Cost and Trust Issues in On-Line Auctions

Lyle H. Ungar[1], David C. Parkes[2], and Dean P. Foster[3]

[1] Computer and Information Science Department, University of Pennsylvania, Philadelphia PA 19104 ungar@central.cis.upenn.edu
[2] Computer and Information Science Department, University of Pennsylvania, Philadelphia PA 19104 dparkes@unagi.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 different costs of determining reservation prices, determining bids and communicating bids, and on whether the auctioneer is trusted. Different auction formats prove to be attractive for agent-mediated on-line auctions than for off-line ones. For example, second-price sealed-bid ("Vickrey") auctions are attractive in encouraging truth revealing behavior, and avoiding the communication costs of the multiple bids of a first-price ascending ("English") auction 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 the English auction is preferred, since it often allows bidders to save the effort of estimating their value for a good. As another example, we show that when an on-line auction is being conducted by a non-trusted auctioneer (e.g. the auctioneer is selling their own items), rational participants will build bidding agents which transform second-price auctions into first-price auctions.

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

1 Introduction

Costs and trust issues drive the design of mechanisms and bidding agents for on-line auctions. Researchers have made the case for market-aware computational agents (Wellman 1997), discussed some of the issues in electronic commerce and automated negotiation (Sandholm and Lesser 1995), and presented applications of economic mechanism design to electronic commerce (Varian 1995). Guttman *et al.* (1998) provide a useful survey of agent-mediated electronic commerce. Auctions provide a simple framework for automated negotiation over price in electronic commerce. This paper addresses the question of why real on-line auctions are the way the are, and how they might be improved.

The most important cost to a human buyer in an agent-mediated electronic auction is not the communication cost of participating in the auction (an agent can cheaply and autonomously monitor auctions and place bids on behalf of the buyer), but the cognitive cost of deciding on the value of an item. The economic literature considers the cognitive cost to a buyer of strategic bidding in an auction, given a reservation price, but the cost of deciding what that reservation price is in the first place is mainly ignored. Just as auctions can be designed to make strategic bidding unnecessary, we identify a new characteristic, termed "reservation-price-not-necessary", that allows a bidding-agent that only has bounds on the value of a good to participate in an auction, often at no disadvantage. We argue that first-price ascending auctions have this characteristic, and also have better trust properties than other theoretically desirable auction types. The analysis of the cost and trust properties of on-line auctions also suggests a new design for bidding agents.

1.1 Design Space for Auctions

Auctions are simple and robust mechanisms for selling items that maybe non-standardized, in short supply, or of uncertain value. Auctions allow a seller to use the marketplace to discover the price of an item. Auctions can be characterized as lying in a design space according to whether they are progressive (ascending-price, descending-price) or sealed-bid, and whether they are first or second-price. In ascending-price auctions the auctioneer posts the highest price bid at any time, and the auction terminates after a period of inactivity. In descending-price auctions the auctioneer lowers the price until a bid is received, whereupon the auction is terminated. In a sealed-bid auction the bid information is concealed, and the auction is closed after a fixed period of time. In all auction formats the good is sold to the highest bidder, but in first-price auctions the bidder pays the amount she bid, while in second-price auctions the bidder pays the highest amount bid by another agent (McAfee and McMillan 1987; Milgrom 1989).

Most combinations of auction types are possible, but many are generally not used; e.g. a second-price ascending auction has the same strategic properties as a first-price ascending auction for small bid increments. Traditional economic theory emphasizes the allocation efficiency and expected revenue properties of different auction types, given different assumptions about how a bidder's value

Table 1. Cost and Trust Properties of Common Auctions. **Key:** IC_B and IC_S Buyer and Seller incentive-compatible, nRP reservation-price-not-necessary, nT trusted-auctioneer-not-necessary, $COMM_L$ low communication cost.

