

Optimal Auction Design for Agents with Hard Valuation Problems

David C. Parkes*

Computer and Information Science Department, University of Pennsylvania,
Philadelphia PA 19104 dparkes@unagi.cis.upenn.edu

Abstract. As traditional commerce moves on-line more business transactions will be mediated by software agents, and the ability of agent-mediated electronic marketplaces to efficiently allocate resources will be highly dependent on the complexity of the decision problems that agents face; determined in part by the structure of the marketplace, resource characteristics, and the nature of agents' local problems. We compare auction performance for agents that have hard local problems, and uncertain values for goods. Perhaps an agent must solve a hard optimization problem to value a good, or interact with a busy and expensive human expert. Although auction design cannot simplify the valuation problem itself, we show that good auction design can simplify *meta*-deliberation – providing incentives for the “right” agents to deliberate for the “right” amount of time. Empirical results for a particular cost-benefit model of deliberation show that an ascending-price auction will often support higher revenue and efficiency than other auction designs. The price provides agents with useful information about the value that other agents hold for the good.

Keywords: agent-mediated electronic commerce, valuation problem, metadeliberation, auction theory.

In *Agent Mediated Electronic Commerce II: Towards Next-Generation Agent-Based Electronic Commerce Systems*, Ed.: A. Moukas, C. Sierra, F. Ygge (LNAI; 1788), Springer, Berlin, 2000. Earlier version, *Proc. IJCAI'99 Workshop on Agent Mediated Electronic Commerce (AmEC-99)*, Stockholm, July 1999.

* This research was funded in part by the National Science Foundation under Grant SBR 97-08965.

1 Introduction

As traditional commerce moves on-line more business transactions will be mediated by software agents, and dynamically negotiated between multiple and fluidly changing partners. The ability of agent-mediated electronic marketplaces to efficiently allocate resources will be highly dependent on the complexity of the decision problems that agents face; determined in part by the structure of the marketplace, resource characteristics, and the nature of agents' local problems.

While many of the costs that are associated with traditional auctions, such as the cost of participation (making bids and watching the progress of an auction), are unimportant in agent-mediated electronic auctions, the cost of valuation remains important [17]. The value of a good is often uncertain, and an accurate valuation can require that an agent solves a hard optimization problem, or interacts with a busy and expensive human expert. In fact, electronic markets may make the valuation problem more difficult, because of mitigating factors such as decreased aggregation, increased product differentiation, and increased dynamics [1, 4, 5]. In this paper we compare auction performance for agents that have hard local problems, and uncertain values for goods.

Just as careful market design can reduce the complexity of the bidding problem, for example by providing incentives for agents to reveal their true value for a good [28], careful market design can also reduce the loss in efficiency that is associated with agents that have hard valuation problems. Unlike the bidding problem, market design can not simplify the valuation problem itself. However market design can improve the quality of an agent's decisions about when to reason about the value of a good. A well structured marketplace can provide information to enable the "right" agents to deliberate for the "right" amount of time. Roughly, agents with high values should deliberate more than agents with low values.

For example, consider a bidding agent that participates in an on-line auction for a flight to Stockholm, initialized by a user with a *lower bound* \underline{v} on value. The user does not know her exact value for the flight, but finds it relatively easy to bound her value. Although the agent can absorb the costs of monitoring the auction and placing bids, the agent cannot easily refine the user's value for the flight. The value of non-standard and short-supply goods is often subjective, and can depend on many factors that an agent cannot know. However, in an ascending-price auction the agent can bid up to \underline{v} , and then prompt the user for a more accurate value. Compare this to a sealed-bid auction where the user needs *a priori* information about the distribution of bids from other agents to make a good decision about how much time to spend deliberating about her value for the flight. The ascending-price auction provides dynamic information on the value of other participants, and can enable the user to avoid deliberation altogether – for example if the price increases above an *upper bound* on value.

We compare the performance of three market designs with agents that have hard valuation problems: a posted-price sequential auction; a second-price sealed-bid auction; and a first-price ascending-price auction [12]. In the posted-price auction the seller offers the good at a fixed price to each agent in turn, and

does not sell the good if no agent accepts the price. The price is set dynamically in the ascending- and sealed-bid auctions, and we allow the seller to optimize the ask price for distributional information about the values of agents in the posted-price auction.

