Addressing the Exposure Problem of Bidding Agents Using Flexibly Priced Options

Valentin Robu[1] and **Ioannis A. Vetsikas**[1] and **Enrico H. Gerding**[1] and **Nicholas R. Jennings**[1]

Abstract. In this paper we introduce a new option pricing mechanism for reducing the exposure problem encountered by bidding agents with complementary valuations when participating in sequential, second-price auction markets. Existing option pricing models have two main drawbacks: they either apply fixed exercise prices, 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. The proposed mechanism involving *flexibly priced options* addresses these problems by calculating the exercise price as well as the option price based on the bids received during an auction. For this new model, which extends and encompasses all the previous models examined, we derive the optimal strategies for a bidding agent with complementary preferences. Finally, we use these strategies to evaluate the proposed option mechanism through Monte-Carlo simutions, and compare it to existing mechanisms, both in terms of the seller revenue and the social welfare. We show that our new mechanism achieves higher market efficiency compared to having no options and free options, while achieving higher revenues for the seller than any existing option mechanism.

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 primarily on settings 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, 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 resources become available over time. Examples include inter-related items sold on eBay by different sellers in auctions with different closing times [4], decentralised transportation logistics, in which part-capacity loads are offered throughout a day by different shippers [9], and electricity grids, where a larger generation capacity needs to be acquired from different sellers [7]. In such settings, an agent desiring a bundle of goods (henceforth called a *synergy bidder*) 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 new option pricing mechanism, and shows its superiority over existing approaches.

In more detail, 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, 8, 11]). In this paper, we base our work on a different approach, that preserves the sequential nature of the allocation problem, and 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 are effective because they can help a synergy bidder 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. Thereby, she is able to limit her loss. Such options work well because 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.

Applying options to reduce the exposure problem in sequential auctions is not new, and there have been a limited number of papers that study this setting. In this context, two main types of pricing mechanisms exist in the current literature. One approach (e.g. by [4]) is to have free options (i.e., an option price of zero), and then let the exercise price be determined by the market (i.e., the submitted bids). However, this approach enables self-interested agents to hoard those options, even if they are unlikely to exercising them, thus considerably reducing both the allocative efficiency of the market and seller revenue. To address this, another approach (proposed in [6]) is to have a fixed exercise price, set by the seller, and then have the market determine the option price. In this case, however, the exercise price can be perceived as a reserve price 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.

To address the shortcomings of existing option models, in this paper we introduce a general pricing mechanism with flexibly priced

[1] Intelligence, Agents, Multimedia Group, University of Southampton, SO17 1BJ Southampton, United Kingdom, email: {vr2,iv,eg,nrj}@ecs.soton.ac.uk

options, in which the exercise price, as well as the option price, are determined by the market, and which generalises the aforementioned option models. We then compare this model to the existing option models and demonstrate the efficacy of our approach for a setting in which a *synergy bidder* requires items from all auctions, whereas other bidders (called *local bidders*) only participate in one of the auctions. We choose this setting since it highlights the exposure problem, but the mechanism can be equally applied to more general settings. In doing so, 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 prices.
- Before presenting the most general model, we start by introducing a new option pricing mechanism, which involves both the option and exercise price of an option being determined by the market. We then extend this basic model and present a generalised flexibly priced options model which encompasses all of the existing models, including the fixed exercise price model, as well as the model with free options.
- We then derive the optimal bidding strategy for the synergy bidder in this general model.
- We experimentally 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 12%-13% improvement in market allocation efficiency (measured in terms of the social welfare of all participating agents), compared to the state of the art fixed price options model. Furthermore, we show that the generalised model can achieve between 4%-12% higher revenue than all existing option pricing models.

The remainder of the paper is structured as follows. We first introduce our setting more formally in Section 2. In Section 3, we then present the various option models and the derivation of the optimal bidding strategies. In Section 4, we present the experiments comparing the various option models, and Section 5 concludes.

2 The Problem Setting

We consider a setting with m sequential, second-price, sealed-bid auctions, each selling an option to buy a single item, where the closing time for auction at time j is strictly prior to the closing time of auction $j+1$ ($\forall j = 1..m-1$). We choose these auctions because local bidders desiring a single item (sold in only one of the m auctions) have a simple, dominant bidding strategy and, furthermore, they are (weakly) 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. We choose this setting to highlight the exposure problem. 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 seek to maximise their expected profits (i.e. expected value minus expected payments).[2]

Given this setting, we are interested in finding the *Bayes-Nash* equilibrium strategies for all of the bidders for different option pricing mechanisms.[3] 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 setting is addressed in Section 3.4.

