

Optimal Bidding Strategies for Simultaneous Vickrey Auctions with Perfect Substitutes

Enrico H. Gerding, Rajdeep K. Dash, David C. K. Yuen and Nicholas R. Jennings

University of Southampton, Southampton, SO17 1BJ, UK.

{eg,rkd02r,dy,nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper, we derive optimal bidding strategies for a global bidder who participates in multiple, simultaneous second-price auctions with perfect substitutes. We first consider a model where all other bidders are local and participate in a single auction. For this case, we prove that, assuming free disposal, the global bidder should always place non-zero bids in all available auctions, irrespective of the local bidders' valuation distribution. Furthermore, for non-decreasing valuation distributions, we prove that the problem of finding the optimal bids reduces to two dimensions. These results hold both in the case where the number of local bidders is known and when this number is determined by a Poisson distribution. In addition, by combining analytical and simulation results, we demonstrate that similar results hold in the case of several global bidders, provided that the market consists of both global and local bidders. Finally, we address the efficiency of the overall market, and show that information about the number of local bidders is an important determinant for the way in which a global bidder affects efficiency.

Keywords

Simultaneous Auctions, Perfect Substitutes, Bidding Strategies, Vickrey Auction, Multiple Sellers, Market Efficiency

1. INTRODUCTION

In recent years, there has been a surge in the application of auctions, both online and within multi-agent systems [3, 8, 13, 14, 15]. As a result, there are an increasing number of auctions offering very similar or even identical goods and services. In eBay alone, for example, there are often hundreds or sometimes even thousands of concurrent auctions running worldwide selling such substitutable items¹. Against this background, it is important to develop bidding strategies that agents can use to operate effectively across a wide number of auctions. To this end, in this paper we devise and analyse optimal bidding strategies for a bidder that participates in multiple, simultaneous second-price auctions for goods that are perfect substitutes.

To date, much of the existing literature on simultaneous auctions focuses either on complementarities, where the value of items together is greater than the sum of the individual items, or on heuristic strategies for simultaneous auctions (see Section 6 for more details). In contrast, here

¹To illustrate, at the time of writing, over one thousand eBay auctions were selling the iPod mini 4GB.

we consider bidding strategies analytically and for the case of perfect substitutes. In particular, our focus is on simultaneous Vickrey or second-price sealed bid auctions. We choose these because they are communication efficient and well known for their capacity to induce truthful bidding [10], which makes them suitable for many multi-agent system settings. Within this setting, we are able to characterise, for the first time, a bidder's utility-maximising strategy for bidding in any number of such auctions and for any type of bidder valuation distribution.

In more detail, we first consider a market where a single bidder, called the *global bidder*, can bid in any number of auctions, whereas the other bidders, called the *local bidders*, are assumed to bid only in a single auction. For this case, we find the following results:

- Whereas in the case of a single second-price auction a bidder's best strategy is to bid its true value, this is generally not the case for a global bidder. As we shall show, its best strategy is in fact to bid below the true value.
- We are able to prove that, even if a global bidder requires only one item and assuming free disposal, the expected utility is maximised by participating in all the auctions that are selling the desired item.
- Finding the optimal bid for each auction can be an arduous task when considering all possible combinations. However, for most common bidder valuation distributions, we are able to significantly reduce this search space.
- Empirically, we find that a bidder's expected utility is maximised by bidding relatively high in one of the auctions, and equal or lower in all other auctions.

We then go on to consider markets with more than one global bidder. Due to the complexity of the problem, we combine analytical results with a discrete simulation in order to numerically derive the optimal bidding strategy. By so doing, we find that, in a market with only global bidders, the dynamics of the best response do not converge to a pure strategy. In fact it fluctuates between two states. If the market consists of both local and global bidders, however, the global bidders' strategy quickly reaches a stable solution and we approximate a symmetric Nash equilibrium outcome.

Finally, we consider the issue of market efficiency when there are such simultaneous auctions. Efficiency is an important system-wide consideration within multi-agent systems since it characterises how well the allocations in the

system maximise the overall utility [5]. Now, efficiency is maximised when the goods are allocated to those who value them the most. However, a certain amount of inefficiency is inherent to a distributed market where the auctions are held separately. In this paper, we measure the inefficiency of markets with local bidders only and consider the impact of global bidders on this inefficiency. In so doing, we find that the presence of a global bidder has a slight, but positive, impact on the efficiency when the number of local bidders is known, but is, in general, negative when there exists uncertainty about the exact number of bidders. Therefore, information about the market, such as the number of bidders, plays an important role in the social welfare of the system.

The remainder of the paper is structured as follows. In Section 2 we describe the bidders and the auctions in more detail. In Section 3 we investigate the case with a single global bidder and characterise the optimal bidding behaviour for it. Section 4 considers the case with multiple global bidders and in Section 5 we address the market efficiency and the impact of a global bidder. Finally, Section 6 discusses related work and Section 7 concludes.

2. BIDDING IN MULTIPLE VICKREY AUCTIONS

The model consists of M sellers, each of whom acts as an auctioneer. Each seller auctions one item; these items are complete substitutes (i.e., they are equal in terms of value and a bidder obtains no additional benefit from winning more than one item). The M auctions are executed simultaneously; that is, they end simultaneously and no information about the outcome of any of the auctions becomes available until the bids are placed². We also assume that all the auctions are symmetric, i.e., a bidder does not prefer one auction over another. Finally, we assume free disposal and risk neutral bidders.

