

Accounting for Cognitive Costs in On-line Auction Design^{*}

David C. Parkes¹, Lyle H. Ungar², and Dean P. Foster³

¹ Computer and Information Science Department, University of Pennsylvania, Philadelphia PA 19104 dparkes@unagi.cis.upenn.edu

² Computer and Information Science Department, University of Pennsylvania, Philadelphia PA 19104 ungar@central.cis.upenn.edu

³ Department of Statistics, University of Pennsylvania, Philadelphia PA 19104 foster@hellspark.wharton.upenn.edu

Abstract. Many auction mechanisms, including first and second price ascending and sealed bid auctions, have been proposed and analyzed in the economics literature. We compare the usefulness of different mechanisms for on-line auctions, focusing on the cognitive costs placed on users (e.g. the cost of determining the value of a good), the possibilities for agent mediation, and the trust properties of the auction. Different auction formats prove to be attractive for agent mediated on-line auctions than for traditional off-line auctions. For example, second price sealed bid auctions are attractive in traditional auctions because they avoid the communication cost of multiple bids in first price ascending auctions, and the “gaming” required to estimate the second highest bid in first price sealed bid auctions. However, when bidding agents are cheap, communication costs cease to be important, and a progressive auction mechanism is preferred over a closed bid auction mechanism, since users with semi-autonomous agents can avoid the cognitive cost of placing an accurate value on a good. As another example, when an on-line auction is being conducted by an untrusted auctioneer (e.g. the auctioneer is selling its own items), rational participants will build bidding agents that transform second price auctions into first price auctions.

Keywords: electronic commerce, auctions, semi-autonomous agents.

^{*} In Noriega P., and Sierra C., eds., *Agent Mediated Electronic Commerce* (LNAI 1571), pp.25–40, Springer-Verlag, 1999.

1 Introduction

Electronic commerce has made possible many market institutions that are too inefficient and expensive in traditional marketplaces [4]. Auctions have typically been reserved for special domains, such as the FCC spectrum auction [7], fine art auction houses (e.g. Sotheby's and Christie's), or selling perishable goods, such as in Spanish fish markets [10]. Auctions are preferred to posted prices in these domains because the goods are of uncertain value, and dynamic price adjustment will often maximize revenue for a seller. In traditional off-line auctions the interested parties must gather in one physical location, or be prepared to participate in a protracted auction by correspondence (e.g. the FCC auction). The infrastructure of electronic commerce significantly reduces the costs of participation, and allows auctions to reach a large and physically distributed audience. Furthermore, artificial bidding agents can represent the preferences of interested parties, bids are electronic, and auctions can be automatically cleared in electronic clearing houses. Auctions provide a useful framework for agent mediated electronic commerce [1].

The design of mechanisms for on-line auctions should be informed by the costs placed on users,¹ the possibilities for effective agent² mediation, and also by the trust relations between users and auctioneers. The most important cost to a user in an agent mediated electronic auction is not the communication cost of participation (an agent can cheaply and autonomously monitor auctions and place bids on behalf of a user), but the cognitive cost of deciding on the value of a good. Traditional auction theory assumes that buyers *know* the value of a good (private values assumption), or that the value of a good will follow from new information (affiliated or common value auctions). In fact, the problem of determining the value of a good can be computationally complex, and a buyer's value might still be uncertain when a bid is required [14].

We identify auction mechanisms that allow a user to avoid the cognitive cost of determining an accurate value for a good whenever possible, while enabling an agent to place the same bids in auctions that would be optimal with an accurate valuation. We argue that progressive auctions, such as the English auction, are preferable to sealed bid auctions when there is a high cost associated with determining the exact value of goods but users can determine approximate values more cheaply. Furthermore, when progressive auctions are coupled with semi-autonomous agents that are able to request more accurate valuations from users when necessary, the high monitoring and communication costs of progressive auctions can be hidden from users. We survey commercial on-line auctions from the perspective of our analysis, and suggest a new design for automatic bidding agents.

¹ A user refers to an interested human participant, either a buyer or a seller.

² An agent refers to a software agent that can participate in an auction on behalf of a user.

2 Auction-based Negotiation for Electronic Commerce

Agent mediated electronic commerce enables cheap negotiation between buyers and sellers on the details of an individual transaction – product features, value-added services, financing and price. Agents that can dynamically negotiate transaction terms with vendors are a natural extension of first-generation shopping agents that compare the posted-price of goods across static on-line vendor sites [5]. While personalized negotiation has traditionally been too expensive, except for high value and highly configurable goods, such as new cars, agent mediated negotiation absorbs many of the costs and inconveniences of negotiation.

