)$ in each auction. Moreover, we note that we consider a seller that sets an exercise price K in advance, and the *same* for both auctions (i.e. $K_1 = K_2 = K$), where the value of K is varied on the abscissa. The optimal price intervals are illustrated in Figure 2.

Figure 2 illustrates that, for both configurations of values/competition setting/price expectations, there is an interval in which the seller can set the exercise price K , such that $op' + \lambda_l \leq op' + \lambda^*$. In these case, the increase in the bids of the synergy bidder compared to direct auctions (i.e. λ^*) is above the minimum threshold increase expected by the seller (i.e. λ_l), to compensate for the risk of remaining with the first item unsold. In our example, note that this interval is considerably narrower in the second case, due to the increased competition. In both cases, adding more local bidders per auction and increasing the mean of the valuation distributions have an effect of narrowing the “window” in which options are beneficial for both parties. Note that we do not claim this happens in every configuration, and there are many value settings in which it always holds that $op' + \lambda_l \geq op' + \lambda_h$, i.e. the window in which sellers have an incentive to offer options - either free or positively priced - may be empty. However, as we discuss in the next section, options can be beneficial for both buyer and seller in a wide variety of settings, and in such settings *both* seller and synergy buyers would benefit, in expectation, from using option contracts.

4. SIMULATION OF A MARKET WITH A SINGLE SYNERGY BUYER

This section presents an experimental examination of a market with one synergy buyer. It introduces the market entry effects in the synergy buyer's behaviour, as well as the threshold effects that may determine which exercise prices the seller chooses for her options. This experimental analysis is performed here for a market with one synergy bidder and several local bidders, while Sect. 5 considers a market with multiple synergy bidders.

The experimental setting is as follows: we consider a simulation where two goods A and B are auctioned n_A and n_B times respectively. The synergy buyer desires one copy of both goods and has zero valuation for the individual goods. That is, each synergy (or global) bidder requires exactly one bundle of $\{A, B\}$ ¹³. In the setting considered in this section, local bidders want only one good and participate in one auction, thus their bids can be modeled as a distribution.

Furthermore, in order to simplify the analysis of the model, we assume there is a single seller who auctions all the goods. This is actually equivalent to studying whether *on average* sellers have an incentive to use options. To explain, on any single sequence of auctions taken in isolation, the sellers of different items may have diverging incentives to use options, based on their position in the auction queue (as noted at the end of section 2.5, sellers with a later position in the auction queue may make more money). However, in a very large setting, where buyers enter the market randomly, it is difficult for any individual seller to strategize about her particular place in the sequence (and, furthermore, in most markets she may simply have no information to do this). Our goal is to study under which conditions, on average, sellers benefit from using options if there are synergy buyers in the market. Here, the average revenue can be also interpreted as the benefits of a typical seller, if her position in the sequential queue were chosen at random (which is realistic in large markets, with repeated interactions). Also, to somewhat reduce the number of test parameters, we further assume that the exercise price is the same for all goods of the same type. So the seller needs to determine which exercise price for A and which for B maximize her expected profit.

Note that, typically a seller has a resale value for the goods that remain unsold, which is usually lower than the value at the start of the auction sequence. The reason for this may be that there is some time discounting associated with waiting for a sequence of auctions to resell her items, or even a listing cost, which is paid per auction (such as in the eBay case). In this paper, we do not explicitly simulate resale, but we use a reservation value, which represents the expected resale value the seller expects to get, if she is forced to resell her items. To summarize, simulations were run in Matlab and had the following parameters:

Name	Description
n	The number of auctions.
$mean$	The mean of price distribution.
std	The standard deviation of price distribution.
res	Reserve prices.
v_{syn}	Valuation of synergy buyer for A and B combined.
k	Number of simulations for each auction run (i.e. how many times a sequence of auctions is repeated for one set of parameters).

¹³An intuitive way to think about this setting is as a sequential sale of individual shoes of exactly the same type, where A is the left shoe, and B is the right shoe, and each synergy buyer requires exactly one pair.

A basic simulation run is as follows. First, all possible auction sequences are determined for the given number of auctions for A and B. The simulation is then run for all these sequences, both for a direct sale setting and for a setting where the items are sold through options with given exercise prices.

For each auction, in each simulation run, there is a set of local bidders, who are assumed not to reason about the bids of the synergy buyer. The bids of these local bidders are therefore, assumed to follow a normal price distribution, with the parameters n , $mean$, std and res consisting out of two values: one for good A and one for good B. For each simulation run, the synergy bidders(s) are asked to determine their profit-maximizing bid for that setting, as described in the next section. The optimization required for determining their optimal bid is done using the Matlab function “fminsearch” from the Optimization Toolbox.

Since there may be considerable variance in the bids of the local bidders each possible auction sequence is run k times (typically, we had $k > 10000$). The average profit of the seller and the synergy buyer which are reported here, for both the case of with and without options, are averages over all these k simulations and also over all possible auction orders of items A and B in the sequence.

4.1 Synergy buyer's bid strategy

This section describes how the synergy buyer determines her bids in the simulation. In order to neutralize the effect that the exact order items are auctioned in plays on the bidding strategy, we add the assumption that the synergy buyer knows the number of remaining auctions, but not the order they will be held in. This remaining number of auctions of each type is common knowledge (i.e. the synergy bidders can always observe how many auctions of each type are left before they have to leave the market, and so does the seller).

The model described here is for a situation without options. But in order to apply it to a situation with options, one merely has to replace the variables: $b_t = op_t - K_t$ and $v_{syn}(A, B) := v_{syn}(A, B) - K_A - K_B$. As in the analytical section, we assume a bidder wants only a complete bundle of $\{A, B\}$. Therefore, $v_{syn}(A) = 0$, $v_{syn}(B) = 0$.