In Section 2 we introduce a simple model for agents with hard valuation problems that allows the derivation of optimal expected-case metadeliberation and bidding strategies for risk-neutral agents in each auction; we describe the optimal strategies in Section 3. Section 4 presents empirical results from simulation, comparing the efficiency and revenue in each auction for different numbers of agents and different levels of local problem complexity. Finally we discuss related work in auction theory, artificial intelligence, and economics, before presenting our conclusions.

2 The Valuation Problem

In standard auction theory agents either know their value for a good (private values) or the value is common across all agents but unknown because of missing information (common-values) [12]. We model an auction with agents that have private independent values for a good, but uncertainty about the value. We believe that this model is especially relevant in on-line auctions, where agents can have hard local problems (e.g. a manufacturing agent that bids for components), or goods are non-standard and difficult to value (e.g. collectibles at www.ebay.com). We assume that agents have an option to refine their value for a cost. The cost is considered to represent the actual cost of consulting an expensive expert, or the cost that results from suboptimal or missed bids because of lost deliberation about the value of goods in other marketplaces.

Conceptually, one can partition the decision problem of a bidding agent into three sub-problems: *metadeliberation*, *valuation*, and *bidding*. The optimal bidding strategy depends on the auction and an agent's (possibly approximate) solution to its valuation problem. The optimal metadeliberation strategy follows from an analysis of an agent's valuation problem and bidding strategy. An agent bids only when it has decided to perform no more deliberation about the value of the good.¹ Previous models of agent-mediated markets have addressed the complexity of the *bidding problem*, that is deciding on an optimal bid given the value of a good, but largely ignored the valuation and metadeliberation problems (although see [24, 25]).

¹ In an alternative model agents do not explicitly solve the valuation problem, but select a bidding strategy directly, based on the payoff from past bids [2]. It is often useful to separate the valuation and bidding phases because: (a) there might be a separation of skills/information – for example when a software agent *bids* for a person that *values* the good; (b) in a business context, separation can enable a bidding agent to leverage existing decision analysis tools and models for the valuation problem; (c) in some markets (for example incentive compatible markets) the bidding problem is trivial.

2.1 A Simple Theoretical Model

We propose a simple model for the valuation problem of an agent, that enables the derivation of optimal metadeliberation and bidding strategies for agents in each auction. The model matches some of the properties of standard algorithmic techniques for solving hard optimization problems, such as Lagrangian relaxation, depth-first search, and branch-and-bound. Furthermore, the model supports a mode of interaction between people and software bidding agents that is provided in some current on-line auctions [17]. We do not expect the valuation problems and decision procedures of real agents (or real experts) to have characteristics that match the precise assumptions (e.g. distributional assumptions) of our model. However, we believe that the general results from our analysis will hold in many real problem domains for agents with hard valuation problems.

Every agent i has an unknown true value v_i for a good, and maintains a lower bound \underline{v} and upper bound \bar{v} on its value, see Fig 1 (a). Agent i believes that its true value is uniformly distributed between its bounds, $v \sim U(\underline{v}, \bar{v})$. Given this belief the expected value for the good is $\hat{v} = (\underline{v} + \bar{v})/2$. As an agent deliberates its bounds are refined and its belief about the value of the good changes, with expected value \hat{v} converging to v over time.

Let $\Delta = \bar{v} - \underline{v}$ denote an agent’s current uncertainty about the value of the good. Agents have a deliberation procedure that adjusts the bounds on value, reducing uncertainty by a multiplicative factor α , where $0 < \alpha < 1$. The new bounds are $\alpha\Delta$ apart, and consistent with the current bounds (but not necessarily adjusted symmetrically), see \underline{v}' and \bar{v}' in Fig 1 (a). For a small α the uncertainty is reduced by a large amount, and we refer to $(1 - \alpha)$ as the “computational effectiveness” of an agent’s deliberation procedure. Furthermore, we model the new expected value \hat{v}' for the good after deliberation as uniformly distributed $\hat{v}' \sim U(\underline{v} + \alpha\Delta/2, \bar{v} - \alpha\Delta/2)$, such that the new bounds are consistent with the current bounds. After deliberation an agent believes the value of the good is uniformly distributed between its new bounds. Agents incur a cost C for each deliberation step, that we assume is constant for all steps, and independent of the final outcome of the auction.