The design of auction servers and market-aware computational agents presents a key challenge for agent technologies, artificial intelligence, and real-time systems [14, 19, 20]. We can make some challenges easier through the design of market mechanisms that enable straightforward and provably optimal agent strategies, at least for simple domains [17].

There are a number of testbeds for agent mediated electronic trading: *AuctionBot* [21] is a general purpose Internet auction server with an Application Programmable Interface (API) that allows software agents to participate in user-configured auctions; *FM 97.6* [11], another multiagent test bed, simulates complex Dutch auctions; and *Kasbah* is a prototype agent-based bilateral exchange market for buying and selling goods [3]. *Kasbah* illustrates a key feature of useful agent mediated electronic trading systems. It reduces the cognitive cost to users of the system by requiring that users provide only a few high-level parameters in order to specify their preferences. The bidding agents within *Kasbah* are given autonomy to pursue strategies consistent with the stated preferences.

Auctions provide a well-defined and simple framework for negotiation between self-interested buyers and sellers in a market. Although auctions traditionally allow negotiation over price alone, auction mechanisms can be extended to include negotiation over product features, warranties, and service policies – still within a framework that assumes self-interested, possibly insincere and strategic agents. For example, in the recent FCC spectrum auction buyers had synergies over combinations of spectrum rights over a geographically consistent area. The FCC implemented a simultaneous ascending auction that allowed bidders to place bids for multiple goods, and dynamically adjust over the course of the auction the precise combination of goods that they tried to acquire [7]. In a combinatorial auction buyers can express synergies directly by submitting bids for combinations of goods (or for one good with a specific combination of features) [13]. Wurman *et al.* [21] are developing languages that allow users to represent their preferences over bundles of interdependent goods efficiently.

Bidding agents need to make *tradeoffs* between features and price within an auction framework. The space of different product specifications is often very large, particularly when goods are nonstandard, and the problem that one must address as an auction designer is how to elicit only *necessary information* from users on their preferences between different product specifications when many feature combinations might prove to be unavailable or too expensive. We as-

sume that this information is costly for a user to provide. We compare standard single-unit auctions over a large and diverse space of possible goods, from the perspective of reducing the cognitive costs involved in providing information to a bidding agent. We assume that a user can specify hard constraints over product features that she will consider, but does not know (or does not wish to specify) her exact value tradeoff within this space.

3 Traditional Auction Theory

Auctions are simple and robust mechanisms for selling nonstandard and short supply items with uncertain market values. Auction mechanisms discover the optimal price for a good through the bidding action of self-interested agents. We characterize price discovery mechanisms as progressive (ascending price, descending price) or sealed bid, and first price or second price. In ascending price auctions the auctioneer continuously reveals the highest bid received, and the *ask price* is a minimum increment above the price of the current highest bid. The auction terminates after a period of inactivity. In descending price auctions the auctioneer lowers the ask price until the first bid is received, when the auction is immediately terminated. In a sealed bid auction all bids are private, and the auctioneer selects the winning bid after a fixed period of time.

In all simple auction mechanisms the good is sold to the bidding agent with the highest bid. In first price auctions an agent pays the price of its bid, while in second price auctions an agent pays the highest amount that was bid by another agent [6, 9]. While there are many combinations of auction characteristics, some auction mechanisms are redundant and do not have general use. For example the second price ascending auction is strategically equivalent to the first price ascending auction for small bid increments, and will have the same revenue and efficiency properties.

Traditional auction theory emphasizes the allocation efficiency and expected revenue properties of auction mechanisms, given assumptions about how buyers assign value to goods [6]. For example, the prevalence of the first price ascending (English) auction in traditional auction environments (e.g. fine art auction houses) is best explained with a model that assumes that buyers' values are affiliated: when one buyer places a high value on a good it is more likely that other buyers will also value the good highly [8]. A progressive price auction maximizes revenue for sellers in this environment because bids reveal information, and the knowledge that one buyer values an item highly can increase the valuations of other buyers [6]. A sealed bid auction does not allow this dynamic information sharing between buyers.

Second price sealed bid (Vickrey) auctions are attractive in traditional auction domains because they avoid the communication cost of multiple bids in ascending auctions, and the "gaming" that is required to estimate the second highest bid in first price sealed bid auctions. Furthermore, privacy and information revelation considerations often favor a sealed bid auction over an (open) progressive auction (e.g. government procurement contracts). However, there are

few real-world examples of second price sealed bid auctions, possibly because the Vickrey auction is vulnerable to manipulation by an untrusted auctioneer that overstates the second highest bid (unless the bid must be verified in some way), and because buyers are often reluctant to reveal their true valuations for goods [12].