Determining the synergy buyer's profit-maximizing bid b_t^* at state t basically involves solving the Markov Decision Process (MDP), where we select the optimal bid b_t^* at time t , subject to the optimal bid b_{t+1}^* being selected for the future time point $t + 1$ (which in this case, is an auction). We can, however, use the valuation function of the bidding agent to significantly reduce the state space of the MDP, as shown below. However, first we introduce some notation.

Let b^* be the immediate best response to the state, which depends on four variables: z_A , z_B , X and I_t . The variables z_A and z_B are the number of remaining auctions for A and B respectively (including the current auction), so $z_A \leq n_A$, $z_B \leq n_B$. The type of good, which is currently sold, is denoted by I_t . The set of goods the synergy buyer owns (i.e. the endowment) is described by X , which can either be \emptyset , $\{A\}$ or $\{B\}$. If X is $\{A, B\}$ then the synergy buyer is done¹⁴. Let $Q(z_A, z_B, X, I_t, b_t)$ be the expected profit of

¹⁴Note that the experimental settings used in the model considered here preclude the possibility of the synergy buyer from acquiring more options than she needs to make up her desired bundle. But it is theoretically possible in our model, especially in settings with very low option prices compared to the synergy valuations, that the synergy buyer is incentivised to hoard options for more items than she really needs, and only choose to exercise some of these in the end. We leave the examination of such cases to further work.

the synergy buyer when bidding b_t . Note that, in these definitions, b_{t+1}^* and $V_{t+1}()$ denote the best available bid, respectively best expected value for the next state (as computed by recursion), while I_{t+1} is the type of the next item in the auction sequence. Therefore, using MDP notation, the profit-maximizing bid b_t^* is determined as follows:

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, z_B, X, I_t, b_t) \quad (12)$$

Where the expected profit is determined via:

$$Q(z_A, z_B, X, I_t = A, b_{t+1}^*) = F_A(b_t)(-b_t + V_{t+1}(z_A - 1, z_B, X \cup A, b_{t+1}^*)) + (1 - F_A(b_t))V_{t+1}(z_A - 1, z_B, X, b_{t+1}^*) \quad (13)$$

$$Q(z_A, z_B, X, I_t = B, b_t) = F_B(b_t)(-b_t + V_{t+1}(z_A, z_B - 1, X \cup B, b_{t+1}^*)) + (1 - F_B(b_t))V_{t+1}(z_A, z_B - 1, X, b_{t+1}^*) \quad (14)$$

Where $V()$ is the value of a state, which simply means the maximum expected profit of that state:

$$V_t(z_A, z_B, X, b_t) = \operatorname{max}_{b_t} Q(z_A, z_B, X, I_t, b_t) \quad (15)$$

Looking at the formula for $Q()$, it basically says that for the probability of winning the auction with her bid, the synergy buyer has to pay a price equal to her bid and the good is included in the endowment X of the next state. If she does not win the auction, then the value of the current state is equal to the value of the next state.

As we mentioned before, in computing its optimal bidding strategy used in the experimental section, we assume the synergy buyer does not know whether the next auction will be for A or B, she knows only the total numbers of auctions for A and B remaining. We acknowledge this is a departure from the formulas in the theoretical analysis, where the exact order of the auctions was taken into account to compute the bidding strategies. There are two reasons to use this assumption here. The first is that it reduces considerably the state space that needs to be modeled when computed the optimization. But the second is that we also find this choice more realistic if this model is to be applied to real-life settings. For example, when bidding on a part-truck order in a logistic scenario, it is more realistic to assume that a carrier can approximate the number of future opportunities to buy a complementary load, but not the exact auction order in which future loads will be offered for auction.

If we assume the synergy buyer only knows the total numbers of auctions for A and B remaining (and not their exact order), then her bidding strategy is based on assuming each future auction has an equal probability to occur. Therefore, the probability of an auction for A occurring next is simply the number of remaining auctions A divided by the total number of remaining auctions. Thus, a weighted average can be used to determine the value of the next auction, while not knowing for which good it will be for.

Apart from this general framework, we can prune the state space with the cases in which we know the synergy buyer's bid is zero:

$$b_t^* = \operatorname{argmax}_{b_t} Q(0, z_B, X, B, b_t) = 0, \text{ with } A \notin X \quad (16)$$

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, 0, X, A, b_t) = 0, \text{ with } B \notin X \quad (17)$$

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, z_B, X, I_t \in X, b_t) = 0 \quad (18)$$

With the first two cases, the synergy buyer can no longer obtain her desired bundle, because she does not own the complementary item and there is no chance left of acquiring

it. The last equation is for the case when the synergy buyer already has a copy of the type of good (and, from her valuation function, she only wants exactly one copy of A and B). The corresponding values of these states are:

$$V(0, z_B, X, b_t^*) = 0, \text{ if } A \notin X \quad (19)$$

$$V(z_A, 0, X, b_t^*) = 0, \text{ if } B \notin X \quad (20)$$

$$V(z_A, z_B, \{A\}, b_t^*) = V(0, z_B, \{A\}, b_t^*) \quad (21)$$

$$V(z_A, z_B, \{B\}, b_t^*) = V(z_A, 0, \{B\}, b_t^*) \quad (22)$$

The first two equations correspond to the case when the buyer can no longer get the complementary-valued item, therefore the sequence of auctions of the same type has no value to her. In both these cases $b_t^* = 0$. The last two equations are important, since they help the most to reduce the state space. Basically, as already mentioned, we assume that a synergy bidder only wants exactly one bundle of $\{A, B\}$. If she already owns a good of one of the two types, she will no longer be interested in the remaining auctions for that type of good. Therefore, the valuation $V()$ of these states is equivalent to a state when no auctions are remaining for the type of good she already owns (as she would not take part in those anyway). All these techniques help reduce the recursive search.

