

Bidding agents for online auctions with hidden bids

Albert Xin Jiang · Kevin Leyton-Brown

Received: 4 September 2005 / Revised: 30 June 2006 / Accepted: 6 September 2006 / Published online: 2 November 2006
Springer Science + Business Media, LLC 2007

Abstract There is much active research into the design of automated bidding agents, particularly for environments that involve multiple decoupled auctions. These settings are complex partly because an agent's strategy depends on information about other bidders' interests. When bidders' valuation distributions are not known *ex ante*, machine learning techniques can be used to approximate them from historical data. It is a characteristic feature of auctions, however, that information about some bidders' valuations is systematically concealed. This occurs in the sense that some bidders may fail to bid at all because the asking price exceeds their valuations, and also in the sense that a high bidder may not be compelled to reveal her valuation. Ignoring these "hidden bids" can introduce bias into the estimation of valuation distributions. To overcome this problem, we propose an EM-based algorithm. We validate the algorithm experimentally using agents that react to their environments both decision-theoretically and game-theoretically, using both synthetic and real-world (eBay) datasets. We show that our approach estimates bidders' valuation distributions and the distribution over the true number of bidders significantly more accurately than more straightforward density estimation techniques.¹

Editors: Amy Greenwald and Michael Littman

A. X. Jiang (✉) · K. Leyton-Brown
Department of Computer Science, University of British Columbia, Vancouver, Canada
e-mail: jiang@cs.ubc.ca

K. Leyton-Brown
e-mail: kevinlb@cs.ubc.ca

¹ An earlier version of this work was presented at the Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT) 2005, Edinburgh, Scotland.

Keywords Bidding agents · Online auctions · eBay auctions · Expectation maximization · Density estimation · Game-theoretic and decision-theoretic approaches

1 Introduction

There has been much research on the study of automated bidding agents for auctions and other market-based environments. The Trading Agent Competitions (TAC-Classic and TAC Supply Chain Management) have attracted much interest (Wellman et al., 2002). There have also been research efforts on bidding agents and bidding strategies in other auction environments (Byde, 2002; Boutilier, Goldszmidt, & Sabata, 1999; Greenwald & Boyan, 2004; Arora et al., 2003; Cai & Wurman, 2003; Anthony et al., 2001). Although this body of work considers many different auction environments, bidding agents always face a similar task: given a valuation function, the bidding agent needs to compute an optimal bidding strategy that maximizes expected surplus. (Some environments such as TAC-SCM also require agents to solve additional, e.g., scheduling tasks.)

The “Wilson Doctrine” in mechanism design argues that mechanisms should be constructed so that they are “detail-free”—that is, so that agents can behave rationally in these mechanisms even without information about the distribution of other agents’ valuations. For example, the VCG mechanism is detail-free because under this mechanism it is a weakly dominant strategy to bid exactly one’s valuation, regardless of other agents’ beliefs, valuations or actions. Under common assumptions (in particular, independent private values) single-item English auctions are similar: an agent should remain in the auction until the bidding reaches the amount of her valuation.

While detail-free mechanisms are desirable, they are not ubiquitous. Very often, agents are faced with the problem of deciding how to behave in games that do not have dominant strategies and where other agents’ preferences are strategically relevant. For example, we may want to participate in a series of auctions run by different sellers at different times. This is a common scenario at online auction sites such as eBay. In Section 4 we consider a sequential auction model of this scenario, and show that information about other bidders’ preferences is very relevant in constructing a bidding strategy.

1.1 Game-theoretic and decision-theoretic approaches

How should a bidding agent be constructed? Depending on the assumptions we choose to make about other bidders, two broad approaches to computing bidding strategies suggest themselves: a game-theoretic approach and a decision-theoretic approach. The game theoretic approach assumes that all agents are perfectly rational and that this rationality is common knowledge; the auction is modeled as a Bayesian game. Under this approach, a bidding agent would compute a Bayes-Nash equilibrium of the auction game, and play the equilibrium bidding strategy. Much of the extensive economics literature on auctions follows this approach (see, e.g., the survey in Klemperer (2000)). For example, in environments with multiple, sequential auctions for identical

items and in which each bidder wants only a single item, the Bayes-Nash equilibrium is well-known (Milgrom & Weber, 2000; Weber, 1983). Such equilibria very often depend on the distribution of agents' valuation functions and the number of bidders. Although this information is rarely available in practice, it *is* usually possible to estimate these distributions from the bidding history of previous auctions of similar items. Note that this involves making the assumption that past and future bidders will share the same valuation distribution.

The game-theoretic approach has received a great deal of study, and is perhaps the dominant paradigm in microeconomics. In particular, there are very good reasons for seeking strategy profiles that are resistant to unilateral deviation. However, this approach is not always useful to agents who need to decide what to do in a particular setting, especially when the rationality of other bidders is in doubt, when the computation of equilibria is intractable, or when the game has multiple equilibria. In such settings, it is sometimes more appropriate to rely on decision theory. A decision-theoretic approach effectively treats other bidders as part of the environment, and ignores the possibility that they may change their behavior in response to the agent's actions. We again make the assumption that the other bidders come from a stationary population; however, in this case we model agents' bid amounts directly rather than modeling their valuations and then seeking an equilibrium strategy. We then solve the resulting single-agent decision problem to find a bidding strategy that maximizes expected payoff. We could also use a reinforcement-learning approach, where we continue to learn the bidding behavior of other bidders while participating in the auctions. Much recent literature on bidding agent design follows the decision-theoretic approach, e.g. (Boutilier, Goldszmidt, & Sabata, 1999; Bye, 2002; Greenwald & Boyan, 2004; Stone et al., 2002; Mackie-Mason et al., 2004; Osepashvili et al., 2005).

In this work we do not attempt to choose between these two approaches; it is our opinion that each has domains for which it is the most appropriate. The important point is that regardless of which approach we elect to take, we are faced with the subproblem of estimating two distributions from the bidding history of past auctions: the distribution on the number of bidders, and the distribution of bid amounts (for decision-theoretic approaches) or of valuations (for game-theoretic approaches). Consequently, this problem has received much discussion in the literature. For example, Athey and Haile (2002) and various others in econometrics have studied the estimation of valuation distributions in various standard auction types given observed bids, assuming that bidders are perfectly rational and follow equilibrium strategies. On the decision-theoretic front, a popular approach is to estimate the distribution of the final prices in auctions based on observed selling prices, and then to use this distribution to compute the optimal bidding strategy. Examples of this include Stone et al.'s (2004) ATTac agent for the Trading Agent Competition, Mackie-Mason et al.'s (2004) study of bidding in simultaneous ascending auctions and the follow-up work in Osepashvili et al. (2005), Bye's (2002) study of bidding in simultaneous online English auctions, and Greenwald and Boyan's (2004) analysis of sequential English auctions. A paper by Boutilier, Goldszmidt, and Sabata (1999) takes a different decision-theoretic approach which is relevant to the approach we propose in this paper. We defer discussion of this work until Section 6, after we have presented our approach.

1.2 Overview

In this paper we consider sequential English auctions in which a full bidding history is revealed, such as the online auctions run by eBay. It might seem that there is very little left to say: we learn the distributions of interest from the bidding history, then compute a bidding strategy based on that information for the current and future auctions. However, we show that under realistic bidding dynamics (described in Section 2) the observed bidding histories omit some relevant information. First, some bidders may come to the auction when it is already in progress, find that the current price already exceeds their valuation, and leave without placing a bid. Second, the amount the winner was willing to pay is never revealed. Ignoring these sources of bias can lead to poor estimates of the underlying valuation distribution and distribution of number of bidders. In Section 3 we propose a learning approach based on the Expectation Maximization (EM) algorithm, which iteratively generates hidden bids consistent with the observed bids, and then computes maximum-likelihood estimations of the valuation distribution based on the completed set of bids. The learned distributions can then be used in computing decision-theoretic or game-theoretic bidding strategies. Section 4 discusses the computation of the optimal strategy for an agent in a repeated English auction setting under the decision-theoretic approach, and in a similar (non-repeated) setting under the game-theoretic approach. In Section 5 we present experimental results on synthetic data sets as well as on data collected from eBay, which show that our EM learning approach makes better estimates of the distributions, gets more payoff under the decision-theoretic model, and makes a better approximation to the Bayes-Nash equilibrium under the game-theoretic model, as compared to the straightforward approach which ignores hidden bids.