2.1 The Auctions

The seller's auction is implemented as a second-price sealed bid auction, where the highest bidder wins but pays the second-highest price. This format has several advantages for an agent-based setting. Firstly, it is communication efficient. Secondly, for the single-auction case (i.e., where a bidder places a bid in at most one auction), the optimal strategy is to bid the true value and thus requires no computation (once the valuation of the item is known). This strategy is also weakly dominant (i.e., it is independent of the other bidders' decisions), and therefore it requires no information about the preferences of other agents (such as the distribution of their valuations).

2.2 Global and Local Bidders

We distinguish between global bidders and local bidders. The former can bid in any number of auctions, whereas the latter only bid in a single auction. Local bidders are assumed to bid according to the weakly dominant strategy and bid their true valuation³. We consider two ways of modelling

local bidders: *static* and *dynamic*. In the first model, the number of local bidders is assumed to be known and equal to N for each auction. In the latter model, on the other hand, the average number of bidders is equal to N , but the exact number is unknown and may vary for each auction. This uncertainty is modelled using a Poisson distribution (more details are provided in Section 3.1).

As we will later show, a global bidder that bids optimally has a higher expected utility compared to a local bidder, even though the items are complete substitutes and a bidder only requires one of them. Nevertheless, we can identify a number of compelling reasons why not all bidders would choose to bid globally:

- **Participation Costs.** Although the bidding itself may be automated by an autonomous agent, it still takes time and/or money, such as entry fees and time to setup an account, to participate in a new auction. Occasional users may not be willing to make such an investment, and they may restrict themselves to sellers or auctions that they are familiar with.
- **Information.** Bidders may simply not be aware of other auctions selling the same type of item. Even if this is known, however, a bidder may not have sufficient information about the distribution of the valuations of other bidders and the number of participating bidders. Whereas this information is not required when bidding in a single auction (because of the dominance property in a second-price auction), it is important when bidding in multiple simultaneous auctions. Such information can be obtained by an expert user or be learned over time, but is often not available to a novice.
- **Risk Attitude.** Although a global bidder obtains a higher utility on average, such a bidder runs a risk of incurring a loss (i.e., a negative utility) when winning multiple auctions. A risk averse bidder may not be willing to take that chance, and so may choose to participate only in a single auction to avoid such a potential loss.
- **Budget Constraints.** Related to the previous point, a budget constrained bidder may not have sufficient funds to make a loss in case it wins more than one auction. In more detail, for a fixed budget b , the sum of bids should not exceed b , thereby limiting the number of auctions a bidder can participate in and/or lowering the actual bids that are placed in those auctions.
- **Bounded Rationality.** As will become clear from this paper, an optimal strategy for a global bidder is harder to compute than a local one. A bidder will therefore only bid globally if the costs of computing the optimal strategy outweigh the benefits of the additional utility.

From the above, we believe it is reasonable to expect a combination of global and local bidders, and for only a few of them to be global bidders. In this paper, we analyse the case of a single global bidder theoretically, and then use a computational approach to address the case with at least two such bidders.

²Although this paper focuses on sealed-bid auctions, where this is the case, the conditions are similar for last-minute bidding in iterative auctions such as eBay [15].

³Note that, since bidding the true value is optimal for local bidders irrespective of what others are bidding, their strategy is not affected by the presence of global bidders.

3. A SINGLE GLOBAL BIDDER

In this section, we provide a theoretical analysis of the optimal bidding strategy for a global bidder, given that all other bidders are local and simply bid their true valuation. After we describe the global bidder's expected utility in Section 3.1, we show in Section 3.2 that it is always optimal for a global bidder to participate in the maximum number of auctions available. Subsequently, in Section 3.3 we discuss how to significantly reduce the complexity of finding the optimal bids for the multi-auction problem, and we then apply these methods to find optimal strategies for specific examples.

3.1 The Global Bidder's Expected Utility

We use the following notation. The number of sellers (or auctions) is $M \geq 2$ and the number of *local* bidders is $N \geq 1$. A bidder's valuation $v \in [0, v_{max}]$ is randomly drawn from a cumulative distribution F with probability density f , where f is continuous, strictly positive and has support $[0, v_{max}]$. F is assumed to be equal and common knowledge for all bidders. A global bid \mathcal{B} is a set containing a bid $b_i \in [0, v_{max}]$ for each auction $1 \leq i \leq M$ (the bids may be different for different auctions). For ease of exposition, we introduce the cumulative distribution function for the first-order statistics $G(b) = F(b)^N \in [0, 1]$, denoting the probability of winning a specific auction conditional on placing bid b in this auction, and its probability density $g(b) = dG(b)/db = NF(b)^{N-1}f(b)$. Now, the expected utility U for a global bidder with global bid \mathcal{B} and valuation v is given by:

$$U(\mathcal{B}, v) = v \left[1 - \prod_{b_i \in \mathcal{B}} (1 - G(b_i)) \right] - \sum_{b_i \in \mathcal{B}} \int_0^{b_i} yg(y)dy \quad (1)$$

Here, the left part of the equation is the valuation multiplied by the probability that the global bidder wins *at least one* of the M auctions and thus corresponds to the expected benefit. In more detail, note that $1 - G(b_i)$ is the probability of *not* winning auction i when bidding b_i , $\prod_{b_i \in \mathcal{B}} (1 - G(b_i))$ is the probability of not winning any auction, and thus $1 - \prod_{b_i \in \mathcal{B}} (1 - G(b_i))$ is the probability of winning at least one auction. The right part of equation 1 corresponds to the total expected costs or payments. To see the latter, note that the expected payment of a single second-price auction when bidding b equals $\int_0^b yg(y)dy$ (see [10]) and is independent of the expected payments for other auctions.