To conclude, to determine the synergy buyer's bids in any situation, the values of b_t^* and $V()$ need to be calculated for the following states:

$$\forall z_B > 0 \quad Q(0, z_B, \{A\}, B, b_t)$$

$$\forall z_A > 0 \quad Q(z_A, 0, \{B\}, A, b_t)$$

$$\forall z_A > 0, z_B > 0 \quad Q(z_A, z_B, \emptyset, A, b_t)$$

$$\forall z_A > 0, z_B > 0 \quad Q(z_A, z_B, \emptyset, B, b_t)$$

Note that, in general, solving for b_t^* involves solving a continuous MDP - except for some cases for which a closed form solution exists (e.g. the case of uniform distributions in Section 3.2). Basically, in the setting considered here with small sequences of auctions, we can treat solving for the optimal bids as a multi-variable optimization problem, which can be solved with standard optimization packages available in Matlab. In larger settings with more auctions, computing the solutions of this MDP may be considerably more involved, and may require additional computational techniques that have not been studied as part of this paper. We note, however, that solving continuous MDPs efficiently, while not trivial, is an active research area, and we provide sufficient details that the solutions developed there could be applied to our framework.

4.2 Experimental results: market entry effect for one synergy buyer

First, we study experimentally the incentives to use options for the sellers and buyers, in the case there is just one synergy bidder present in the market. In order to study different dimensions of such markets, we considered several combinations of parameter settings.

The first setting has $n_A = 2$ and $n_B = 2$. As mentioned above, the local bidders only bid in one local auction, without considering the bids placed by the synergy bidder. Therefore, their bids can be modeled as a distribution $\sim N(10, 4)$ for both goods. The goods A and B are, in this model, of equal rarity and attract an equal amount of independent competition during bidding. This choice is not random, as having a certain degree of symmetry in the experimental model allows us to reduce the number of parameter settings we need to

consider. More specifically, we assume the same exercise prices are set for both goods of type A and B. This is a reasonable assumption, because A and B are of symmetric value and because bidders do not know in advance the exact order goods will be sold in.

Furthermore, for each good, the seller has a reservation value $res = 8$, which gives its estimate resell value in the case the synergy buyer acquires an option for the item, but fails to exercise it. Since, on average, local bidders bid have an expected mean of 10 for an item, 20% is a reasonably safe estimate of a resell value.

The value of a bundle of $\{A,B\}$ for the synergy buyer is an important choice, especially in relation to the mean expectation μ of the bids placed by single-item bidders. We considered two settings: $v(A, B) = 24$ (thus 20% more, on average, than local competition) - with results shown in Fig. 3, and $v(A, B) = 21$ (which is only 5% more on average than local competition) - with results shown in Fig. 4.

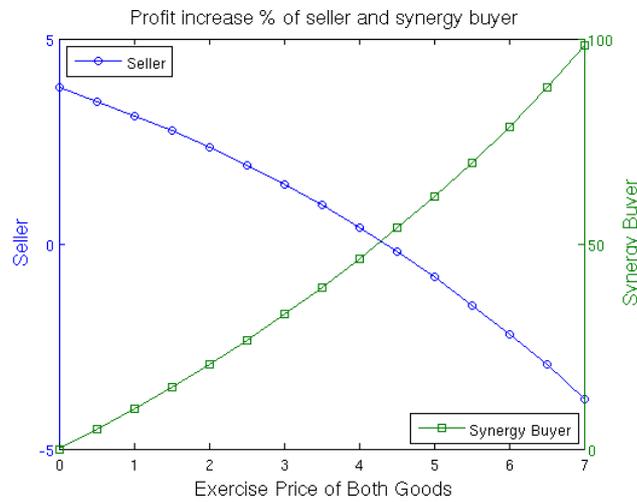


Fig. 3. Percentage increase in profit for a model using options with respect to direct sale, for the case there is one synergy buyer is present in the market. In the setting, there are two items of type A sold and two items of type B. For all 4 items, the bids of the local bidders follow the distribution $N(10, 4)$, while the valuation of the synergy buyer is $v(A, B) = 24$ (thus 20% more, on average, than the local bidders). What is varied on the horizontal axis is the exercise price with which the items are sold (assuming they are set the same for all items, being of equal rarity). Note that the figure is super-imposed: the left-hand side axis refers exclusively to the seller, while the right-hand side axis refers exclusively to the synergy bidder. From this picture, one can already see the important effect: synergy buyer prefers, on average, higher exercise prices, while seller prefers lower ones. Note that there is a sudden increase in profit, on the seller side, for the options case with $k = \epsilon > 0$, with respect to direct auctioning. This is simply because, with options, the seller gets to keep the item (for which it has a non-residual value), rather than the buyer, who disposes of it (as in the direct sale case).

Looking at these two figures, some important effect can be observed. First, we mention that the seller has an immediately higher expected profit with options compared to direct sale. This is because an option is sometimes not exercised and then the seller gets to keep the good (for which she has a positive valuation), while the synergy buyer still pays the option price.