3 The Metadeliberation Problem

The *metadeliberation problem* is to determine how much deliberation to perform before placing a bid. The decision is a tradeoff between reducing uncertainty about the value of the good so that the bid is accurate, and avoiding the cost of deliberation. An agent’s optimal metadeliberation strategy *does* depend on the bids that other agents will make, even in incentive compatible auctions (unlike an agent’s optimal bidding strategy).

Given our model of an agent’s valuation problem and decision procedure we derive optimal metadeliberation strategies for agents within the general framework of Russell and Wefald [22]. The key observation is that the value of deliberation is derived from the effect of deliberation on an agent’s bid. Deliberation

can only be worthwhile when: (1) it changes an agent’s bid; (2) the new bid has greater expected utility than the old bid. For example, an agent should never deliberate about its value for a good if its current upper bound on value is less than the ask price, because further deliberation can never cause the agent to accept the price. Metadeliberation is hard because of uncertainty, about: the effect of placing a bid b in an auction (this can depend on the bids of other agents); the outcome of further deliberation (otherwise deliberation is unnecessary!), and the value of goods. We describe normative metadeliberation strategies for agents in each auction below, but see [19] for derivations. We assume risk-neutral agents, who receive utility $v_i - p$ for purchasing a good at price p .²

3.1 Second-Price Sealed-Bid

In a second-price sealed-bid (Vickrey) auction agents need distributional information about the bids from other agents to make good metadeliberation decisions. For example, if an agent does not have any information about the bids from other agents it cannot know the probability of winning the auction with a bid b , or its expected surplus if it wins (this depends on the second-highest bid received). Uninformed agents are left to either follow a worst-case metadeliberation strategy (don’t deliberate), a best-case strategy (that recommends too much deliberation), or an ad-hoc strategy that makes implicit assumptions about the bids from other agents. We do not consider the mechanism (for example Bayesian learning) with which agents become informed, but provide agents with approximately correct distributional information about the bids from other agents.

Informed agents can follow expected-utility maximizing metadeliberation strategies; agents compare the expected utility of placing an optimal bid after deliberation with the expected utility of placing an optimal bid before deliberation, given their beliefs about deliberation, the value of the good, and the bids of other agents. The mapping from expected value \hat{v} to expected utility for an agent’s optimal bid ($b^* = \hat{v}$) is non-linear, because an agent is more likely to win the auction with a higher bid. Although an agent’s mean expected value \hat{v}' after deliberation is equal to an agent’s expected value \hat{v} before deliberation, deliberation can have positive utility because of this non-linear mapping. The number of deliberation steps that an agent performs depends on the number of agents in the auction, an agent’s current beliefs $[\underline{v}, \bar{v}]$, and the computational effectiveness and cost of its deliberation procedure. Agents with large uncertainty and high expected values tend to deliberate more than other agents.

² The optimal bidding strategy for a risk-neutral agent with an uncertain value for the good is “expected-value equivalent” to the optimal bidding strategy for an agent that knows its value for the good. For example, when offered a good at a fixed price p , an agent with beliefs $[\underline{v}, \bar{v}]$ should buy the good for a price $p < \hat{v}$. Similarly, in a second-price sealed-bid auction an agent should bid \hat{v} .

3.2 Posted-Price Sequential

In a posted-price sequential auction an agent that receives an ask-price p for the good holds an exclusive offer until it accepts or rejects the price. The only uncertainty in an agent’s metadeliberation problem is due to the agent’s own uncertainty about the value of the good. There is no uncertainty from the actions of other agents in the auction. Although agents that receive the offer can make good metadeliberation decisions, the revenue and efficiency is often less than for auctions that set the price dynamically.

The optimal expected-value metadeliberation and bidding strategy for an agent that faces a fixed price p and has beliefs $[\underline{v}, \bar{v}]$ is to deliberate while its expected value \hat{v} is close to the ask price, and then accept a price $p < \hat{v}$, and reject the price otherwise, see Fig 1 (b). An agent deliberates while its expected value \hat{v} is within $\gamma^* \Delta / 2$ of the price, for a threshold $\gamma^*(\alpha, C, \Delta)$ that depends on the computational effectiveness $(1 - \alpha)$ of its deliberation procedure, its cost C for deliberation, and its current uncertainty Δ in value. The threshold decreases as an agent deliberates, and eventually an agent will not deliberate for any ask price (when $\gamma^* = 0$).