Clearly, equation 1 applies to the model with *static* local bidders, i.e., where the number of bidders is known and equal for each auction (see Section 2.2). However, we can use the same equation to model *dynamic* local bidders in the following way:

LEMMA 1. *By replacing the first-order statistic $G(y)$ with*

$$\hat{G}(y) = e^{N(F(y)-1)}, \quad (2)$$

and the corresponding density function $g(y)$ with $\hat{g}(y) = d\hat{G}(y)/dy = Nf(y)e^{N(F(y)-1)}$, equation 1 becomes the expected utility where the number of local bidders in each auction is described by a Poisson distribution with average N , i.e., where the probability that n local bidders participate is given by $P(n) = N^n e^{-N}/n!$.

PROOF. To prove this, we first show that $G(\cdot)$ and $F(\cdot)$ can be modified such that the number of bidders per auction

is given by a *binomial* distribution (where a bidder's decision to participate is given by a Bernoulli trial) as follows:

$$G'(y) = F'(y)^N = (1 - p + pF(y))^N, \quad (3)$$

where p is the probability that a bidder participates in the auction, and N is the total number of bidders. To see this, note that not participating is equivalent to bidding zero. As a result, $F'(0) = 1 - p$ since there is a $1 - p$ probability that a bidder bids zero at a specific auction, and $F'(y) = F'(0) + pF(y)$ since there is a probability p that a bidder bids according to the original distribution $F(y)$. Now, the average number of participating bidders is given by $N = pN$. By replacing p with N/N , equation 3 becomes $G'(y) = (1 - N/N + (N/N)F(y))^N$. Note that a Poisson distribution is given by the limit of a binomial distribution. By keeping N constant and taking the limit $N \rightarrow \infty$, we then obtain $G'(y) = e^{N(F(y)-1)} = \hat{G}(y)$. This concludes our proof. \square

The results that follow apply to both the static and dynamic model unless stated otherwise.

3.2 Participation in Multiple Auctions

We now show that, for any valuation $0 < v < v_{max}$, a utility-maximising global bidder should always place non-zero bids in all available auctions. To prove this, we show that the expected utility increases when placing an arbitrarily small bid compared to not participating in an auction. More formally,

THEOREM 1. *Consider a global bidder with valuation $0 < v < v_{max}$ and global bid \mathcal{B} , where $b_i \leq v$ for all $b_i \in \mathcal{B}$. Suppose $b_j \notin \mathcal{B}$ for $j \in \{1, 2, \dots, M\}$, then there exists a $b_j > 0$ such that $U(\mathcal{B} \cup \{b_j\}, v) > U(\mathcal{B}, v)$.*

PROOF. Using equation 1, the marginal expected utility for participating in an additional auction can be written as:

$$U(\mathcal{B} \cup \{b_j\}, v) - U(\mathcal{B}, v) = vG(b_j) \prod_{b_i \in \mathcal{B}} (1 - G(b_i)) - \int_0^{b_j} yg(y)dy$$

Now, using integration by parts, we have $\int_0^{b_j} yg(y)dy = b_jG(b_j) - \int_0^{b_j} G(y)dy$ and the above equation can be rewritten as:

$$U(\mathcal{B} \cup \{b_j\}, v) - U(\mathcal{B}, v) = G(b_j) \left[v \prod_{b_i \in \mathcal{B}} (1 - G(b_i)) - b_j \right] + \int_0^{b_j} G(y)dy \quad (4)$$

Let $b_j = \epsilon$, where ϵ is an arbitrarily small strictly positive value. Clearly, $G(b_j)$ and $\int_0^{b_j} G(y)dy$ are then both strictly positive (since $f(y) > 0$). Moreover, given that $b_i \leq v < v_{max}$ for $b_i \in \mathcal{B}$ and that $v > 0$, it follows that $v \prod_{b_i \in \mathcal{B}} (1 - G(b_i)) > 0$. Now, suppose $b_j = \frac{1}{2}v \prod_{b_i \in \mathcal{B}} (1 - G(b_i))$, then $U(\mathcal{B} \cup \{b_j\}, v) - U(\mathcal{B}, v) = G(b_j) \left[\frac{1}{2}v \prod_{b_i \in \mathcal{B}} (1 - G(b_i)) \right] + \int_0^{b_j} G(y)dy > 0$ and thus $U(\mathcal{B} \cup \{b_j\}, v) > U(\mathcal{B}, v)$. This completes our proof. \square

3.3 The Optimal Global Bid

A general solution to the optimal global bid requires the maximisation of equation 1 in M dimensions, an arduous task, even when applying numerical methods. In this section, however, we show how to reduce the entire bid space to two dimensions in most cases (one continuous, and one discrete), thereby significantly simplifying the problem at

hand. First, however, in order to find the optimal solutions to equation 1, we set the partial derivatives to zero:

$$\frac{\partial U}{\partial b_i} = g(b_i) \left[v \prod_{b_j \in \mathcal{B} \setminus \{b_i\}} (1 - G(b_j)) - b_i \right] = 0 \quad (5)$$