2 A model of online auctions

Consider a (possibly repeated) auction held online, e.g., on eBay. There are m potential bidders interested in a certain auction for a single item. We assume that bidders are risk-neutral, have independent private values (IPV), and that utilities are quasilinear. The number of bidders m is drawn from a discrete distribution $g(m)$ with support $[2, \infty)$. Bidders' potential valuations (in the game-theoretic context) or bids (in the decision-theoretic context) are independently drawn from a continuous distribution $f(x)$.

2.1 Bidding dynamics

The m potential bidders arrive at the auction site sequentially. When each bidder arrives, she observes the bids that have been accepted in the auction so far, places a single proxy bid² and then leaves the system. The auctioneer processes new bids as follows:

² A proxy bidding system asks bidders for their maximum willingness to pay, and then bids up to this amount on the bidder's behalf. Such systems are common on online auction sites such as eBay; see e.g., <http://pages.ebay.com/help/buy/proxy-bidding.html>.

1. When a proxy bid is submitted, the auctioneer compares it to the current price level, which is the second-highest proxy bid so far plus a small bid increment.³
 - (a) If the submitted bid is not greater than the current price level the bid is dropped and no record of the bid is recorded in the auction's history.
 - (b) If the submitted bid is higher than the current price level but lower than the highest proxy bid so far, then the submitted bid is accepted, the bid amount of the currently winning bid is increased to equal the submitted bid (i.e., this bid remains the winning bid), and the submitted bid loses.
 - (c) If the submitted bid is higher than the previously winning bid then the price level is increased to equal the previously winning bid and the submitted bid is made the current winner.
2. At the end of the auction, the item is awarded to the bidder who placed the highest bid, and the final price level is the amount of the second highest bid.

2.2 Hidden bids

According to our model, some bidders' proxy bid amounts will be revealed. This includes any bid that was higher than the current price level at the time it was placed, even if that bid was immediately outbid by another proxy bid. However, other bids will not be observed. There are two types of hidden bids:

1. The highest bid of each auction x_{hi} .
2. Dropped bids x_d that were lower than the current price when they were placed.

Let us denote the set of visible bids as x_v and the set of hidden bids as x_h . Let n denote the number of bidders that appears in the bidding history. This means that x_v will always contain $(n - 1)$ bids, since the winning bidder's proxy bid is never observed. Since there are m potential bidders in total, $(n - 1)$ bids are visible, and one bid is the highest bid x_{hi} , there are $(m - n)$ dropped bids in x_d .

Figure 1 shows an example of the bidding process for an auction according to our online auction model, and illustrates how bids are dropped. Bids arrive sequentially from left to right. Bids 3, 4 and 7 (the grey bars) will be dropped because they are lower than the current bid amount at their respective time steps. The amount of the winning bid (6) will not be revealed, although an observer would know that it was at least the amount of bid 2, which would be the amount paid by bidder 6.

2.3 Discussion

Our model of the bidding process is quite general. Notice that when a bidder observes that the price level is higher than her potential bid, she may decide not to bid in this auction. This is equivalent to our model in which she always submits the bid, because dropped bids do not appear in the bidding history. (Indeed, this is the motivation for our model.) Also our model covers the case of last-minute bidding, which happens

³ For simplicity in this paper we ignore this small increment and assume that the current price level is the second-highest proxy bid so far.




Fig. 1 An example of the bidding process of an auction with 7 potential bidders

quite often in eBay auctions (Roth & Ockenfels, 2002): even though last-minute bids may be submitted almost simultaneously, eBay processes the bids in sequence.

Observe that with a proxy bidding system, and when agents have independent private values, bidders would not benefit from bidding more than once in an auction. However, in practice eBay bidders quite often make multiple bids in one auction. One possible motivation for these bids is a desire to learn more information about the proxy bid of the current high bidder (Shah et al., 2003). However, only the last bid of the bidder represents her willingness to pay. Given a bidder's last bid, her earlier bids carry no extra information. Therefore, we will be interested in only the last bid from each bidder.⁴ We can preprocess the bidding histories by removing all bids except the last bids from each bidder, without losing any relevant information.

3 Learning the distributions

Given the model of the bidding process, the first task of our bidding agent is to estimate the distributions $f(x)$ and $g(m)$ from the bidding history. Suppose we have access to the bidding history of K auctions for identical items.

3.1 The simple approach

A simple approach is to ignore the hidden bids, and to directly estimate $f(x)$ and $g(m)$ from observed data. The observed number of bidders, n , is used to estimate $g(m)$. To

⁴ In a common value model, the earlier bids would carry some information, and we would not be able to simply ignore those bids.

estimate $f(x)$ we use the observed bids x_v , which consists of $(n - 1)$ bids for each auction. Any standard density estimation technique may be used.

Because it ignores hidden bids, this approach can be expected to produce biased estimates of f and g :

- $g(m)$ will be skewed towards small values because $n \leq m$.
- $f(v)$ may be skewed towards small values because it ignores the winning bid x_{hi} , or it may be skewed towards large values because it ignores the dropped bids x_d .

A popular variation of this approach (mentioned in Section 1) is to directly estimate the distribution of final prices from the selling prices of previous auctions, then to use this distribution to compute a decision-theoretic bidding strategy. The problem with this approach is that for English auctions, the selling prices are the second-highest bids. As we will show in Section 4, to compute a decision-theoretic strategy we really need the distribution of the other bidders' highest bids. Using the distribution of final prices introduces bias, as this distribution is skewed towards small values compared to the distribution of highest bids.

3.2 EM learning approach

We would like to have an estimation strategy that accounts for the hidden bids x_h and any bias introduced by their absence. Suppose $f(x)$ belongs to a class of distributions parameterized by θ , $f(x | \theta)$. Further suppose that $g(m)$ belongs to a class of distributions parameterized by λ , $g(m | \lambda)$. For example, $f(x)$ could be a Normal distribution parameterized by its mean μ and variance σ^2 , whereas $g(m)$ could be a Poisson distribution parameterized by its mean, λ . We want to find the maximum likelihood estimates of θ and λ given the observed data x_v .

Suppose that we could actually observe the hidden bids x_h in addition to x_v . Then estimating θ and λ from the completed data set (x_v, x_h) would be easy. Unfortunately we do not have x_h . Given x_v , and with the knowledge of the bidding process, we could generate x_h if we knew θ and λ . Unfortunately we do not know θ and λ .

A popular strategy for learning this kind of model with missing data is the Expectation Maximization (EM) algorithm (Dempster et al., 1977). EM is an iterative procedure that alternates between E steps which generate the missing data given current estimates for the parameters and M steps which compute the maximum likelihood (or maximum a posteriori) estimates for the parameters based on the completed data, which consists of the observed data and current estimates for the missing data. Formally, the E step computes

$$Q(\theta) = \int \log(p(x_h, x_v | \theta)) p(x_h | x_v, \theta^{(old)}, \lambda^{(old)}) dx_h, \quad (1)$$

and the M step solves the optimization

$$\theta^{(new)} = \arg \max_{\theta} (Q(\theta)). \quad (2)$$

Analogous computations are done to estimate λ , the parameter for $g(m \mid \lambda)$. The EM algorithm terminates when λ and θ converge. It can be shown that EM will converge to a local maximum of the observed data likelihood function $p(x_v \mid \theta, \lambda)$.

EM is a particularly appropriate algorithm in our auction setting with hidden bids, because

1. EM was designed for finding maximum likelihood (ML) estimates for probabilistic models with unobserved latent variables.
2. EM is especially helpful when the observed data likelihood function $p(x_v \mid \theta, \lambda)$ (which we want to maximize) is hard to evaluate, but the ML estimation given complete data (x_v, x_h) is relatively easy (because the M step corresponds to maximizing the expected log-likelihood of (x_v, x_h)). This is exactly the case in our auction setting.

3.2.1 The E step

The integral in Eq. (1) is generally intractable for our complex bidding process. However, we *can* compute a Monte Carlo approximation of this integral by drawing N samples from the distribution $p(x_h \mid x_v, \theta^{(old)}, \lambda^{(old)})$, and approximating the integral by a small sum over the samples (see e.g. (Andrieu et al., 2003)). Applied to our model, in each E step our task is therefore to generate samples from the distribution $p(x_h \mid x_v, \theta^{(old)}, \lambda^{(old)})$. Recall that x_h consists of the highest bid x_{hi} and the dropped bids x_d .

Given $\theta^{(old)}$ and the second highest bid (which is observed), the highest bid x_{hi} can easily be sampled. According to the bidding process described in Section 2, the highest bid was generated by repeatedly drawing from $f(x \mid \theta^{(old)})$ until we get a bid higher than the previously highest (now second-highest) bid. This is exactly the rejection-sampling procedure for the distribution $f(x \mid \theta^{(old)})$ truncated at the second highest bid and renormalized. For distributions with a simple functional form (e.g. normal distributions), it may be easier to sample directly from the truncated distribution by reversing the CDF (see e.g. (West, 1994)).