Fig. 1. (a) Valuation problem. Upper and lower bounds \bar{v} and \underline{v} on value before deliberation, with uncertainty $\Delta = \bar{v} - \underline{v}$; new bounds after deliberation \bar{v}' and \underline{v}' are $\alpha\Delta$ apart and consistent with the initial bounds. Optimal Metadeliberation and Bidding Strategies: (b) Posted-price auction; (c) Ascending-price auction.

3.3 Ascending-Price

The auctioneer in an ascending-price auction announces an initial ask price, p , and increases the price a minimum bid increment ϵ whenever a bid is received. The auction closes when no bids are received, with the good sold to the highest bidder for the price that it bid. The optimal metadeliberation and bidding strategy in the ascending-price auction is different than in the posted-price auction

because: (1) the price of the good can increase over time; (2) an agent that bids for the good at price p cannot be sure that it will win the good. In addition to choosing to `deliberate` or `bid`, it can also be useful for agents to `wait` because the price can increase as the result of bids from other agents. Every agent hopes that another agent will deliberate and place a new bid that increases the price. Agents are locked into a deliberation “waiting game”. If the price increases above an agent’s upper bound \bar{v} on value the agent can avoid deliberation completely.

Agents can be in one of three states, depending on the relative position of the ask price with respect to their beliefs $[\underline{v}, \bar{v}]$ about the value of the good, see Fig 1 (c). Agents in state $s1$ always bid, and therefore the ask price in the auction is always ϵ (the minimum bid increment) above the second-highest lower threshold on deliberation for all agents (after a bidding war between all agents in state $s1$). Agents in state $s3$ leave the auction because the ask price is greater than their upper thresholds on deliberation, and can only increase. Agents in state $s2$ remain *active*, and are locked into a waiting game. Every active agent will deliberate to prevent the auction closing, but prefers to wait for another agent to deliberate. The unique symmetrical Nash equilibrium of the waiting game has all active agents that are not currently winning the auction play a mixed strategy: deliberate with probability $1/(N_a - 1)$ when the auction is about to close, for N_a active agents (the active agent that is winning the good will not deliberate because it is happy for the auction to terminate).³ When there is a single active agent left the auction terminates, with that agent winning the auction.

4 Empirical Results

We model a simple market for a single good, with agents that have true values v_i for the good drawn from a uniform distribution, such that $v_i \sim U(0, 10)$. We implement the optimal bidding and metadeliberation strategies for agents in each auction, and compare the performance of each auction, in terms of efficiency, revenue and average utility from participation. Every agent has initial beliefs $\underline{v} = 0$ and $\bar{v} = 10$, and a deliberation procedure with computational effectiveness $(1 - \alpha)$ and cost C . We simulate deliberation to match the agents’ simple model. After assigning a true value for the good to each agent we use a stochastic procedure to generate new bounds after deliberation such that: (1) the true value remains between the bounds; (2) the true value is uniformly distributed between the bounds over all possible stochastic sequences of deliberations.⁴

Efficiency is computed as the (true) value for the good of the agent that wins the auction, as a fraction of the maximum value over all agents. Revenue is

³ The agents need to know how many agents remain active to implement the Nash equilibrium. We could ask the agents to pay a small “participation fee” in each round of the auction to remain active, so that the auctioneer can report this information in each round.

⁴ Actually, this requires that the new expected value for the good in simulation is not uniformly distributed with respect to the current bounds, but favors more central values.

computed as the price paid for the good, as a fraction of the maximum value over all agents. Finally, the average utility to an agent for participation in the auction is computed as the surplus $(v_i - p)$ to agent i that wins the auction for price p , minus the total cost of deliberation for all agents, and divided by the number of agents in the auction. Agents can lose utility from participation because of assumptions made in deriving metadeliberation strategies, for example about the bids of other agents and an agent’s own deliberation procedure.⁵ We check that the utility for participation is positive to validate agent strategies.

We compare the performance of each auction as we vary the number of agents, N , and the computational effectiveness $(1 - \alpha)$ and cost C of agents’ deliberation procedures. All results are averaged over at least 1000 trials. Efficiency is often the primary performance measure of mechanism design, but if the auctioneer is also the seller then revenue can be important. We write $M_1 \succ M_2$ if mechanism M_1 dominates M_2 in terms of efficiency and revenue, or $M_1 \approx M_2$ if the mechanisms have conflicting ordering for revenue and efficiency, or very similar performance; A denotes the ascending-price auction, S the sealed-bid auction, and P the posted-price auction.

