Flexibly Priced Options: A New Mechanism for Sequential Auctions with Complementary Goods

Valentin Robu, Ioannis A. Vetsikas, Enrico H. Gerding, and Nicholas R. Jennings

University of Southampton
School of Electronics and Computer Science
Highfield Campus, Southampton, UK
{vr2, iv, eg, nrj}@ecs.soton.ac.uk

Abstract. In this paper we present a novel option pricing mechanism for reducing the exposure problem encountered by bidders with complementary valuations when participating in sequential, second-price auction markets. Existing option pricing models have two main drawbacks: they either apply a fixed exercise price, which may deter bidders with low valuations, thereby decreasing allocative efficiency, or options are offered for free, in which case bidders are less likely to exercise them, thereby reducing seller revenues. Our novel mechanism with *flexibly priced options* addresses these problems by calculating the exercise price as well as the option price based on the bids in an auction. For this novel setting we derive the optimal strategies for a bidding agent with complementary preferences. Furthermore, to compare our approach to existing ones, we derive, for the first time, the bidding strategies for a fixed price mechanism, in which exercise prices for options are fixed by the seller. Finally, we use these strategies to empirically evaluate the proposed option mechanism and compare it to existing ones, both in terms of the seller revenue and the social welfare. We show that our new mechanism achieves higher market efficiency, while still ensuring higher revenues for the seller than direct sale auctions (without options).

1 Introduction

Auctions are an efficient method for allocating resources or tasks between self-interested agents and, as a result, have been an important area of research in the multi-agent community. In recent years, research has focused on complex auctions where agents have combinatorial preferences and are interested in purchasing bundles of resources. Most of the solutions designed to address this problem involve one-shot, combinatorial auctions, where all parties declare their preferences to a center, which then computes the optimal allocation and corresponding payments [2]. Although such auctions have many desirable properties, such as incentive compatibility or efficiency of the allocation, in practice many settings are inherently decentralized and sequential. Often, the resources to be allocated are offered by different sellers, sometimes in different markets, or in auctions with different closing times. In such settings, however, a buyer of a bundle is faced with the so-called *exposure* problem when it has purchased a number of items from the bundle, but is unable to obtain the remaining items. In order to address this important problem, this paper proposes a novel option pricing mechanism, and shows its superiority over existing approaches.

In detail, sequential auctions form an important part of many application settings in which agents could be realistically applied, and such settings cannot usually be mapped to one-shot, combinatorial mechanisms. Examples include inter-related items sold on eBay by different sellers in auctions with different closing times [6], or decentralised transportation logistics, in which part-capacity loads are offered throughout a day by different shippers [10]. In these examples, a buyer is often faced with complementary preferences, i.e., where bundles of resources have more value than the sum of individual resources through synergies between these goods. In sequential auctions, this can result in the *exposure problem*. This occurs whenever an agent is faced with placing a bid for an item which is higher than its immediate marginal value, in the expectation of obtaining extra value through a synergy with another item sold later. However, if she fails to get the other item for a profitable price, she risks making a loss. In this paper, we refer to such a bidder as a *synergy bidder*. Due to the exposure problem, synergy bidders will shade their bids when entering a sequential auction market, bidding less than their bundle valuation for the items they desire [1, 3]. This reduces not only their own expected profits, but also seller revenues and the allocative efficiency of the market.

Although the exposure problem is well known, it has mostly been studied from the perspective of designing efficient bidding strategies that would help agents act in such market settings (e.g. [1, 3, 9, 11]). In this paper, we consider a different approach, that preserves the sequential nature of the allocation problem, and we propose a novel mechanism which involves auctioning *options* for the goods, instead of the goods themselves. An option is essentially a contract between the buyer and the seller of a good, where (1) the writer or seller of the option has the *obligation* to sell the good for the *exercise price*, but not the right, and (2) the buyer of the option has the *right* to buy the good for a pre-agreed *exercise price*, but not the obligation. Since the buyer gains the right to choose in the future whether or not she wants to buy the good, she pays for this right through an *option price* which she has to pay in advance, regardless of whether she chooses to exercise the option or not.

Options can help a synergy buyer reduce the exposure problem she faces since, even though she has to pay the option price, if she fails to complete her desired bundle, she does not have to pay the exercise price as well and thereby is able to limit her loss. 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 otherwise have stayed away, due to the risk of exposure to a potential loss.

1.1 Related Work

Options have a long history of research in finance (see [4] for an overview). However, the underlying assumption for all financial option pricing models is their dependence on an underlying asset, which has a known, public value that moves independently of the actions of individual agents This type of assumption does not hold for the online, sequential auctions setting we consider, as each individual synergy buyer has its own, private value for the goods/bundles on offer.

When applying options to the problem of reducing the exposure problem in sequential auctions, previous literature considers two main types of pricing mechanisms. One approach, proposed by Juda & Parkes [5, 6]) is to offer options for free (i.e., without an option price or an advance payment), and then let the exercise price be determined by the submitted bids in the market. However, this approach enables self-interested agents to hoard those options, even if they are highly unlikely to exercise them, thus considerably reducing both the allocative efficiency of the market and seller revenue. The second main approach (proposed by Mous et. al. [8]) is to have a fixed exercise price set by the seller, and then have the market determine the option price through an open auction. In this case, however, this preset exercise price can be perceived as a reserve value, since no bidder with a valuation below that price has an incentive to participate. This negatively effects the market efficiency, and may also affect the seller's profits by excluding some bidders from the market.