Sampling the dropped bids x_d is a more difficult task. We use the following procedure, which is based on simulating the bidding process:

1. Sample m from $g(m \mid \lambda^{(old)})$.
2. If $m < n$, reject the sample and go back to step 1.
3. Simulate the bidding process using x_v and $m - n$ dropped bids:
 - (a) Repeatedly draw a sample bid from $f(x \mid \theta^{(old)})$, and compare it to the current price level. If it is lower than the price level, add the bid to the set of dropped bids x_d . Otherwise, the current price level is increased to the next visible bid from x_v .
 - (b) If the number of bids in x_d exceeds $m - n$, or if we used up all the bids in x_v before we have $m - n$ dropped bids in x_d , we reject this sample and go back to step 1. Only when we used up all bids in x_v and we have $m - n$ bids in x_d , do we accept the sample of x_d .
4. Repeat until we have generated N samples of x_d .

3.2.2 The M step

Our task at each M step is to compute the maximum likelihood (ML) estimates of λ and θ from x_v and the generated samples of x_h . For many standard parametric families of distributions, there are analytical solutions for the ML estimates. For example, if $f(x)$ is a normal distribution $N(\mu, \sigma)$, then given the complete set of bids (x_v, x_h) , the ML estimate of μ is the sample mean, and the ML estimate of σ is the sample standard deviation. If $g(m)$ is a Poisson distribution, then the ML estimate of the mean parameter λ is the mean of the number of bidders per auction. If analytical solutions do not exist we can use numerical optimization methods such as simulated annealing. If prior distributions on λ and θ are available, we may instead compute maximum a posteriori (MAP) estimates, which are point estimates of λ and θ that maximize their posterior probabilities.

3.3 Learning distributions in a game-theoretic setting

The approach we just described is decision-theoretic because we estimate the distribution of bid amounts without considering how other agents would react to our behavior. What if we want to take a game-theoretic approach? Athey and Haile (2002) discussed estimation in the game-theoretic setting, however they generally assume that the number of bidders is known (there is a brief discussion of unknown number of bidders, but it is not relevant to our online auction setting).

Let $f(v)$ be the distribution of bidder's *valuations* (instead of bid amounts), and let $g(m)$ denote the distribution of number of bidders, as before. Given a bidder's valuation v , her bid amount in the game-theoretic setting is given by the Bayes-Nash equilibrium of the auction game. Many (but not all) auction games have symmetric pure strategy equilibria. Assume there is a symmetric pure strategy equilibrium given by the bid function $b(v | f, g)$. Our task is to estimate $f(v)$ and $g(m)$ given the observed bids.

We can use an EM learning algorithm similar to the one in Section 3.2 to estimate $f(v | \theta)$ and $g(m | \lambda)$, where θ and λ are parameters for f and g :

- E step: for each auction with observed bids x_v :
 - Compute observed bidders' valuations v_v from x_v by inverting the bid function.⁵
 - Generate bidders with valuations v_h who place hidden bids $x_h = b(v_h | f^{(old)}, g^{(old)})$. This is done by a similar procedure to the one in Section 3.2.1 that simulates the auction's bidding process.
- M step: update θ and λ to maximize the likelihood of the valuations (v_v, v_h) .

4 Constructing a bidding agent

In Sections 2 and 3 we presented a model of the bidding process for a single auction, and proposed methods to estimate the distributions of bids and number of bidders in

⁵ If the bidding function does not have a single-valued inverse function, or if the equilibrium has mixed strategies, we just generate bidders with valuation v_v that would likely bid x_v under the equilibrium.

an auction. But our work is not done yet: how do we make use of these estimated distributions to compute a bidding strategy?

If we only participate in one English auction, bidding is simple: under the IPV model it is a dominant strategy to bid up to our valuation for the item, and we do not even need to estimate the distributions. However if we are interested in more complex environments, good estimates of the distributions $f(x)$ and $g(m)$ are essential for computing a good bidding strategy. In this section we discuss two such environments: finitely-repeated online auctions and unrepeated online auctions without proxy bidding. We consider the first problem from a decision-theoretic point of view, and take a game-theoretic approach to the second problem.

4.1 A decision-theoretic approach to repeated auctions

In this section we develop a decision-theoretic bidding agent for finitely repeated auctions. We choose this setting because it is a reasonable model of the decision-theoretic problem we would face if we wanted to buy one item from an online auction site. Our estimation algorithm can straightforwardly be applied to more complex decision-theoretic models such as infinite horizon models with discount factors and combinatorial valuation models.

4.1.1 The auction environment

Suppose we only want to buy one item (say a Playstation 2) in an environment (say eBay) where multiple auctions for similar, but not necessarily identical, Playstation 2 systems are held regularly. Recall that we have assumed that utility is quasilinear: thus if we successfully win one item, our utility will be equal to our valuation for the item minus the price we paid. So our bidding agent's task is to compute a bidding strategy that will maximize this utility. Assume that we are only interested in the first k auctions that will be held after we arrive at the auction site. One motivation for such a restriction is that we prefer to have the item within a bounded amount of time. If we fail to win an item from the k auctions, we lose interest in the item and leave the auction site, and our utility is 0. (Alternatively, we could say that if we fail to win an auction then we could buy the item from a store after leaving the auction site, in which case we would get some other known and constant amount of utility.)

Some of the k auctions may overlap in time, but since eBay auctions have strict closing times, this can be modeled as a sequential decision problem, where our agent makes bidding decisions right before each auction closes. Number the auctions $1 \dots k$ according to their closing times. Let v_j denote our valuation for the item from auction j . Note that this allows the items in the auctions to be non-identical. Let b_j denote our agent's bid for auction j . Let U_j denote our agent's expected payoff from participating in auctions $j \dots k$, assuming we did not win before auction j . Let U_{k+1} be our payoff if we fail to win any of the auctions. For simplicity we define $U_{k+1} = 0$, though the analysis would be similar for any constant value. Suppose that for each auction j the number of other bidders is drawn from $g_j(m)$ and each bidder's bid is drawn from $f_j(x)$. Since each auction j is an English auction, only the highest bid from other

bidders affects our payoff. Let $f_j^1(x)$ and $F_j^1(x)$ respectively denote the probability density function and cumulative density function (CDF) of the highest bid from other bidders in the j th auction. Then

$$F_j^1(x) = \sum_{m=2}^{\infty} g_j(m)(F_j(x))^m,$$

where $F_j(x)$ is the CDF of $f_j(x)$. Now U_j can be expressed as the following function of the future bids $b_{j:k} = (b_j, \dots, b_k)$ and valuations $v_{j:k} = (v_j, \dots, v_k)$:

$$U_j(b_{j:k}, v_{j:k}) = \int_{-\infty}^{b_j} (v_j - x)f_j^1(x)dx + (1 - F_j^1(b_j))U_{j+1}(b_{j+1:k}, v_{j+1:k}). \quad (3)$$

The first term in Eq. (3) is the expected payoff from the j th auction; the second term is the expected payoff from the later auctions.

4.1.2 The optimal strategy

Greenwald and Boyan (2004) and Arora et al. (2003) have analyzed similar auction models. Following similar reasoning, we can derive the optimal bidding strategy for our auction model. Let $b_{j:k}^*$ be the optimal bidding strategy for auctions j, \dots, k . Let $U_j^*(v_{j:k})$ denote the expected payoff under the optimal strategy, i.e. $U_j^*(v_{j:k}) = U_j(b_{j:k}^*, v_{j:k})$. We can optimize U_j by working from the k th auction to the first one in a manner similar to backward induction. By solving the first-order conditions of U_j , we obtain the optimal bidding strategy:

$$b_j^* = v_j - U_{j+1}^*(v_{j+1:k}). \quad (4)$$

In other words, our agent should shade her bids by the “option value”, i.e. the expected payoff of participating in future auctions. The exception is of course the k th auction; in this case there are no future auctions and the optimal bid is $b_k^* = v_k$. Thus we see that the standard result that honest bidding is optimal in an unrepeated auction is recovered as the special case $k = 1$.

