

# Bidding Optimally in Concurrent Second-Price Auctions of Perfectly Substitutable Goods

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

University of Southampton, Southampton, SO17 1BJ, UK.  
{eg,rkd,dy,nrj}@ecs.soton.ac.uk

## ABSTRACT

We derive optimal bidding strategies for a global bidding agent that 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. This analysis extends to online markets where, typically, auctions occur both concurrently and sequentially. 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.

## Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems; J.4 [Social and Behavioral Sciences]: Economics

## General Terms

Economics

## Keywords

Vickrey Auctions, Simultaneous Auctions, Market Efficiency

## 1. INTRODUCTION

The recent surge of interest in online auctions has resulted in an increasing number of auctions offering very similar or

even identical goods and services [9, 10]. In eBay alone, for example, there are often hundreds or sometimes even thousands of concurrent auctions running worldwide selling such substitutable items<sup>1</sup>. Against this background, it is essential to develop bidding strategies that autonomous 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 an important yet barely studied setting — namely, an agent that participates in multiple, concurrent (i.e., simultaneous) second-price auctions for goods that are perfect substitutes. As we will show, however, this analysis is also relevant to a wider context where auctions are conducted sequentially, as well as concurrently.

To date, much of the existing literature on multiple auctions focuses either on sequential auctions [6] or on simultaneous auctions for complementary goods, where the value of items together is greater than the sum of the individual items (see Section 2 for related research on simultaneous auctions). In contrast, here we consider bidding strategies for markets with multiple concurrent auctions and perfect substitutes. In particular, our focus is on Vickrey or second-price sealed bid auctions. We choose these because they require little communication and are well known for their capacity to induce truthful bidding, which makes them suitable for many multi-agent system settings. However, our results generalise to settings with English auctions since these are strategically equivalent to second-price auctions. Within this setting, we are able to characterise, for the first time, a bidder's utility-maximising strategy for bidding simultaneously 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, the best strategy for a global bidder is to bid below it.
- We are able to prove that, even if a global bidder requires only one item, 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 distribu-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS '07 Honolulu, Hawaii  
Copyright 2007 IFAAMAS.

<sup>1</sup>To illustrate, at the time of writing, within eBay alone over 1300 auctions were selling the Apple iPod Video 30GB.

tions, we are able to significantly reduce this search space and thus the computation required.

- 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.

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

## 2. 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 and decision-theoretic analysis of simultaneous auctions which concentrates on studying the equilibrium strategy of rational agents [3, 7, 8, 9, 12, 11]. Such analyses are typically used when the auction format employed in the concurrent auctions is the same (e.g. there are  $M$  Vickrey auctions or  $M$  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, 13, 5]. This paper adopts the former approach in studying a market of  $M$  simultaneous Vickrey auctions since this approach yields provably optimal bidding strategies.

In this case, the seminal paper by Engelbrecht-Wiggans and Weber 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. Our analysis differs from theirs in that we study concurrent auctions in which bidders have different valuations and the global bidder can bid in all the auctions concurrently (which is entirely possible given autonomous agents).

Following this, [7] 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 [7] is further extended to the case of common values in [9]. However, neither of these works extend easily to the case of substitutable goods which we consider. This case is studied in [12], 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. Finally, [11] considers the case of concurrent *English* auctions, in which he develops bidding algorithms for buyers with different risk attitudes. However, he forces the bids to be the same across auctions, which we show in this paper not always to be optimal.

## 3. BIDDING IN MULTIPLE 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 concurrently; that is, they end simultaneously and no information about the outcome of any of the auctions becomes available until the bids are placed<sup>2</sup>. However, we briefly address markets with both sequential and concurrent auctions in Section 4.4. We also assume that all the auctions are equivalent (i.e., a bidder does not prefer one auction over another). Finally, we assume free disposal (i.e., a winner of multiple items incurs no additional costs by discarding unwanted ones) and risk neutral bidders.

### 3.1 The Auctions

The seller's auction is implemented as a Vickrey 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).

### 3.2 Global and Local Bidders

We distinguish between global and local bidders. The former can bid in any number of auctions, whereas the latter only bid in a single one. Local bidders are assumed to bid according to the weakly dominant strategy and bid their true valuation<sup>3</sup>. 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 4.1).