[2] This implies that bidders are assumed to be risk neutral.

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

3 Option Pricing Models

In this section we present the three option models that we examine in this paper. In Section 3.1 we start by formally presenting the option model with fixed exercise prices, followed, in Section 3.2, by our new model with flexibly priced options. Although this model improves efficiency, this comes at the expense of seller revenue compared to the fixed exercise model (as we show in Section 4). We therefore introduce a generalised flexibly priced options model in Section 3.3 which extends and encompasses all of the option models discussed in this paper. We then derive the equilibrium bidding strategies for this general model. Finally, in Section 3.4 we address the setting with multiple synergy bidders.

3.1 Fixed Exercise Price

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.[4] 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. 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 difference is positive, and zero otherwise. The following theorem then specifies the optimal bidding strategy of the synergy bidder (we omit the proof since we later present the proof of Theorem 2, which is a generalisation of this theorem):

Theorem 1. *Consider the setting presented above, 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:*

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

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

3.2 Flexibly Priced Options

In the above fixed exercise price options model, the existence of the exercise prices creates a secondary effect similar to having a reserve price in the auction. 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

[4] This option protocol is similar to the one in [6], but their bidding strategy relies on using a heuristic, and they do not derive analytical expressions for the equilibrium strategies w.r.t. the local bidders.

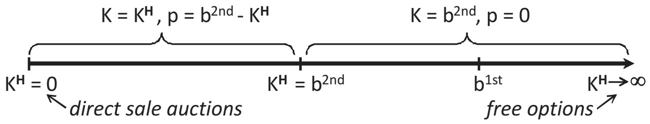

Figure 1. The relationship between the maximum exercise price, K^H, and the second-highest bid, b^{2nd}, and its effect in determining the option price, p, and actual exercise price, K, for flexibly priced options model.

the synergy bidder compared to direct sales, 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, we introduce a model involving flexibly priced options, i.e., that have a flexible exercise price (first presented, as an extended abstract, in [10]). 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, this is set to:

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

where b_j^{2nd} is the second-highest bid. Furthermore, the price paid by the winning bidder in order to purchase the option (the option price) is set to:

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

To understand how this mechanism works, Figure 1 illustrates its features 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 to the second highest bid. Crucially, this means that, unlike the option mechanism with fixed exercise prices, 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 $K_j^H \rightarrow \infty$), then the exercise price is always equal to the second highest bid, and the option is always purchased for free.

The advantage of the flexibly priced options model is that it removes the reserve price effect and thereby increases market efficiency. However, as we will show in the experiments (Section 4), a seller may still prefer using the fixed exercise price model since this can generate higher revenue when the fixed exercise price is appropriately selected. Therefore, in the next section we propose a generalised model that can obtain both high efficiency and seller revenue.

3.3 Generalised Flexibly Priced Options Model

In this section, we extend the flexibly priced model of section 3.2 to include a *minimum exercise price*[5], which enables the seller to control the revenue level of the auction. This more general model encompasses a wide range of existing option mechanisms and can take advantage of all the benefits of the models discussed so far. This means that it can achieve both high efficiency and high seller revenue at the same time. More formally:

Definition 1. ***Generalised Flexibly Priced Options*** *Let K_j^L and K_j^H denote the minimum and maximum exercise prices respectively for auction j. These parameters are set by the seller in advance of the auction and are known to the bidders. Furthermore, b_j^{2nd} is the second highest bid. Then the actual exercise price, K_j, of the auction is given by:*

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

Furthermore, the option price is given by:

$$p_j = \max\{0, b_j^{2nd} - K_j\}$$

Note that by setting $K_j^L = K_j^H = K_j$, this corresponds to the the model with fixed exercise prices.Furthermore, the flexibly priced options model is given by $K_j^L = 0$. We now derive the equilibrium strategy of a synergy bidder for the generalised model:

