
Modelling Bidders in Sequential Automated Auctions

KUMAARA VELAN

*Intelligent Systems and Networks Group, Department of Electrical and Electronic Engineering, Imperial College, London SW7 2BT, UK
Email: kumaara.velan@imperial.ac.uk*

Auctions are mechanisms that formalise the rules with which automated trading schemes can be conducted, and in this paper we model the interaction of bidder and seller agents in sequential computerised auctions. We study the outcome of strategies that a designated “special bidder” (SB) may follow in the presence of a collection of other bidders in an English auction, under the assumption that the SB can make bids based on its observation of the ongoing auction as a collective system. In our model, bidding and sale events are continuous time random processes with discrete state-space, where the state-space represents the current value of the most recent bid. We obtain analytical solutions which allow the evaluation of measures of interest to the SB such as the probability of winning, the savings with respect to the maximum payable price in the event of a win, and the expected waiting time to win. We examine the effects of the SB’s time to bid, and study how its decisions may be selected so as to optimise the SB’s measures of interest.

Keywords: Automated Auctions; Analytical Models; Autonomic Systems; Internet Technologies; E-Commerce; Queueing Systems

Received 00 January 2009; revised 00 Month 2009

1. INTRODUCTION

Computer systems and the Internet have enabled a wide variety of automated trading schemes which are in use for stocks, commodities, derivatives and other financial instruments. More recently, the Internet has also allowed individuals to buy and sell various items in an open and easily accessible manner. Thus we can envisage a future when large parts of the economy will be driven by sequences of automated and interconnected trading patterns.

Auctions are a convenient mechanism for formalising the rules with which such automating trading schemes can be conducted, and they have been widely used for many centuries in human based trading and commerce. In recent work the stochastic behaviour of collections of bidders acting in an auction has been analysed through the use of discrete state-space and continuous time probability models [?]. The approach constructs stochastic dynamical models of auctions and bidders, and then obtains the steady-state behaviour to compute significant properties of both the “one-shot” outcome, and the long-run repetitive outcome of auctions. One-shot properties include the probability distribution and the expected sale price, while long-term properties include the income per unit time obtained by the seller over a large number of transactions. In [?] the model

has been extended to study a network of interconnected auctions, where bidders are allowed to move freely between auctions, and an analytical solution has been obtained. These mathematical techniques, and their variants, have also been successfully deployed in other diverse range of problem areas: from biological applications in modelling populations of viruses and agents [?], to communication systems in minimising packet travel time across a wireless network [?, ?], and choosing adaptive routing decisions [?].

In reality, knowledge of the existence of opportunity and availability to purchase similar goods in the future may influence a bidder to choose to forgo the opportunity of procuring a good at a high cost, and instead wait in the hope of securing a better deal later. We can also imagine situations where goods are reusable resources that are not sold per se, but rented out and returned to the seller at the end of some period, so that bidder agents who are not in urgent need to obtain a good immediately may defer a purchase. Interesting work along these lines, where bidders exhibit rational forward-looking behaviour when deciding strategies for the current auction can be found in [?]. A possible application for auctions of reusable goods in allocating computing resources is given in [?].

When we discuss automated auctions, we will invariably imagine software agents representing human

counterparts in the digital marketplace [?], acting autonomously and yet guided by its design objectives to fulfill the interests of their owners to the best possible extent. In such instances it is crucial that the underlying communications infrastructure is designed to allow and support the user, i.e. the agent, to set the criteria for its service requirements. For example, a bidder agent may have certain specific needs with respect to its connectivity to the seller, and so may request for a network path with the minimum overall delay to the seller be established, or an agent physically located on a mobile node may prefer a connection that consumes the minimum power, or any weighted combination of its other goals. In this regard, auctions, or any digital marketplace activity for that matter, will benefit from autonomic communications [?] where emphasis is placed to insulate the user experience and services from changes, whether predicted or not, to the underlying infrastructure. In particular, a self-aware network, as proposed and implemented in the *cognitive packet network* (CPN) [?, ?, ?], accomplishes this by online internal probing and measurement mechanisms that are used for self-management, so that it adapts itself to provide the user the best effort quality of service (QoS). The CPN performs these corrective actions by using random neural networks with reinforcement learning [?, ?, ?], and finds newer routes with improved QoS using genetic algorithms [?, ?]. In a wireless mobile ad hoc network environment, where power efficiency is an overriding concern, the CPN has been extended to incorporate power-awareness [?], and the problem of controlling the admission of new users into the network while preserving the QoS of all users has been addressed in [?].

Beyond ensuring QoS for users, at a different level of abstraction, the ongoing research in autonomic communications has a broader objective of providing an intelligent platform for efficient interaction between digital objects such as users and services [?], or as in our context between buyer and seller agents. An important aspect of applying the “intelligence”, as discussed in [?], is in creating new knowledge based on collected raw sensitive data, through a knowledge network which, in turn, can be used to enhance economic efficiency. In this example, it is shown that a good allocation efficiency can be achieved in a trading application where the aggregated knowledge is used to create new markets so that sellers can respond to buyers’ needs as they arise.