4.1 Adjusting the Number of Agents

Fig 2 (a) compares the performance of each auction mechanism as the number of agents is varied between 5 and 100 (log scale). All agents have deliberation procedures with $\alpha = 0.7, C = 0.1$. We plot efficiency (i), revenue (ii), average computational-cost (iii), and average utility (iv). The bid-increment in the ascending-price auction is set to enable agents to achieve positive expected utility from participation in the auction, while maximizing performance.

For $N < 10$ we have $A \succ S \succ P$, with agents achieving positive expected utility for participation in all auctions. For $10 < N < 35$ it at first appears that $S \succ A$, however the agents in the sealed-bid auction now have negative utility for participation (iv).⁶ So, discounting the performance of the sealed-bid auction, for $10 < N < 70$ agents we have $A \succ P \succ S$. Finally, for large numbers of agents, $N > 70$ we have $P \succ A \succ S$. Therefore, the ascending-price auction performs best for small to medium numbers of agents ($N < 70$), and the posted-price auction performs better than the sealed-bid auction for medium to large numbers of agents ($N > 10$), and better than the ascending-price auction for large numbers of agents ($N > 70$).

Although the surplus to the winning agent remains approximately constant in all auctions as the number of agents increases (unlike in regular auctions,

⁵ Agents that deliberate in the ascending-price auction assume that they will win the good for current ask price p if they bid. The auction actually remains open and other agents can bid.

⁶ This is because there are no homogeneous beliefs that agents can hold about the distribution of bids from agents that are consistent with the actual distribution of bids that occurs when agents hold the beliefs. The inaccuracy in agents’ models leads to a loss in utility.

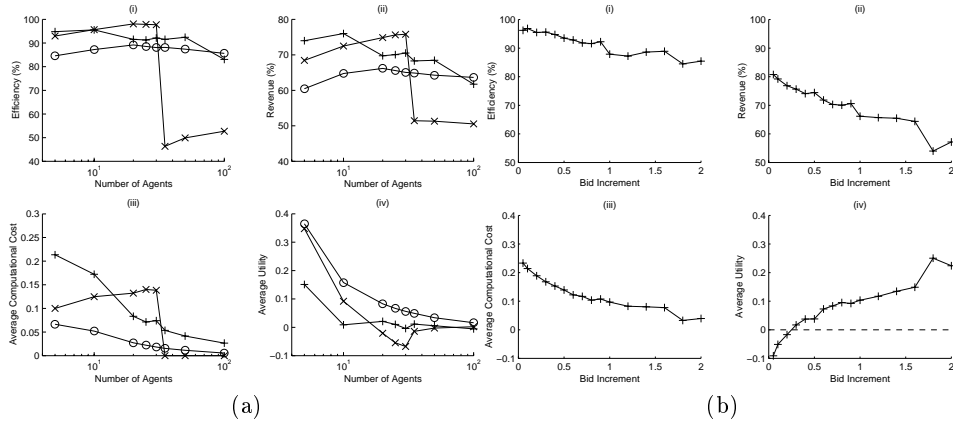


Fig. 2. (a) Auction performance: ‘x’ sealed-bid; ‘o’ posted-price sequential; ‘+’ ascending-price. Agents with $\alpha = 0.7, C = 0.1$. Bid increment in the ascending-price auction is adjusted to provide agents with positive expected utility from participation. Averaged over 1000 trials. (b) Performance of ascending-price auction. $N = 10$ agents, $\alpha = 0.7, C = 0.1$.

where it falls), the average utility for participation can only remain positive if the total computational cost also remains approximately constant (or increases to no more than the surplus). The only auction that can sustain a fixed but positive amount of deliberation as the number of agents increases is the posted-price mechanism, because it isolates agents from the effect of more agents by offering the good sequentially to each agent. Efficiency and revenue tend to decrease as the number of agents increases, again contrary to the performance markets with agents that have easy valuation problems. Performance decreases because the surplus from participation is less able to support the amount of deliberation that is necessary for high efficiency and revenue with large numbers of agents.

Fig 2 (b) plots the performance of the ascending-price auction for $\alpha = 0.7, C = 0.1$ and $N = 10$, as the bid-increment ϵ is increased. Small bid increments support more efficient allocations and higher revenue, but also lead to more deliberation and can result in agents receiving negative utility from participation in the auction. In this example the agents have positive utility for participation with a bid increment $\epsilon > 0.25$, see Fig 2 (b) iv, and the seller is able to achieve a revenue that is almost as good as that possible with very small bid increments.