Now, equality 5 holds either when $g(b_i) = 0$ or when $\prod_{b_j \in \mathcal{B} \setminus \{b_i\}} (1 - G(b_j)) v - b_i = 0$. In the dynamic model, $g(b_i)$ is always greater than zero, and can therefore be ignored (since $g(0) = Nf(0)e^{-N}$ and we assume $f(y) > 0$). In case of the static model, $g(b_i) = 0$ only when $b_i = 0$. However, theorem 1 shows that the optimal bid is non-zero for $0 < v < v_{max}$. Therefore, we can ignore the first part, and the second part yields:

$$b_i = v \prod_{b_j \in \mathcal{B} \setminus \{b_i\}} (1 - G(b_j)) \quad (6)$$

In other words, the optimal bid in auction i is equal to the bidder's valuation multiplied by the *probability of not winning any of the other auctions*. It is straightforward to show that the second partial derivative is negative, confirming that the solution is indeed a maximum when keeping all other bids constant. Thus, equation 6 provides a means to derive the optimal bid for auction i , given the bids in all other auctions.

3.3.1 Reducing the Search Space

In what follows, we show that, for non-decreasing probability density functions, such as the uniform and logarithmic distributions, the optimal global bid consists of at most two different values for any $M \geq 2$. That is, the search space for finding the optimal bid can then be reduced to two continuous values. Let these values be b_{high} and b_{low} , where $b_{high} \geq b_{low}$. More formally:

THEOREM 2. *Suppose the probability density function f is non-decreasing within the range $[0, v_{max}]$, then the following proposition holds: given $v > 0$, for any $b_i \in \mathcal{B}$, either $b_i = b_{high}$, $b_i = b_{low}$, or $b_i = b_{high} = b_{low}$.*

PROOF. Using equation 6, we can produce M equations, one for each auction, with M unknowns. Now, by combining these equations, we obtain the following relationship: $b_1(1 - G(b_1)) = b_2(1 - G(b_2)) = \dots = b_m(1 - G(b_m))$. By defining $H(b) = b(1 - G(b))$ we can rewrite the equation to:

$$H(b_1) = H(b_2) = \dots = H(b_m) = v \prod_{b_j \in \mathcal{B}} (1 - G(b_j)) \quad (7)$$

In order to prove that there exist at most two different bids, it is sufficient to show that $b = H^{-1}(y)$ has at most two solutions that satisfy $0 \leq b \leq v_{max}$ for any y . To see this, suppose $H^{-1}(y)$ has two solutions but there exists a third bid $b_j \neq b_{low} \neq b_{high}$. From equation 7 it then follows that there exists a y such that $H(b_j) = H(b_{low}) = H(b_{high}) = y$. Therefore, $H^{-1}(y)$ must have at least three solutions, which is a contradiction.

Now, note that, in order to prove that $H^{-1}(y)$ has at most two solutions, it is sufficient to show that $H(b)$ is strictly concave⁴ for $0 \leq b \leq v_{max}$. The function H is strictly

⁴More precisely, $H(b)$ can be either strictly convex or strictly concave. However, it is easy to see that H is not convex since $H(0) = H(v_{max}) = 0$, and $H(b) \geq 0$ for $0 < b < v_{max}$.

concave if and only if the following holds:

$$\frac{d^2 H}{db^2} = \frac{d}{db} (1 - b \cdot g(b) - G(b)) = -b \frac{dg}{db} + 2g(b) < 0$$

By performing standard calculations, we obtain the following condition for the static model:

$$b (N - 1) \frac{f(b)^N}{F(b)} + N \frac{f'(b)}{f(b)} > -2 \text{ for } 0 \leq b \leq v_{max}, \quad (8)$$

and similarly for the dynamic model we have:

$$b N f(b) + \frac{f'(b)}{f(b)} > -2 \text{ for } 0 \leq b \leq v_{max}, \quad (9)$$

where $f'(b) = df/db$. Since both f and F are positive, conditions 8 and 9 clearly hold for $f'(b) \geq 0$. In other words, conditions 8 and 9 show that $H(b)$ is strictly concave when the probability density function is non-decreasing for $0 \leq b \leq v_{max}$, completing our proof. \square

Note from conditions 8 and 9 that the requirement of non-decreasing density functions is sufficient, but far from necessary. Although we are as yet not able to make a more precise formal characterisation, in practice even most density functions with decreasing parts satisfy these conditions. Moreover, the requirement for $H(b)$ to be strictly concave is also stronger than necessary in order to guarantee only two solutions. As a result, for practical purposes, we expect the reduction of the search space to apply in most cases.

Given there are at most 2 possible bids, b_{low} and b_{high} , we can further reduce the search space by expressing one bid in terms of the other. Suppose the buyer places a bid of b_{low} in M_{low} auctions and b_{high} for the remaining $M_{high} = M - M_{low}$ auctions, equation 6 then becomes:

$$b_{low} = v(1 - G(b_{low}))^{M_{low}-1} (1 - G(b_{high}))^{M_{high}},$$

and can be rearranged to give:

$$b_{high} = G^{-1} \left(1 - \frac{b_{low}}{v(1 - G(b_{low}))^{M_{low}-1}} \right)^{\frac{1}{M_{high}}} \quad (10)$$

Here, the inverse function $G^{-1}(\cdot)$ can usually be obtained quite easily. Furthermore, note that, if $M_{low} = 1$ or $M_{high} = 1$, equation 6 can be used directly to find the desired value.