In our analysis of mechanisms for on-line auctions we make standard assumptions about how buyers assign value to goods: (a) private value assumption – the value of a good only depends on a buyer’s own preferences; (b) independence of bids – the values placed on a good by buyers are statistically independent. Given these assumptions the *revenue equivalence theorem* [18] states that the four most common auctions (English, first price descending, first price sealed bid and Vickrey) all yield the same expected revenue for risk-neutral agents. However good auction design can still increase expected revenue by lowering the costs of participation. Theoretical analysis of revenue equivalence assumes that a fixed set of buyers participate in the auction, irrespective of the mechanism. In reality the costs (cognitive and communication) and conditions of trust within a system will affect the number of buyers that participate. *A seller can expect to receive a higher sale price for a good if she is able to attract bids from more buyers.*

4 Auction Cost and Trust Properties

The design of an auction mechanism for on-line agent mediated auctions should be informed by the costs placed on human participants in the system, and the conditions of trust that exist between the buyers, the sellers and the auctioneer. Auctions for agent mediated on-line electronic commerce require somewhat different cost structures than more traditional off-line auctions, and conditions of trust are often more critical. In this section we introduce some key auction *properties* that influence both the cost of participation in an auction, and the trust requirements that are placed on parties. We claim that the most important barrier to user participation in an agent mediated on-line auction is the cognitive cost of deciding values for new, previously unvalued, goods.

Standard economic theory assumes that a buyer knows the value of a good when the good is for self-consumption (private value assumption), or that the value of a good is set exogenously and cannot be known by an agent without further information (common value assumption). However, in many real situations a good can only be valued by careful reasoning. This is a cognitive cost that cannot be easily transferred to an electronic bidding agent, indeed buyers often cannot even formulate the tradeoffs that determine their exact reservation price for a particular good. That there is a nontrivial cost to decide what the reservation price is: “how much am I willing to pay for this item?” may not be obvious, but note that people often take a lot of time deciding whether they want to buy a given dress, bike or house, even when the price is posted. In business settings determining the value of a given item, such as a quantity of

crude oil, or a shipment of electronic components, may require solving a complex optimization problem.

We summarize the cost and trust properties over six different auction mechanisms in Table 1 and Table 2.³ We will introduce two new auction properties: *bounded-rational compatible* and *untrusted-auctioneer compatible*. The properties listed are all beneficial, they either reduce the cost of participation, or reduce the degree of trust that is required between agents within the auction for faithful implementation.

4.1 Cost Properties Relevant to On-Line Auction Design

The costs to participation in an auction are *cognitive* (information gathering and processing) and *communicative*, and depend on the mechanism, the goods that are being auctioned, and the complexity of the local problems of buyers. The costs that are important in an on-line agent mediated auction differ from the costs in a more traditional off-line auction.

Cognitive costs can be separated into two processes. First a buyer must estimate her *reservation price* for a good, the maximum price that she will pay for a good. The reservation price is equal to the perceived value of a good to the buyer, and depends on her preferences. Although a buyer can place a bid in an auction on the basis of her reservation price alone, in some auctions counterspeculation is also useful – and a buyer might choose to reason about the preferences and strategies of other buyers in the system when placing a bid. This is the second cognitive process. We can design auction mechanisms and agents that reduce and remove this second cost, either through mechanisms that make counterspeculation redundant, or through bidding agents that implement complex bidding strategies for a user. The premise of our work is that the cost of the first cognitive process, that of determining a reservation price for a new, previously unvalued, good is nontrivial, and will often dominate any other costs to participation in an agent mediated on-line auction.

In an auction that is *buyer incentive compatible* (IC_B , Table 1) the dominant strategy for a buyer is to bid her true value for a good [9]. It is not necessary for a buyer to incur the cognitive cost of strategic counterspeculation that can be useful in some auction mechanisms. In buyer incentive compatible ascending price auctions, such as the first price ascending auction, a buyer will maximize her expected utility by placing a new bid whenever the current highest bid is below her reservation price and held by another buyer.

³ The second price descending auction is not a standard auction format. The auctioneer starts with a high price, and lowers the price until two bids have been received. The agents are not informed of the first bid until the second bid has also been made. The item is then sold to the highest bidder at the price of the second highest bid. The auction is strategically equivalent to the English and Vickrey auctions. Furthermore, the auction is “bounded-rational compatible” because a good upper bound on a buyer’s reservation price determines an optimal strategy when the auction terminates before the upper bound is reached.