Security is another vital feature in a communication network, especially when transactions involving e-commerce activities such as auctions are conducted; the servers running these applications are easy targets for attackers, either as a malicious act of sabotage or for profitable gains. The denial of service (DoS) attack is a particularly critical threat, since it is easy to launch, and by its usually distributed nature, difficult to protect against. Thus an autonomic approach to defending the network, based on self-monitoring and

adaptive measures, has been suggested in [?], and several biologically inspired DoS detectors have been evaluated in [?, ?, ?].

In this paper we extend the work in [?] to study the outcome of strategies that a designated bidder may follow in an English auction, in the presence of a collection of other bidders, under the assumption that this “special bidder” (SB) observes the parameters resulting from the auction as a collective (many bidders and the seller) system. Note that in [?] bidders are lumped together in a pool, where everyone shares a similar behaviour; whereas in this work we propose a generalisation in that the SB is allowed to have its own activity (bidding) rate which may differ from the other bidders’, and examine how the SB should select its bidding rate in a self-serving manner.

We first sketch the model to be studied, and then in Section 2 we analyse it in detail. The manner in which the model provides performance measures of interest to the SB and to the seller, is discussed in Section 3 where we first discuss how the SB can behave in order to optimise outcomes that are in its best interest, and provide numerical examples to illustrate the approach and the model predictions. We then explore how the SB can try to achieve balance and compete with the other bidders in Section 4. Finally Section 5 generalises the analysis to the case where the bidding rates depend on the current price attained in the auction. Conclusions are drawn in Section 6 where we also suggest further work.

1.1. An English auction with $n + 1$ competing bidders

Consider an English auction in which the SB participates with n other bidders. We assume that a single good is being sold, and that it has a maximum valuation $v > 0$, so that none of the bidders will bid beyond the sum v . Initially we assume that v is a fixed and identical valuation for all of the bidders, but it is easy (see [?]) to generalise the results to the case where v is a random variable so that the actual valuation that buyers associate with a good is known in terms of the probability distribution of a random valuation V .

During the auction, each bidder may take some time to consider the current highest offer before deciding to place a new counter-offer. We assume that these thinking times are exponentially distributed random variables with parameters β and λ , for the SB and the other bidders, respectively, so that we may distinguish the behaviour of the SB from all the other bidders which have a common statistical behaviour. All the bidders can participate in submitting bids, except obviously for the bidder who owns the current highest offer. Furthermore we assume that all bids proceed with unit increments with respect to the previous bid, in order to minimally surpass previous highest bid.

It should be noted that, by allowing the SB to have its

own bidding rate, our model generalises the previous [?], and enables us to characterise the system outcomes as a function of the “divergence” of the SB’s behaviour from the other bidders. We examine these outcomes from both the seller’s and the bidder’s interests. As we would expect, if there is no divergence, i.e. we fix $\beta = \lambda$, then this model reduces to the previous.

Once a bid is received, the seller considers the offer for some time before accepting. If a higher bid is submitted before this time expires, then the earlier offer is rejected (and the previous highest bidder rejoins the bidder pool), while the seller waits for another random time, represented by an exponentially distributed random variable with rate parameter δ .

On the other hand, if no new bid arrives by the end of the seller’s waiting time, the auction concludes with a sale to the current and hence highest bidder, and as in [?], the seller is indifferent to the identity of the bidder.

After the sale is successfully concluded, the seller “rests” for some random time and then the auction repeats itself as a statistically independent replica with a population of $n + 1$ bidders. The rest time can be thought of as the time spent in declaring the winner, eliciting payment and allocating the item, followed by the time spent in preparing the next item for sale. We assume the rest times are exponentially distributed with expected value r^{-1} , and that successive rest times are independent of all past events. Note that all of our results will hold if we assume that the rest times obey some general (i.e. not necessarily exponential) distribution function.

2. THE MATHEMATICAL MODEL

The system that we have described is modelled as a continuous time Markov chain $\{X_t : t \geq 0\}$ with state-space

$$X_t \in Y = \{0, O(l), R(l), A(O, l), A(R, l) : 1 \leq l \leq v\}. \quad (1)$$

Initially we have $X_0 = 0$ and the state valuations are described as follows for $t \geq 0$:

- $X_t = 0$, if no bid is placed at time t . Note that this may occur after any one of the instants $t_{i+1} = \inf\{t : t > t_i \text{ and } X_{t_{i+1}} = 0\}$ when the seller accepts a bid, and the auction restarts. We set $t_0 = 0$.
- $X_t = O(l)$ where $0 < l \leq v$, if at time t the current valuation of the bid is l and the current bidder is not the SB, regardless of who placed the previous $l - 1$ bids.
- $X_t = R(l)$ where $0 < l \leq v$, if at time t the current bidder is the SB and the valuation of his bid, i.e. the current highest bid, is l .
- $X_t = A(O, l)$, if at time t the auction has concluded with a sale at price l to one of the “other” n bidders, i.e. other than the SB, and the next auction has

not yet restarted.