As we will later show, a global bidder who 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. However, we can identify a number of compelling reasons why not all bidders would choose to bid globally. Firstly, participation costs such as entry fees and time to set up an account may encourage occasional users to

<sup>2</sup>Although this paper focuses on sealed-bid auctions, where this is the case, the conditions are similar for last-minute bidding in English auctions such as eBay [10].

<sup>3</sup>Note 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.

participate in auctions that they are already familiar with. Secondly, bidders may simply not be aware of other auctions selling the same type of item. Even if this is known, however, additional information such as the distribution of the valuations of other bidders and the number of participating bidders is required for bidding optimally across multiple auctions. This lack of expert information often drives a novice to bid locally. Thirdly, an optimal global strategy is harder to compute than a local one. An agent with bounded rationality may therefore not have the resources to compute such a strategy. Lastly, even though a global bidder profits on *average*, such a bidder may incur a loss when inadvertently winning multiple auctions. This deters bidders who are either risk averse or have budget constraints from participating in multiple auction. As a result, in most market places we expect a combination of global and local bidders.

In view of the above considerations, human buyers are more likely to bid locally. The global strategy, however, can be effectively executed by autonomous agents since they can gather data from many auctions and perform the required calculations within the desired time frame.

## 4. 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 4.1, we show in Section 4.2 that it is always optimal for a global bidder to participate in the maximum number of auctions available. In Section 4.3 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. Finally, in Section 4.4 we extend our analysis to sequential auctions.

### 4.1 The Global Bidder's Expected Utility

In what follows, the number of sellers (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 [6]) 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 3.2). 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:

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

where  $p$  is the probability that a bidder participates in the auction, and  $\mathcal{N}$  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  $N = p\mathcal{N}$ . By replacing  $p$  with  $N/\mathcal{N}$ , equation 3 becomes  $G'(y) = (1 - N/\mathcal{N} + (N/\mathcal{N})F(y))^{\mathcal{N}}$ . Note that a Poisson distribution is given by the limit of a binomial distribution. By keeping  $N$  constant and taking the limit  $\mathcal{N} \rightarrow \infty$ , 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.

### 4.2 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  $\mathcal{B}$  contains no bid for auction  $j \in \{1, 2, \dots, M\}$ , 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  $\int_0^{b_j} yg(y)dy = b_j G(b_j) - \int_0^{b_j} G(y)dy$  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  $\epsilon$  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.  $\square$

### 4.3 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 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.

#### 4.3.1 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 necessary and sufficient to show that  $H(b)$  has at most one local maximum for  $0 \leq b \leq v_{max}$ . A sufficient conditions, however, is for  $H(b)$  to be strictly concave<sup>4</sup>. The function  $H$  is strictly concave if and only if the following condition holds:

$$H''(b) = \frac{d}{db} (1 - b \cdot g(b) - G(b)) = - \left( b \frac{dg}{db} + 2g(b) \right) < 0 \quad (8)$$

where  $H''(b) = d^2 H/db^2$ . By performing standard calculations, we obtain the following condition for the static model:

$$b \left( (N-1) \frac{f(b)}{F(b)} + N \frac{f'(b)}{f(b)} \right) > -2 \text{ for } 0 \leq b \leq v_{max}, \quad (9)$$

and similarly for the dynamic model we have:

$$b \left( N f(b) + \frac{f'(b)}{f(b)} \right) > -2 \text{ for } 0 \leq b \leq v_{max}, \quad (10)$$

where  $f'(b) = df/db$ . Since both  $f$  and  $F$  are positive, conditions 9 and 10 clearly hold for  $f'(b) \geq 0$ . In other words, conditions 9 and 10 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 9 and 10 that the requirement of non-decreasing density functions is sufficient, but far from necessary. Moreover, condition 8 requiring  $H(b)$  to be strictly concave is also stronger than necessary to guarantee only two solutions. As a result, in practice we find that the reduction of the search space applies to 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:

$$b_{high} = G^{-1} \left( 1 - \left[ \frac{b_{low}}{v(1 - G(b_{low}))^{M_{low}-1}} \right]^{\frac{1}{M_{high}}} \right) \quad (11)$$

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

<sup>4</sup>More 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  $0 < b < v_{max}$ .