Using the above, we are able to reduce the bid search space to a single continuous dimension, given M_{low} or M_{high} . However, we do not know the number of auctions in which to bid b_{low} and b_{high} , and thus we need to search M different combinations to find the optimal global bid. Moreover, for each combination, the optimal b_{low} and b_{high} can vary. Therefore, in order to find the optimal bid for a bidder with valuation v , it is sufficient to search along one continuous variable $b_{low} \in [0, v]$, and a discrete variable $M_{low} = M - M_{high} \in \{1, 2, \dots, M\}$.

3.3.2 Empirical Evaluation

In this section, we present results from an empirical study and characterise the optimal global bid for specific cases. Furthermore, we measure the actual utility improvement that can be obtained when using the global strategy. The results presented here are based on a uniform distribution of the valuations with $v_{max} = 1$, and the static local bidder model, but they generalise to the dynamic model and other distributions (not shown due to space limitations). Figure 1

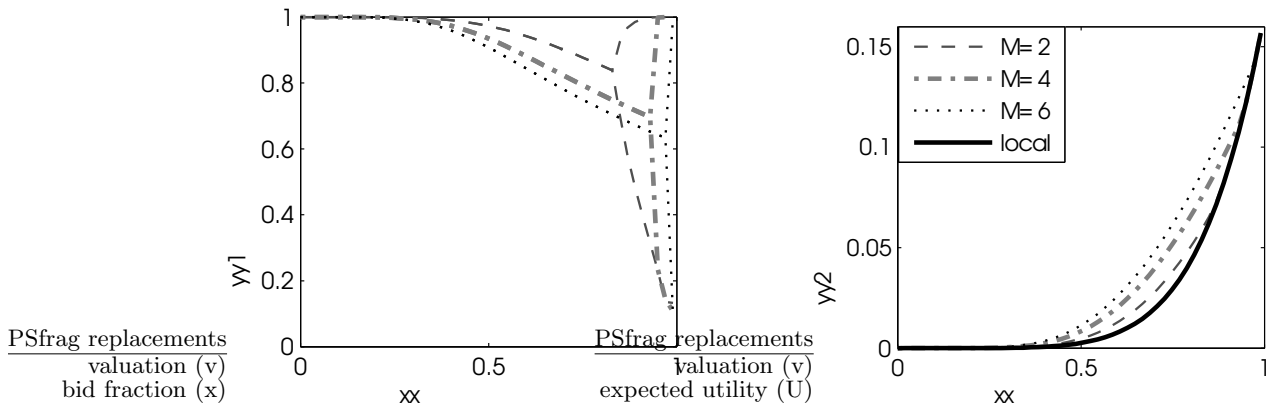


Figure 1: The optimal bid fractions $x = b/v$ and corresponding expected utility for a single global bidder with $N = 5$ static local bidders and varying number of auctions (M). In addition, for comparison, the solid line in the right figure depicts the expected utility when bidding locally in a randomly selected auction, given there are no global bidders (note that, in case of local bidders only, the expected utility is not affected by M).

illustrates the optimal global bids and the corresponding expected utility for various M and $N = 5$, but again the bid curves for different values of M and N follow a very similar pattern. Here, the bid is normalised by the valuation v to give the bid fraction $x = b/v$. Note that, when $x = 1$, a bidder bids its true value.

As shown in Figure 1, for bidders with a relatively low valuation, the optimal strategy is to submit M equal bids at, or very close to, the true value. The optimal bid fraction then gradually decreases for higher valuations. Interestingly, in most cases, placing equal bids is no longer the optimal strategy after the valuation reaches a certain point. At this point, a so-called pitchfork bifurcation is observed and the optimal bids split into two values: a single high bid and $M - 1$ low ones. This transition is smooth for $M = 2$, but exhibits an abrupt jump for $M \geq 3$. In all experiments, however, we consistently observe that the optimal strategy is always to place a high bid in one auction, and an equal or lower bid in all others. In case of a bifurcation and when the valuation approaches v_{max} , the optimal high bid becomes very close to the true value and the low bids go to almost zero⁵.

As illustrated in Figure 1, the utility of a global bidder becomes progressively higher with more auctions. In absolute terms, the improvement is especially high for bidders that have an above average valuation, but not too close to v_{max} . The bidders in this range thus benefit most from bidding globally. This is because bidders with very low valuations have a very small chance of winning any auction, whereas bidders with a very high valuation have a high probability of winning a single auction and benefit less from participating in more auctions. In contrast, if we consider the utility *relative* to bidding in a single auction, this is much higher for bidders with relatively low valuations (this effect cannot be seen clearly in Figure 1 due to the scale). In particular, we notice that a global bidder with a low valuation can improve its utility by up to M times the expected utility

⁵Note in Figure 1 that the low bids are significantly higher than zero at this point. This is because as v approaches v_{max} , the low bids have very little impact on the utility and finding the optimum numerically at this point requires an extremely high precision.

of bidding locally. Intuitively, this is because the chance of winning one of the auctions increases by up to a factor M , whereas the increase in the expected cost is negligible. For high valuation buyers, however, the benefit is not that obvious because the chances of winning are relatively high even in case of a single auction.