- $X_t = A(R, l)$, if at time t the auction has concluded with a sale at price l to the SB, and the next auction has not yet restarted.

Any bidder which is not the current highest bidder can place a bid at rate β and λ , respectively, for SB and the other bidders, as long as the current bid valuation has not attained v . When the valuation v has been attained, no further bids will be placed. Also, the transition rate that denotes the start of a new auction, from either state $A(O, l)$ or $A(R, l)$ to state 0 is r , and the transition rate (denoting the seller’s decision to sell) from state $O(l)$ to $A(O, l)$ and from $R(l)$ to $A(R, l)$ is δ . Note that the seller cannot tell the difference between the SB and the other bidders, and the transition rates in this first model do not depend on the current valuation of the highest bid.

For any state $x \in Y$, let the stationary probability of the state be denoted by $P(x) = \lim_{t \rightarrow \infty} P\{X_t = x\}$; then the balance equations satisfied by the stationary probabilities are

$$P(O(1))((n-1)\lambda + \beta + \delta) = n\lambda P(0), \quad (2)$$

$$P(O(l))((n-1)\lambda + \beta + \delta) = (n-1)\lambda P(O(l-1)) + n\lambda P(R(l-1)), \quad 2 \leq l \leq v-1,$$

$$P(O(v))\delta = (n-1)\lambda P(O(v-1)) + n\lambda P(R(v-1)),$$

$$P(A(O, l))r = \delta P(O(l)), \quad 1 \leq l \leq v,$$

$$P(R(1))(n\lambda + \delta) = \beta P(0),$$

$$P(R(l))(n\lambda + \delta) = \beta P(O(l-1)), \quad 2 \leq l \leq v-1,$$

$$P(R(v))\delta = \beta P(O(v-1)),$$

$$P(A(R, l))r = \delta P(R(l)), \quad 1 \leq l \leq v,$$

$$P(0)(n\lambda + \beta) = r \sum_{U=O,R} \sum_{l=1}^v P(A(U, l)),$$

$$1 = P(0) + \sum_{U=O,R} \sum_{l=1}^v [P(U(l)) + P(A(U, l))].$$

After some algebra we can write

$$P(O(l)) = H(l)P(0), \quad (3)$$

$$P(R(l)) = G(l)P(0),$$

$$P(A(O, l)) = \frac{\delta}{r} H(l)P(0),$$

$$P(A(R, l)) = \frac{\delta}{r} G(l)P(0),$$

where

$$H(l) = \begin{cases} \frac{n\lambda}{(n-1)\lambda + \beta + \delta}, & l = 1 \\ \frac{(n-1)\lambda}{(n-1)\lambda + \beta + \delta} H(l-1) \\ \quad + \frac{n\lambda}{(n-1)\lambda + \beta + \delta} G(l-1), & 2 \leq l \leq v-1 \\ \frac{(n-1)\lambda}{\delta} H(l-1) + \frac{n\lambda}{\delta} G(l-1), & l = v \end{cases}$$

$$G(l) = \begin{cases} \frac{\beta}{n\lambda + \delta}, & l = 1 \\ \frac{\beta}{n\lambda + \delta} H(l-1), & 2 \leq l \leq v-1 \\ \frac{\beta}{\delta} H(l-1), & l = v \end{cases}$$

and

$$P(0) = \frac{r\delta}{r\delta + (r + \delta)(n\lambda + \beta)}. \quad (4)$$

In the following we will obtain the closed form expression for $H(l)$ where $1 \leq l \leq v-1$. Let us define the constants

$$\begin{aligned} \alpha_1 &= \frac{n\lambda}{(n-1)\lambda + \beta + \delta}, & (5) \\ \alpha_2 &= \frac{(n-1)\lambda}{(n-1)\lambda + \beta + \delta}, \\ \alpha_3 &= \frac{(n-1)\lambda}{\delta}, \\ \alpha_4 &= \frac{n\lambda}{\delta}, \\ \alpha_5 &= \frac{\beta}{n\lambda + \delta}, \\ \alpha_6 &= \frac{\beta}{\delta}. \end{aligned}$$

We can immediately identify the recurrence relation in $H(l)$ by substituting $G(l)$ with its valuation as a function of $H(l-1)$:

$$H(l) = \alpha_2 H(l-1) + \alpha_1 \alpha_5 H(l-2), \quad 3 \leq l \leq v-1, \quad (6)$$

with initial values $H(1) = \alpha_1$ and $H(2) = \alpha_1(\alpha_2 + \alpha_5)$. Let R_1, R_2 be the roots of this recurrence equation, then

$$R_{1,2} = \frac{1}{2} \left[\alpha_2 \pm \sqrt{\alpha_2^2 + 4\alpha_1\alpha_5} \right].$$