The computation of the optimal bidding strategies requires the computation of the expected payoffs $U_j^* = U_j(b_{j:k}^*, v_{j:k})$, which involves an integral over the distribution $f_j^1(x)$ (Eq. (3)). In general this integral cannot be solved analytically, but we can compute its Monte Carlo approximation if we can sample from $f_j^1(x)$. If we can sample from $f_j(x)$ and $g_j(m)$, we can use the following straightforward procedure to generate a sample from $f_j^1(x)$:

1. draw m from $g_j(m)$;
2. draw m samples from $f_j(x)$;
3. keep the maximum of these m samples.

The bidding strategy $b_{1:k}^*$ computed using Eqs. (4) and (3) is optimal, provided that the distributions $f_j(x)$ and $g_j(m)$ are the correct distributions of bids and number of bidders for all $j \in 1 \dots k$. Of course in general we do not know the true $f_j(x)$ and $g_j(m)$ and

the focus of this paper is to estimate the distributions from the bidding history and use the estimated distributions to compute the bidding strategy. As a result, the computed bidding strategy should be expected to achieve less than the optimal expected payoff. However, it is reasonable to think that better estimates of $f(x)$ and $g(m)$ should give bidding strategies with higher expected payoffs. This is confirmed in our experiments across a wide range of data sets, which we discuss in Section 5.

4.1.3 Auctions that overlap in time: Exploiting early bidding

We observe that while the optimal bid in auction j does not depend on f_j^1 , it does depend on f_l^1 for $l > j$. So far we have been estimating $f_l^1(x)$ using $f_l(x)$ and $g_l(m)$. In practice, auctions overlap in time, and we often observe some early bidding activity by other bidders in auctions $j + 1, \dots, k$ before we have to make a bid on auction j . This extra information allows us to make even more informed (posterior) estimates on $f_l^1(x)$, $l > j$, based on $f_l(x)$, $g_l(m)$ and the observed bids for auction l , which leads to a better bid for auction j .

Suppose we have observed $n - 1$ early bids, denoted by x_v ; the current highest bid x_{hi} is not revealed (but can be sampled from $f(x)$ truncated at the current price). Since the auction is not over, there will be some set of future bids x_{future} (possibly empty). When the auction closes, the highest bid from the other bidders will be $\max\{x_{hi}, x_{future}\}$. We can generate x_{future} if we know the number of future bids. We know the total number of bids m is drawn from $g(m)$, and the number of bids made so far is $n + |x_d|$, where x_d are the dropped bids so far, so the number of future bids is $m - n - |x_d|$. Now we have a procedure that samples from $f^1(x)$:

1. Simulate the auction using our model in Section 2 to generate x_d , the dropped bids so far.
2. Sample the total number of bids m from $g(m)$.
3. Compute the number of future bids, $m - n - |x_d|$.
 - (a) If this quantity is negative, reject the sample.
 - (b) Otherwise generate x_{future} as $(m - n - |x_d|)$ bids drawn from $f(x)$.
4. Generate x_{hi} and take the maximum of x_{future} and x_{hi} .

4.2 A game-theoretic approach to bidding in online auctions without proxies

In Section 4.1 we discussed building a decision-theoretic bidding agent for repeated English auctions. What happens if we try to use the game-theoretic approach on our auction model to build an agent that plays a Bayes-Nash equilibrium of the auction game?

The difficulty turns out to be identifying the equilibrium. If each bidder other than our agent only participates in one auction, then the situation is easy. In this case the other bidders' dominant strategies will be to bid truthfully up to the amount of their valuations. Assuming all bidders follow this dominant strategy as a game-theoretic approach would have us do, we find that the distribution of bids is the same as the distribution of valuations. Thus, our agent's strategy should be the same as the decision-theoretic optimal strategy described in Section 4.1.2. It is easy to verify

that this strategy together with the other players' truthful bidding forms a Bayes-Nash equilibrium.

If the other bidders participate in more than one auction then the equilibrium strategy gets more complex, both strategically and computationally. Milgrom and Weber (2000) have derived equilibrium strategies for sequential auctions under some strong simplifying assumptions, such as identical items, a fixed number of bidders, and all bidders entering the auctions at the same time. In an online auction setting, these assumptions are not reasonable: the items are often heterogeneous, the number of bidders is unknown *ex ante*, and bidders have different entry times and exit policies. The equilibrium strategy in the general case is not known. Therefore, it is an open problem to find a game-theoretic model for repeated online auctions that is realistic and yet still tractable.

4.2.1 Online auctions without proxies

Because of the difficulties described above, we turn to a slightly different auction setting in order to show how our distribution learning techniques can be used to build a game-theoretic bidding agent. Specifically, we will consider an online auction setting which differs in one crucial respect from the setting defined in Section 2. Bidders still participate in an online ascending auction against an unknown number of opponents and still arrive and bid sequentially; however, now we add the restriction that bidders cannot place proxy bids. Instead, bidders can now place only a single bid before they leave the auction forever; the game now takes on a flavor similar to a first-price auction, but with different information disclosure rules. We chose this game because it has hidden-bid characteristics similar to the online auctions we discussed earlier, but at the same time it also has a known, computationally tractable—and yet non-trivial—Bayes-Nash equilibrium.

More formally, suppose that there are m potential bidders interested in a single item. The number m is drawn from a distribution $g(m)$, and is observed by the auctioneer but not the bidders. Bidders have independent private valuations, drawn from the distribution $f(v)$. A bid history records every bid which is placed along with the bid amount, and is observable by all bidders. The auction has the following rules:

- The auctioneer determines a random sequential order to use in approaching bidders.
- Each bidder is approached and asked to make a bid. Each bidder is asked once and only once.
- When asked, a bidder has to either make a bid which is at least the current highest bid plus a constant minimum increment δ , or decide not to bid.
 1. If the bidder places a bid, it is recorded in the bidding history.
 2. Otherwise, no record is made of the fact that the auctioneer approached the bidder.
- The auction ends after the last bidder makes her decision. The highest submitted bid wins, and the winner pays her bid amount. The winner's utility is her valuation minus her payment; other bidders get zero utility.

4.2.2 Computing equilibrium strategies

We now turn to the construction of a game-theoretic bidding agent for the auction described in Section 4.2.1. We begin by supposing that the distributions $f(v)$ and $g(m)$ are common knowledge, which allows us to treat the auction as a Bayesian game and find its Bayes-Nash equilibrium.

Suppose there exists a pure-strategy equilibrium of this game. Consider bidder i with valuation v , who is asked to make the next bid when the bidding history is x_v . Denote $b(v | x_v)$ the equilibrium bidding strategy. Before we compute $b(v | x_v)$, let us first eliminate dominated strategies. As in classical sealed-bid first-price auctions, it is obvious that we should never bid higher than our valuation v . Let b_o denote the highest submitted bid so far, thus $b_o + \delta$ is the minimum bid amount. It immediately follows that if $v < b_o + \delta$ then we should not bid. On the other hand, if $v \geq b_o + \delta$, then we have non-negative expected payoffs by making a bid (as long as we bid below v).

Define as P the class of strategies which take the following form:

- if $v < b_o + \delta$, then do not bid
- if $v \geq b_o + \delta$, then bid an amount in the interval $[b_o + \delta, v]$

Following the above reasoning, any strategy not in P is weakly dominated. Thus, in equilibrium all bidders will play some strategy in P . Suppose bidder i 's valuation is $v \geq b_o + \delta$. Thus she places a bid $b \in [b_o + \delta, v]$. Note that bidder i 's expected payoff depends on her valuation, her bid b , and what future bidders do. Since future bidders play strategies in P , we know i will be outbid if and only if at least one future bidder has valuation greater than or equal to $b + \delta$. In other words, as long as everyone plays a strategy in P , a bidder's expected payoff depends on future bidders' valuations but not their exact bidding functions.

Denote by m_f the number of future bidders, and denote by $p_o = \Pr(m_f = 0 | x_v)$ the probability that there will be no future bidders. Let $F^1(v)$ be the CDF of the (posterior) distribution of the highest future bid given x_v , conditioned on having at least one future bidder. In other words,

$$F^1(v) = E_{m_f | m_f > 0} [F(v)^{m_f}] \quad (5)$$

where $F(v)$ is the CDF of the value distribution. F^1 and p_o depend on the posterior distribution of the number of future bidders, $\Pr(m_f | x_v)$. If we can generate samples of m_f , then F^1 and p_o can be straightforwardly approximated. Note that the set of bidders consists of past bidders (observed and hidden), the current bidder, and future bidders. Thus we can generate m_f using essentially the same procedure as described in Section 4.1.3, by drawing m from $g(m)$, simulating the auction so far to get the number of already dropped bids $|x_d|$, then letting $m_f = m - |x_v| - |x_d| - 1$. Note that we do not need to know the exact equilibrium strategy to generate the hidden bidders; the knowledge that the equilibrium strategy is in P is enough to determine whether a bidder should bid or not.