4. MULTIPLE GLOBAL BIDDERS

As argued in section 2.2, we expect a real-world market to exhibit a mix of global and local bidders. Whereas so far we assumed a single global bidder, in this section we consider a setting where multiple global bidders interact with one another and with local bidders as well. The analysis of this problem is complex, however, as the optimal bidding strategy of a global bidder depends on the strategy of other global bidders. A typical analytical approach is to find the symmetric Nash equilibrium solution [6, 8, 14, 16], which occurs when all global bidders use the same strategy to produce their bids, and no (global) bidder has any incentive to unilaterally deviate from the chosen strategy. Due to the complexity of the problem, however, here we combine a computational simulation approach with the analytical results from section 3. The simulation works by iteratively finding the best response to the optimal bidding strategies in the previous iteration. If this should result in a stable outcome (i.e., when the current and previous optimal bidding strategies are the same), the solution is by definition a (symmetric) Nash equilibrium.

In more detail, the simulation is based on the observation that the valuation distribution F of the local bidders corresponds to the distribution of bids (since local bidders bid their true valuation). Therefore, by maximising equation 1 we find the best response given the current distribution of bids. Now, we first discretize the space of possible valuations and bids. Then, by performing this maximisation for each bidder *type*, where a bidder type is defined by its (discrete) valuation v , we find a new distribution of bids. Note that this distribution can include bids from any number of both global and local bidders, where the latter simply bid their true valuation. This distribution of bids can then be used to find a new best response, resulting in a new distribution of bids, and so on, for a fixed number of iterations or until

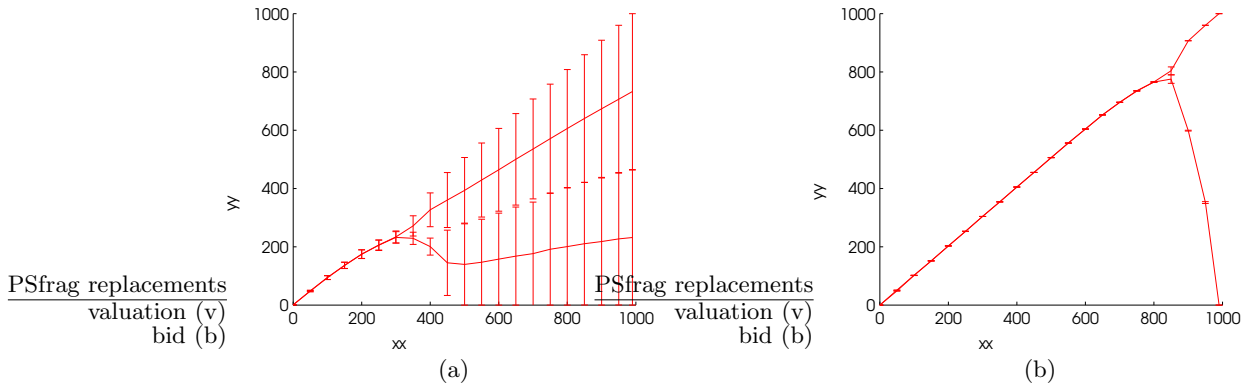


Figure 2: Best response strategy for 2 auctions and 3 global bidders without local bidders (a), and with 10 local bidders (b), averaged over 10 iterations and 20 runs with different initial conditions. The measurements are taken after an initialisation period of 10 iterations. The error-bars indicate the standard deviation.

a stable solution has been found⁶.

In what follows, we first describe the simulation settings, and then apply the simulation to settings with global bidders only, followed by settings with both global and local bidders.

4.1 The Setting

The simulation is based on discrete valuations and bids. The valuations are natural numbers ranging from 1 to $v_{max} \in \mathbb{N}$, where v_{max} is set to 1000. Each valuation $v \in \{1, 2, \dots, v_{max}\}$ occurs with equal probability, equivalent to a uniform valuation distribution in the continuous case. Note, however, that even though the bidder valuations are distributed uniformly, the resulting distribution of bids is typically not uniform (since global bidders typically bid below their valuation). The number of different bid levels that a bidder is allowed is set to $\mathcal{L} \in \mathbb{N}$. Thus, a bidder with valuation v can place the bids $b \in \{v/\mathcal{L}, 2v/\mathcal{L}, \dots, v\}$. For the results reported here, we use $\mathcal{L} = 300$. The initial state can play an important role in the experiments. Therefore, to ensure our results are robust, experiments are repeated with different random initial bid distributions. In the following, we assume the number of local bidders to be static and use N_G and N_L to denote the number of global and local bidders respectively.

4.2 The Results

First, we describe the results with no local bidders (i.e., $N_L = 0$). For this case, we find that the simulation does not converge to a stable state. That is, when the number of (global) bidders is at least 2, the best response strategy keeps fluctuating, irrespective of the number of iterations, and of the initial state. The fluctuations, however, show a distinct pattern and more or less alternate between two states. Figure 2a depicts the average best response strategy for $N_G = 3$ and $M = 2$. Here, the standard deviation is a gauge for the amount of fluctuation and thus the instability of the strategy. In general, we find that the best response for low valuations remain stable, whereas the strategy for bidders with high valuations fluctuates heavily, as is shown in Figure 2a. These results are robust for different initial

⁶This approach is similar to an alternating-move best-response process with pure strategies [7], although here we consider symmetric strategies within a setting where an opponent's best response depends on its type.

conditions and simulation parameters.

If we include local bidders, on the other hand, we observe that the strategies stabilise. Figure 2b shows the simulation results for the same settings as before except with both local and global bidders. As can be seen from this figure, the variation is very slight and only around the bifurcation point. We note that these outcomes are obtained after only a few iterations of the simulation. The results show that the principal conclusions in case of a single global bidder carry over to the case of multiple global bidders. That is, the optimal strategy is to participate in all auctions and to bid high in one auction, and equal or lower in the others. A similar bifurcation point is also observed. These results are also obtained for other values of M , N_L , and N_G . Moreover, the results are very robust to changes to the parameters of the simulation.