We then have

$$H(l) = \frac{1}{2(R_1 - R_2)} \left[(-\alpha_2 + 2\alpha_1 + R_1 - R_2) R_1^l \right. \\ \left. + (\alpha_2 - 2\alpha_1 + R_1 - R_2) R_2^l \right], \quad (7)$$

$$1 \leq l \leq v-1,$$

and at the boundary $l = v$, the solution involves a different set of coefficients:

$$H(v) = \alpha_3 H(v-1) + \alpha_4 \alpha_5 H(v-2). \quad (8)$$

Since $G(l)$ is defined as a function of $H(l-1)$, we also have

$$G(1) = \alpha_5, \quad (9)$$

$$G(l) = \frac{\alpha_5}{2(R_1 - R_2)} \left[(-\alpha_2 + 2\alpha_1 + R_1 - R_2) R_1^{l-1} \right. \\ \left. + (\alpha_2 - 2\alpha_1 + R_1 - R_2) R_2^{l-1} \right], \quad 2 \leq l \leq v-1,$$

$$G(v) = \alpha_6 H(v-1).$$

Notice that because all the $\alpha_i > 0$, $\alpha_2^2 + 4\alpha_1\alpha_5 > 0$ and $R_1 - R_2 = \sqrt{\alpha_2^2 + 4\alpha_1\alpha_5} > 0$, we are assured of the solution of these equations. Also, if the valuation v is replaced by a random variable V with some general distribution function $Prob[V = v] = p(v)$ both for the SB and the other bidders, then the analysis follows directly from the previous discussion by computing expectations with respect to the random variable V .

3. PERFORMANCE MEASURES OF INTEREST TO THE SB AND TO THE SELLER

Some measures of interest to the SB are:

1. whether the SB is actually able to purchase the item it is seeking,
2. how quickly it can purchase the item,
3. whether it is able to minimise the cost of its purchase or equivalently how much it saves with respect to the maximum price that it is willing to pay, and what is its savings per unit time with respect to the maximum price v that it might have paid.

On the other hand, the seller's interest may be to maximise its income from a sale, or to maximise its *income per unit time* for a sequence of sales.

Note that $P(0)$ is the ratio of the average time elapsing from when the auction starts until the first bid arrives, to the total average time τ an auction cycle lasts (including the "rest time" of average valuation r^{-1} after an auction ends). Since the system leaves state 0 only when the first bid in an auction is made, the average time spent in this state is simply the inverse of the rate at which the first bid is made, i.e. $[n\lambda + \beta]$, and

$$P(0) = \frac{\text{Average time in state 0}}{\tau} \quad (10)$$

$$\tau = \frac{P(0)^{-1}}{n\lambda + \beta}$$

$$= \frac{r\delta + (r + \delta)(n\lambda + \beta)}{r\delta(n\lambda + \beta)}.$$

When a sale is made, the *expected income of the seller* is

$$I = \frac{\sum_{l=1}^v l [P(A(R, l)) + P(A(O, l))]}{\sum_{l=1}^v [P(A(R, l)) + P(A(O, l))]}, \quad (11)$$

and the *seller's income per unit time* is

$$\iota = \frac{I}{\tau}. \quad (12)$$

Concerning (1), the *probability that the SB is the bidder that makes the purchase* at an auction, rather than one of the other bidders, which we denote by π , it is given by

$$\begin{aligned}\pi &= \frac{\sum_{l=1}^v P(A(R, l))}{\sum_{l=1}^v [P(A(R, l)) + P(A(O, l))]} \\ &= \left[\sum_{l=1}^v P(A(R, l)) \right] \cdot \left[\frac{r}{n\lambda + \beta} [P(0)]^{-1} \right] \\ &= \left[\sum_{l=1}^v P(A(R, l)) \right] \cdot \left[\frac{r\delta + (n\lambda + \beta)(r + \delta)}{\delta(n\lambda + \beta)} \right].\end{aligned}\quad (13)$$

Hence regarding (2) *the average time ψ that the SB waits to win* an auction is the inverse of its winning rate or

$$\psi(v) = \frac{\tau}{\pi} = \frac{1}{r \sum_{l=1}^v P(A(R, l))}.$$

Concerning (3) *the average difference between the valuation v for the good, and the price at which the auction concludes given that the SB makes the purchase*, is denoted by

$$\phi(v) = \frac{\sum_{l=1}^v (v - l)P(A(R, l))}{\sum_{l=1}^v P(A(R, l))}.\quad (14)$$

3.1. Optimisation on the part of the SB

All that the SB can do, without reverting to deceit, is to adjust its bidding rate β to the situation it is observing, including the bid rate it observes concerning other bidders, so as to optimise the performance measures that it is selfishly and legitimately interested in.