Now we can write the expected payoff $U(b, v)$ as:

$$U(b, v) = (1 - p_o)F^1(b + \delta)(v - b) + p_o(v - b). \quad (6)$$

Since v is fixed for each bidder, $U(b, v)$ is then a function of b . We can then use any standard numerical technique for function maximization in order to find the optimal $b^* \in [b_o + \delta, v]$ that maximizes U .

Theorem 1. *It is a Bayes-Nash equilibrium of our auction setting for all bidders to follow the strategy:*

- if $v < b_o + \delta$, do not bid.
- if $v \geq b_o + \delta$, bid $b^* = \arg \max_{b \in [b_o + \delta, v]} U(b, v)$.

The proof is straightforward: first observe that this strategy is in P . We have showed above that if everyone else is playing a strategy in P , then b^* as defined will maximize the bidder's utility. It follows that if everyone is playing the above strategy, then each bidder is playing a best response to the other bidders' strategies and so we have an equilibrium.

4.2.3 Learning the distributions

At the beginning of Section 4.2.2 we assumed that the distributions $f(v)$ and $g(m)$ were commonly known. Now we relax this assumption. We assume that we have access to the bidding histories of K past auctions, and consider the problem of learning the distributions.

As in our online English auctions, some information is hidden from the bidding history: the number and valuations of the bidders who decided not to bid. As discussed in Section 3.1, the simple approach of ignoring these hidden bidders could be expected to introduce bias to the estimates of f and g . Furthermore, the observed bids are not equal to the bidders' valuations; as shown above a bidder's equilibrium bid is always lower than her valuation.

To correctly account for the hidden information, we use our EM algorithm for game-theoretic settings, as described in Section 3.3. To implement the EM algorithm, an important subproblem is reversing the equilibrium bidding strategy. Formally, given a complete bidding history x_v of an auction and the current estimates $f(v | \theta)$ and $g(m | \lambda)$, compute the corresponding valuations v_v such that bidders with valuations v_v playing the equilibrium strategy would have placed the bids x_v .

Consider each bid $b_i \in x_v$. Denote by b_o the previous observed bid. From Section 4.2.2 we know that $b_i = \arg \max_{b \in [b_o + \delta, v]} U(b, v)$. There are two cases:

1. If $b_i > b_o + \delta$, then b_i must be a local maximum in $[b_o + \delta, v]$ satisfying the first-order condition $\frac{\partial U(b, v)}{\partial b} = 0$. Solving for v , we get

$$v = b_i + \frac{F^1(b_i + \delta) + p_o / (1 - p_o)}{f^1(b_i + \delta)} \quad (7)$$

where f^1 is the derivative of F^1 . We can numerically approximate F^1 , f^1 and p_o by generating samples of m_f in the same way as in Section 4.2.2.

2. Otherwise, $b_i = b_o + \delta$. Intuitively, this means that the bidder's valuation is high enough to make a bid (i.e. $v \geq b_o + \delta$), but not high enough to fall into Case 1. Any

bidder with a valuation in the interval

$$\left[b_o + \delta, \quad b_o + \delta + \frac{F^1(b_o + \delta) + p_o/(1 - p_o)}{f^1(b_o + \delta)} \right]$$

would bid in this way. Thus we generate samples of v by drawing from $f(v \mid \theta)$ truncated to the above interval.

Once the EM algorithm has estimated the distributions f and g , our agent can use these distributions to compute the Bayes-Nash equilibrium strategy as described in Section 4.2.2.

5 Experiments

We evaluated both our EM learning approach and the simple approach⁶ on several synthetic data sets and on real world data collected from eBay. We compared the approaches in three ways:

1. Which approach gives better estimates of the distributions $f(x)$, $g(m)$ and $f^1(x)$? This is important because better estimation of these distributions should give better results, regardless of whether agents take a decision-theoretic approach or a game-theoretic approach to bidding. We measure the closeness of an estimated distribution to the true distribution using the Kullback-Leibler (KL) Divergence from the true distribution to the estimated distribution. The smaller the KL Divergence, the closer the estimated distribution to the true one.
2. For repeated online auction data, which approach gives better expected payoff under the decision-theoretic bidding model, as described in Section 4.1?
3. For data from online auctions without proxies, which approach gives closer estimates to the Bayes-Nash equilibrium under the game-theoretic bidding model (i.e., computes ϵ -equilibria for smaller values of ϵ), as described in Section 4.2?

Our experiments show that the EM learning approach outperforms the simple approach in all three ways, across a wide range of data sets.

5.1 Repeated online auctions

In this section we consider repeated online auctions, and thus attempt to answer the first two questions above. We present results from four data sets:

- Data Set 1 has auctions of identical items, and we know the family of distributions that $f(x)$ and $g(m)$ belong to.
- Data Set 2 has auctions of non-identical items, but we know the bid distribution $f(x)$ is influenced linearly by an attribute a .

⁶ We also looked at the variant of the simple approach that directly estimates the distribution $f^1(x)$ using the selling prices. Our results show that this approach consistently underestimates $f^1(x)$ as expected, and its performance is much worse than both the EM approach and the simple approach. To avoid clutter, detailed results on this approach are not shown.




Fig. 2 Results for data set 1, distribution estimation: Distribution of bids $f(x)$ (left); distribution of highest bids $f^1(x)$ (right)

- Data Set 3 has auctions of identical items, but we do not know what kind of distributions $f(x)$ and $g(m)$ are. We use nonparametric estimation techniques to estimate the distributions.
- Data Set 4 is real-world auction data on identical items, collected from eBay.

5.1.1 Synthetic data set 1: Identical items

In this data set, the items for sale in all auctions are identical, so the number of bidders and bid amounts come from stationary distributions $g(m)$ and $f(x)$. $f(x)$ is a Normal distribution $N(4, 3.5)$. $g(m)$ is a Poisson distribution shifted to the right: $g(m - 2) = P(40)$, i.e. the number of bidders is always at least 2. The bidding history is generated using our model of the bidding process as described in Section 2. Each instance of the data set consists of bidding history from 40 auctions. We generated 15 instances of the data set.

Estimating the distributions. Both estimation approaches are informed of the parametric families from which $f(x)$ and $g(m)$ are drawn; their task is to estimate the parameters of the distributions, (μ, σ) for $f(x)$ and λ for $g(m)$. At the M step of the EM algorithm, standard ML estimates for μ , σ , and λ are computed, i.e. sample mean of the bid amounts for μ , standard deviation of the bid amounts for σ , and the mean of the number of bidders minus 2 (due to the shifting) for the Poisson parameter λ .

Our results show that the EM approach outperforms the simple approach in the quality of its estimates for the distributions $f(x)$, $g(m)$ and $f^1(x)$. Figure 2 shows typical estimated distributions⁷ and the true distributions. We observe that the plot of the estimated $f(x)$ by the simple approach is significantly shifted to the right of the true distribution, i.e. the simple approach overestimated $f(x)$. We have also calculated KL Divergences from the true distributions to the estimated distributions, and the EM

⁷ The distributions shown were randomly chosen from the 15 instances of the data set. We have verified that the plots of the other distributions are qualitatively similar.

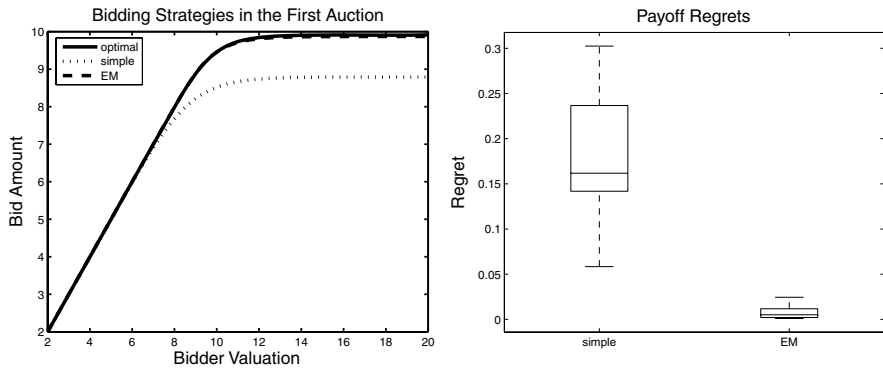


Fig. 3 Results for data set 1, bidding: Bidding strategies in the first auction (left); box plot of payoff regrets of the two approaches (right)

estimations have consistently lower KL Divergences. This difference was verified to be significant, using the non-parametric Wilcoxon sign-rank test.