To conclude, even though a theoretical analysis proves difficult in case of several global bidders, we can approximate a (symmetric) Nash equilibrium for specific settings using a discrete simulation in case the system consists of both local and global bidders. Our experiments show that, even in the case of multiple global bidders, the best strategy is to bid in multiple auctions. Thus, our simulation can be used as a tool to predict the market equilibrium and to find the optimal bidding strategy for practical settings where we expect a combination of local and global bidders.

5. MARKET EFFICIENCY

Efficiency is an important system-wide property since it characterises to what extent the market maximises social welfare (i.e. the sum of utilities of all agents in the market). To this end, in this section we study the efficiency of markets with either static or dynamic local bidders, and the impact that a global bidder has on the efficiency in these markets. Specifically, efficiency in this context is maximised when the bidders with the M highest valuations in the entire market obtain a single item each. More formally, we define the efficiency of an allocation as:

DEFINITION 1. Efficiency of Allocation. *The efficiency η_K of an allocation K is the obtained social welfare proportional to the maximum social welfare that can be achieved in*




Figure 3: Average efficiency for different market settings as shown in the legend. The error-bars indicate the standard deviation over the 10 runs.

the market and is given by:

$$\eta_K = \frac{\sum_{i=1}^{N_T} v_i(K)}{\sum_{i=1}^{N_T} v_i(K^*)}, \quad (11)$$

where $K^* = \arg \max_{K \in \mathcal{K}} \sum_{i=1}^{N_T} v_i(K)$ is an efficient allocation, \mathcal{K} is the set of all possible allocations, $v_i(K)$ is bidder i 's utility for the allocation $K \in \mathcal{K}$, and N_T is the total number of bidders participating across all auctions (including any global bidders).

Now, in order to measure the efficiency of the market and the impact of a global bidder, we run simulations for the markets with the different types of local bidders. The experiments are carried out as follows. Each bidder's valuation is drawn from a uniform distribution with support $[0, 1]$. The local bidders bid their true valuations, whereas the global bidder bids optimally in each auction as described in Section 3.3. The experiments are repeated 5000 times for each run to obtain an accurate mean value, and the final average results and standard deviations are taken over 10 runs in order to get statistically significant results.

The results of these experiments are shown in Figure 3. Note that a degree of inefficiency is inherent to a multi-auction market with only local bidders [4].⁷ For example, if there are two auctions selling one item each, and the two bidders with the highest valuations both bid locally in the same auction, then the bidder with the second-highest value does not obtain the good. Thus, the allocation of items to bidders is inefficient. As can be observed from Figure 3, however, the efficiency increases when N becomes larger. This is because the differences between the bidders with the highest valuations become smaller, thereby decreasing the loss of efficiency.

Furthermore, Figure 3 shows that the presence of a global bidder has a slightly positive effect on the efficiency in case the local bidders are static. In the case of dynamic bidders, however, the effect of a global bidder depends on the number of sellers. If M is low (i.e., for $M = 2$), a global bidder

⁷An exception is when $N = 1$ and bidders are static, since the market is then completely efficient without a global bidder. However, since this is a very special case and does not apply to other settings, we do not discuss it further here.

significantly increases the efficiency, especially for low values of N . For $M = 6$, on the other hand, the presence of a global bidder has a negative effect on the efficiency (this effect becomes even more pronounced for higher values of M). This result is explained as follows. The introduction of a global bidder potentially leads to a decrease of efficiency since this bidder can unwittingly win more than one item. However, as the number of local bidders increase, this is less likely to happen. Rather, since the global bidder increases the number of bidders, its presence makes an overall positive (albeit small) contribution in case of static bidders. In a market with dynamic bidders, however, the market efficiency depends on two other factors. On the one hand, the efficiency increases since items no longer remain unsold (this situation can occur in the dynamic model when no bidder turns up at an auction). On the other hand, as a result of the uncertainty concerning the actual number of bidders, a global bidder is more likely to win multiple items (we confirmed this analytically). As M increases, the first effect becomes negligible whereas the second one becomes more prominent, reducing the efficiency on average.

To conclude, the impact of a global bidder on the efficiency clearly depends on the information that is available. In case of static local bidders, the number of bidders is known and the global bidder can bid more accurately. In case of uncertainty, however, the global bidder is more likely to win more than one item, decreasing the overall efficiency.

6. RELATED WORK

Research in the area of simultaneous auctions can be segmented along two broad lines. On the one hand, there is the game-theoretic analysis of simultaneous auctions which concentrates on studying the equilibrium strategy of rational agents [6, 11, 12, 14, 16]. Such analyses are typically used when the auction format employed in the simultaneous auctions is the same (e.g. there are N second-price auctions or N first-price auctions). On the other hand, heuristic strategies have been developed for more complex settings when the sellers offer different types of auctions or the buyers need to buy bundles of goods over distributed auctions [1, 2, 9]. This paper adopts the former approach in studying a market of N second-price simultaneous auctions since this approach yields provably optimal bidding strategies.