4.2 Adjusting the Complexity of the Valuation Problem

Fig 3 compares the efficiency (a), and revenue (b), for markets with $N = 10$ agents as the effectiveness $(1 - \alpha)$ and cost C of deliberation changes. In a single experiment all agents have the same α and C parameters. Equivalently, we model agents with the same bounded resources, but easy local problems for large $(1 - \alpha)$ and relaxed time-constraints for small costs C .




Fig. 3. Auction performance for $N = 10$ agents with different levels of deliberation effectiveness and cost. ‘+’ ascending-price; ‘x’ sealed-bid; ‘o’ posted-price sequential auctions. Bid increment $\epsilon = 0.4$, 1000 trials. (a) Efficiency; (b) Revenue.

In this market, with $N = 10$ agents, when agents have small deliberation costs ($C = 0.05$) then $A \succ S \succ P$, and $A \approx S \succ P$ for $C = 0.1$. Auctions with dynamic prices, such as the ascending-price or sealed-bid auctions, perform better than the posted-price auction for small C . For medium deliberation costs ($C = 0.5$), the sealed-bid auction fails and $A \approx P \succ S$. For large deliberation costs (e.g. $C = 1$) there is no deliberation in any auction, and auction design does not matter. Efficiency and revenue increase as deliberation effectiveness increases, and decrease as deliberation cost increases. The ascending-price and sealed-bid auctions become approximately *revenue equivalent* for small deliberation costs and high computational effectiveness (c.f. easy local problems), see Fig 3 (b) i, $C = 0.05, 1 - \alpha = 0.7$.

Finally, consider a market with a *mixture* of agents; some with hard valuation problems (“inexperienced agents”), and some with easy valuation problems (“experienced agents”). We assume a fraction f of inexperienced agents; and a fraction $1 - f$ of experienced agents that know their value v_i for the good. Fig 4 plots results for $N = 10$, and inexperienced agents with $\alpha = 0.3$ and $C = 0.5$. The ascending-price auction has the best performance, but the sealed-bid auction often outperforms the posted-price auction, even though the efficiency and revenue comes only from experienced agents. For a small to medium fraction of inexperienced agents, $f < 0.4$, the sealed-bid and ascending-price auctions are approximately *revenue equivalent*, and $A \approx S \succ P$. For a medium to large fraction of inexperienced agents, $0.4 < f < 0.8$, we have $A \succ S \succ P$, and with many inexperienced agents the sealed-bid auction fails and $A \succ P \succ S$. The posted-price auction performs badly for heterogeneous markets, e.g. $f = 0.6$, because the seller must choose an ask price to target one type of agent. The other auctions are better able to involve both agent types. Also, we see that when there are only a few experienced agents they are able to use their informational

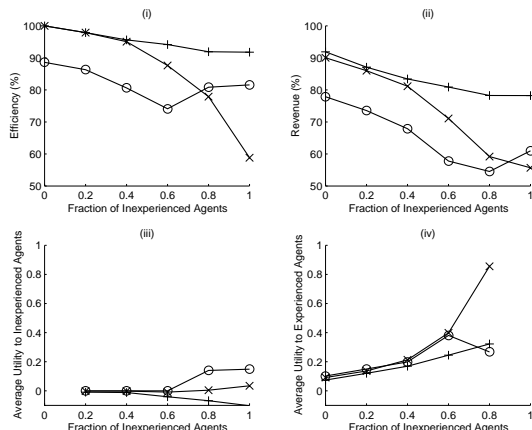


Fig. 4. Performance of the sealed-bid ‘x’, ascending-price ‘+’, and posted-price ‘o’ auctions as the fraction of inexperienced agents in the system increases. $N = 10$, $\alpha = 0.3$, $C = 0.5$. The results are computed over 2000 trials. Bid increment $\epsilon = 0.2$.

advantage, and achieve high average utility from participation in the auction, see Fig 4 (iv).

5 Related Work

In Parkes and Ungar [15] we introduce a formal definition for *bounded-rational compatible* (BRC) auctions, to describe auctions in which agents can bid optimally with approximate values. Iterative auctions are a special class of BRC auctions which can compute optimal allocations with enough agent computation. Empirical results for a model of agents with *limited computation* highlight the importance of BRC auctions in combinatorial allocation problems, with agents that need *bundles* of items. *Iterative* bundle auctions, such as *iBundle* [16, 18], that allow agents to perform incremental value computation and adjust their bids in response to bids from other agents, are particularly important in applications to hard distributed optimization problems.