1.2 Contributions

To address the shortcomings of existing option models, in this paper we introduce a novel pricing mechanism in which the exercise price, as well as the option price are determined by the open market, and we compare this model to existing options models. In more detail, we extend the state-of-the-art in the following ways:

– To compare our new approach with existing ones, we derive, for the first time, the optimal bidding strategy for a synergy bidder in an options model with fixed exercise price (Section 3).
– We introduce a novel option pricing mechanism where the exercise price, as well as the options price are determined by the bids in an auction, and we show that both direct auctions (without options), and offering free options to the bidders appear in this model as particular subcases. Furthermore, we derive the optimal bidding strategy for the synergy bidder in this new model (Section 4).
– We empirically compare our new pricing model to existing option models from the literature, as well as to using direct auctions. We show that our flexible options approach achieves much better market allocation efficiency (measured in terms of the social welfare of all participating agents) than the state of the art fixed price options model. Furthermore, we show that sellers do not stand to lose any revenue by using this option model, by comparison to auctioning their items directly, without using options (Section 5).

The remainder of the paper is structured as follows. Section 2 formally defines our sequential auction setting, while Section 3 presents the bidding solution for the option model using fixed exercise price options. Section 4 introduces our new, flexible options approach, together with the optimal bidding strategies for this case. Section 5 empirically compares the two approaches, and Section 6 concludes.

2 The Problem Setting

In this section we formally describe the auction setting and introduce the notation used. We consider a setting with m second-price, sealed-bid auctions, each selling an option to buy a single item. We choose these auctions because bidders without synergies have a simple, dominant bidding strategy and, furthermore, they are strategically equivalent to the widely-used English auction. We assume that there exists a single *synergy* bidder who is interested in purchasing all of the items and receives a value of v if it succeeds, and 0 otherwise. Furthermore, every auction $j \in \{1, \ldots, m\}$ has N_j *local* bidders. These bidders only participate in their own auction, and are only interested in acquiring a single item. The values of this item for local bidders in auction j are i.i.d. drawn from a cumulative distribution function F_j. Finally, we assume that all bidders are expected utility maximisers.

Given this setting, we are interested in finding the *Bayes-Nash* equilibrium strategies for all of the bidders for different option pricing mechanisms.[1] However, even with options, due to the second-price auction, the local bidders have a dominant bidding strategy. Therefore, the main problem is finding the optimal strategy for the synergy bidder and this is largely decision-theoretic in nature.

We furthermore note that, although we focus largely on a single participating synergy bidder when presenting the strategies and results, this analysis can be easily extended to multiple synergy bidders. This is because synergy bidders are assumed to only be interested in either winning all of the auctions or none at all, and therefore, after the first auction all but one synergy bidder (with the highest valuation) will leave the market. Therefore, having multiple synergy bidders only affects the bid distribution in the first auction. We address this setting in more detail in Section 4.1.

3 Optimal Bidding Strategies for Fixed Exercise Price Options

To compare our new approach with existing option pricing mechanisms, we first derive the optimal bidding strategies for a synergy bidder in a fixed exercise price setting, where the exercise price for the options to be acquired is set by the seller before the start of the auctions.[2] While the different exercise prices for the auctions are fixed in advance, the option prices are determined by the second-highest bid in the auction. In the following, we let $\overrightarrow{K}$ denote the vector of fixed exercise prices, where K_j is the exercise price of the j^{th} auction, i.e. the price that the winner will have to pay in order to purchase the item in question. Note that, if $K_j = 0$, this is equivalent to a *direct sale* auction, i.e., without any options. Furthermore, note that local bidders have a dominant strategy to bid their value minus the exercise price if this is positive, and zero otherwise. We denote by $b_1^* \ldots b_m^*$ denote the optimal bids of the synergy bidder in the m auctions, and by $p_1 \ldots p_m$ the prices paid in these auctions. The expiry time for the options is set after the auctions for all m items close. The following theorem then specifies the optimal bidding strategy of the synergy bidder:

[1] The Bayes-Nash equilibrium is the standard solution concept used in game theory to analyze games with imperfect information, such as auctions.

[2] This option protocol is similar to the one proposed by Mous et al. [8], but that work relies on using a heuristic bidding strategies, and they do not derive analytical expressions for the equilibrium strategies w.r.t. the local bidders.

Theorem 1. *Consider the setting from Section 2, with a pre-specified exercise price vector $\overrightarrow{K}$. If $v \leq \sum_{j=1}^{m} K_j$, then $b_r^* = 0, r \in \{1, \ldots, m\}$ constitutes a Bayes-Nash equilibrium for the synergy bidder. Otherwise, the equilibrium is given by:*

$$b_r^* = \begin{cases} v - \sum_{j=1}^{m} K_j, & \textit{if } r = m \\ \\ \int_{K_{r+1}}^{b_{r+1}^* + K_{r+1}} H(\omega) d\omega, & \textit{if } 1 \leq r < m \end{cases} \tag{1}$$

where $H_j(x) = (F_j(x))^{N_j}$.