There are two main effects to be observed from Fig. 3 and 4:

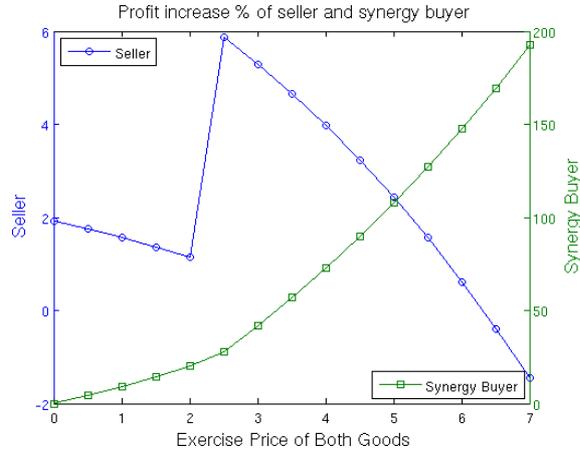


Fig. 4. Percentage increase in profit for a model using options with respect to direct sale, for the case there is one synergy buyer is present in the market. The settings are exactly the same as those is in Fig. 3 above: 2 auctions for A and 2 for B, with local bidders following $N(10, 4)$. However, now the valuation of the synergy buyer is $v(A, B) = 21$ (thus only 5% more, on average, than the local bidders). One can see, however, that there is an important difference by comparison to Fig. 3: the threshold effect in the profit increase for the seller when the exercise price $K \geq 2.5$. Intuitively, the reason this effect occurs is the market-entry effect on the part of the synergy buyer, who would otherwise stay out for this lower valuation



Fig. 5. Percentage increase in profit for the case of one synergy buyer, for longer auction sequences. The settings in terms of valuations are exactly the same as those is in Fig. 4 above: the synergy buyer has a value $v(A, B) = 21$, while single-item bidders bid according to $N(10, 4)$. One change is that now there are 4 auctions available for each type, i.e. 4 auctions for an item of type A and 4 for B. Notice that now there are multiple thresholds, since there are multiple points when the market entry effect of the synergy buyers appears. However, on average, the percentage increases in expected profits for the synergy buyers are lower, when compared to the direct auctions case. The reason for this is that, with multiple future buying opportunities, the exposure problems that synergy bidder faces decreases.

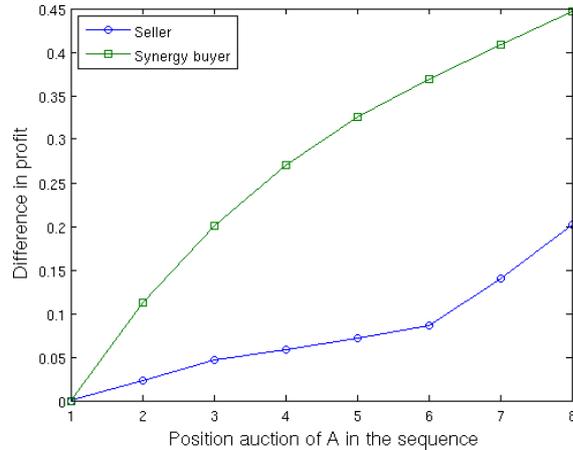


Fig. 6. Influence of the position in an auction queue of an item on the seller's expected profit. Settings are the same as in Fig. 3, but with one important difference: the rarity of the goods is no longer symmetric. There is now only 1 auction for a good of type A, but 7 auctions for a good of type B. What is varied along the horizontal axis is the position in the auction queue of the sale of the rarer item (of type A). The graph shows the absolute difference in profit for a seller of an item of type B and for the synergy buyer (i.e. the difference in profit between an options and direct auctions model). Note that, if the rare item of type A is sold at the end of the auction sequence, the benefit of selling item B through an option increases, because the exposure risk of not acquiring item of type A increases.

- First, the synergy buyer in such a market always prefers *higher* exercise prices (an effect clearly seen in both Figs. 3 and 4). This may be counter-intuitive at first, but is a rational expectation. If the option for an item is sold with a higher exercise price, then the synergy buyer can bid more aggressively on the option price to get the item, since she is “covered” for the loss represented by the exercise price. The local bidders extract no advantage from being offered the good as an options vs. a direct sale, because, if they acquire the option, they would always exercise it regardless. Therefore, they will simply lower their bid for the option with the amount represented by the exercise price.
- Second, the expected profit of the seller seems to decrease between intervals if she has to sell the option with a higher exercise price. The main reason for this is that there is some chance that she or she would remain with her item unsold (because the option is not exercised), and thus extract only her reservation value for that item. There is, however, an important difference between the cases shown in Fig. 3 and 4, which is the participation thresholds (that appear as “peaks” in the picture), where the expected profit of the seller seems to “jump” at a new level. These can be explained by the synergy buyer joining the market, as the expected profit becomes non-negative. The threshold nature is determined by the discrete nature of the auction sequence, as is explained below.

Such a participation threshold is illustrated in Fig. 4 is the increase in the seller's expected profit when the exercise price is set above a certain level ($K \geq 2.5$, for the settings in Fig. 4). Such thresholds can be explained as follows. If the synergy buyer currently owns nothing, then she will only bid on a good if the number of remaining auctions and their exercise prices give her a prior expectation of a positive profit. Conversely, if the synergy buyer is not offered a sequence of option sales from which she derives a positive

expected profit, she has the incentive to leave the market altogether. There are two main factors that increase a synergy buyer's expected profit in a sequence of auctions (sold as options):

- The number of remaining future auctions of the other good, necessary to complete her bundle.
- The exercise price of the options (that only needs to be paid at the end). This should be high enough to cover the risk, given her valuation for the bundle.

Note that in some market setting (such as the one in Fig. 3), no participation effects (i.e. thresholds) occur, because the value the synergy buyer assigns to her desired bundle is already high enough, so she would participate in the market anyway (i.e. regardless of whether she gets offered options or not), and at any point in the sequence that there is still a chance of completing her bundle.

However, in the valuation settings in Fig. 4, the synergy buyer will only bid on a good if there are two remaining auctions for the other good. So she places a bid for A if the auctions are $[A, B, B]$, but not if they are $[A, B]$. This is because with a single auction for B, the risk of ending up with only a worthless A is too great. But in a market with exercise prices of at least 2.5, the risk is reduced and one remaining auction is already enough for the synergy buyer to stay in the market. So a higher exercise price enables the synergy buyer to stay the market, even if she owns nothing and there are only a few auctions left, which increases the seller's expected profit. This increase in participation is beneficial to the seller, who thus has an incentive to fix the exercise prices $K_A = K_B = 2.5$.