In order to minimise $\psi(v)$ it would suffice to take $\beta \gg n\lambda$. Then the SB raises its bid to the valuation v very quickly so that it is always the winner, and $\pi(v)$ tends to 1. However this means that the SB would be buying the good at its maximum price, rather than driving a good bargain.

Thus a reasonable approach would be to choose a valuation of β which maximises the SB's return on the auction, such as $\gamma(v)$, the average savings per unit time that the SB makes with respect to the maximum price that it would pay, or specifically

$$\gamma(v) = \frac{\phi(v)}{\psi(v)} = r \sum_{l=1}^v (v - l)P(A(R, l)).\quad (15)$$

If v is replaced by the random variable V , the function of interest to the SB is

$$\Gamma = E[\gamma(V)] = r \sum_{v=1}^{\infty} \sum_{l=1}^v (v - l)p(v)P(A(R, l)),\quad (16)$$

and with the previous analysis we have

$$\Gamma = \delta \sum_{v=1}^{\infty} \sum_{l=1}^v (v - l)p(v)G(l)P(0).\quad (17)$$

Hence the SB could choose a valuation of β that maximises Γ .

FIGURE 1. SB's expected time to win with $\delta = 0.5$, $r = 1$, $n = 10$, $V \sim U(80, 100)$.

FIGURE 2. SB's expected payoff with $\delta = 0.5$, $r = 1$, $n = 10$, $V \sim U(80, 100)$.

3.2. Numerical examples

We will now provide some numerical examples that illustrate the predictions of the model. In all the numerical results that are shown, we provide curves for the case when all bidders including the SB follow a symmetric bidding strategy, i.e. $\lambda = \beta$, and for the more interesting case when the SB varies its bidding rate while the rest have a fixed bidding rate. The topic of mutual adaptation of all bidders to each other is yet another important subject which is not discussed in this paper.

Comparisons of the asymmetrical bidders case where λ is constant, against the case with identical bidders with $\lambda = \beta$ are also shown in Figures 1, 2, 3, 4. In Figure 1, as we would expect, we see that in the asymmetric case it suffices for the SB to bid at a sufficiently high rate (the x -axis) in order to reduce its time until it can make a purchase (the y -axis).

In Figure 2 we study the quantity $\phi(v)$. It is interesting to see that for fixed δ and λ , even if the SB increases β to very high valuations (and hence wins the bid), the "expected payoff" $\phi(v)$ does not tend to zero and only drops slowly with β . However, if $\lambda = \beta$ and they increase, then the pay-off will tend to zero.

On the other hand, when buyers are interested in purchasing multiple goods from the auction or when they have a long term view of things, the expected payoff per unit time $\gamma(v)$ can be a good criterion for decision making. Figure 3, with the y -axis in logarithmic scale, shows that the expected pay-off per unit time increases very rapidly with β , and furthermore this effect is accentuated, and the pay-off is greater, when the other bidders are relatively slower, i.e. have smaller valuations of λ . Figure 3 shows that bidding at a high rate increases the payoff rate for the bidder and that it leads to diminishing returns of payoff per time beyond some valuation of β . In other words, when the SB's actions do not impact the other bidders' behaviour, it should bid quickly. This is contrary to what we observe in online auctions such as eBay, in which "sniping" is often used regardless of other bidders' strategies. Bidders wait until the last possible moment before the auction expires to place their true bids.

Indeed, it has been suggested that sniping is a good

FIGURE 3. SB's expected payoff per unit time with $\delta = 0.5$, $r = 1$, $n = 10$, $V \sim U(80, 100)$.

FIGURE 4. Expected income per unit time for the seller, with $\delta = 0.5$, $r = 1$, $n = 10$, $V \sim U(80, 100)$.

strategy [?, ?, ?] if the information on the closing time of the auction is made public by the seller. But this strategy has its shortcomings: in balancing the benefits of submitting the very last bid against the risk of bid being rejected for arriving after the auction has ended, the bidders can misjudge. Technical issues such as communication delay can aggravate the problem, causing the item to sell at a lower price than what it can fetch. It is desirable that the time spent in waiting for the auction to close be shortened, thus saving time for both seller and bidders; this is especially so when the seller has many items to sell and time is of the essence. Interestingly enough, a variant of the auction protocol that was until recently used by Amazon tackles sniping behaviour by automatic deadline extension: if any bid is submitted within the last 10 minutes of the scheduled closing time, the deadline is automatically extended for another 10 minutes. This process continues until 10 minutes have passed since the last received bid, at which time the auction concludes. While it may succeed in discouraging sniping [?], this approach is not always time effective: the scheduled deadline is the best-case time within which the seller can hope to make a sale, and in general it is likely that it will take longer. This may not be suitable if the seller is pressed for time. Note also that Amazon has stopped running auctions as indicated in their Changes to the Participation Agreement of April 14 2008¹.