Proof. The synergy bidder cares about the value of the highest bid among the N_j local bidders which participate in the j^{th} auction. The latter would place bids which would be truthful (see [7] for 2^{nd} price auctions), if it were not for the existence of the exercise price K_j of each auction; given that K_j is an additional cost for the winner, the actual bid placed by a non-synergy bidder is equal to his valuation minus K_j. Given that each valuation is drawn from a distribution with c.d.f. $F_j(x)$, the corresponding bid is drawn from $F_j(x + K_j)$, and therefore the c.d.f. of the highest bid among local bidders is the highest order statistic, $(F_j(x + K_j))^{N_j}$.

We now compute the bidding strategy of the synergy bidder using backward induction, starting from the last (m^{th}) auction. If the synergy bidder has not won all the auctions up to the last one, then it will bid $b_m^* = 0$, as it needs to obtain all items in order to make a profit. On the other hand, assuming that the synergy bidder has won the first round, it will make a profit equal to $v - \sum_{j=1}^{m} K_j - p_m$ if it wins the last item at a price equal to p_m. This price p_m is equal to the highest opponent bid, which is drawn from $(F_m(x + K_m))^{N_m}$ when the synergy bidder is the winner. Let $H_j(x) = (F_j(x))^{N_j}$. Then the expected utility of the synergy bidder when bidding b_m is:

$$EP_m(v, b_m, \overrightarrow{K}) = \left(v - \sum_{j=1}^{m} K_j \right) H_m(K_m) + \int_0^{b_m} \left(v - \sum_{j=1}^{m} K_j - \omega \right) H_m^{'}(\omega + K_m) d\omega \tag{2}$$

The bid which maximizes this utility is found by setting:

$$\frac{dEP_m(v, b_m, \overrightarrow{K})}{db_m} = 0 \Leftrightarrow v - \sum_{j=1}^{m} K_j - b_m = 0,$$

which gives Equation 1, for the case of $r = m$. We can furthermore compute the optimal expected utility of the synergy bidder in this round by using Equation 2.

Now, we can compute the bid b_r placed in auction round r, assuming that the bid and expected utility for the next $(r + 1)$ round has been computed. The synergy bidder will make expected profit equal to $EP_{r+1}^*(v, \overrightarrow{K}) - p_r$ if it wins the r^{th} item at a price equal to p_r and it has won all auctions up to that point, where $EP_j^*(v, \overrightarrow{K})$ denotes the expected profit of the synergy bidder by bidding optimally from round j onwards. The price p_r is equal to the highest opponent bid, which is again drawn from $(F_r(x+K_r))^{N_r}$

when the synergy bidder is the winner. The expected utility of the synergy bidder from bidding b_r is thus:

$$EP_r(v, b_r, \overrightarrow{K}) = EP^*_{r+1}(v, \overrightarrow{K})H_r(K_r) + \int_0^{b_r} \left(EP^*_{r+1}(v, \overrightarrow{K}) - \omega\right) H_r^{'}(\omega + K_r)d\omega \tag{3}$$

The bid which maximizes this utility is found by setting:

$$\frac{dEP_r(v, b_r, \overrightarrow{K})}{db_r} = 0 \Leftrightarrow EP^*_{r+1}(v, \overrightarrow{K}) - b_r = 0. \tag{4}$$

Now, we need to compute the optimal expected profit EP^*_r. The expected utility of the synergy bidder when bidding b_r is given by Equation 3. Replacing the solution $b^*_r = EP^*_{r+1}(v, \overrightarrow{K})$ from Equation 4, this gives the optimal utility:

$$EP^*_r(v, \overrightarrow{K}) = b^*_r H_r(K_r) + \int_0^{b^*_r} (b^*_r - \omega)H_r^{'}(\omega + K_r)d\omega \tag{5}$$

We can then substitute the subscripts r in Equation 5 by $r + 1$, and then since $EP^*_{r+1}(v, \overrightarrow{K}) = b^*_r$ we get:

$$b^*_r = b^*_{r+1}H_{r+1}(K_{r+1}) + \int_0^{b^*_{r+1}} (b^*_{r+1} - \omega)H_{r+1}^{'}(\omega + K_{r+1})d\omega$$

Which, after integration by parts and substitution gives Equation 1:

$$b^*_r = \int_{K_{r+1}}^{b^*_{r+1}+K_{r+1}} H(\omega)d\omega$$

□

4 Optimal Bidding Strategies for Flexibly Priced Options

In the fixed exercise price options model from the previous section, the existence of the exercise prices created a secondary effect similar to having a reserve price in the auction.[3] This is because any bidder with a private valuation lower than K_j will not participate in the auction and the same will happen if the synergy bidder has a valuation lower than the sum of the exercise prices.Although this reduces the exposure problem of the synergy bidder, at the same time it may significantly reduce the market efficiency, and also negatively effects seller revenue if this value is set too high.