Table 1. Cost properties of common auctions. **Key:** IC_B Buyer incentive compatible, BRC Bounded-rational compatible, $COMM_L$ Low communication cost.

	Ascending Price	Descending Price	Sealed bid
First price	IC_B, BRC	BRC	$COMM_L$
Second price	IC_B, BRC	IC_B, BRC	$IC_B, COMM_L$

A buyer in an incentive compatible sealed bid auction should simply bid her reservation price. For example, the second price sealed bid (Vickrey) auction is incentive compatible because the price that a buyer pays for a good is determined by the bids received from other buyers, and does not depend on her own bid. Instead, the price that a buyer bids defines a range of prices that she is willing to accept (any price up to and including the price of her bid). It is a dominant strategy to submit a bid equal to her reservation price [9]. When all buyers adopts this dominant strategy the good is sold to the buyer with the highest reservation price, for a price equal to the second highest reservation price.⁴ In comparison, the optimal strategy for a buyer in a *first price* sealed bid auction is to place a bid that is just above the highest bid of another buyer, so long as the price is below her own reservation price. In this case a buyer can use information about the valuations and strategies of other buyers in the system to increase her expected utility.

Different auction structures require vastly different expected effort by buyers to determine a value for a good that is being auctioned. We introduce a new auction property, *bounded-rational compatible* (BRC , Table 1). A bounded-rational compatible auction, such as the English auction, will often allow a bidding agent to follow an optimal bidding strategy with only approximate information on the value of a good. We assume that it takes less effort for a buyer to place bounds on her value for a good than it does to compute her exact value for a good. For example, consider the behavior of a bidding agent in an English auction that has a lower and upper bound on the value that a buyer places on a good. The bidding agent will place a bid whenever the ask price is below the lower bound on value, and drop out of the auction when the ask price is above the upper bound on value. The buyer avoids the cognitive cost of refining its valuation when the auction terminates with an ask price that is either below its lower bound or above its upper bound. The buyer only needs to compute an accurate valuation when other buyers have similar reservation prices. Compare this with a sealed bid auction, where an agent that has an estimate of the value of a good risks

⁴ This is the same outcome as is achieved with an English auction, but without the overhead of multiple bids.

winning the good for a price that is above the buyer's true value, or missing a price that is below the buyer's true value.

Finally, we also consider whether an auction mechanism has low communication costs ($COMM_L$, Table 1). The communication cost of an auction mechanism depends on the size of a single bid, and the expected number of auction rounds. Sealed bid auctions have low communication costs in comparison with progressive auctions because they require only one bid from each agent in each auction while progressive auctions may require several rounds of bidding.

The importance of communication cost depends on the infrastructure of the auction. For example, the amortized per-message communication cost in a traditional off-line auction is much greater than the per-message cost in an electronic on-line auction. In a traditional auction the cost includes the need to physically meet in one location for a period of time, while the per-message communication cost in an Internet-based electronic auction is minimal [1]. Furthermore, users in an on-line auction can be physically distributed and use bidding agents to automatically monitor auctions and place bids.

4.2 Trust Properties Relevant to On-Line Auction Design

Trust is another important consideration that drives auction design – will users have any reason not to trust that an auctioneer will implement a bidding mechanism truthfully? It is not unusual in on-line commerce for the seller and the auctioneer to be the same firm (e.g. www.onsale.com). There is a clear conflict of interest. The auctioneer has an incentive to inflate the price of a good. We can design auction mechanisms that are not vulnerable to this kind of direct price manipulation.⁵

We introduce a new auction property, *untrusted-auctioneer compatible* (UAC , Table 2), for auctions that are not vulnerable to *direct* price manipulation by a strategic auctioneer. Direct price manipulation is strategic action by the auctioneer, or parties acting on behalf of the seller, that is risk-free (while undetected) and can only increase revenue. An auction is secure to this type of price manipulation when the buyer that wins the auction pays an amount equal to the price that she bid. For example, the first price sealed bid auction is untrusted-auctioneer compatible because the auctioneer cannot inflate the price that the winning buyer pays. However, an untrusted auctioneer in a second price sealed bid auction can receive all the bids, create a false bid just below the highest outside bid received, and then charge the buyer with the highest bid a price just below her bid [12]. This is possible because the price paid by the highest bidder is different from the price that is bid, and revealed only to the auctioneer.

A rational participant in a sealed bid auction with an untrusted auctioneer will build a bidding agent that transforms second price auctions into first price auctions because buyer incentive compatibility is lost. This will occur when there

⁵ An alternative (institutional) solution to untrusted auctioneers is to provide trusted third parties that are certified to manage auctions.