Theorem 2. *Considering the same setting as before, the following bids $b_r^*, r \in \{1, \dots, m\}$ constitute a Bayes-Nash equilibrium for the synergy bidder in auctions with generalised flexibly priced options. If $r = m$:*

$$b_m^* = v - \sum_{i=1}^{m-1} K_i \quad (2)$$

If $r < m$ and $v \leq \sum_{i=1}^{r-1} K_i + K_r^H + \sum_{i=r+1}^{m} K_i^L$:

$$b_r^* = v - \sum_{i=1}^{r-1} K_i - \sum_{i=r+1}^{m} K_i^L \quad (3)$$

Otherwise:
$$b_r^* = K_r^H + EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H)) \quad (4)$$

where EP_j^* denotes the equilibrium expected profit of the synergy bidder at time j. EP_j^* can be calculated recursively as follows. First, if $j = m$:

$$EP_m^*(v, \overrightarrow{K}_m^*) = (v - \sum_{i=1}^{m-1} K_i - K_m^L) H_m(K_m^L) + \int_{K_m^L}^{v - \sum_{i=1}^{m-1} K_i} (v - \sum_{i=1}^{m-1} K_i - \omega) H_m'(\omega) d\omega \quad (5)$$

If $j < m$ and $v \leq \sum_{i=1}^{j-1} K_i + K_j^H + \sum_{i=j+1}^{m} K_i^L$:

$$EP_j^*(v, \overrightarrow{K}_j^*) = EP_{j+1}^*(v, \overrightarrow{K}_j^*(K_j^L)) H_j(K_j^L) + \int_{K_j^L}^{v - \sum_{i=1}^{j-1} K_i - \sum_{i=j+1}^{m} K_i^L} EP_{j+1}^*(v, \overrightarrow{K}_j^*(\omega)) H_j'(\omega) d\omega \quad (6)$$

Otherwise:

$$EP_j^*(v, \overrightarrow{K}_j^*) = EP_{j+1}^*(v, \overrightarrow{K}_j^*(K_j^L)) H_j(K_j^L) + \int_{K_j^L}^{K_j^H} EP_{j+1}^*(v, \overrightarrow{K}_j^*(\omega)) H_j'(\omega) d\omega + \int_{K_j^H}^{K_j^H + EP_{j+1}^*(v, \overrightarrow{K}_j^*(K_j^H))} (EP_{j+1}^*(v, \overrightarrow{K}_j^*(K_j^H)) - \omega + K_j^H) H_j'(\omega) d\omega \quad (7)$$

where $H_j(x) = (F_j(x))^{N_j}$. Here, K_j denotes the actual exercise price for an option that the synergy bidder has purchased. Furthermore, the vector $\overrightarrow{K}_j^*$ contains known exercise prices from all past auctions when participating in the j^{th} auction, as well as the parameters for auctions still to come. Specifically: $\overrightarrow{K}_j^* = \langle \{K_1\}, \dots, \{K_{j-1}\}, \{K_j^L, K_j^H\}, \dots, \{K_m^L, K_m^H\} \rangle$, for $j = 1, \dots, m$. Furthemore, we let $\overrightarrow{K}_j^*(x) = \langle \{K_1\}, \dots, \{K_{j-1}\}, \{x\}, \{K_{j+1}^L, K_{j+1}^H\} \dots, \{K_m^L, K_m^H\} \rangle$.

[5] An alternative approach is to set a minimum option or reserve price, which is equivalent from a local bidder's perspective, but is less effective in reducing the exposure problem for a synergy bidder.

Proof. Note that the local bidders would have to pay either $p_j + K_j = K_j^L$ if the second highest bid $b_j^{2nd} \leq K_j^L$, or $p_j + K_j = b_j^{2nd}$ otherwise. Therefore each auction is strategically equivalent to a second price auction with reserve K_j^L, from the point of view of a local bidder. Thus, these bidders would participate and bid truthfully, if their valuation is at least K_j^L (see [5]).

The synergy bidder cares about the value of the highest bid (and thus the highest valuation) among the N_j local bidders participating in the j^{th} auction. These valuations are drawn from a distribution with cdf $F_j(x)$, thus the top bid of the local bidders is a random variable with cdf $H_j(x) = (F_j(x))^{N_j}$.[6]

To compute the bidding strategy of the synergy bidder, we start from the last (m^{th}) auction. She bids $b_m = 0$, if she has not won all the previous auctions, as she needs all the items to make a profit. Otherwise, if the exercise prices for the options purchased are K_j, her profit equals $v - \sum_{i=1}^{m-1} K_i - \max\{K_m^L, \omega\}$ by winning the last item when the highest opponent bid is equal to ω,[7] where ω is drawn from $H_m(x)$. Her expected utility (when she bids b_m) is:

$$EP_m(v, b_m, \overrightarrow{K}_m^*) = (v - \sum_{i=1}^{m-1} K_i - K_m^L)H_m^{'}(K_m^L) + \int_{K_m^L}^{b_m} (v - \sum_{i=1}^{m-1} K_i - \omega)H_m^{'}(\omega)d\omega$$