Finally Figure 4 looks at things from the perspective of the seller; for fixed λ we see that the SB's bid rate β affects the seller's income per unit time, but only in a moderate way. This is to be expected because after the SB makes a bid, it must pause and the remaining bidders then have a chance to bid. Since there are many other bidders (in this example $n = 10$) they will have a significant impact on the outcome, while the SB's effect remains limited.

4. WHEN THE SB TRIES TO KEEP UP WITH OTHER BIDDERS

An interesting question arises if the SB adjusts its bidding rate β in a manner proportional to the bidding rate of all other bidders. From the state equations (2) we can set a value μ representing the relative rate at which both SBs and other bidders are bidding, with respect to the other bidding and decision rates. Thus the quantity μ illustrates the "similar" behaviour of SBs and of the other bidders, and we have

$$\mu \equiv \frac{\beta}{n\lambda + \delta} = \frac{(n-1)\lambda}{(n-1)\lambda + \beta + \delta}.$$

¹See <http://www.amazon.co.uk/gp/help/customer/display.html?ie=UTF8&nodeId=200239030> that we have accessed on 22-04-2008

FIGURE 5. Expected time to win for SB when keeping up with the other bidders for various λ and n . Other parameters: $\delta = 0.5$, $r = 1$.

FIGURE 6. Expected payoff per time for SB when keeping up with the other bidders for various λ and n . Other parameters: $\delta = 0.5$, $r = 1$.

Then, the outcome of the auction for the SB will be equivalent to that for the other bidders taken together. In fact if n is large enough, this simplifies to

$$0 = \beta^2 + \beta[n\lambda + \delta] - n\lambda[n\lambda + \delta], \quad (18)$$

which yields

$$\beta \approx \frac{n\lambda + \delta}{2} \left[\sqrt{1 + 4 \frac{n\lambda}{n\lambda + \delta}} - 1 \right], \quad (19)$$

or

$$\mu \approx \frac{1}{2} \left[\sqrt{1 + 4 \frac{n\lambda}{n\lambda + \delta}} - 1 \right]. \quad (20)$$

Figure 5 shows our model's predictions on the expected time to win for SB, while in figures 6 and 7 we show the payoff and income rates, respectively, as functions of varying λ and δ , when SB follows this policy in keeping up with the other bidders. In Figure 5, for each case of n , there exists a minimum expected time to win that occurs at some λ , and for increasing n this minimal point occurs at smaller λ . Likewise, the highest payoff per time is obtained at a distinct valuation of λ and this decreases with n . These observations correspond to increasing competition with n , and the penalty suffered by SB for an increase in λ is larger for systems with large n ; the drop in payoff per time is increasingly steeper with n (see Figure 6). On the other hand, the seller benefits from large n and its income per time has higher peaks, as shown in Figure 7.

5. PRICE DEPENDENT BIDDING

In many cases the current price attained by a good offers useful information about its valuation, and about the situation of other bidders. Thus a model with bidding rates dependent on price was analysed in [?]. Here we extend this approach to the behaviour of both the SB and the other bidders.

We use $\beta(l)$ and $\lambda(l)$ to denote the bidding rates when the price is at level l for the SB and the other bidders, respectively. Likewise, $\delta(l)$ will be the seller's decision rate when price is at level l . By a simple extension of the previous model, the steady state probabilities for

FIGURE 7. Expected income per time for the seller when SB keeps up with the other bidders for various δ and n . Other parameters: $\lambda = 1$, $r = 1$.

the system satisfy

$$\begin{aligned}
 P(O(1)) &= \frac{n\lambda(0)}{(n-1)\lambda(1) + \beta(1) + \delta(1)} P(0), \quad (21) \\
 P(R(1)) &= \frac{\beta(0)}{n\lambda(1) + \delta(1)} P(0), \\
 P(O(l)) &= \frac{(n-1)\lambda(l-1)}{(n-1)\lambda(l) + \beta(l) + \delta(l)} P(O(l-1)) \\
 &\quad + \frac{n\lambda(l-1)}{(n-1)\lambda(l) + \beta(l) + \delta(l)} P(R(l-1)), \\
 &\quad 2 \leq l \leq v-1, \\
 P(R(l)) &= \frac{\beta(l-1)}{n\lambda(l) + \delta(l)} P(O(l-1)), \quad 2 \leq l \leq v-1, \\
 P(O(l)) &= \frac{(n-1)\lambda(l-1)}{\delta(l)} P(O(l-1)) \\
 &\quad + \frac{n\lambda(l-1)}{\delta(l)} P(R(l-1)), \quad l = v, \\
 P(R(l)) &= \frac{\beta(l-1)}{\delta(l)} P(O(l-1)), \quad l = v, \\
 P(A(O, l)) &= \frac{\delta(l)}{r} P(O(l)), \quad 1 \leq l \leq v, \\
 P(A(R, l)) &= \frac{\delta(l)}{r} P(R(l)), \quad 1 \leq l \leq v, \\
 P(0) &= \frac{r}{n\lambda(0) + \beta(0)} \sum_{U=O,R} \sum_{l=1}^v P(A(U, l)), \\
 1 &= P(0) + \sum_{U=O,R} \sum_{l=1}^v [P(U(l)) + P(A(U, l))].
 \end{aligned}$$