Table 2. Trust properties of common auctions. **Key:** IC_S Seller incentive compatible, UAC untrusted-auctioneer compatible.

	Ascending Price	Descending Price	Sealed bid
First price	UAC	IC_S, UAC	IC_S, UAC
Second price	UAC		

is reason to believe that the auctioneer might profit from the final sale price of the good.

An auction that is also *seller incentive compatible* (IC_S , Table 2) is secure from *indirect* price manipulation. Indirect price manipulation is strategic action by the auctioneer, or parties acting on behalf of the seller, that includes some risk of decreasing revenue, for example when information about the preferences and strategies of buyers is not fully accurate. In a seller incentive compatible auction it is optimal for a seller to bid her true valuation for a good, because her ask price does not affect the price that she receives, but defines the range of prices that she is willing to receive. Seller incentive compatibility is necessary and sufficient for preventing the type of indirect price manipulation that we consider here.

First price descending and sealed bid auctions are seller incentive compatible because the ask price does not directly determine the price that a seller receives. The optimal strategy for a seller is to set the ask price equal to her reservation price. Second price auctions are not seller incentive compatible because the seller can influence the price that she receives for a good, for example by setting the ask price just below the highest bid price. Similarly, the first price ascending auction is not seller incentive compatible because a seller can increase her expected revenue by setting her ask price just below the highest reservation price of the buyers in the system.

When an auction mechanism is not seller incentive compatible the price that the seller receives can be manipulated through a third party that places bids on behalf of the seller, called a *shill*. For example, a shill in a first price ascending auction can compete with the buyer with the highest reservation price, and try to drive the ask price to just below that buyer's reservation price. This type of price manipulation is not risk free without perfect information about the preferences and strategies of the buyers in an auction. The seller might be left with a good that she could have sold for a profit.⁶

⁶ In traditional English auctions all participants are present in the same room, and a shill can reduce his risk with visual clues about when a bidder has been pushed close to her reservation price.

We can also compare the auction mechanisms with regard to manipulation through buyer collusion. A group of buyers can form a coalition, and elect one member to participate in the auction, bidding up to the highest reservation price of the members of the coalition. The buyers can then re-auction the good among themselves, and share the surplus that is extracted from the seller. Possible in any auction mechanism, this type of static collusion that occurs prior to an auction is hard to detect and prevent. Collusion may also occur dynamically *during* a progressive auction, if the active bidders can reach a self enforcing agreement to stop bidding and re-auction the item amongst themselves [15]. One key property that is necessary for dynamic collusion is the identification of bidders. A progressive on-line auction can offer some protection from buyer collusion by posting all bids anonymously.

5 Agent Mediated Progressive Auctions

In this section we use the earlier analysis of the cost and trust properties of auctions (Section 4) to identify an appropriate auction mechanism for agent mediated electronic commerce, paying particular attention to the need to delay and avoid the cost to a buyer of valuing a good.

Aside from cognitive costs, we can simplify the design of optimal bidding strategies for software agents through good auction design. With well-designed auctions the set of optimal strategies that might be useful to an agent is small, and we can program agents with a complete set of strategies – there is no need for agents to learn better strategies from success or failure [2]. In particular, agents in *buyer incentive compatible* auctions can maximize expected utility by bidding according to their own reservation price, and without consideration of the preferences or strategies of other agents in the system. For example, the optimal strategy of an agent in a first price ascending auction is to bid whenever the current highest bid is below its reservation price, and held by another agent.

Furthermore, because the space of interesting goods that a bidding agent might encounter is potentially very large, we would like *bounded-rational compatible* auctions that allow optimal bidding strategies with only approximate reservation prices. This allows buyers to defray, and if possible avoid, the cognitive cost of placing an accurate valuation on a good.

The auction mechanisms that we have considered that are both buyer incentive compatible and bounded-rational compatible are the first price ascending auction and the second price descending auction (see Table 1). We rule out the second price sealed bid (Vickrey) auction because it is not bounded-rational compatible and we want to avoid the high cognitive cost of computing the reservation price for every auction.⁷

We prefer the first price ascending auction to the second price descending auction for a number of reasons. Firstly, an agent in an ascending price auction can drop out of an auction before it terminates, when the ask price is greater

⁷ Sandholm [16] has noted some other limitations of Vickrey auctions in computational multiagent systems.

than its reservation price upper bound. In a descending price auction an agent must monitor the auction as long as it is open, and accept an ask price that is below its reservation price lower bound. Secondly, in an ascending price auction bidding agents control the rate at which the price of a good is increased as they refine their valuations for goods. This is preferable to a descending price auction where the auctioneer must lower the price of a good at a rate that is appropriate to the bounded rationality of the bidding agents, in order to avoid missing a price that an agent would accept given enough time to reason about the true value of its reservation price.