The bid which maximizes this utility is found by setting:

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

which gives Equation 2. The expected profit of synergy bidder when bidding b_m is therefore: $EP_m^*(v, \overrightarrow{K}_m^*) = EP_m(v, b_m, \overrightarrow{K}_m^*)$, which gives Equation 5, when substituting b_m from Equation 2.

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 now compute the bid and expected profit for the r^{th} auction. We examine two cases, depending on whether $b_r \leq K_r^H$ or not:

Case I: $b_r \leq K_r^H$. Then $\forall \omega \leq b_r \Rightarrow \omega \leq K_r^H$. Hence, if the synergy bidder wins with a bid of b_r, the second bid in the auction (which is the highest opponent bid) must be smaller than K_r^H and therefore the exercise price is equal to this bid (or K_r^L if this bid is lower than K_r^L) and it is given for free. Thus, the expected profit of the synergy bidder is:

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

To find the bid b_r that maximizes this expected utility we use Lagrange multipliers. After we introduce a factor δ to convert the inequality, it can be rewritten as: $b_r - K_r^H + \delta^2 = 0$. Then, the Lagrange equation becomes:

$$\Lambda(b_r, \lambda, \delta) = -EP_r(v, b_r, \overrightarrow{K}_r^*) + \lambda(b_r - K_r^H + \delta^2)$$

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

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

The last equation can mean that: either (i) $\delta = 0$, thus $b_r = K_r^H$, or (ii) $\lambda = 0$, and by substituting into the first equation, we get that $EP_{r+1}^*(v, \overrightarrow{K}_r^*(b_r)) = 0$. The expected utility EP_{r+1}^* is given either by Equation 6 or 7, depending on whether the valuation discounted by the exercise prices up to that point and the minimum exercise prices of future rounds $v - \sum_{i=1}^{i=r} K_i - \sum_{i=r+2}^{m} K_i^L$ is less or greater (respectively) than K_{r+1}^H. The second case cannot be, if $EP_{r+1}^*(v, \overrightarrow{K}_r^*(b_r)) = 0$, because the first integral of Equation 7 is greater than 0, unless $b_r = K_r^L = K_r^H$, which would yield an expected utility of 0 and thus cannot be the optimal bid (in general). On the other hand, if the expected utility EP_{r+1}^* is given by Equation 6, then it is $EP_{r+1}^*(v, \overrightarrow{K}_r^*(b_r)) = 0$ exactly when the upper bound of the integral is equal to K_{r+1}^L, which gives Equation 3. This happens because K_r is set to bid b_r of Equation 3, and thus, by placing this bid b_r, both parts of Equation 6, for round $r + 1$ become equal to 0.

Note that of the two possible maxima $b_r = v - \sum_{i=1}^{r-1} K_i - \sum_{i=r+1}^{m} K_i^L$ (Eq. 3) and $b_r = K_r^H$, given by this analysis, the first one yields a higher revenue, as $EP_{r+1}^*(v, \overrightarrow{K}_r^*(\omega))H_r^{'}(\omega) < 0, \forall \omega > v - \sum_{i=1}^{r-1} K_i - \sum_{i=r+1}^{m} K_i^L$. This means that the optimal bid is given by Eq. 3. It should be $b_r \leq K_r^H \Leftrightarrow v - \sum_{i=1}^{r-1} K_i - \sum_{i=r+1}^{m} K_i^L \leq K_r^H$, which gives the bound of the theorem.

Case II: $b_r > K_r^H$. If the synergy bidder wins with bid b_r, the second bid in the auction could be smaller than K_r^H, thus the exercise price is equal to this bid (or K_r^L if the bid is even smaller than this) 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:

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

To find the bid b_r that maximizes this expected utility we use Lagrange multipliers again:

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

The last equation can mean that: either (i) $\delta = 0$, thus $b_r = K_r^H$, or (ii) $\lambda = 0$, and by substituting into the first equation, we get that $H_r^{'}(b_r)\big(EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H)) - b_r + K_r^H\big) \Leftrightarrow b_r = EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H)) + K_r^H$, which is indeed Equation 4. It should be $b_r \geq K_r^H \Leftrightarrow EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H)) \geq 0$. This happens exactly when $v - \sum_{i=1}^{r} K_i > \sum_{i=r+1}^{m} K_i^L$. But since $K_r = K_r^H$, this means that $v > \sum_{i=1}^{r-1} K_i + K_r^H + \sum_{i=r+1}^{m} K_i^L$.