In order to remove this effect, in this section we introduce a novel model with flexibly priced options, i.e., that have a flexible exercise price. In more detail, in this model, we set a *maximum* exercise price K_j^H for the auction, but the actual exercise price K_j depends on the bids placed by the bidders so as to eliminate the reserve price effect. Specifically:

[3] Note, however, that there is a subtle but important difference between having a reserve price and selling options with a fixed exercise price; whereas the two auctions are equivalent from the perspective of a local bidder, the same is clearly not true for a synergy bidder.

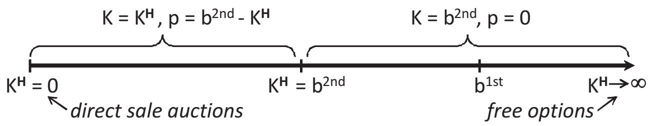

Fig. 1. The relationship between the maximum exercise price parameter, K^H, and the second-highest bid, b^{2nd}, and its effect in determining the option price, p, and actual exercise price, K, for the mechanism with flexibly priced options.

Definition 1. ***Flexible Exercise Price Options Mechanism*** *Let each seller in a sequence of m second-price auctions select a parameter K_j^H, which is the maximum exercise price she is willing to offer for the item sold in auction j. Let b_j^{2nd} denote the second highest bid placed in this auction by the participating bidders. Then, the actual exercise price of the auction is given by:*

$$K_j = \min\{K_j^H, b_j^{2nd}\}$$

Furthermore, the price paid by the winning bidder in order to purchase the option is set to:

$$p_j = b_j^{2nd} - K_j$$

Figure 1 illustrates the features of this mechanism and how it compares to some existing approaches. As shown, depending on the values of K_j^H and b^{2nd}, one of two situations can occur. Either $K_j^H < b^{2nd}$, in which case the actual exercise price is set to K_j^H, and the winner pays $b^{2nd} - K_j^H$. Otherwise, if $K_j^H \geq b^{2nd}$, then the actual exercise price is set to the second highest bid and the option is given to the winning bidder for free. In both cases, however, the total payment of the winner (if she decides to exercise the purchased option) will be equal the second highest bid. Crucially, this means that, unlike the option mechanism with fixed exercise price, from a local bidder's perspective, this auction is identical to a regular second-price auction, and there are no secondary effects on these bidders. Therefore, this options model only affects bidders with synergies.

Moreover, note from Figure 1 that this approach is a generalization of two other auction mechanisms. If the seller sets $K_j^H = 0$, then the auction becomes identical to a direct sales auction (without options). Furthermore, if K_j^H is set at a sufficiently high value (i.e. as $K_j^H \to \infty$), then the exercise price is always equal to the second highest bid, and the option is always purchased for free.

We now proceed with deriving the optimal bidding strategy for a synergy bidder.

Theorem 2. *Consider the setting from Section 2, using auctions with flexibly priced options with pre-specified maximal exercise prices K_r^H, for $r \in \{1..m\}$. The following bids b_r^* constitute a Bayes-Nash equilibrium for the synergy bidder:*

$$b_r^* = v - \sum_{i=1}^{r-1} K_i \text{ if } v \leq \sum_{i=1}^{r-1} K_i + K_r^H \tag{6}$$

$$b_r^* = K_r^H + EP_{r+1}^*(v, \overrightarrow{K}_r^H) \textit{ if } v > \sum_{i=1}^{r-1} K_i + K_r^H \tag{7}$$

$$b_m^* = v - \sum_{i=1}^{m-1} K_i \tag{8}$$

with the expected profit EP_r^* when bidding b_r^* being:

$$EP_r^*(v, \overrightarrow{K}_r^H) = \int_0^{v-\sum_{i=1}^{r-1} K_i} EP_{r+1}^*(v, \overrightarrow{K}_r^H(\omega))H_r^{'}(\omega)d\omega \tag{9}$$

if $v \leq \sum_{i=1}^{r-1} K_i + K_r^H$, otherwise it is:

$$EP_r^*(v, \overrightarrow{K}_r^H) = \int_0^{K_r^H} EP_{r+1}^*(v, \overrightarrow{K}_r^H(\omega))H_r^{'}(\omega)d\omega + \tag{10}$$

$$\int_{K_r^H}^{K_r^H + EP_{r+1}^*\left(v, \overrightarrow{K}_r^H\right)} \left(EP_{r+1}^*(v, \overrightarrow{K}_r^H) - \omega + K_r^H\right) H_r^{'}(\omega)d\omega$$

for $r \in \{1, \ldots, m-1\}$, and

$$EP_m^*(v, \overrightarrow{K}_m^H) = \int_0^{v-\sum_{i=1}^{m-1} K_i} \left(v - \sum_{i=1}^{m-1} K_i - \omega\right) H_m^{'}(\omega)d\omega \tag{11}$$

where $H_r(x) = (F_r(x))^{N_r}$ and K_r are the exercise prices for the options that have been purchased in the previous auctions. The vector $\overrightarrow{K}_r^H$ is defined as:
$\overrightarrow{K}_r^H = \{K_1, \ldots, K_{r-1}, K_r^H, \ldots, K_{m-1}^H\}$, for all rounds $r = 1, \ldots, m$.
Furthermore $\overrightarrow{K}_r^H(x)$ is defined as the vector $\overrightarrow{K}_r^H$ where the element K_r^H is replaced by value x, thus $\overrightarrow{K}_r^H(x) = \{K_1, \ldots, K_{r-1}, x, K_{r+1}^H, \ldots, K_{m-1}^H\}$.