Furthermore, when we consider the trust properties of auctions, we see that the descending second price auction is susceptible to manipulation by an untrusted auctioneer in the same way as the Vickrey auction, while the first price ascending auction is untrusted-auctioneer compatible because the winning agent pays the price that it bid (*UAC*, Table 2). Both progressive auctions are however susceptible to indirect price manipulation (*IC_S*, Table 2). It is interesting to note that it is impossible for an auction mechanism to be both buyer and seller incentive compatible (compare Table 1 and Table 2) [22], and that we must accept indirect price manipulation in return for simple optimal bidding strategies.

We conclude that the most suitable auction for agent mediated on-line auctions, when the space of possible goods is large and diverse, and users do not know their reservation prices for all goods, is the first price ascending auction. In on-line auctions it is worth having agents submit many bids if this enables some users to avoid computing accurate valuations in order to participate. Prices in progressive auctions provide feedback that allow users to avoid and delay the cognitive cost of determining an accurate value for a good.

5.1 Bidding Agent Design

When progressive auctions are coupled with semi-autonomous bidding agents users can achieve significant cost savings. While a fully autonomous agent requires a complete set of preferences in order to represent a user correctly in all situations that it might encounter, a semi-autonomous agent will bid on behalf of the user when it has enough knowledge to proceed, and query the user when its best action is ill-defined given its current information.

The design of bidding agents should be appropriate to the auction mechanism, and sensitive to the cost characteristics of users. An autonomous agent that places bids up to the value of a fixed reservation price, a “reservation-price agent”, is appropriate for users that have precise valuations for goods that are auctioned. A reservation-price agent competes in relevant auctions while the current ask price is above its reservation price or until the agent is sold the good.

When users have approximate valuations for goods, and would prefer to avoid the cost of providing accurate valuations, then an agent that is semi-autonomous and requests a refined valuation as necessary, a “progressive-price agent”, is more appropriate. A progressive-price agent is initialized with a *lower bound* and an optional *upper bound* on the true reservation prices of the user. The agent always

places a bid when the current ask price is below its lower bound and the agent does not hold the highest bid, and will leave an auction when the ask price is above its upper bound. Given this strategy, the agent is sold a good whenever the auction closes at a price below its current lower bound. The agent behaves autonomously while current ask prices are below the lower bound or above the upper bound. The agent will request more information from the buyer, in the form of refined lower and upper bounds on its valuation, whenever the current ask price lies within the range of uncertain value between the two bounds. The agent will leave the auction if the new upper bound is below the current ask price, place a bid if the new lower bound is above the current ask price, and request further refinement if the ask price remains between the bounds.

For example, consider a user that is interested in purchasing an automobile. The user initializes a progressive-price agent that will make bids on any convertible VW Beetle that was made after 1975, has done less than 40000 miles, and is in mint condition. The user provides a lower bound of \$3000, but no upper bound. The bidding agent will behave autonomously while no such vehicles are for sale (`search`), or while the current ask price in an auction is less than \$3000 (`bid`). The agent will only request further information from the user when the price of a suitable car increases above \$3000. With updated reservation price bounds the agent will continue to bid, or drop out the auction – as appropriate. This is the sense in which the bidding agent is semi-autonomous. The user provides the bidding agent with some initial information, and is ready to provide new information when necessary.

As another example, consider a user who wants to buy a notebook computer with features that satisfy a set of hard constraints, but does not know her exact price tradeoff, for example between a Pentium-166 and a Pentium-233 processor. Instead she initializes a progressive-price agent with upper and lower valuation bounds over the feature space of satisficing goods. When her agent finds an auction for a notebook with suitable features it will bid while the price is below the lower bound for that particular set of features. Should the auction end with the agent winning the notebook for a price below the reservation price, then the agent can complete the transaction. However, if the price in the auction rises above the upper bound, then the agent can stop monitoring the auction, and move to another auction. The agent only needs to request a more accurate value for the notebook (with this particular set of specifications) if the price stops within the region of uncertain value.

6 Commercial On-Line Auctions

There are many on-line auctions that are currently in use [23]. Some sites are run by discounters selling multiple goods, often from odd lots (Major Vendor auctions), while others are open to the public and anyone can sell (Person-to-Person auctions), see Table 3. It is illuminating to look at the design of real on-line auctions in the light of the above discussion.