Bidding in repeated auctions. Estimates from both approaches were used to compute bidding strategies for an auction environment with 8 sequentially held auctions of the same kind of items, using the decision-theoretic model introduced in Section 4.1. The agent’s “actual” expected payoffs $U_1(b, v)$ under these bidding strategies were then computed, using the true distributions. The optimal bidding strategy and its expected payoff were also computed.

Our results show that the EM approach gives rise to bidding strategies closer to the optimal strategy, and achieves higher expected payoffs, as compared to the simple approach. Figure 3 shows a plot of the bidding strategies in the first auction, and a box plot of the mean regrets, which is the differences between optimal expected payoffs and actual expected payoffs. Formally, let b^* denote the optimal strategy and \hat{b} the strategy computed using estimated distributions, then the regret given our agent’s valuation v is

$$R(v) = U_1(b^*, v) - U_1(\hat{b}, v)$$

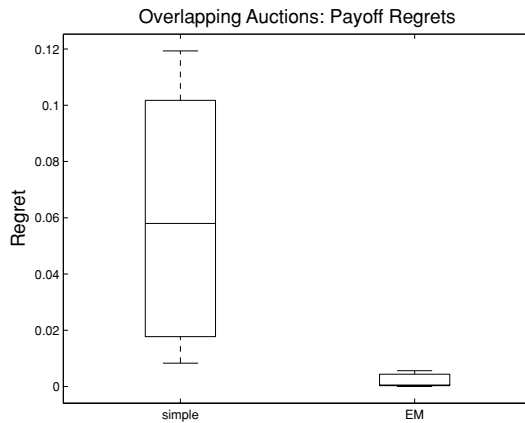
and the mean regret is the expected value of $R(v)$ over the distribution of v , which we set to be the same as the other bidders’ bid distribution f :

$$\bar{R} = \int_{-\infty}^{\infty} R(v) f(v) dv.$$

From the box plot we observe that the mean regret of the EM approach is much smaller than that of the simple approach. The ratio between the mean regrets of the EM and simple approaches is 1 : 56.

We also used the estimated distributions on the decision-theoretic model with overlapping auctions, as described in Section 4.1.3. We again tested both approaches on 8 sequentially held auctions, but now some early bidding activity on each auction was

Fig. 4 Box plot of expected payoff regrets for overlapping auctions



generated. These results are shown in Fig. 4. Again we see that the EM approach achieves higher expected payoffs (and thus less regret) compared to the simple approach. The EM approach seemed to benefit more from the extra information of early bids than the simple approach: the ratio between the mean regrets of the EM and simple approaches increased to 1 : 390.

5.1.2 Synthetic data set 2: Non-identical items

In our second data set, the items on sale are not identical; instead the distribution of valuations are influenced by an observable attribute a . In this data set the dependence is linear: $f(x | a) = N(1.1a + 1.0, 3.5)$. $g(m)$ is a Poisson distribution as before: $g(m - 2) = P(35)$. For each auction, a is sampled uniformly from the interval $[3, 9]$. In other words, this data set is similar to Data Set 1, except that the bid distribution $f(x)$ is drawn from a different parametric family. Both approaches now use linear regression to estimate the linear coefficients.

Estimating the distributions. Again, our results show that the EM approach outperforms the simple approach for this data set, in terms of its estimates for $f(x)$ and $g(m)$. Figure 5 (left) shows the estimated linear relation between the mean of $f(x | a)$ and a . From the figure we can see that the EM approach gives a much better estimate to the linear function. The simple approach again significantly overestimates the bid amounts. In fact the simple approach has consistently overestimated $f(x)$ for all the synthetic data sets we tested.

Bidding in repeated auctions. We then used the estimated distributions to compute a decision-theoretic agent's bidding strategies and expected payoffs of an auction environment with 8 sequential auctions, where the attribute a of each item is observed. The EM approach also gives better expected payoff, the statistical significance of which is confirmed by Wilcoxon's sign-rank test. Figure 5 (right) shows a box plot of regrets from different instances of data sets, which shows that the EM approach achieved consistently higher payoffs.




Fig. 5 Results for Data Set 2: Linear relationship between the mean of $f(x | a)$ and a (left). Box plot of payoff regrets (right)

5.1.3 Synthetic data set 3: Unknown distributions

We go back to the identical items model with stationary distributions $f(x)$ and $g(m)$. For this data set, $f(x)$ is a Gamma distribution with shape parameter 2 and scale parameter 3. $g(m)$ is a mixture of two Poisson distributions: $P(4)$ with probability 0.6 and $P(60)$ with probability 0.4. But now the estimation approaches does not know the types of the true distributions. Instead, both use kernel density estimation (kernel smoothing), a nonparametric estimation strategy. Essentially, given N samples from a distribution $p(x)$, we estimate $p(x)$ by a mixture of N kernel functions centered at the N samples.

Estimating the distributions. A Gaussian kernel is used for estimating $f(x)$ and a uniform kernel is used for estimating $g(m)$. At each M step of the EM algorithm, the bandwidth parameters of the two kernel estimations need to be selected. We use the simple “rule of thumb” strategy (Silverman, 1986) for bandwidth selection. We used the same kernel estimation and bandwidth selection technique for the simple approach.

Our results show that the EM approach gives better estimates than the simple approach. Figure 6 shows typical estimated distributions and true distributions. From the figure we can observe that the EM estimates of $f(x)$, $g(m)$ and $f^1(x)$ are much closer to the true distributions than the simple estimates. The EM estimates have significantly smaller KL Divergences compared to the simple estimates, verified by Wilcoxon’s sign-rank test.

Bidding in repeated auctions. We then computed the expected payoffs under the decision-theoretic model with 8 sequential auctions. The expected payoffs of the EM approach were not significantly better than that of the simple approach, as shown by the box plot in Fig. 6. One possible explanation is that although KL divergence is a good measure of similarity of distributions in general, under this particular sequential auction decision-theoretic model KL divergence might not be the best measure of quality. The bidding strategy as defined in (4) and (3) is optimal if the distributions f and g are the true underlying distributions. Using our estimated distributions, the

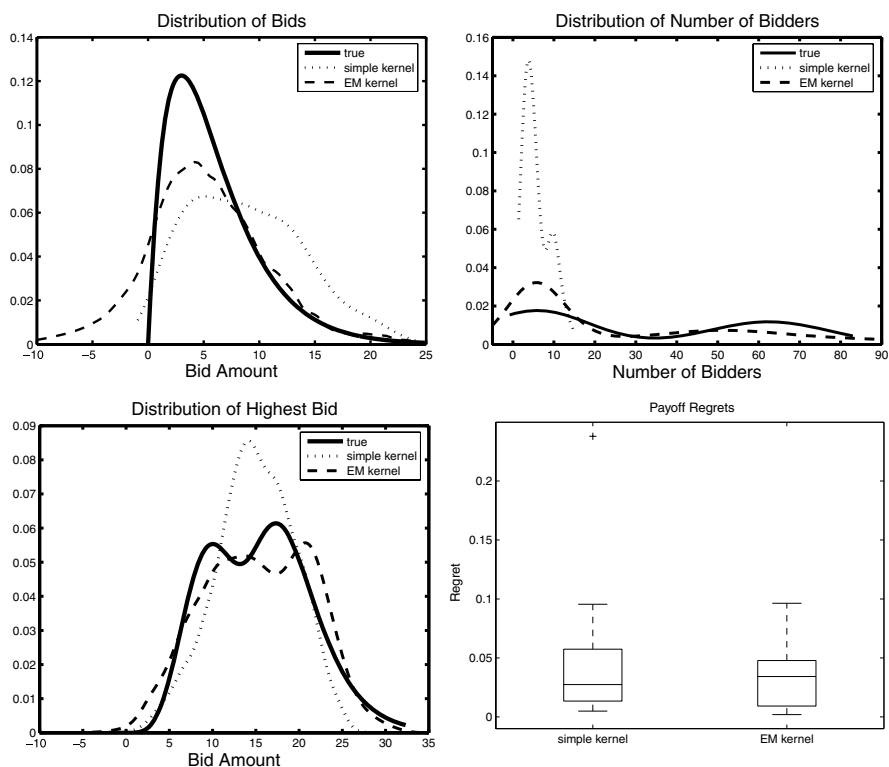


Fig. 6 Results for Data Set 3: Distribution $f(x)$ (top-left). Distribution $g(m)$ (top-right). Distribution $f^1(x)$ (bottom-left). Box plot of payoff regrets (bottom-right)

resulting bidding strategy may be higher or lower than the optimal bids. However from our experience in these experiments, bidding too high is more costly than bidding too low. This is not taken into account in the KL divergence computation as well as our ML estimation procedure. This suggests that a possible future research direction is to identify a loss function suited for the sequential auction bidding model, and compute estimated distributions by minimizing that loss function. Another possible approach is to compute a posterior distribution instead of point estimates of the parameters in each iteration. For example, West (1994) used Gibbs sampling techniques to compute the posterior distribution of parameters in a relatively simple model of incomplete data. Bidding strategies computed using the distribution of parameters should be free of the problem mentioned above. However, this kind of approach would be more computationally expensive.