Proof. The synergy bidder cares about the value of the highest bid among the N_r non-synergy bidders which participate in the r^{th} auction. The latter place truthful bids (as we discussed). Given that each of their valuations is drawn from a distribution with cdf $F_r(x)$, therefore the cdf of the highest bid among the local, one-item bidders is: $H_r(x) = (F_r(x))^{N_r}$, as it is the maximum of N_r random variables drawn from $F_r(x)$.

To compute the bidding strategy of the synergy bidder, we start from the last (m^{th}) auction. If the synergy bidder has not won all the previous auctions, then it will bid $b_m = 0$, as it needs to obtain all items in order to make a profit. If it has won all these auctions and the exercise prices for the options purchased are K_r, then it will make a profit equal to $v - \sum_{i=1}^{m-1} K_i - \omega$ if it wins the last item when the highest opponents bid is equal to ω.[4] Note that whether $\omega > K_m^H$ or the other way round, in both cases

[4] Note that it should be $v \geq \sum_{i=1}^{m-1} K_i$ when the synergy bidder has won all the previous auctions and the exercise prices are K_i. This follows from the fact that the bidder will not bid more than $v - \sum_{i=1}^{j-1} K_i$ in any auction j and the exercise price cannot be higher than his bid if he has won.

the winning synergy bidder makes a total payment equal to ω, where ω is drawn from $H_m(x)$ when the synergy bidder is the winner. We compute the expected utility of the synergy bidder when he bids b_m as:

$$EP_m(v, b_m, \overrightarrow{K}^H_m) = \int_0^{b_m} \left(v - \sum_{i=1}^{m-1} K_i - \omega \right) H^{'}_m(\omega) d\omega$$

The bid which maximizes this utility is found by setting:

$$\frac{dEP_m(v, b_m, \overrightarrow{K}^H_m)}{db_m} = 0 \Leftrightarrow v - \sum_{i=1}^{m-1} K_i - b^*_m = 0,$$

which gives Equation 8.

The expected profit of synergy bidder when bidding b^*_m is therefore: $EP^*_m(v, \overrightarrow{K}^H_m) = EP_m(v, b^*_m, \overrightarrow{K}^H_m)$, which gives Equation 11, when substituting b_m from Equation 8.

Assume that we have computed the bids b^*_j and expected profit EP^*_j for $\forall j > r$ and that these are given by the equations of this theorem. To complete the proof, we will now show how to compute the bid and the expected utility for the r^{th} auction (round). Depending on whether $b_r \leq K^H_r$ or not, we need to distinguish two different cases:

Case 1: $b_r \leq K^H_r$. In this case, $\forall \omega \leq b_r \Rightarrow \omega \leq K^H_r$. Therefore if the synergy bidder wins with a bid of b_r, then the second bid in the auction (which is the highest opponent bid) must be smaller than K^H_r and therefore the exercise price is equal to this bid and it is given for free. Hence, the expected profit of the synergy bidder is:

$$EP_r(v, b_r, \overrightarrow{K}^H_r) = \int_0^{b_r} EP^*_{r+1}(v, \overrightarrow{K}^H_r(\omega)) H^{'}_r(\omega) d\omega$$

To find the bid b^*_r that maximizes this expected utility we use Lagrange multipliers. The inequality is rewritten as $b_r - K^H_r + \delta^2 = 0$, and the Lagrange equation for this problem becomes:

$$\Lambda(b_r, \lambda, \delta) = - \int_0^{b_r} EP^*_{r+1}(v, \overrightarrow{K}^H_r(\omega)) H^{'}_r(\omega) d\omega + \lambda(b_r - K^H_r + \delta^2)$$

The possible variables which maximize this function are found by setting the partial derivatives for dependent variables b_r, λ, δ to 0:

$$\frac{\vartheta \Lambda(b_r, \lambda, \delta)}{\vartheta b_r} = 0 \Leftrightarrow EP^*_{r+1}(v, \overrightarrow{K}^H_r(b_r)) H^{'}_r(b_r) = \lambda$$
$$\frac{\vartheta \Lambda(b_r, \lambda, \delta)}{\vartheta \lambda} = 0 \Leftrightarrow b_r - K^H_r + \delta^2 = 0$$
$$\frac{\vartheta \Lambda(b_r, \lambda, \delta)}{\vartheta \delta} = 0 \Leftrightarrow \lambda\delta = 0$$

The last equation can mean that either: (i) $\delta = 0$, thus $b^*_r = K^H_r$, or (ii) $\lambda = 0$, and by substituting into the first equation, we get $EP^*_{r+1}(v, \overrightarrow{K}^H_r(b_r)) H^{'}_r(b_r) = 0$, thus $EP^*_{r+1}(v, \overrightarrow{K}^H_r(b_r)) = 0$. Here, the expected optimal utility EP^*_{r+1} is given either by