Table 3. Commercial on-line auctions. **Key:** **M** – Major Vendor Auction, **P** – Person-to-Person Auction

Company	M / P	Mechanism	Bidding Agent	Automatic notification
auctionworks.com	P	English	“proxy bidding system”	yes
auctionuniverse.com	P	English	“RoboBid”	yes
ebay.com	P	English	no	yes
haggle.com	P	English	“proxy robot”	no
onsale.com	M	English	“BidAgent”	yes (“BidWatch”)
zauction.com	M	English	no	no
surplusauction.com	M	English	no	no

Most commercial on-line auction mechanisms are a variant of the English auction. The auctions often include a deadline for the auction (whether or not there is any bidding activity), and the final stage is sometimes a sealed bid auction to avoid strategic bids that seek to take advantage of Internet communication latencies. In some auctions the deadline is extended with a “going going gone” phase that continues while there is bidding activity.

Rudimentary bidding agents are provided at many Person-to-Person auction sites.⁸ For example, `onsale.com` provides a “BidAgent” that is initialized with a reservation price, and will monitor an auction and place a bid when the current highest bid is below the agent’s reservation price and held by another agent. The agent will also provide automatic e-mail notification when the ask price rises above the reservation price of the agent, the price that the agent is authorized to bid up to. This enables a user to set a new maximum bid price if appropriate. Most on-line auctions encourage a user to initialize a bidding agent with an accurate reservation price, although they do (implicitly) allow a user to refine a lower bound on her reservation price through this “e-mail and update” mechanism.

The bidding agents actually implement a second price auction, but do so in a way that allows the current second highest bid to be publicly posted so

⁸ A user of an untrusted auction might prefer to implement and execute her own bidding agents, given that the bidding agents implemented and executing at the auction site will hold information on her reservation price – information that could be used to the advantage of an untrusted auctioneer [20].

that other potential buyers can judge whether it is worth their time to enter the bidding. The bidding agents also reveal the incentive compatibility of a second price sealed bid auction to the users [20].

In another complication of the simple auction, sellers are often permitted to start the auction at a price below their reservation price in order to generate interest. An auction is sometimes marked as a “reservation price auction” in this case. The seller does not have to sell the good if the highest bid falls below her reservation price. A better way to generate interest, that avoids the spectacle of auctions closing without the highest bidder receiving the good, is to provide a “seller agent” that will bid for the *seller* if an outside bid is received that is lower than her reservation price.

7 Conclusions

Real auctions differ from those typically analyzed by economists and computer scientists. Issues of cost and trust drive these differences. On-line electronic commerce provides a cheap and readily available communication infrastructure (the Internet), that makes differences in the communication efficiency of sealed bid and progressive auctions unimportant. The largest overhead to user participation in an on-line auction is deciding what to buy, and how much it is worth. This is a much greater problem on-line than off-line because there are many different goods available across many different auctions. In traditional off-line auctions goods are typically sold in special lots to a group of invited experts.

We have argued that first price ascending (English) auctions have many advantages over other mechanisms for on-line auctions. In particular English auctions enable semi-autonomous agents that can act on behalf of users, and follow optimal bidding strategies with only approximate valuations for goods that they encounter. A request is made to a user for more accurate valuations only when necessary – when it will affect the outcome of an auction. This frees users of the need to provide accurate valuations for every single good that the agent might encounter (or equivalently for every possible combination of features for a particular type of good).

Semi-autonomous agents may not always be desirable. For example, if it is important to a user to choose *when* to provide information to her agent, then she might prefer to decide on an accurate reservation price when she has time to do so, and request that her agent behaves completely autonomously until a later predetermined time.

Although theoretical analysis of the revenue equivalence of private-value auctions assumes risk-neutral agents, and human buyers tend to be risk-averse, we believe that our conclusions are still valid. The first price sealed bid and first price descending bid auctions have a greater expected revenue than English auctions for risk-averse agents [6]. However, we consider an *open system* with voluntary participation, where the number of buyers that choose to participate in an auction is determined by the costs (cognitive and communicative) and conditions of trust within a system. A seller can expect to receive a higher sale

price for a good if she is able to attract bids from more buyers. We believe that the positive effect on increased participation from an English auction will more than outweigh any negative effect on expected revenue for the set of buyers that choose to participate.

Finally, we emphasize that when designing a mechanism for on-line commerce it is critical not only to consider the allocation efficiency and revenue maximization properties of a mechanism, but also the costs that are relevant to participants in the auction.