4.3 Settings with longer sequences of auctions and effect of auction order

In the previous section, we examined a sequence of auctions of a specific length of $n_A = 2, n_B = 2$. We now look at whether we can observe similar effects in the case when the number of opportunities to buy goods A and B increases. With the exception of auction lengths, the parameters are kept the same as in the previous case. First, we keep the relative rarity of both goods symmetrical, but increase the number of auctions available for each to 4, i.e. $n_A = n_B = 4$. Results are shown in Fig. 5.

Basically, there are two main effects to observe here. First, the benefits to the buyer of having options mechanism decreases (seen from comparing the percentage increases shown in the right-hand vertical axis of Figs. 4 and 5). The reason for this is that, in sequential auctions, the number of available future opportunities plays a big role in how big the exposure problem the synergy buyer faces is. If there is less exposure, then the relative benefits of using options becomes smaller (although it is still quite considerable). The second effect to be observed from Fig. 5 is that there are more participation thresholds (denoted by peaks), but they are smaller. The reason is that, for a longer sequence of auctions, there are more possible sequences of remaining auction combinations. The synergy bidder will join in the bidding in some, but not in others, leading to multiple participation thresholds.

The second problem we look in this subsection at is what happens if the relative frequency of the two goods is more asymmetric. We keep the same total number of auctions in the sequence (8), but the relative frequency is highly asymmetric: $n_A = 1, n_B = 7$. As mentioned, in the previous graphs, results were averaged over all possible auction orders - while here, by contrast, we look at auction orders one by one.

For this setting, there are exactly 8 possible auction orders, corresponding to the point where the rarer good (type A) can be inserted in the auction queue. What is varied on the horizontal axis is this position of the type A good. The reason why we look at whether a seller of items of type B would use options is that the exposure of the synergy buyer exists for the other good in the sequence. For the single item of type A, the benefits of using options are limited, because the synergy buyer has 7 other auctions in which to acquire the second item anyway, hence she has much less of an exposure problem.

Clearly, we can see an important effect of the position of the rarer good in the auction queue, from the perspective of both parties. If the item of type A is sold at the very beginning of the auction sequence, then the synergy bidder has no exposure problem left for the rest of the sequence, hence there is no incentive to use options, for either party. However, it is at the very end of the auction sequence, the synergy buyer will not know whether she would need the item acquired until all auctions end. For this case, the benefits of using options are considerably greater.

5. MULTIPLE SYNERGY BUYERS

Finally, we consider market settings in which multiple synergy buyers are active simultaneously. Much of the experimental set-up and parameter choices are the same as described in the above sections, for the case of one for the single synergy buyer. The only difference is that now multiple synergy buyers may enter and leave the market at different times and they have different valuations for the combination of A and B.

We have to emphasize that the results from this section are still rather preliminary and are based on some restrictions on the reasoning capability of the synergy buyers in the market. Specifically, as in the single-bidder case, we assume the synergy bidders have some prior expectations about the closing prices in future auctions and compute their optimal strategy with respect to this expectation. In these results, this expectation is assumed the same for all synergy bidders, which is a reasonable choice in comparing their strategies, but assuming the sequence of auctions considered is too short for other synergy buyers to learn about existing competition and adapt their bids. In a more realistic market, however, synergy bidders could be expected to be able to learn and adjust their expectations based on past interactions, as well as reason game-theoretically about the fact that another synergy bidder may present in the market at the same time. At this point, these more sophisticated forms of reasoning are left to future work.

As in the previous section all simulations of this section have reserve prices of 8 and local bidders following $\sim N(10, 2.5)$. The first two experiments also have two synergy buyers syn_1 and syn_2 with valuations for both goods of 21.5 and 22.5 respectively. The order the synergy bidders enter the market (and the number of auctions they can stay in) are given in Figs. 9 and 10, while results for all settings are shown in Fig. 7, respectively 8. In the following, we will discuss these in separate subsections.

5.1 Two synergy buyers interacting indirectly through the exercise price level

In the setting examined here, the two synergy buyers each have $n_A = 3$ and $n_B = 3$, without the other agent participating in these auctions. An example of such an auction sequence is shown in Fig. 9. However, these two synergy bidders do interact indirectly as follows. Since options are sold through open auctions based on the option price, the seller has to fix the exercise prices for the whole market (i.e. for all auctions in the sequence). So while synergy buyers may not participate in the same auctions, their presence does

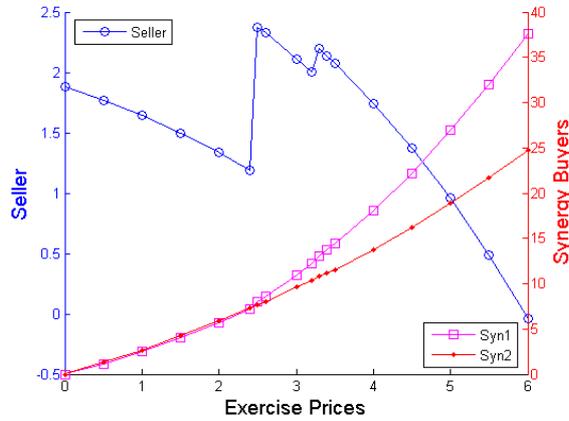


Fig. 7. Percentage increase in profits for a market with with 2 synergy bidders. There are 3 auctions for A and 3 for B, and for each one the bids from the competition formed by local bidders follows the distribution $N(10, 2.5)$. The valuations of the two synergy bidders for a bundle $\{A,B\}$ are 21.1 for *syn1*, respectively 22.5 for *syn2*. The order the agents enter the market is described by Fig. 9 below (so the two agents do not compete directly against each other in this setting). Notice that, in this case, the average profit of *syn2* does not decrease with the entry of *syn1* in the market.