Equation 9 or 10, depending on whether the valuation discounted by the exercise prices up to that point $v - \sum_{i=1}^{i=r} K_i$ is less or greater (respectively) than K_{r+1}^H. The second case cannot occur if $EP_{r+1}(v, \overrightarrow{K}_r^H(b_r)) = 0$, because the first integral of Equation 10 is greater than 0, unless $b_r^* = 0$, which would yield an expected utility of 0 and thus cannot be the optimal bid (in general). On the other hand, if the optimal expected utility EP_{r+1}^* is given by Equation 9, then $EP_{r+1}^*(v, \overrightarrow{K}_r^H(b_r)) = 0$ exactly when the upper bound of the integral is 0, i.e. $v - \sum_{i=1}^{r-1} K_i - b_r = 0$, which gives Equation 6.

Note that of the two possible maxima $b_r^* = v - \sum_{i=1}^{r-1} K_i$ and $b_r^* = K_r^H$, given by this analysis, the first one yields a higher revenue, as $EP_{r+1}^*(v, \overrightarrow{K}_r^H(\omega))H_r^{'}(\omega) < 0$, $\forall \omega > v - \sum_{i=1}^{r-1} K_i$ in this case. This means that the optimal bid is $b_r^* = v - \sum_{i=1}^{r-1} K_i$. By adding the case condition $b_r^* \leq K_r^H \Leftrightarrow v - \sum_{i=1}^{r-1} K_i \leq K_r^H$, we obtain the bound in the theorem.

Case 2: $b_r > K_r^H$. In this case, if the synergy bidder wins with a bid of b_r, then the second bid in the auction could be smaller than K_r^H and therefore the exercise price is equal to this bid and it is given for free, like in the previous case. However, it could also be higher than K_r^H, and thus the exercise price is equal to K_r^H, whereas the payment for getting the options would be equal to the second bid minus K_r^H. Hence, the expected profit of the synergy bidder is:

$$EP_r(v, b_r, \overrightarrow{K}_r^H) = \int_0^{K_r^H} EP_{r+1}^*(v, \overrightarrow{K}_r^H(\omega))H_r^{'}(\omega)d\omega$$
$$+ \int_{K_r^H}^{b_r} (EP_{r+1}^*(v, \overrightarrow{K}_r^H) - \omega + K_r^H)H_r^{'}(\omega)d\omega$$

The optimal bid b_r^* that maximizes this expected utility is derived in the same way as in Case 1, by applying Lagrange multipliers to the above equation. We skip this part of the proof due to space constraints. □

4.1 Multiple Synergy Bidders

We finish the theoretical analysis by showing that in all the theorems that we presented the bidding strategies would remain unchanged, even if multiple (n) synergy bidders participate. We prove this for the flexible auction model from Theorem 2, but the same argument applies to the fixed option model in Theorem 1.

Proposition 1. *A setting with n synergy bidders with a valuation v_i for synergy bidder i when it obtains m items, and 0 otherwise, is strategically equivalent (in the case of a sequence of m auctions with flexibly priced options) to a setting where only a single synergy bidder participates and his valuation is equal to $\max_i\{v_i\}$.*

Proof. In the model considered in this paper, each synergy bidder would need to win all items in order to make a profit (i.e. there are no substitutable items available). Hence, it follows that only the synergy bidder (if any) who won the first auction would participate in the remaining ones, and the bidding strategies and the expected profit in the remaining rounds would remain unchanged. Now, in the first auction, we notice from

Equations 6 and 7 that the optimal bids of any synergy bidder in any auction (round) are not affected by the bids placed by its opponents in that particular auction, but only in the remaining ones. Therefore, the bid in the first auction is affected only by the opponent bidding in the remaining auctions, which remains unchanged as there is only one synergy bidder participating in all remaining auctions. Finally, from Equations 6 and 7 it follows that the bids placed (assuming that all the parameters K_i^L and K_i^H remain unchanged) increase as the valuation v increases, therefore the synergy bidder with the highest valuation $v = \max_i\{v_i\}$ is the only one who might win the first auction. From this, it follows that this scenario is thus strategically equivalent to one where this bidder with the highest valuation is the only synergy bidder participating.

5 Empirical Analysis

In this section, we evaluate and compare experimentally the fixed and flexible exercise price option models discussed above, as well as the direct auctions case which, as discussed above, appears as a special case in both models when $K = K_H = 0$.

5.1 Experimental Setup

The settings used for our experiments are as follows. In each run, we simulate a market consisting of $m = 3$ sequential auctions. Each auction involves $N = 5$ local bidders, and one synergy bidder. The valuations of the local bidders are i.i.d. drawn from normal distributions $\mathcal{N}(\mu = 2, \sigma = 4)$. This high variance makes the distribution almost flat (i.e. close to uniform). This uncertain valuation setting makes options more desirable.