8 Acknowledgments

This research was funded in part by National Science Foundation Grant SBR 96-02053.

References

1. Beam, C., Segev, A., Shanthikumar, J. G.: Electronic negotiation through Internet-based auctions. CITM Working Paper 96-WP-1019 (1996) <<http://haas.berkeley.edu/~citm/wp-1019-summary.html>>
2. Beam, C., Segev, A.: Automated negotiations: A survey of the state of the art. CITM Working Paper 96-WP-1022 (1997) <<http://haas.berkeley.edu/~citm/wp-1022-summary.html>>
3. Chavez, A., Maes, P.: Kasbah: An agent marketplace for buying and selling goods. Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96) (1996)
4. Guttman, R., Moukas, A., Maes, P.: Agent-mediated electronic commerce: A survey. Knowledge Engineering Review **13**(3) (1998) <<http://ecommerce.media.mit.edu/papers/ker98.ps>>
5. Guttman, R., Maes, P.: Agent-mediated integrative negotiation for retail electronic commerce. Workshop on Agent Mediated Electronic Trading (AMET-98) (1998)
6. McAfee, R. P., McMillan, J.: Auctions and bidding. Journal of Economic Literature **25** (1987) 699–738
7. McAfee, R. P., McMillan, J.: Analyzing the airwaves auction. Journal of Economic Perspectives **10**(1) (1996) 159–175
8. Milgrom, P., Weber, R. J.: A theory of auctions and competitive bidding. Econometrica **50**(5) (1982) 1089–1122
9. Milgrom, P.: Auctions and bidding: A primer. Journal of Economic Perspectives **3**(3) (1989) 3–32
10. Rodríguez-Aguilar, J. A., Noriega, P., Sierra, C., Padget, J.: FM96.5 A Java-based electronic auction house. Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97) (1997)
11. Rodríguez-Aguilar, J. A., Martín, F. J., Noriega, P., Garcia, P., Sierra, C.: Competitive scenarios for heterogeneous trading agents. Proceedings of the Second International Conference on Autonomous Agents (Agents'98) (1998) 293–300
12. Rothkopf, M. H., Teisberg, T. J., Kahn, E. P.: Why are Vickrey auctions rare? Journal of Political Economy **98**(1) (1990) 94–109

13. Rothkopf, M. H., Pekeć, A., Harstad, R.M.: Computationally manageable combinatorial auctions. Rutgers University, RUTCOR research report **13-95** (1995)
14. Sandholm, T. W., Lesser, V.: Issues in automated negotiation and electronic commerce: Extending the Contract Net framework. First International Conference on Multiagent Systems (ICMAS-95) (1995) 328–335 <<ftp://ftp.cs.umass.edu/pub/lesser/sandholm-icmas95-issues.ps>>
15. Sandholm, T. W., Lesser, V.: Coalition formation among bounded rational agents. Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95) (1995) 662–669
16. Sandholm, T. W.: Limitations of the Vickrey auction in computational multiagent systems. Second International Conference on Multiagent Systems (ICMAS-96) (1996) 299–306 <<http://www.cs.wustl.edu/~sandholm/vickrey.ICMAS96.ps>>
17. Varian, H.: Economic mechanism design for computerized agents. First USENIX Workshop on Electronic Commerce (1995) 13–21 <<ftp://alfred.sims.berkeley.edu/pub/Papers/mechanism-design.ps.Z>>
18. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16** (1961) 8–37
19. Wellman, M. P., Wurman, P. R.: Market-aware agents for a multiagent world. Forthcoming, *Robotics and Autonomous Systems* (1998) <<ftp://ftp.eecs.umich.edu/people/wellman/maamaw.ps>>
20. Wellman, M. P., Wurman, P. R.: Real time issues for Internet auctions. First IEEE Workshop on Dependable and Real-Time E-Commerce Systems (DARE-98) (1998) <<ftp://ftp.eecs.umich.edu/people/wellman/dare98.ps>>
21. Wurman, P. R., Wellman, M. P., Walsh, W. E.: The Michigan Internet AuctionBot: A configurable auction server for human and software agents. *Proceedings of the Second International Conference on Autonomous Agents (Agents'98)* (1998) 301–308
22. Wurman, P. R., Walsh, W. E., Wellman, M. P.: Flexible double auctions for electronic commerce: Theory and implementation. Forthcoming, *Decision Support Systems* (1998) <<ftp://ftp.eecs.umich.edu/people/wellman/dss98.ps>>
23. Yahoo! Online auction <http://www.yahoo.com/Business_and_Economy/Companies/Auctions/Online_Auctions>