Note that of the two possible maxima $b_r = K_r^H$ and $b_r = K_r^H + EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H))$, given by this analysis, the first one yields a higher revenue, as $(EP_{r+1}^*(v, \overrightarrow{K}_r^*(K_r^H)) - \omega + K_r^H)H_r^{'}(\omega) > 0, \forall \omega \in [K_r^H, b_r]$. Hence, the optimal bid is given by Eq. 4, when $v > \sum_{i=1}^{r-1} K_i + K_r^H + \sum_{i=r+1}^{m} K_i^L$. □

The intuition behind these equations, in simple terms, is that in the last auction the bidder bids according to its valuation v minus the exercise prices K_i which it will have to pay in order to buy all the previous items; the last auction is essentially strategically equivalent to a standard second price auction with a reserve price, so the bidder bids its true (discounted) value. In any previous auction $r < m$,

[6] To be entirely precise, we assume that if this highest bid is lower than the minimum exercise price K_j^L, then the bids are actually placed but ignored eventually.

[7] Note that it should be $v \geq \sum_{i=1}^{m-1} K_i + K_m^L$ when she has won all the previous auctions, which follows from the bids placed.

the discounted value v minus the exercise prices K_i up to that point give the discounted value of the bidder; this is further discounted by the minimum exercise prices K_i^L of the future auctions. Now, if this value is higher than the maximum exercise price limit K_r^H for that auction, then it bids this maximum exercise price limit K_r^H plus the profit it expects from the remaining rounds conditional on the fact that $K_r = K_r^H$; if the exercise price K_r is indeed smaller then its profit can only increase. On the other hand, since the agent cannot bid higher than the discounted value, if this is smaller than K_r^H, then the agent hopes to get the option for free, therefore it bids the discounted value; if it wins it will have to pay 0 and the exercise price K_r will be smaller than the discounted value $v-\sum_{i=1}^{r-1} K_i-\sum_{i=r+1}^{m} K_i^L$, thus in the next round the discounted value $v-\sum_{i=1}^{r} K_i-\sum_{i=r+2}^{m} K_i^L$ at that point will be higher than K_{r+1}^L, and the synergy bidder will always be able to participate. This is also true, when the discounted value is greater than K_r^H.

3.4 Multiple Synergy Bidders

We finish the theoretical analysis by showing that in all the theorems that we have presented, the bidding strategies will remain unchanged, even if multiple (i.e. 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 generalised flexibly priced options) to a setting where only a single synergy bidder participates and his valuation is equal to $\max_i\{v_i\}$.*

Proof. As Eq. 3 and 4 show, the optimal bids of synergy bidders in any auction are not affected by their opponents' bids placed in that particular auction, but only in the remaining ones. Furthermore, each synergy bidder would need to win all items in order to make a profit, thus 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. Thus, only the synergy bidder with the highest valuation $v = \max_i\{v_i\}$ matters in the analysis, as he is the only one who can win the first auction. This scenario is strategically equivalent to one where only the highest valuation synergy bidder participates. □

4 Experimental Analysis

In this section, we use the the optimal bidding strategies derived in Section 3 to evaluate experimentally the flexible option pricing models and compare them to existing option models as well as direct sales (which appears as a special case for $K = K_L = K_H = 0$).

4.1 Experimental Setup

The setting used for our experiments is 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 $\mathcal{N}(\mu = 2, \sigma = 4)$, bounded at 0. The reason for choosing such a high dispersion is that this increases the uncertainty in the market, and hence the benefits of using options. The valuation for the synergy bidder, v_{syn}, is drawn from the normal distribution $\mathcal{N}(\mu_{syn} = 20, \sigma_{syn} = 2)$. This setting was chosen as it demonstrates the effect of the exposure problem; if the value of the synergy bidder is set too high, then the synergy bidder 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.