There is a growing literature on auctions for e-commerce, see [7] for an introduction. Hard valuation problems are relevant in many online consumer auctions, for example for collectibles and refurbished electronic goods. This provides an appealing explanation for why many on-line auctions are ascending-price, while few on-line auctions are sealed-bid. This explanation seems more probable than alternative explanations in terms of risk-seeking agents that enjoy risk-taking in iterative auctions [13].

We believe that this work is the first to compare the performance of auctions with a normative deliberation model for agents with limited or costly computation. Early work in market-oriented programming [29] assumed that agents were provided with closed-form solutions to their local valuation problems, so

minimizing agent valuation work was not important. Problems associated with limited or costly computation in the *auctioneer* have received recent attention, in particular with respect to the generalized Vickrey auction [27]. Interesting recent work explores methods to introduce approximate solutions but retain the incentive-compatibility property (such that truth-telling remains optimal for self-interested agents) [8, 9, 14]. However, these methods still require that agents compute their values for all possible outcomes, which is often impossible. Shoham and Tennenholtz [26] consider the communication complexity of auction mechanisms.

Sandholm [25] proposed TRACONET, a decentralized system for dynamic task reallocation via negotiation between agents, with an application to a transport scheduling problem in which agents have hard local valuation problems. Safe and opportunistic strategies are proposed for contracting strategies with approximate valuations. Sandholm and Lesser [23] introduce decommitment to allow agents that make early mistakes because of approximate values and uncertainty to back-track, and increase allocative efficiency. The technique allows agents to integrate local deliberation with negotiation between many other agents. The focus is on a *decentralized* system, while our work considers techniques for auction systems—with a centralized auctioneer.

Parunak *et al.* [20] describe a market-based constraint system for interactive decision support between cooperative agents. The MARCON system aims to reduce agents’ work; agents can express broad initial preferences and must only refine preferences as necessary to compute a good system-wide solution.

Models of economic search with one-time participation costs have demonstrated that *posted-price* sequential auctions can reduce agent uncertainty and improve performance when agents are uninformed about the bids of other agents [3, 10]. Rothkopf and Harstad [21] argue that auction design should be considered in the context of a marketplace, with agents that take active participation decisions. Although some auction models consider agents with *one-time* participation costs [6, 11], the models cannot distinguish between sealed-bid and iterative auctions.

6 Conclusions

We have presented the results of an empirical comparison of the performance of first-price ascending, second-price sealed and posted-price sequential auctions, for a simple model of agents with hard valuation problems. The empirical results show clear support for the bounded-rational compatible [15] properties of ascending-price and other iterative auctions, first identified by Parkes *et al.* [17] as important for on-line auctions. It is interesting that the *revenue-equivalence theorem* [12] fails when agents have hard valuation problems and limited or costly computation.

The ascending-price auction dominates the sealed-bid auction, in terms of allocative efficiency and revenue, when there are more than a few agents in the market (e.g. $N > 5$). This holds even when agents are uninformed about

the values of other agents in the ascending-price auction but informed in the sealed-bid auction. Iterative auctions do more than provide agents with useful information about the values of other agents, they can also reduce an agent's *uncertainty* about the outcome, i.e. about the final prices and allocation.

Posted-price auctions are often necessary with many agents (e.g. $N > 50$) because they eliminate a bidder's uncertainty about the outcome of the auction. An agent knows that if it accepts the price of an item it will definitely receive the item, and at that price. This simplifies an agent's metadeliberation problem. However, with small numbers of agents posted-price auctions are typically less efficient than auctions with dynamic pricing unless the seller is very well informed about agents' values.

Agents with information about the likely outcome of an auction can deliberate more efficiently about the values of different outcomes, for example quickly eliminating from their possibility set all bundles of items that are very expensive. Iterative auctions provide agents with this information dynamically during the auction as bids are received. The "right" agents can deliberate for the "right" amount of time, shifting deliberation away from agents with low values and towards agents with high values, and towards values for bundles that fit into good global solutions.

Another approach to improve deliberation efficiency could provide agents with historical information about clearing prices from previous problem instances, to allow agents to predict likely outcomes even before an auction begins. This is a useful role for *middleware* in agent-mediated auctions.