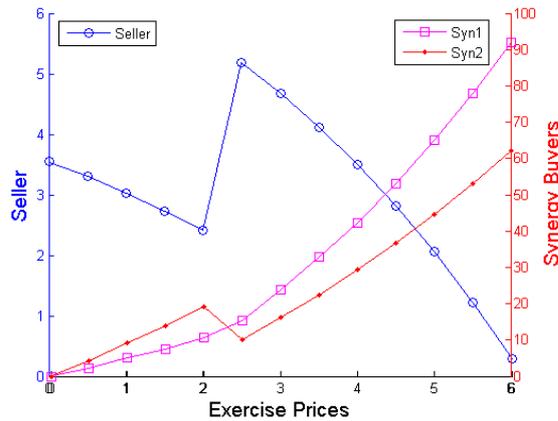


Fig. 8. Percentage increase in profits for a market with with 2 synergy bidders. The setting and valuations are the same as in Fig. 7 above. However, the order the agents enter the market is now described by Fig. 10 below (so the two agents *do* compete directly for the same goods). Notice that, in this case, the average profit of *syn2* decreases due to the additional competition from *syn1*.

influence the competition through the exercise prices set by the seller.

This effect can be seen in Fig. 7, in which the seller maximizes her expected profit at $K = K_A = K_B = 2.4$. In this case *syn2* is better off, because without the presence of *syn1* she would be offered options with lower exercise prices. But *syn1* is worse off, because if she were alone in the market the seller would choose $K = 3.2$, which gives her a higher expected profit. Yet, due to *syn2*, the seller sets $K = 2.4$. In this case, due to the

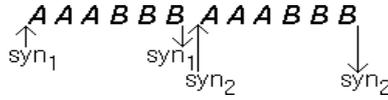


Fig. 9. An auction sequence for the case shown in Fig 7.

seller's choice of exercise prices, one synergy buyer (syn_1) gains, while syn_2 loses.

5.2 Direct synergy buyer competition in the same market

Next, we considered a setting in which synergy buyers compete directly for some of the goods. The entry points for such a setting are shown in Fig. 10, while simulation results are given in Fig. 8.

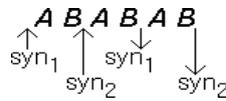


Fig. 10. An auction sequence for the case shown in Fig. 8.

As can be seen in Figure 8, the profit of syn_2 drops at 2.5. In previous figures the synergy buyers' profits were monotonically increasing in the exercise prices, because they then have a smaller loss when they fail to complete their bundle. But now this effect cannot immediately compensate the extra competition coming from syn_1 , who participates in the same auctions more often after this threshold at 2.5. So, in this case, both synergy buyers lose from the presence of additional bidders. While one synergy buyer (i.e. syn_2) should benefit because she is offered better (higher) exercise prices than if she were alone in the market, this effect cannot immediately compensate the additional competition.

5.3 Larger simulation with random synergy buyers' market entry

In the final results we report in this paper, we conducted a larger scale simulation with multiple synergy buyers, which can enter the market randomly, with a certain probability.

The experimental setup implies that each sequence of auctions (forming a test case) has 10 items of each type (i.e. $n_A = 10$ and $n_B = 10$). What differs from previous settings is the random entry of synergy buyers. For each auction, there is a 25% chance that a synergy buyer will enter the market. If she does, then her valuation is drawn from a uniform distribution between 20 and 22 and she will stay in the market for exactly four auctions. To simplify matters, the auction sequence is fixed at first selling A, then B, then A etc. so that each synergy buyer will face exactly two auctions for an item of type A and two for an item of type B. However, the general result of this section is also true for a random auction sequence, since the basic effects remain the same.

As shown in Figure 11, the seller's profit now only has one maximum at 5, because initially each increase in exercise prices causes, with some probability, a synergy buyer to participate more often. So each point is a threshold and the profit graph smooths out over those many local maxima, corresponding to a steady increase (on average) of the expected profit. This result shows why it can be rational for the seller to have the same exercise prices for all goods of the same type (e.g. the same K_A). In a market with random

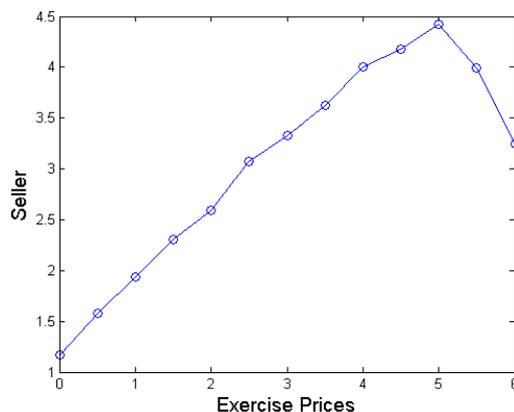


Fig. 11. Percentage increase in seller's profits in a larger experimental setting, with synergy buyers randomly entering the market.

entry of synergy buyers, the seller does not know which buyers are participating in any particular auction. Her optimal policy is to set her exercise prices which maximize her overall expected profit (in this case, $K = 5$).

6. DISCUSSION AND FURTHER WORK

This paper examined, from a decision-theoretic perspective, the use of priced options as a solution to the exposure problem in sequential auctions. We consider a model in which the seller is free to fix the exercise price for options on the goods she has to offer, and then sell these options in the open market, through a first price auction mechanism.