The exercise prices for the options are set as follows. For the fixed exercise options model, the value of K was varied between $K = 0$ (a setting equivalent to direct sale) and $K = 20$ (the average valuation of the synergy bidder), for all the 3 items offered in the auction sequence. For the generalized flexible case, we consider 2 values of the parameter K_L: $K_L = 0$ and $K_L = 6$. The rationale for this choice is that $K_L = 0$ represents the fully flexible auction model (introduced in Section 3.2), while $K_L = 6$ represents the value which was found experimentally to provide the best revenues for the seller. In both of these generalised cases, the value of the upper limit K_H was varied from 0 to 20, as for the fixed case.

4.2 Efficiency Criteria

We compare the three option models above based on two main comparison criteria, which are standard in auction theory: the allocative efficiency of the market and the revenue obtained by the auctioneer.

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:

$$\eta = \frac{\sum_{i=1}^{n}\sum_{k=1}^{m} x_i^k v_i^k + x_{syn} v_{syn}}{\max\left(v_{syn}, \sum_{k=1}^{m} \max_{i\in\{1,\ldots,N\}}(v_i^k)\right)} \tag{8}$$

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).

4.3 Experimental Results

Figure 2 compares the allocative efficiency of the market (left), and the seller revenues (right) for the three option mechanisms. Here, *flexible K* corresponds to the general model where $K_L = 0$; *fixed K* corresponds to $K = K_L = K_H$; the third line shows the generalised option model which starts at $K = K_L = K_H = 6$ after which we keep K_L fixed at 6 while K_H increases.

Note that, for the *fixed exercise price* option model, both the efficiency and the seller revenue decrease sharply when K becomes larger than around 6. 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.

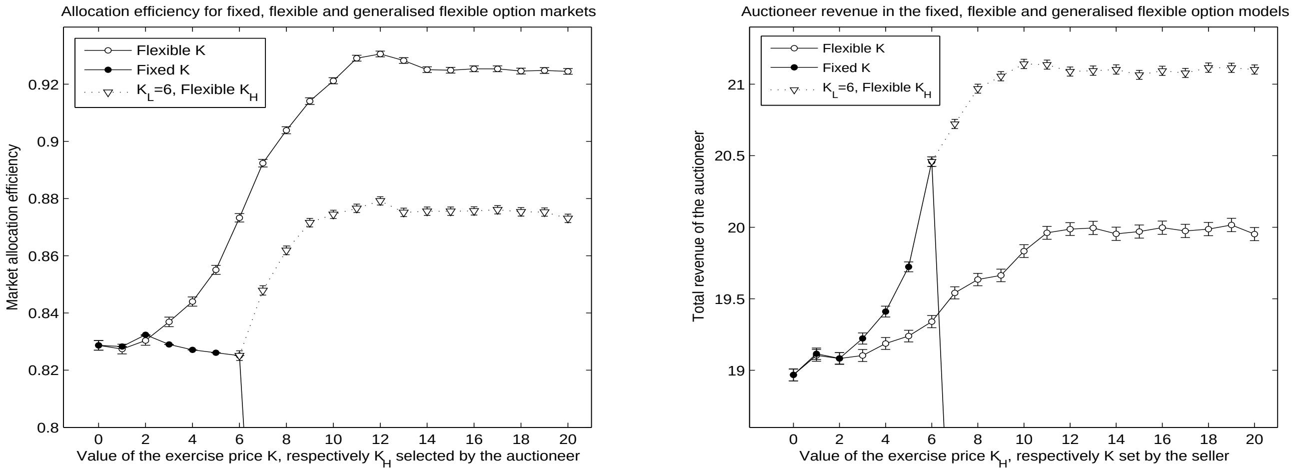

Figure 2. Allocative efficiency (left) and seller revenue (right) using the options mechanism with fixed and flexible exercise price, for different values of K, respectively and K_H. The graph shows averages and standard error bars over 10000 runs, for each data point.

As shown by the results, this reservation effect can be avoided by using the flexibly priced option mechanism. In this mechanism, the level of K_H has a different effect than K in fixed-price options. In particular, the options become effectively free when K_H becomes sufficiently high. This is because, when K_H is set 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 Section 3.2). However, free options can also be sub-optimal in terms of market efficiency. This is because a synergy bidder will always bid to acquire the options, even if she has a very low probability of exercising them. In such a case, if the seller receives a high bid from one of the local bidders, it may be better to allocate the item to this local bidder, rather than the synergy bidder.