	Progressive		**Sealed-bid**
	Ascending	**Descending**	
First-price	IC_B, nT, nRP	IC_S, nT, nRP	$IC_S, COMM_L, nT$
Second-price	IC_B, nT, nRP	IC_B, nRP	$IC_B, COMM_L$

of a good is formed (McAfee and McMillan 1987). We make the standard assumptions that the value of a good only depends on the bidder's own preferences (*private values assumption*), and that the bids made by competitors are statistically independent (*independence assumption*). The *revenue equivalence theorem* (Vickrey 1961) states that the four most common auction types (English, Dutch, first-price sealed-bid and Vickrey) all yield the same expected revenue for the seller in independent private value auctions. The appropriate auction type then depends on the costs (cognitive and communication) and conditions of trust (or lack of trust) of the parties in the particular situation. Auction properties that reduce **costs** for bidders include:

Buyer-Incentive-Compatible, IC_B. An auction is buyer-incentive-compatible if a buyer can do no better than to truthfully reveal her value for a good. In traditional auction theory this is a desirable property because an auctioneer can determine an optimal allocation of resources if she can extract truthful information on the private value of goods from participants. Computationally, incentive-compatibility is desirable because a bidder needs only to consider her own value for a good, not the values and strategies of the other bidders (Milgrom 1989). In ascending-price auctions, buyer incentive-compatibility implies that the buyer should place a new, highest bid, whenever the existing high bid is below her *reservation price*[1] and held by another buyer. In sealed-bid auctions buyer incentive-compatibility implies that the buyer should bid her reservation price.

Reservation-Price-Not-Necessary, nRP. An auction is reservation-price-not-necessary if a bidder can participate with an estimate of the value of a good, often at no disadvantage. This property holds for progressive auctions, but not sealed-bid auctions when a bidder with a poor estimate of the value for a good must always risk buying the good above value, or missing a price below value. While traditional auction theory assumes that the bidder *knows* the value of a good (private values assumption), we consider the cognitive cost to the bidder of determining that value. The problem of determining the expected value of a good can be computationally complex, and a bidder's valuation might still be uncertain when a bid is required (Sandholm 1995).

Low Communication Cost, $COMM_L$. The communication cost reflects the total number and size of messages sent during an auction, and the associated

[1] The reservation price is a bidder's best estimate of her value for a good.

transmission-costs. In a traditional off-line auction this includes the cost of getting all the bidders together in the same place on the same day.

The auction properties that improve **trust conditions** for bidders in an auction include:

Trusted-Auctioneer-Not-Necessary, nT**.** An auction is trusted-auctioneer-not-necessary when the price that the highest bidder must pay is the price that she bid, or is revealed to all participants through the auction mechanism.

Seller-Incentive-Compatible, IC_S**.** An auction is seller-incentive-compatible if the seller can do no better than by truthfully revealing her value for an auctioned good. Seller incentive-compatibility implies that a seller cannot benefit from manipulating the price of any item she is selling, either by the use of a shill (a third party) or by placing a bid herself.

We present these properties over a space of auction types in Table 1. The incentive-compatibility properties are standard (see McAfee and McMillan (1987), for example). The descending second-price auction[2] and the sealed-bid second-price (Vickrey) auction require a trusted auctioneer because the price paid by the highest bidder is different from the price bid, and revealed only to the auctioneer. The progressive auctions are reservation-price-not-necessary because it is often sufficient to have an estimate of the reservation price for a good. Further refinements can be made during the auction as necessary. For example, when the current price for a good in an English auction is well above the buyer's true valuation and the buyer can place an upper-bound on her valuation that is below that price, then the bidding-agent does not need to know the buyer's true value. Finally, the sealed-bid auctions have low communication costs in comparison with the progressive auctions because they require only one bid from each agent in each auction.

2 Off-line Auctions

Although second-price sealed-bid auctions achieve the same revenue and resource allocation as English auctions (Vickrey 1961) with lower communication costs, the most widely used non-electronic auction is the ascending first-price (English) auction (e.g. fine art auction houses). Both auction mechanisms are incentive-compatible, the Vickrey auction because the price that a winning bidder pays is determined by its competitors' bids alone and does not depend on its actual bid.