5.1.4 eBay data on Sony Playstation 2 systems

Our experiments on synthetic data sets showed that our EM approach gave good estimates of the true distributions in several different settings. However, the synthetic data sets are generated using our model for the bidding process. Thus, the above ex-

periments do not tell us whether our model for the bidding process is an accurate description of what happens in real world online auctions. To answer this question, we wanted to test our approach on real world bid data. On eBay, the bidding histories of completed auctions are available for 30 days. Unfortunately, information on the hidden bids, especially the proxy bids of the winners of the auctions, is not publicly available. So unlike in the synthetic data experiments, we cannot compare our estimated distributions with the “true” distributions.⁸

To get around this problem, we used the following approach. First we collected bidding histories from a set of eBay auctions. Next we pretended that those highest bids were not placed, and the previously second highest bids were the highest bids. We then “hid” these new highest bids of each auction. This gave us a “ground truth” for the hidden bids which allowed us to evaluate our different approaches. It is true that this approach of hiding bids changed the underlying distribution of bids; however, this did not worry us as our goal was not to learn the true distribution of bids in our eBay auctions. Instead, our goal was to evaluate our algorithms’ performance on a realistic, non-synthetic data set. We believe that despite the removal of the high bid our data set preserves qualitative characteristics of the original bidding history data. If our model of the bidding process is correct, then our EM approach should be able to correctly account for the hidden bids in this data set and produce good estimates of $f^1(x)$.

We collected bidding history of eBay auctions on brand new Sony Playstation 2 (Slim Model) consoles, over the month of March 2005. We chose to study these auctions because they had been previously studied by Shah et al. (2003), who argued that bidders’ valuations on Playstations tend to be close to the private value model. We considered only auctions that lasted one day and had at least 3 bidders.⁹ We found 60 auctions that satisfied these requirements. We then randomly divided our data into a training set of 45 auctions and a testing set of 15 auctions.


Estimating the distributions. We tested four learning approaches: the EM and simple approaches that estimate a Normal distribution for $f(x)$ and a Poisson distribution for $g(m)$, and the EM and simple approaches that use kernel density estimation to estimate $f(x)$ and $g(m)$. Of course we did not have a ground truth for these distributions, so it was not possible to compare the accuracy of the predictions. However, we could use both approaches’ estimates to estimate $f^1(x)$ based on the training set, and compare these estimates against the highest bids from the test set. We did 8 runs of this experiment with different random partitions of training set and testing set, and aggregated the results. The KL Divergences of $f^1(x)$ of the approaches were similar, and no one approach was significantly better than the others.

Bidding in repeated auctions. We then computed the expected payoffs under the decision-theoretic model. The EM approaches achieved significantly higher payoffs than the simple approaches, as shown in Fig. 7. The approaches using parametric models achieved similar payoffs to the corresponding approaches with kernels. The good performance of the parametric estimation EM approach for the eBay data set

⁸ Of course we could have used our techniques to generate values for the missing bids; however, this would have been unfair when the goal was to test these techniques!

⁹ We needed at least 3 bidders so that we could drop one and still have one observed bid.

Fig. 7 Box plot of payoff regrets on the eBay Data Set



indicates that the Normal and Poisson models for $f(x)$ and $g(m)$ may be adequate models for modeling bidding on eBay.

The EM approaches did not have better KL divergence than the simple approaches, but nevertheless outperformed the simple approaches in the repeated auction experiments. This is similar to our situation in Data Set 3. Our experiments have shown that there is a positive correlation between better KL divergence and better performance in repeated auctions, but that this correlation is not perfect.

5.2 Online auctions without proxies

We built a data set modeling online auctions without proxies in order to compare the EM approach against the simple approach in the game-theoretic setting described in Section 4.2. Thus, in this section we consider questions 1 and 3 from the list at the beginning of Section 5.

Similar to Section 5.1.1, we use a normal distribution $N(5.0, 2.0)$ for $f(v)$ and a shifted Poisson distribution (with $\lambda = 10$) for $g(m)$. Each instance of the data set consists of bidding histories from 30 auctions.

Estimating the distributions. As might be expected from the results in the previous section, we found that the EM approach gave much better estimates of $f(v)$ and $g(m)$ than the simple approach. Figure 8 shows typical estimated distributions and the true distributions.

Estimating the Bayes-Nash equilibrium. We then compared the game-theoretic bidding agents that would have been built using the distributions estimated by the two learning approaches. Unlike in Section 5.1, we cannot simply use expected payoff as a measure of performance, since in a game-theoretic setting our payoff depends other agents' strategies; for example, a bad approximation to an equilibrium could actually lead to increased payoffs for all agents. (Consider e.g., the prisoner's dilemma.) Instead, what we can measure is the amount that each player could gain by unilaterally deviating from the current strategy at the "equilibrium" strategy profile computed using each learning approach. One way of thinking of this value relates to the concept of




Fig. 8 Results for online auctions without proxies: the value distributions $f(v)$ (left); the distributions of number of bidders $g(m)$ (right)

ϵ -equilibrium. (Recall that a strategy profile is an ϵ -equilibrium if each agent can gain at most ϵ by unilaterally deviating from the strategy profile.) Every strategy profile is an ϵ -equilibrium for some ϵ ; what we compute is the smallest ϵ for which the learned strategies form an ϵ -equilibrium.

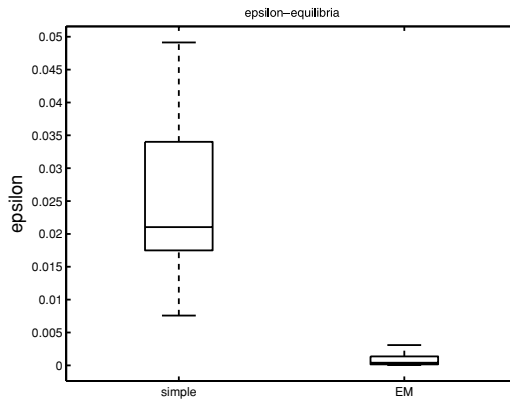
For our auction setting, we observe (from Theorem 1) that no matter what distributions we use for f and g , our agents will play a strategy in P . As a result, the true equilibrium strategy (computed from Theorem 1 using the true distributions) is always a best response to the strategies our agents play. Given our agent's valuation v and observed bidding history so far x_v , we can compute the difference between the expected payoff of the best response and the expected payoff of our agent's strategy. The ϵ for the game is then the expectation of this difference over v , m , and bidding history x_v . We approximate this expectation by sampling v , m , x_v and taking the mean of the payoff differences. We generated bidding histories for 20 auctions using each approach, and for each bidder in these auctions we computed expected payoff differences for 50 valuations v drawn from the true $f(v)$.

Our results show that strategy profiles computed using the EM approach are ϵ -equilibria for much smaller values of ϵ than the strategy profiles computed using the simple approach. This means that the EM bidding agent's strategy is much closer to the true equilibrium. We computed ϵ for 15 instances of the training data set, and Fig. 9 gives a box plot of the resulting ϵ 's.

6 Related work

Our EM approach is similar in spirit to an approach by Boutilier, Goldszmidt, and Sabata (1999). This work concerns a decision-theoretic MDP approach to bidding in sequential first-price auctions for complementary goods. For the case when these sequential auctions are repeated, this paper discusses learning a distribution of other agents' highest bids for each good, based on the winning bids in past auctions. If the agent's own bid wins in an auction, the highest bid by the other agents is hidden because only the winning bid is revealed. To overcome this problem of hidden bids, the paper uses an EM approach to learn the distribution of the highest bid.

Fig. 9 Results for online auctions without proxies: box plot of epsilon-equilibria using the simple and EM approaches



In our English auction setting the highest bids are *always* hidden, and thus cannot be directly estimated. Instead we have to estimate the distributions of bids $f(x)$ and of number of bidders $g(m)$ —which requires us to take into account bids other than the highest—and then compute the distribution of highest bids using $f(x)$ and $g(m)$. As a result the learning task in our domain is more complicated than the problem considered by Boutilier, Goldszmidt, and Sabata (1999).