While the flexibly priced option model outperforms the other models in terms of efficiency (as is shown in Figure 2), the same cannot be said for seller revenue. As Figure 2 (right) shows, the seller can achieve higher revenues by using fixed exercise prices. Here, the fixed K in our model acts effectively as 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 [5].

Now, both revenue and efficiency can be addressed using our more generalised flexible model, which admits both a lower and an upper limit of the exercise price, and thus can take care of both of these aspects. In the case shown in the figure, we fix the lower threshold to $K_L = 6$, which was found to be the level of K_L for which the best auctioneer revenue can be achieved, for this setting. Beyond this value, bidding agents would start to drop out of the market. The upper limit K_H is allowed to rise, however, thus capturing some of the benefits of both the fixed and flexible options. Somewhat surprisingly, as shown in Figure 2 (right), the generalised option pricing outperforms both of the other option models (as well as direct auctions, i.e. for $K = K_H = 0$) in terms of revenue which can be extracted by the seller. On closer analysis, this is because it can take advantage of the reserve price effect like the fixed option pricing model, but at the same time it is able to maintain a higher level of allocation efficiency. Naturally, the allocative efficiency is higher for $K_L = 0$ since this removes the reserve price effect, but this comes at a cost in terms of revenue. Nonetheless, the setting which generates the highest revenue for the seller still results in a higher allocative efficiency compared to the fixed pricing model and direct auctions (denoted Fig. 2 by the point in which $K = K_L = K_H = 0$).

5 Conclusions

In this paper, we introduce an option pricing mechanism which addresses the exposure problem of a synergy bidder. Here, 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 for a combined mechanism which encompasses our new approach, as well as existing option models. We show that our combined pricing model can significantly increase the efficiency of the resulting allocations compared to fixed-priced options and having no options, while at the same time obtaining higher revenues than any of the other option models, including the flexibly priced one.

The results in this paper are based on the assumption that the synergy bidder is interested in winning all the auctions, and derives no utility from winning less. This setting best captures the exposure problem, and our main goal here was to demonstrate the effectiveness of the new option pricing model in addressing this issue. Now, while there are a number of settings where such an assumption is realistic[8], there are also settings where bidders are interested in more complex or partial subsets of the available goods. Thus, our future work will seek to extend our analysis to deal with such cases.

REFERENCES

[1] C. Boutilier, M. Goldszmidt, and B. Sabata, ‘Sequential auctions for the allocation of resources with complementarities’, in *IJCAI’99*, pp. 527–523, (1999).

[2] P. Cramton, Y. Shoham, R. Steinberg, *Combinatorial auctions*, MIT Press 2006

[3] A. Greenwald and J. Boyan, ‘Bidding under uncertainty: theory and experiments’, in *UAI ’04*, pp. 209–216, (2004).

[4] A. I. Juda and D. C. Parkes, ‘An options-based solution to the sequential auction problem’, *Artificial Intelligence*, **173**(7-8), 876–899, (2009).

[5] V. Krishna, *Auction theory*, Academic Press, 2002.

[6] L. Mous, V. Robu, and J. A. La Poutré, ‘Using priced options to solve the exposure problem in sequential auctions’, in *AMEC’08*, Springer LNBIP, (2008).

[7] J. Nicolaisen, V. Petrov, and L. Tesfatsion, ‘Market power and efficiency in a computational electricity market with discriminatory double-auction pricing’, *IEEE Trans. Evolutionary Computation*, **5**(5), 504–523, (2001).

[8] A. Osepayshvili, M. Wellman, D. Reeves, and J. MacKie-Mason, ‘Self-confirming price prediction for bidding in simultaneous ascending auctions’, in *UAI’05*, pp. 441–449, (2005).

[9] 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 *AAMAS’08, Industry Track*, pp. 3–10, (2008).

[10] V. Robu, E.H. Gerding, I.A. Vetsikas, N.R. Jennings, ‘Flexibly Priced Options: A New Mechanism for Sequential Auctions with Complementary Goods (Extended Abstract)’, in *AAMAS’10, Toronto*, pp. 1485-1486, (2010).

[11] I. A. Vetsikas and N. R. Jennings, ‘Bidding strategies for realistic multi-unit sealed-bid auctions’, in *AAAI’08*, (2008).

[8] One such example is transportation logistics with part-truck or return loads. For example, an agent acquiring orders to fill one truck may need to get both an outgoing and a return order for the combination to be profitable.