We will first give the general solutions for this system, and then look at a plausible example of forms that the dependent functions λ , β and δ might assume. Suppose, following similar approach in (3), we let

$$\begin{aligned}
 P(O(l)) &= H(l)P(0), \quad 1 \leq l \leq v, \\
 P(R(l)) &= G(l)P(0), \quad 1 \leq l \leq v.
 \end{aligned}$$

We can then express the second order recurrence relations in $H(l)$:

$$H(l) = c_1(l)H(l-1) + c_2(l)H(l-2), \quad 3 \leq l \leq v-1, \quad (22)$$

where the coefficients c_1 and c_2 are

$$\begin{aligned}
 c_1(l) &= \frac{(n-1)\lambda(l-1)}{(n-1)\lambda(l) + \beta(l) + \delta(l)}, \quad (23) \\
 c_2(l) &= \frac{n\lambda(l-1)\beta(l-2)}{((n-1)\lambda(l) + \beta(l) + \delta(l))} \\
 &\quad \times \frac{1}{(n\lambda(l-1) + \delta(l-1))},
 \end{aligned}$$

and, the initial valuations will satisfy

$$\begin{aligned}
 H(1) &= \frac{n\lambda(0)}{(n-1)\lambda(1) + \beta(1) + \delta(1)} \quad (24) \\
 H(2) &= \frac{1}{(n-1)\lambda(2) + \beta(2) + \delta(2)} \\
 &\quad \times \left[\frac{n(n-1)\lambda(0)\lambda(1)}{(n-1)\lambda(1) + \beta(1) + \delta(1)} + \frac{n\lambda(1)\beta(0)}{n\lambda(1) + \delta(1)} \right].
 \end{aligned}$$

Clearly, the difference equations (22) are linear homogeneous with variable coefficients (23), and, hence, the solution for $H(l)$ can be expressed in closed form, purely in terms of the coefficients[?, ?, ?]. First, define a matrix:

$$\mathbf{M}_l \equiv \begin{bmatrix} c_2(l) & c_1(l) \\ c_1(l+1)c_2(l) & c_2(l+1) + c_1(l+1)c_1(l) \end{bmatrix}. \quad (25)$$

Then, the solution sequence $\{H(l) : 1 \leq l \leq v-1\}$, can be represented as a product of matrices $\{\mathbf{M}_l\}$ and the initial valuations:

$$\begin{aligned}
 \begin{bmatrix} H(2j+1) \\ H(2j+2) \end{bmatrix} &= \mathbf{M}_{2j+1} \mathbf{M}_{2j-1} \cdots \mathbf{M}_3 \begin{bmatrix} H(1) \\ H(2) \end{bmatrix} \quad (26) \\
 &= \prod_{i=1}^j \mathbf{M}_{2i+1} \begin{bmatrix} H(1) \\ H(2) \end{bmatrix}, \quad 0 \leq j \leq \left\lfloor \frac{v-2}{2} \right\rfloor.
 \end{aligned}$$

Solving the above equation yields a set of two $H(l)$, one corresponding to an odd l and another to an even, for every j . However, the solutions will not hold at the boundary $l = v$, because it involves a different set of coefficients as given in (21). Thus, the boundary solution will be distinct and dependent on the previous two valuations:

$$\begin{aligned}
 H(v) &= \frac{(n-1)\lambda(v-1)}{\delta(v)} H(v-1) \\
 &\quad + \frac{n\lambda(v-1)\beta(v-2)}{\delta(v)\delta(v-1)} H(v-2). \quad (27)
 \end{aligned}$$

Similarly, the solutions for $G(l)$ will be

$$\begin{aligned}
 \begin{bmatrix} G(2j+2) \\ G(2j+3) \end{bmatrix} &= \mathbf{N}_{2j+2} \prod_{i=1}^j \mathbf{M}_{2i+1} \begin{bmatrix} H(1) \\ H(2) \end{bmatrix}, \quad (28) \\
 0 \leq j &\leq \left\lfloor \frac{v-3}{2} \right\rfloor,
 \end{aligned}$$

where the matrix and the coefficients are

$$\mathbf{N}_l \equiv \begin{bmatrix} d(l) & 0 \\ 0 & d(l+1) \end{bmatrix}, \quad \text{and } d(l) = \frac{\beta(l-1)}{n\lambda(l) + \delta(l)}. \quad (29)$$

Again, at the boundaries $l = 1$ and $l = v$, the solutions will be different:

$$\begin{aligned}
 G(1) &= \frac{\beta(0)}{n\lambda(1) + \delta(1)}, \quad (30) \\
 G(v) &= \frac{\beta(v-1)}{\delta(v)} H(v-1).
 \end{aligned}$$

FIGURE 8. Payoff per unit time in the price dependent bidding model, against nominal bid rate β_0 for various pressure coefficients. Here $n = 10$, $\lambda_0 = 1.0$, $r = 1$, $\delta = 0.5$, and $\sigma = 0$.