**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 dark 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  $M$ ).**

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  $M_{low} = M - M_{high} \in \{1, 2, \dots, M\}$ .

#### 4.3.2 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 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. A so-called pitchfork bifurcation is then 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 goes to the true value and the low bids go to zero.

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 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.

#### 4.4 Sequential and Concurrent Auctions

In this section we extend our analysis of the optimal bidding strategy to sequential auctions. Specifically, the auction process consists of  $R$  rounds, and in each round any number of auctions are running simultaneously. Such a combination of sequential and concurrent auctions is very common in practice, especially online<sup>5</sup>. It turns out that the analysis for the case of simultaneous auctions is quite general and can be easily extended to include sequential auctions. In the following, the number of simultaneous auctions in round  $r$  is denoted by  $M_r$ , and the set of bids in that round by  $\mathcal{B}_r$ . As before, the analysis assumes that all other bidders are local and bid in a single auction. Furthermore, we assume that the global bidders have complete knowledge about the number of rounds and the number of auctions in each round.

The expected utility in round  $r$ , denoted by  $U_r$ , is similar to before (equation 1 in Section 4.1) except that now additional benefit can be obtained from future auctions if the desired item is not won in one of the current set of simultaneous auctions. For convenience,  $U_r(\mathcal{B}_r, \mathcal{M}_r)$  is abbreviated to  $U_r$  in the following. The expected utility thus becomes:

$$U_r = v \cdot P_r(\mathcal{B}_r) - \sum_{b_{ri} \in \mathcal{B}_r} \int_0^{b_{ri}} yg(y)dy + U_{r+1} \cdot (1 - P_r(\mathcal{B}_r)) \\ = U_{r+1} + (v - U_{r+1})P_r(\mathcal{B}_r) - \sum_{b_{ri} \in \mathcal{B}_r} \int_0^{b_{ri}} yg(y)dy, \quad (12)$$

where  $P_r(\mathcal{B}_r) = 1 - \prod_{b_{ri} \in \mathcal{B}_r} (1 - G(b_{ri}))$  is the probability of winning at least one auction in round  $r$ . Now, we take the partial derivative of equation 12 in order to find the optimal bid  $b_{rj}$  for auction  $j$  in round  $r$ :

$$\frac{\partial U_r}{\partial b_{rj}} = g(b_{rj}) \left[ (v - U_{r+1}) \prod_{b_{ri} \in \mathcal{B}_r \setminus \{b_{rj}\}} (1 - G(b_{ri})) - b_{rj} \right] \quad (13)$$

<sup>5</sup>Rather than being purely sequential in nature, online auctions also often overlap (i.e., new auctions can start while others are still ongoing). In that case, however, it is optimal to wait and bid in the new auctions only after the outcome of the earlier auctions is known, thereby reducing the chance of unwittingly winning multiple items. Using this strategy, overlapping auctions effectively become sequential and can thus be analysed using the results in this section.

Note that equation 13 is almost identical to equation 5 in Section 4.3, except that the valuation  $v$  is now replaced by  $v - U_{r+1}$ . The optimal bidding strategy can thus be found by backward induction (where  $U_{R+1} = 0$ ) using the procedure outlined in Section 4.3.

## 5. MULTIPLE GLOBAL BIDDERS

As argued in section 3.2, 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 [9, 12], 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 analytical results. 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 optimal bidding strategies remains unchanged for two subsequent iterations), the solution is by definition a (symmetric) Nash equilibrium.