The valuation for the synergy bidder, v_{syn}, is drawn the normal distribution $\mathcal{N}(\mu_{syn} = 20, \sigma = 2)$. We choose this setting as it demonstrates the effect of the exposure problem; if the value of the synergy bidder is set too high, then the it would win all of the auctions, even in the case of direct sale. On the other hand, if the value is set too low, then the exposure problem disappears since local bidders will win all of the auctions. Here the value of the synergy bidder is in between these extremes and is representative of a setting in which the exposure problem plays an important role.

For this setting, we compare the allocative efficiency as well as the seller revenue of the fixed and flexible exercise priced options mechanism, and for different values of the auction parameters (the exercise prices K_i in the first model, and maximum exercise prices K_i^H in the second). Formally, the allocative efficiency is calculated as follows. Let v_i^k denote the valuation of local bidder i (where $i \in \{1..N\}$) in the k^{th} auction (where $k \in \{1..m\}$) and let v_{syn} be the valuation of the synergy bidder. Furthermore, let $x_i^k, x_{syn} \in \{0, 1\}$ denote the actual allocation of the options in a certain run of the simulation. That is, $x_i^1 = 1$ means that local bidder i acquired the option in the $1st$ auction, and $x_{syn} = 1$ means that the synergy bidder won *all* the auctions. Given this, the allocative efficiency, η, of the entire market in a given run is defined as:

$$\eta = \frac{\sum_{i=1}^{n} \sum_{k=1}^{m} x_i^k v_i^k + x_{syn} v_{syn}}{\max\left(v_{syn}, \sum_{k=1}^{m} \max_{i\in\{1,\ldots,N\}}(v_i^k)\right)} \tag{12}$$

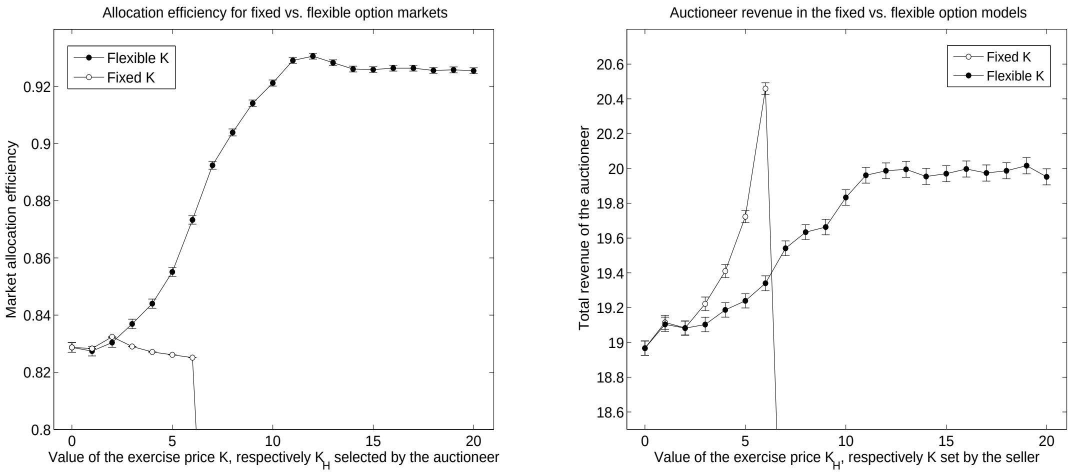

Fig. 2. Allocative efficiency (left) and seller revenue (right) using the options mechanism with fixed and flexible exercise price, and with identical parameters K and K^H respectively in all auctions. Result are averaged over 3500 runs. The error bars indicate the standard error.

By calculating the efficiency of the market in this way, we implicitly assume that local bidders will always exercise their options, and that the synergy bidder will exercise its option if and only if it wins all auctions. We can safely make this assumption because we consider optimal bidding strategies, and a rational bidder will never place a bid such that the combined exercise and option price will exceed the (marginal) value of the item. Therefore, it is optimal for a bidder who has acquired options for all of its desired items to exercise them. Thus an *in*efficient outcome occurs in two situations. Either the local bidders have won the items, but the value of the synergy bidder exceeds the sum of the values of the local bidders; or, the synergy bidder has won some auctions but not all, and will therefore not exercise its option(s).

5.2 Discussion of Numerical Results

To reduce the number of parameters, we consider a setting in which the parameters, K_i and K_i^H for the two options mechanism respectively, are set to the same value in all of the 3 sequential auctions, denoted by K and K^H respectively in the following. This models a setting in which the seller has to pre-specify the protocol for selling her item without any knowledge about a synergy bidder's endowment state[5] (i.e., how many options it has won so far).

To this end, Figure 2 compares the allocation efficiency of the market (left), and the seller revenues (right) for the two option mechanisms. Note that the direct auctions case appears, in both option models, as a particular subcase, for $K = K_H = 0$.

[5] This is especially relevant in a large open environment where synergy bidders can enter and exit the market dynamically.