In this case, the seminal paper by Engelbrecht-Wiggans and Weber [6] provides one of the starting points for the game-theoretic analysis of distributed markets where buyers have substitutable goods. Their work analyses a market consisting of couples having equal valuations that want to bid for a dresser. Thus, the couple's bid space can at most contain two bids since the husband and wife can be at most at two geographically distributed auctions simultaneously. They derive a mixed strategy Nash equilibrium for the special case where the number of buyers is large and also study the efficiency of such a market and show that for local bidders, the market efficiency is $1 - 1/e$. Our analysis differs from theirs in that we study simultaneous auctions in which bidders have different valuations and the global bidder can bid in all the auctions simultaneously (which is entirely possible for online auctions).

Following this, Krishna and Rosenthal [11] then studied the case of simultaneous auctions with complementary goods. They analyse the case of both local and global bidders and characterise the bidding of the buyers and resultant

market efficiency. The setting provided in [11] is further extended to the case of common values by Rosenthal and Wang [14]. However, neither of these works extend easily to the case of substitutable goods which we consider. This case is studied in [16], but the scenario considered is restricted to three sellers and two global bidders and with each bidder having the same value (and thereby knowing the value of other bidders). The space of symmetric mixed equilibrium strategies is derived for this special case, but again our result is more general.

7. CONCLUSIONS

In this paper, we derive utility-maximising strategies for bidding in multiple, simultaneous second-price auctions. We first analyse the case where a single global bidder bids in all auctions, whereas all other bidders are local and bid in a single auction. For this setting, we find the counter-intuitive result that it is optimal to place non-zero bids in all auctions that sell the desired item, even when a bidder only requires a single item and derives no additional benefit from having more. Thus, a potential buyer can considerably benefit by participating in multiple auctions and employing an optimal bidding strategy. For most common valuation distributions, we show analytically that the problem of finding optimal bids reduces to two dimensions. This considerably simplifies the original optimisation problem and can thus be used in practice to compute the optimal bids for any number of auctions.

Furthermore, we investigate a setting with multiple global bidders by combining analytical solutions with a simulation approach. We find that a global bidder's strategy does not stabilise when only global bidders are present in the market, but only converges when there are local bidders as well. We argue, however, that real-world markets are likely to contain both local and global bidders. The converged results are then very similar to the setting with a single global bidder, and we find that a bidder benefits by bidding optimally in multiple auctions. For the more complex setting with multiple global bidders, the simulation can thus be used to find these bids for specific cases.

Finally, we compare the efficiency of a market with multiple simultaneous auctions with and without a global bidder. We show that, if the bidder can accurately predict the number of local bidders in each auction, the efficiency slightly increases. In contrast, if there is much uncertainty, the efficiency significantly decreases as the number of auctions increases due to the increased probability that a global bidder wins more than two items. These results show that the way in which the efficiency, and thus social welfare, is affected by a global bidder depends on the information that is available to that global bidder.

In future work, we intend to expand our analysis for markets with more than one global bidder, and extend the model to investigate optimal strategies for purchasing multiple units of an item, and when the auctions are no longer symmetric. The latter arises, for example, when the number of (average) local bidders differs per auction or the auctions have different settings for parameters such as the reserve price.

Acknowledgements

This research was undertaken as part of the EPSRC funded project on Market-Based Control (GR/T10664/01). This is a collaborative project involving the Universities of Bir-

ingham, Liverpool and Southampton and BAE Systems, BT and HP. This research was also undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Systems) project and is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership. In addition, we would like to thank Alex Rogers, Adam Prügel-Bennett, and the anonymous referees for their insightful comments and suggestions.

8. REFERENCES

- [1] S. Airiau and S. Sen. Strategic bidding for multiple units in simultaneous and sequential auctions. *Group Decision and Negotiation*, 12(5):397 – 413, October 2003.
- [2] A. Byde, C. Preist, and N.R. Jennings. Decision procedures for multiple auctions. In *Proc. 1st Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS2002)*, pages 613–620, 2002.
- [3] S.H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing, 1996.
- [4] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
- [5] R. K. Dash, D. C. Parkes, and N. R. Jennings. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- [6] R. Engelbrecht-Wiggans and R. Weber. An example of a multiobject auction game. *Management Science*, 25:1272–1277, 1979.
- [7] D. Fudenberg and D.K. Levine. *The Theory of Learning in Games*. MIT Press, 1999.
- [8] E. H. Gerding, A. Rogers, R. K. Dash, and N. R. Jennings. Competing sellers in online markets: Reserve prices, shill bidding, and auction fees. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*, 2006.
- [9] A. Greenwald, R.M. Kirby, J. Reiter, and J. Boyan. Bid determination in simultaneous auctions: A case study. In *Proc. of the Third ACM Conference on Electronic Commerce*, pages 115–124, 2001.
- [10] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [11] V. Krishna and R. Rosenthal. Simultaneous auctions with synergies. *Games and Economic Behaviour*, 17:1–31, 1996.
- [12] K. Land and R. Rosenthal. The contractor's game. *RAND J. Econ*, 22:329–338, 1991.
- [13] A. Rogers, E. David, and N.R. Jennings. Self organised routing for wireless micro-sensor networks. *IEEE Trans. on Systems, Man and Cybernetics: Part A*, 35(3):349–359, 2005.
- [14] R. Rosenthal and R. Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behaviour*, 17:32–55, 1996.
- [15] A.E. Roth and A. Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *The American Economic Review*, 92(4):1093–1103, 2002.
- [16] B. Szentes and R. Rosenthal. Three-object two-bidder simultaneous auctions: chopsticks and tetrahedra. *Games and Economic Behaviour*, 44:114–133, 2003.