### 5.1 The Global Bidder's Expected Utility

In order to find a global bidder's best response, we first need to calculate the expected utility given the global bid  $\mathcal{B}$  and the strategies of both the other global bidders as well as the local bidders. In the following, let  $N_g$  denote the number of other global bidders. Furthermore, let the strategies of the other global bidders be represented by the set of functions  $\beta_k(v)$ ,  $1 \leq k \leq M$ , producing a bid for each auction given a bidder's valuation  $v$ . Note that all other global bidders use the same set of functions since we consider symmetric equilibria. However, we assume that the assignment of functions to auctions by each global bidder occurs in a random fashion without replacement (i.e., each function is assigned exactly once by each global bidder). Let  $\Omega$  denote the set of all possible assignments. Each such assignment  $\omega \in \Omega$  is a  $(M, N_g)$  matrix, where each entry  $\omega_{i,j}$  identifies the function used by global bidder  $j$  in auction  $i$ . Note that the cardinality of  $\Omega$ , denoted by  $|\Omega|$ , is equal to  $M^{N_g}$ . Now, the expected utility is the average expected utility over all possible assignments and is given by:

$$U(\mathcal{B}, v) = \frac{1}{|\Omega|} \sum_{\omega \in \Omega} v \left( 1 - \prod_{b_i \in \mathcal{B}} (1 - \tilde{G}_{\omega_i}(b_i)) \right) - \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{b_i \in \mathcal{B}} \int_0^{b_i} y \tilde{g}_{\omega_i}(y) dy, \quad (14)$$

where  $\tilde{G}_{\omega_i}(b) = G(b) \cdot \prod_{j=1}^{N_g} \int_0^b \beta_{\omega_{i,j}}(y) f(y) dy$  denotes the probability of winning auction  $i$ , given that each global bidder  $1 \leq j \leq N_g$  bids according to the function  $\beta_{\omega_{i,j}}$ , and  $\tilde{g}_{\omega_i}(y) = d\tilde{G}_{\omega_i}(y)/dy$ . Here,  $G(b)$  is the probability of winning an auction with only local bidders as described in Section 4.1, and  $f(y)$  is the probability density of the bidder valuations as before.

### 5.2 The Simulation

The simulation works by discretising the space of possible valuations and bids and then finding a best response to an

initial set of bidding functions. The best response is found by maximising equation 14 for each discrete valuation, which, in turn, results in a new set of bidding functions. These functions then affect the probabilities of winning in the next iteration for which the new best response strategy is calculated. This process is then repeated for a fixed number of iterations or until a stable solution has been found<sup>6</sup>.

Clearly, due to the large search space, finding the utility-maximising global bid quickly becomes infeasible as the number of auctions and global bidders increases. Therefore, we reduce the search space by limiting the global bid to two dimensions where a global bidder bids high in one of the auctions and low in all the others<sup>7</sup>. This simplification is justified by the results in Section 4.3.1 which show that, for a large number of commonly used distributions, the optimal global bid consist of at most two different values.

The results reported here are based on the following settings.<sup>8</sup> In order to emphasize that the valuations are discrete, we use integer values ranging from 1 to 1000. Each valuation occurs with equal probability, equivalent to a uniform valuation distribution in the continuous case. A bidder can select between 300 different equally-spaced bid levels. Thus, a bidder with valuation  $v$  can place bids  $b \in \{0, v/300, 2v/300, \dots, v\}$ . The local bidders are static and bid their valuation as before. The initial set of functions can play an important role in the experiments. Therefore, to ensure our results are robust, experiments are repeated with different random initial functions.


### 5.3 The Results

First, we describe the results with no local bidders. For this case, we find that the simulation does not converge to a stable state. That is, when there is at least one other global bidder, 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 alternate between two states. Figure 2 depicts these two states for  $N_g = 10$  and  $M = 5$ . The two states vary most when there are at least as many auctions as there are global bidders. In that case, one of the best response states is to bid truthfully in one auction and zero in all others. The best response to that, however, is to bid an equal positive amount close to zero in all auctions; this strategy guarantees at least one object at a very low payment. The best response is then again to bid truthfully in a single auction since this appropriates the object in that particular auction. As a result, there exists no stable solution. The same result is observed when the number of global bidders is *less* than the number of auctions. This oc-


<sup>6</sup>This approach is similar to an alternating-move best-response process with pure strategies [4], although here we consider symmetric strategies within a setting where an opponent's best response depends on its valuation.

<sup>7</sup>Note that the number of possible allocations still increases with the number of auctions and global bids. However, by merging all utility-equivalent permutations, we significantly increase computation speed, allowing experiments with relatively large numbers of auctions and bidders to be performed (e.g., a single iteration with 50 auctions and 10 global bidders takes roughly 30 seconds on a 3.00 Ghz PC).