Figure 2 shows that, for the *fixed exercise price* option model, both the efficiency and the seller revenue start to decrease sharply when K becomes larger than around 6.2. This is because, at this point, the synergy bidder is likely to leave the market due to the reservation price effect of this mechanism. Specifically, this occurs when the sum of the exercise prices of the auctions exceeds the valuation of the synergy bidder, in which case the bidder no longer has an incentive to participate. This also holds for many of the local bidders. This is precisely the outcome that we would like to avoid using our flexibly priced option mechanism. Note that, using our new mechanism, having a very high value for K^H has a very different effect, and the options become effectively free. This is because, when K_H is very high, it will almost certainly exceed the second-highest bid. If this happens, the exercise price becomes equal to this bid, and the option price becomes zero (see also Section 4).

As is shown in Figure 2 (left), since the flexibly priced option mechanism removes the reserve price effect, this mechanism outperforms a fixed exercise price in terms of efficiency. Furthermore, both mechanisms outperform direct auctions (i.e. without options, which is when $K = 0$ or $K^H = 0$) for an appropriately set parameter. The results also show that having free options is suboptimal in terms of efficiency. This is because there is a small chance that the synergy bidder will win the one or more auctions in the sequence, but loose the second or the third, and hence not exercise her options. If this occurs, some goods remain unallocated. This gives rise to inefficiency since the goods could have been allocated to a local bidder instead.

While the flexibly priced options outperforms other mechanisms in terms of efficiency, the same cannot be said for seller revenue. In this context, Figure 2 (right) shows that a seller can achieve significantly higher revenues by using a fixed exercise price. This is not entirely surprising, since the fixed exercise price has a secondary effect which is similar to setting a reserve price and standard auction theory shows that, even in a single second-price auction, the seller can increase its revenues by using reserve prices [7]. Nevertheless, we find that the flexibly priced option mechanism achieves a higher revenue than regular, direct sale auctions. Furthermore, it is important to point out that, if the aim is to maximise revenue, the flexibly priced option mechanism can be used in combination with a reserve price, and our mechanism enables the separation of the two effects: reducing the exposure problem of synergy bidders and increasing seller revenue. We intend to investigate models in which both of these effects are jointly captured, but through separate parameters, in future work.

6 Conclusions

The exposure problem faced by bidders with valuation synergies in sequential auctions is a difficult, but an important one, with considerable implications for both theory and practice, for a wide range of multi-agent systems. Due to the risk of not acquiring all the desired items in future auctions, bidders with valuation synergies often shade their bids, or do not participate in such markets, which considerably reduces both allocation efficiency and auctioneer revenue. Options have been identified before [5, 6, 8] as a promising solution to address this problem, but existing mechanisms in the literature either prescribe free options, or options in which the seller fixes a minimum exercise price (which can potentially deter many bidders from entering the market).

To this end, in this paper, we propose a novel option mechanism, in which the exercise price is set flexibly, as a minimum between the second highest bid and a seller-prescribed maximum level, while the option price is determined by the open market. We derive the optimal bidding policies of the synergy bidder in this new model and show that this mechanism can significantly increase the social welfare of the resulting allocations, while at the same time outperform sequential auction with direct sales in terms of seller revenues.

While the mechanism proposed in this paper makes a significant step forward in addressing this problem, several aspects are still left open to future work. One of these is combining the two options models by allowing sellers to fix a minimum exercise price for their options, as well as a maximum. Such a mechanism would, on the one hand, be able to reduce the exposure problem for synergy bidders, and therefore result in a corresponding increase in market efficiency, but would also allow the seller to extract more profits, as they do in a fixed-price options model. Another important area where further work is needed, is the derivation of optimal bidding strategies in more general market settings, such as those involving substitutable items as well as complementary ones.

References

1. C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. *Proc. of IJCAI'99*, pp. 527–523, Morgan Kaufmann, 1999.
2. P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial auctions*. MIT Press, 2006.
3. A. Greenwald and J. Boyan. Bidding under uncertainty: theory and experiments. In *UAI '04: Proc. of the 20th Conference on Uncertainty in AI*, pages 209–216, 2004.
4. J. C. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, fifth edition, 2003.
5. A. I. Juda and D. C. Parkes. An options-based method to solve the composability problem in sequential auctions. *Agent-Mediated El. Comm. VI*, pp. 44–58. Springer LNCS 3435, 2006.
6. A. I. Juda and D. C. Parkes. An Options-Based Solution to the Sequential Auction Problem. In *Artificial Intelligence*, vol. 173(7-8), pages 876–899, 2009.
7. V. Krishna. *Auction theory*. Academic Press, 2002.
8. L. Mous, V. Robu, and J. A. La Poutré. Using priced options to solve the exposure problem in sequential auctions. In *Agent-Mediated El. Commerce (AMEC'08)*, Springer LNBIP, 2009.
9. A. Osepayshvili, M. Wellman, D. Reeves, J. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Proc. of UAI'05*, pp. 441–449, 2005.
10. V. Robu, H. Noot, J. A. La Poutré, and W.-J. van Schijndel. An agent platform for auction-based allocation of loads in transportation logistics. In *Proc. of AAMAS'08, Industry Track*, pages 3–10, 2008.
11. I. A. Vetsikas and N. R. Jennings. Bidding strategies for realistic multi-unit sealed-bid auctions. In *23rd Conference on Artificial Intelligence (AAAI'08), Chicago*, 2008.