For this setting, we derived analytically, for a market with a synergy buyer and under some assumptions, expressions that provide the bounds on the option prices between which both synergy buyers and sellers have an incentive to use option contracts over direct auctions. Next, we performed an experimental analysis of several settings, where either one or multiple synergy bidders are active simultaneously in the market. We show that, if the exercise price is chosen appropriately, selling items through priced options rather than directly can increase the expected profits of both parties.

The overall conclusion of our study is that the proposed priced options mechanism can considerably reduce the exposure problem that synergy bidders face when taking part in sequential auctions. Furthermore, and most important, *both* parties in the market have an incentive to prefer and use such a mechanism. We show that in many realistic market scenarios, sellers can fix the exercise prices at a level that both provides sufficient incentive for buyers to take part in the auctions, as well as cover their risk of remaining with the items unsold.

It is important to note, however, that sequential auction allocation is a highly complex and still under-researched area, for which few exact analytical solutions are known to exist. To our knowledge, this study provides a first decision-theoretic analysis for the use of priced options to solve this problem. The analysis and results for the several fundamental cases studied here can serve as a basis for future work in more complex and realistic settings. These include more complex market scenarios, as well as more sophisticated rea-

soning abilities on the part of participating synergy bidders and sellers. For example, in a large market, synergy bidders could be expected to use learning strategies to adapt to changing market conditions, as well as the presence of other synergy bidders who want similar item combinations. However, the sellers of the items could also use learning to choose better levels of the exercise prices K with which to sell the options for their goods.

Other possible issues open to future research include: markets where bidders have imperfect or asymmetric information about other participants, more complex preferences over bundles, or different attitudes to risk. In order to study markets involving a variety of such heterogeneous agents, a promising approach may be to use evolutionary game theory techniques. Such an approach has already been considered for continuous double auctions (CDAs) by [Cai et al. 2007], but to our knowledge this has not been attempted before for sequential auctions with complementarities.

To conclude, sequential auction bidding with complementary valuations is a problem that appears in many real-life settings, although no dominant strategies exist and bidders face a severe exposure problem. The main intuition of this work is that a simple options mechanism, where sellers auction options for their goods (with a pre-set exercise price), instead of the goods themselves can go a long way in solving the exposure problem, and can be beneficial to both sides of such a market.

In practical terms, the potential impact of having a working solution to the exposure problem in sequential auctions is considerable. One example, which was used to illustrate some aspects of the model in this paper is decentralised transportation logistics [Robu et al. 2008; Robu et al. 2011], where loads appear sequentially, over time, and a bidding agent has to acquire a combination of these to fill her transportation capacity (i.e. truck). In decentralised electricity markets, much of the available electricity supply (especially that generated by renewable sources, such as wind or solar energy) comes online with a certain probability. In allocating this intermittent, “green” electricity through an electronic market, options could be a promising solution to deal with the inherent uncertainty. Other potential applications include retail electronic commerce (such as those discussed in [Juda and Parkes 2006] or keyword markets in sponsored search [Jordan et al. 2010; Borgs et al. 2007; Robu et al. 2009]). In our future work, we plan to explore the application of priced option mechanisms to some of these areas.

Acknowledgements

The authors would like to thank Enrico Gerding and Ioannis Vetsikas (University of Southampton, UK) and David Parkes (Harvard University) for many useful discussions. Moreover, we would like thank the editor and reviewers of ACM TOIT for their many useful suggestions for improving both the presentation and technical content of the paper.

REFERENCES

- AMRAM, M. AND KULATILAKA, N. 1998. *Real Options: Managing Strategic Investment in an Uncertain World*. Oxford University Press.
- BAJARI, P. AND HORTACSU, A. 2003. The winner’s curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *The RAND Journal of Economics* 34, 329–155.
- BORGS, C., CHAYES, J., IMMORLICA, N., JAIN, K., ETESAMI, O., AND MAHDIAN, M. 2007. Dynamics of bid optimization in online advertisement auctions. In *WWW ’07: Proceedings of the 16th international conference on World Wide Web*. ACM Press, 531–540.
- BOUTILIER, C., GOLDSZMIDT, M., AND SABATA, B. 1999. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI’99)*. 527–523.