References

1. Bakos, Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science* 43(12).
2. Boutilier, C.; Goldszmidt, M.; and Sabata, B. 1999. Sequential auctions for the allocation of resources with complementarities. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, 527–534.
3. Ehrman, C., and Peters, M. 1994. Sequential selling mechanisms. *Economic Theory* 4:237–253.
4. Greenwald, A., and Kephart, J. O. 1999. Shopbots and pricebots. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, 506–511.
5. Guttman, R. H., and Maes, P. 1998. Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *Proc. 2nd International Workshop on Cooperative Information Agents (CIA'98)*.
6. Hausch, D. B., and Li, L. 1993. A common value auction model with endogeneous entry and information acquisition. *Economic Theory* 3:315–334.
7. Huhns, M. H., and Vidal, J. M. 1999. Online auctions. *IEEE Internet Computing* 3(3):103–105.
8. Kfir-Dahav, N. E.; Monderer, D.; and Tennenholtz, M. 1998. Mechanism design for resource bounded agents. Technical report, Technion.
9. Lehmann, D.; O'Callaghan, L.; and Shoham, Y. 1999. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proc. ACM Conference on Electronic Commerce (EC-99)*.

10. Lippman, S. A., and McCall, J. J. 1976. The economics of job search: A survey. *Economic Inquiry* 14:155–189.
11. McAfee, R. P., and McMillan, J. 1987a. Auctions with entry. *Economic Letters* 23:343–347.
12. McAfee, R. P., and McMillan, J. 1987b. Auctions and bidding. *Journal of Economic Literature* 25:699–738.
13. Monderer, D., and Tennenholtz, M. 1998. Internet auctions - are they gamblers' attraction? Technical report, Technion.
14. Nisan, N., and Ronen, A. 1999. Algorithmic mechanism design (extended abstract). In *Proc. 31st Annual Symposium on Theory of Computing (STOC99)*.
15. Parkes, D. C., and Ungar, L. H. 2000a. Bounded rational compatible auctions. Submitted for publication.
16. Parkes, D. C., and Ungar, L. H. 2000b. Iterative combinatorial auctions: Theory and practice. In *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)*, 74–81.
17. Parkes, D. C.; Ungar, L. H.; and Foster, D. P. 1999. Accounting for cognitive costs in on-line auction design. In Noriega, P., and Sierra, C., eds., *Agent Mediated Electronic Commerce (LNAI 1571)*. Springer-Verlag. 25–40. Earlier version appeared at the Agents'98 Workshop on Agent Mediated Electronic Trading, 1998.
18. Parkes, D. C. 1999a. \mathcal{B} Bundle: An efficient ascending price bundle auction. In *Proc. ACM Conference on Electronic Commerce (EC-99)*, 148–157.
19. Parkes, D. C. 1999b. On agent metadeliberation strategies in auctions. Technical report, University of Pennsylvania.
20. Parunak, V.; Ward, A.; and Sauter, J. 1998. A systematic market approach to distributed constraint problems. Poster session, *Proc. 3rd International Conference on Multi-Agent Systems (ICMAS-98)*, 455–456.
21. Rothkopf, M. H., and Harstad, R. M. 1994. Modeling competitive bidding: A critical essay. *Management Science* 40(3):364–384.
22. Russell, S., and Wefald, E. 1991. Principles of metareasoning. *Artificial Intelligence* 49:361–395.
23. Sandholm, T. W., and Lesser, V. R. 1996. Advantages of a leveled commitment contracting protocol. In *Proc. 14th National Conference on Artificial Intelligence (AAAI-96)*, 126–133.
24. Sandholm, T. W., and Lesser, V. R. 1997. Coalitions among computationally bounded agents. *Artificial Intelligence* 94(1–2):99–137.
25. Sandholm, T. 1993. An implementation of the Contract Net Protocol based on marginal-cost calculations. In *Proc. 11th National Conference on Artificial Intelligence (AAAI-93)*, 256–262.
26. Shoham, Y., and Tennenholtz, M. 1999. What can a market compute, and at what expense? Technical report, Technion.
27. Varian, H., and MacKie-Mason, J. K. 1995. Generalized Vickrey auctions. Technical report, University of Michigan.
28. Varian, H. R. 1995. Economic mechanism design for computerized agents. In *Proc. USENIX Workshop on Electronic Commerce*. Minor update, 2000.
29. Wellman, M. P. 1993. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research* 1:1–23.