The solutions above are general, and will hold for all price dependent functions. Suppose now, that the dependencies are such that λ and β will decrease while δ will increase, with the price level l . Specifically, let a *pressure coefficient* $\kappa \geq 0$ [?] to represent the degree to which the attained price discourages bidding, while $\sigma \geq 0$ represents the effect of higher prices on the seller's tendency to sell:

$$\begin{aligned}\beta(l) &= \frac{\beta_0}{(l+1)^\kappa}, \quad l \geq 0, \\ \lambda(l) &= \frac{\lambda_0}{(l+1)^\kappa}, \quad l \geq 0, \\ \delta(l) &= l^\sigma \delta_0, \quad l \geq 1,\end{aligned}\quad (31)$$

where β_0 , λ_0 and δ_0 are fixed nominal rates. Although we use the same κ for the SB and other bidders, it is easy to relax this restriction. When $\kappa = 1$, we have the case of ‘‘harmonic discouragement’’, and if $\kappa = 0$ the bidders are insensitive to price, and consequently, the whole system reduces to the previously solved model (2).

Now, for functionals of form (31), the explicit solutions for $H(l)$ will follow (26) and (27), where the coefficients c_1 and c_2 are

$$\begin{aligned}c_1(l) &= \left(\frac{l+1}{l}\right)^\kappa \frac{(n-1)\lambda_0}{(n-1)\lambda_0 + \beta_0 + l^\sigma(l+1)^\kappa\delta_0}, \\ c_2(l) &= \left(\frac{l+1}{l-1}\right)^\kappa \frac{n\lambda_0\beta_0}{((n-1)\lambda_0 + \beta_0 + l^\sigma(l+1)^\kappa\delta_0)} \\ &\quad \times \frac{1}{(n\lambda_0 + l^\kappa(l-1)^\sigma\delta_0)},\end{aligned}\quad (32)$$

and the initial valuations become

$$\begin{aligned}H(1) &= \frac{n\lambda_0 2^\kappa}{(n-1)\lambda_0 + \beta_0 + 2^\kappa\delta_0}, \\ H(2) &= \frac{3^\kappa}{(n-1)\lambda_0 + \beta_0 + 2^\sigma 3^\kappa\delta_0} \\ &\quad \times \left[\frac{n(n-1)\lambda_0^2}{(n-1)\lambda_0 + \beta_0 + 2^\kappa\delta_0} + \frac{n\lambda_0\beta_0}{n\lambda_0 + 2^\kappa\delta_0} \right].\end{aligned}\quad (33)$$

For $G(l)$, the solutions will follow the general forms (28) and (30), where the coefficient

$$d(l) = \left(\frac{l+1}{l}\right)^\kappa \frac{\beta_0}{n\lambda_0 + l^\sigma(l+1)^\kappa\delta_0}, \quad (34)$$

and the initial valuation is $G(1) = \frac{2^\kappa\beta_0}{n\lambda_0 + 2^\kappa\delta_0}$.

The examples in Figure 8 illustrate the effect of κ on the expected payoff per unit time for the SB. We see that the pressure coefficient does not make a difference for relatively small bid rates β_0 , and that a higher

coefficient fetches a better payoff rate at higher bid rates. Also, a small increase in κ from 0 to 0.2 yields a bigger difference in payoff rates, than an equal-sized increase from 0.8 to 1.0.

6. CONCLUSIONS

In this paper we have considered auctions in which bidders make offers that are sequentially increasing in value by a unit price in order to minimally surpass the previous highest bid, and modelled them as discrete state-space random processes in continuous time. Analytical solutions are obtained and measures that are of interest to the SB are derived.

The measures that can be computed in this way include the SB's probability of winning the auction, its expected savings with respect to the maximum sum it is willing to pay, and the average time that the SB spends before it can make a purchase. An extension of the model that incorporates price-dependent behaviours of the agents has also been presented.

The model allows us to quantitatively characterise intuitive and useful trade-offs between improving the SB's chances of buying a good quickly, and the price that it has to pay, in the presence of different levels of competition from the other bidders.

There are interesting extensions and applications of these models that can be considered, such as the behaviour of bidders and sellers that may have time constraints for making a purchase, and the possibility of the SB's moving among different auctions so as to optimise measures which represent its self-interest. Another interesting area of study may be to examine bidders who are ‘‘rich’’ and are willing to drive away rivals at any cost, and who may create different auction environments for bidders that have significantly different levels of wealth. Yet another area of interest concerns auctions where items are sold in batches of varying sizes, with prices which depend on the number of items that are being bought.

ACKNOWLEDGEMENTS

This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) project and is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1).