- CAI, K., NIU, J., AND PARSONS, S. 2007. Using evolutionary game-theory to analyse the performance of trading strategies in a continuous double auction market. In *Adaptive Agents and Multi-Agent Systems III*. Springer LNCS, volume 4865, 44–59.
- CRAMTON, P., SHOHAM, Y., AND STEINBERG, R. 2006. *Combinatorial auctions*. MIT Press.
- CZAJKOWSKI, K., FOSTER, I., AND KESSELMAN, C. 1999. Resource co-allocation in computational grids. In *Proceedings of the 8th IEEE Symposium on High Performance Distributed Computing (HPDC-8)*. 219–228.
- ELYAKIME, B., LAFFONT, J.-J., LOISEL, P., AND VUONG, Q. 1994. First-price sealed-bid auctions with secret reservation prices. *Annales d'Economie et Statistique* 34, 115–141.
- FATIMA, S. 2006. A comparative study of sequential and simultaneous auctions. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-06)*. 1205–1207.
- FRIEDMAN, E. J. AND PARKES, D. C. 2003. Pricing wifi at starbucks: issues in online mechanism design. In *Proceedings the 4th ACM conference on Electronic Commerce (EC'03)*. ACM, 240–241.
- GERDING, E., ROBU, V., STEIN, S., PARKES, D., ROGERS, A., AND JENNINGS, N. 2011. Online mechanism design for electric vehicle charging. In *Proceedings of the 10th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'11)*. 811–818.
- GOPAL, R., THOMPSON, S., TUNG, Y. A., AND WHINSTON, A. B. 2005. Managing risks in multiple online auctions: An Options approach. *Decision Sciences* 36(3), 397–425.
- GREENWALD, A. AND BOYAN, J. 2004. Bidding under uncertainty: theory and experiments. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI'04)*. AAAI Press, 209–216.
- HULL, J. C. 2003. *Options, Futures, and Other Derivatives*, Fifth ed. Prentice Hall.
- JORDAN, P. R., WELLMAN, M. P., AND BALAKRISHNAN, G. 2010. Strategy and mechanism lessons from the first ad auctions trading agent competition. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*. 287–296.
- JUDA, A. AND PARKES, D. 2009. An options-based solution to the sequential auction problem. *Artificial Intelligence* 173(7-8), 876–899.
- JUDA, A. I. AND PARKES, D. C. 2006. The sequential auction problem on ebay: An empirical analysis and a solution. In *Proceedings of the 7th ACM Conference on Electronic commerce*. ACM Press, 180–189.
- KLEMPERER, P. 1999. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13, 3, 227–286.
- MOUS, L., ROBU, V., AND POUTRÉ, H. L. 2008. Can priced options solve the exposure problem in sequential auctions? In *ACM SIGEcom Exchanges* 7(2). ACM Press.
- MOUS, L., ROBU, V., AND POUTRÉ, H. L. 2010. Using priced options to solve the exposure problem in sequential auctions. In *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*. Lecture Notes in Business Information Processing, vol. 44. Springer Berlin Heidelberg, 29–45.
- OSEPAYSHVILI, A., WELLMAN, M., REEVES, D., AND MACKIE-MASON, J. 2005. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Proceedings of Uncertainty in Artificial Intelligence (UAI'05)*. 441–449.
- PARKES, D. C. 2007. *Online Mechanisms*. Chapter in Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Eds.), “Algorithmic Game Theory”, Cambridge University Press.
- REEVES, D., WELLMAN, M., MACKIE-MASON, J. K., AND OSEPAYSHVILI, A. 2005. Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39, 1, 67–85.
- ROBU, V. AND LA POUTRÉ, H. 2007. Designing bidding strategies in sequential auctions for risk averse agents. In *Agent-Mediated Electronic Commerce (AMEC'07)*. Vol. 13. Springer LNBIP, 76–89.
- ROBU, V. AND LA POUTRÉ, H. 2010. Designing bidding strategies in sequential auctions for risk averse agents. *Multiagent and Grid Systems* 6, 5–6, 437–457.
- ROBU, V., LA POUTRÉ, H., AND BOHTE, S. 2009. The complex dynamics of sponsored search markets. In *Proceedings of the 4th International Workshop on Agents and Data Mining Interaction (ADMI'09)*. Springer LNAI volume 5680, 183–198.
- ROBU, V., NOOT, H., LA POUTRÉ, J. A., AND VAN SCHIJNDEL, W. 2011. A multi-agent platform for auction-based allocation of loads in transportation logistics. *Expert Systems with Applications* 38, 4, 3483–3491.
- ROBU, V., NOOT, H., POUTRÉ, H. L., AND VAN SCHIJNDEL, W.-J. 2008. An agent platform for auction-based allocation of loads in transportation logistics. In *Proceedings of the 7th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'08), Industry Track*. IFAAMAS Press, 3–10.

- ROBU, V., STEIN, S., GERDING, E., PARKES, D., ROGERS, A., AND JENNINGS, N. 2011. An online mechanism for multi-speed electric vehicle charging. In *Proceedings of the The Second Conference on Auctions, Market Mechanisms and Their Applications (AMMA'11)*. Springer LNICST (to appear).
- ROBU, V., VETSIKAS, I. A., GERDING, E. H., AND JENNINGS, N. R. 2010a. Addressing the exposure problem of bidding agents using flexibly priced options. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*. IOS Press, 581–586.
- ROBU, V., VETSIKAS, I. A., GERDING, E. H., AND JENNINGS, N. R. 2010b. Flexibly priced options: A new mechanism for sequential auctions with complementary goods. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. IFAAMAS, 1485–1486.
- SANDHOLM, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135, 1-2, 1–54.
- SANDHOLM, T. AND LESSER, V. 2001. Leveled commitment contracts and strategic breach. *Games and Economic Behavior* 35, 1-2, 212–270.
- SANDHOLM, T. AND LESSER, V. 2002. Leveled-commitment contracting: a backtracking instrument for multi-agent systems. *AI Magazine* 23, 3, 89–100.
- SMITH, J. E. AND MCCARDLE, K. F. 1999. Options in the real world: Some lessons learned in evaluating oil and gas investments. *Operations Research* 47(1), 1–15.
- STEIN, S., PAYNE, T. R., AND JENNINGS, N. R. 2009. Flexible provisioning of web service workflows. *ACM Transactions on Internet Technology* 9(1), 2–45.
- 'T HOEN, P., ROBU, V., AND LA POUTRÉ, J. A. 2005. Decommitment in a competitive multi-agent transportation setting. In *Software Agent-Based Applications, Platforms and Development Kits*, R. Unland, M. Klusch, and M. Calisti, Eds. Whitestein Series in Software Agent Technologies. Springer, 409–431.
- VETSIKAS, I. A. AND JENNINGS, N. R. 2008. Bidding strategies for realistic multi-unit sealed-bid auctions. In *Proceedings of the 23rd Conference on Artificial Intelligence (AAAI'08)*. 182–189.
- WELLMAN, M., GREENWALD, A., AND STONE, P. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press.
- WELLMAN, M. P., OSEPAYSHVILI, A. V., MACKIE-MASON, J. K., AND REEVES, D. M. 2008. Bidding strategies for simultaneous ascending auctions. *B.E. Journal of Theoretical Economics* 8(1).