Several other papers have tried to solve the hidden bid problem in various auction settings. Rogers et al. (2005) studied English auctions with discrete bid levels. They discussed estimating the distributions of valuations and numbers of bidders (f and g) from data, then using the estimated distributions to compute an optimal set of bid levels (including the reserve price) for the auctioneer. Their learning approach is to look at only the final prices of the auctions, and then to use Bayesian inference to compute posterior distributions of the parameters of f and g given the final prices of previous auctions. We note that this approach ignores all the earlier bids, which carry information about f and g , while our approach uses all bids observed. Furthermore, Rogers et al. (2005) approach works only for parametric distributions. If the true underlying distributions are not in the chosen parametric family of distributions, parametric estimation tends to give poor results. In Section 5.1.3 we showed that our approach can use nonparametric estimation techniques (kernel density estimation). Finally, the method of Rogers et al. (2005) needs to compute the joint distribution of multiple parameters, which takes exponential time in the number of parameters. Our EM approach only tries to compute ML/MAP estimates of the parameters, so each iteration of EM scales roughly linearly with the number of parameters.

Another relevant paper is Song (2004). Like us, Song studies the estimation problem in online English auctions in eBay-like environments. However she used a different approach, based on the theoretical result that the second- and third-highest valuations uniquely identify the underlying value distribution. Unfortunately, applying this fact to the eBay model presents a problem: while the highest observed bid (which is the second-highest actual bid) corresponds to the second-highest valuations, the second-highest observed bid is not necessarily the third-highest valuation. This is because the first- or second- highest bidders may have submitted bids higher than the third-highest value before the third-highest valued bidder had a chance to submit her final bid, which means her bid can be hidden from the bidding history. Thus the second-highest

observed bid is sometimes less than the third-highest valuation, and ignoring this fact would introduce bias to the estimation. Song recognizes this problem, and her solution is to use data from auctions in which the first- or second-highest bidder submitted bids higher than the second-highest observed bid *late* in the auction, because these auctions have relatively higher probability that their second-highest observed bids are third-highest valuations. However, this approach merely reduces expected bias rather than avoiding it entirely. Indeed, we note that there is empirical evidence (Roth & Ockenfels, 2002) that bidding activity late in auctions has much higher density than earlier bidding activity. This suggests that even for the selected auctions of Song's approach, there may be significant probability that second-highest observed bids are less than third-highest valuations. Our approach models the hidden bidders, so does not suffer from the bias introduced in her approach, and does not need to drop any auctions from consideration. Finally, Song (2004) did not discuss how to estimate the distribution of the number of bidders, which is essential for computing optimal bidding strategies in our repeated auctions setting (see Section 4).

Haile and Tamer's (2003) analyzed a different form of the hidden bid problem: a bidder may have submitted an earlier bid below her valuation, then the current price rises above her valuation, so she does not bid again. As a result, bidders' final observed bids may be below their valuations.¹⁰ Haile and Tamer's (2003) proposed a method to compute bounds on the value distributions given the observed bids. In our model, bidders' final observed bids are assumed to coincide with their willingness to pay, so this issue does not arise in our current model. On the other hand, Haile and Tamer's (2003) approach was intended for physical (as opposed to online) auctions. As a result they assume that the number of potential bidders in each past auction is known. This could cause their technique to overestimate the value distribution in the online setting because it does not account for potential bidders who arrive after the auction's current price exceeds their valuation. Since in practice, the *hidden bidders* problem addressed in our paper and the *bidding below valuation* problem studied by Haile and Tamer may both be present, an interesting future research direction is to combine our approach that deal with hidden bidders with Haile and Tamer's (2003) technique for bounding the valuation distribution.

7 Conclusion

In this paper we have described techniques for building bidding agents in online auction settings that include hidden bids. In particular, we have addressed the issue of estimating the distributions of the number of bidders and bid amounts from incomplete auction data. We proposed a learning approach based on the EM algorithm that takes into account the missing bids by iteratively generating missing bids and doing maximum likelihood estimates on the completed set of bids. We applied our approach to both decision-theoretic and game-theoretic settings, and conducted experiments on both synthetic data as well as on eBay data. Our results show that our approach never did worse and often did much better than the the straightforward approach of ignoring

¹⁰ Like us (although in a different way), Haile and Tamer's (2003) also address the problem that the highest bid is not the actual highest value.

the missing data, both in terms of the quality of the estimates and in terms of expected payoffs under a decision theoretic bidding model.

Acknowledgment We would like to thank three anonymous reviewers for their many helpful suggestions on improving this manuscript.

References

- Andrieu, C., de Freitas, N., Doucet, A., & Jordan, M.I. (2003). An introduction to MCMC for machine learning. *Machine Learning*.
- Anthony, P., Hall, W., Dang, V. D., & Jennings, N. (2001). Autonomous agents for participating in multiple online auctions. In *Proc. of the IJCAI workshop on e-business and the intelligent web*.
- Arora, A., Xu, H., Padman, R., & Vogt, W. (2003). Optimal bidding in sequential online auctions. Working paper.
- Athey, S., & Haile, P. (2002). Identification in standard auction models. *Econometrica*, 70(6), 2107–2140.
- Boutillier, C., Goldszmidt, M., & Sabata, B. (1999). Sequential auctions for the allocation of resources with complementarities. In *Proceedings of 16th international joint conference on artificial intelligence*.
- Byde, A. (2002). A comparison among bidding algorithms for multiple auctions. *AAMAS workshop on agent-mediated electronic commerce IV*.
- Cai, G., & Wurman, P. R. (2003). *Monte Carlo approximation in incomplete-information, sequential-auction games* (Technical report). North Carolina State University.
- Dempster, A., Laird, N., & Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39(1), 1–38.
- Greenwald, A., & Boyan, J. (2004). Bidding under uncertainty: Theory and experiments. In *Proc. Uncertainty in Artificial Intelligence-04*.
- Haile, P. A. & Tamer, E. (2003). Inferences with an incomplete model of English auctions. *Journal of Political Economy*, 111(1), 1–51.
- Klemperer, P. (2000). Auction theory: A guide to the literature. In P. Klemperer (Ed.), *The economic theory of auctions*. Edward Elgar.
- Mackie-Mason, J. K., Osepayshvili, A., Reeves, D. M., & Wellman, M. P. (2004). Price prediction strategies for market-based scheduling. In *Fourteenth international conference on automated planning and scheduling* (pp. 244–252).
- Milgrom, P., & Weber, R. (2000). A theory of auctions and competitive bidding, II. In P. Klemperer (Ed.), *The economic theory of auctions*. Edward Elgar.
- Osepayshvili, A., Wellman, M. P., Reeves, D. M., & Mackie-Mason, J. K. (2005). Self-confirming price prediction for simultaneous ascending auctions. In *Proc. uncertainty in artificial intelligence*.
- Rogers, A., David, E., Schiff, J., Kraus, S., & Jennings, N.R. (2005). Learning environmental parameters for the design of optimal English auctions with discrete bid levels. In *Proceedings of 7th international workshop on agent-mediated e-commerce* (pp. 81–94). Utrecht, Netherlands.
- Roth, A. E., & Ockenfels, A. (2002) Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *American Economic Review*.
- Shah, H. S., Joshi, N. R., Sureka, A., & Wurman, P. R. (2003). Mining for bidding strategies on eBay. *Lecture Notes on Artificial Intelligence*.
- Silverman, B.W. (1986) *Density estimation*. London: Chapman and Hall.
- Song, U. (2004). *Nonparametric estimation of an eBay auction model with an unknown number of bidders*. Working paper.
- Stone, P., Schapire, R. E., Csirik, J. A., Littman, M. L., & McAllester, D. (2002) ATTac-2001: A learning, autonomous bidding agent. In *AAMAS workshop on agent-mediated electronic commerce IV* (pp. 143–160).
- Weber, R. (1983). Multi-object auctions. In R. Engelbercht-Wiggans, M. Shubik & R. Stark (Eds.), *Auctions, bidding, and contracting: Uses and theory* (pp. 165–191). New York, University Press.
- Wellman, M. P., Greenwald, A., Stone, P., & Wurman, P. R. (2002). The 2001 trading agent competition. In *Proceedings of innovative applications of artificial intelligence*.
- West, M. (1994). Discovery sampling and selection models. In J. O. Berger & S. S. Gupta (Eds.), *Decision theory and related topics IV*. New York: Springer Verlag.