<sup>8</sup>We also performed experiments with different precision, other valuation distributions, and dynamic local bidders. We find that the principal conclusions generalise to these different settings, and therefore we omit the results to avoid repetitiveness.



**Figure 2: The two states of the best response strategy for  $M = 5$  and  $N_g = 10$  without local bidders.**



**Figure 3: The variance of the best response strategy over 10 iterations and 10 experiments with different initial settings and  $M = 5$ . The errorbars show the (small) standard deviations.**

curs since global bidders randomise over auctions, and thus they cannot coordinate and choose to bid high in different auctions.

As shown in Figure 2, a similar fluctuation is observed when the number of global bidders increases relative to the number of auctions. However, the bids in the equal-bid state (state 2 in Figure 2), as well as the low bids of the other state, increase. Moreover, if the number of global bidders is increased even further, a bifurcation occurs in the equal-bid state similar to the case without local bidders.

We now consider the best response strategies when both local and global bidders participate and each auction contains the same number of local bidders. To this end, Figure 3 shows the average variance of the best response strategies. This is measured as the variance of an actual best-response bid over different iterations, and then taking the average over the discrete bidder valuations. Here, the variance is a gauge for the amount of fluctuation and thus the instability of the strategy. As can be seen from this figure, local bidders have a large stabilising effect on the global bidder strategies. As a result, the best response strategy approximates a pure symmetric Nash equilibrium. We note that the results converge after only a few iterations.

The results show that the principal conclusions in the case of a single global bidder carry over to the case of multiple global bidders. That is, the optimal strategy is to bid positive in all auctions (as long as there are at least as many bidders as auctions). Furthermore, a similar bifurcation point is observed. These results are very robust to changes to the auction settings and 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. 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.

## 6. 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  $\eta_K$  of an allocation  $K$  is the obtained social welfare proportional to the maximum social welfare that can be achieved in the market and is given by:

$$\eta_K = \frac{\sum_{i=1}^{N_T} v_i(K)}{\sum_{i=1}^{N_T} v_i(K^*)}, \quad (15)$$


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 4.3. 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 4. Note that a degree of inefficiency is inherent to a multi-auction market with only local bidders [2].<sup>9</sup> 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 4, 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 4 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 significantly increases the efficiency, especially for low values of

<sup>9</sup>Trivial exceptions are when either  $M = 1$  or  $N = 1$  and bidders are static, since the market is then completely efficient without a global bidder.



**Figure 4: Average efficiency for different market settings as shown in the legend. The error-bars indicate the standard deviation over the 10 runs.**

$N$ . 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.

## 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 achieve considerable benefit by participating in multiple auctions and employing an optimal bidding strategy. For a number of 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 concurrent 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 diminishes 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 extend the results to imperfect substitutes (i.e., when a global bidder gains from winning additional items), and to settings where the auctions are no longer identical. 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.

## 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, 2003.
- [2] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
- [3] R. Engelbrecht-Wiggans and R. Weber. An example of a multiobject auction game. *Management Science*, 25:1272–1277, 1979.
- [4] D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1999.
- [5] A. Greenwald, R. 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.
- [6] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [7] V. Krishna and R. Rosenthal. Simultaneous auctions with synergies. *Games and Economic Behaviour*, 17:1–31, 1996.
- [8] K. Lang and R. Rosenthal. The contractor's game. *RAND J. Econ.*, 22:329–338, 1991.
- [9] R. Rosenthal and R. Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behaviour*, 17:32–55, 1996.
- [10] A. 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.
- [11] O. Shehory. Optimal bidding in multiple concurrent auctions. *Int. Journal of Cooperative Information Systems*, 11:315–327, 2002.
- [12] B. Szentes and R. Rosenthal. Three-object two-bidder simultaneous auctions: chopsticks and tetrahedra. *Games and Economic Behaviour*, 44:114–133, 2003.
- [13] D. Yuen, A. Hyde, and N. R. Jennings. Heuristic bidding strategies for multiple heterogeneous auctions. In *Proc. 17th European Conference on AI (ECAI)*, pages 300–304, 2006.