[2] The descending second-price auction is not a standard auction format. The auctioneer starts with a high price, and lowers the price until two accepts have been received. The bidders are not informed of the first "accept-bid" until the second accept-bid has also been made. The item is then sold to the highest bidder at the second highest price bid. The auction is strategically equivalent to the English auction and the Vickrey auction. A good upper bound on a bidder's reservation price is sufficient when the auction terminates before the upper bound is reached – so the auction is reservation-price-not-necessary.

A bidder will accept any price up to its bid, and none above. It is therefore a dominant strategy for a bidder to submit a bid equal to its true true reservation price (Milgrom 1989). When each bidder adopts this dominant strategy the good is sold to the bidder with the highest valuation, for a price equal to the second highest valuation – the same result as for English auctions. This outcome is achieved without the overhead of multiple bids.

The prevalence of the English auction in traditional off-line auctions is best explained with a model that assumes that bidders' valuations are *affiliated*: when one bidder perceives a high value for a good it is more likely that other bidders will perceive the value to be high (Milgrom and Weber 1982). Under this model the English auction maximizes revenue for sellers because bids reveal information (McAfee and McMillan 1987). The knowledge that one bidder values the item highly may increase the valuations of the other bidders.

Sealed-bid auctions are also common off-line (e.g. government procurement), but not second-price sealed-bid (Vickrey) auctions. The first-price sealed-bid auction is useful when communication costs are important and there are additional *privacy* concerns that favor a sealed-bid auction over an (open) ascending-bid auction (Rothkopf *et al.* 1990). Rothkopf *et al.* (1990) argue that the lack of Vickrey auctions in practice is because of *trust* issues. The Vickrey auction is vulnerable to manipulation by a non-trusted auctioneer. The auctioneer can overstate the second highest bid unless the bid can be verified, and the auction is therefore vulnerable to manipulation.

3 On-Line Auctions: What and Why

The choice of the auction format to be used should be derived by the costs and conditions of trust (or lack of trust) of the parties in a particular situation. As we will see, agent-mediated on-line auctions offer somewhat different cost and trust structures from more traditional off-line auctions, and hence favor English auctions.

3.1 Agent-mediated On-line Auctions

We consider on-line auctions where artificial bidding agents may represent the preferences of human buyers. Communication in an on-line auction is electronic, and auction clearing is fully automatic. Auctions are a particularly simple form of distributed negotiation, but they are a first step towards more sophisticated automated negotiation for electronic commerce (Beam *et al.* 1996). Beam and Segev (1997) identify two forms of automated-negotiation agents, those that are programmed with a complete set of strategies, and those that learn better strategies from success or failure in previous negotiations. In auctions the strategy is generally simple, especially in incentive-compatible auctions, and the issue is not whether the strategy should be programmed up-front (it should), but whether a complete set of preferences for every good that a bidding-agent might discover should be programmed up-front. A fully-autonomous agent would require a complete set of preferences for the human buyer in order to represent the buyer correctly in all situations that it might encounter. We promote the idea of a *semi-autonomous agent* that will bid on behalf of a human buyer when it has enough knowledge to proceed, but will query the buyer for further information and instruction when its best action is not well-defined.

For example, consider a car-buying agent that has been instructed that the buyer is interested in a convertible VW Beetle that was made after 1975, has done less than 40000 miles, and is in mint condition – and that the agent should definitely bid up to $3000. The bidding agent will behave autonomously while no such cars are for sale, or while the current highest bid (which it will be holding) is less than $3000. The agent will only seek further advice (either to continue bidding, or to stop bidding) when the price of a suitable car increases above $3000. This is the sense in which the bidding agent is semi-autonomous. A human buyer programs a bidding agent with some initial knowledge, and then provides more information as necessary.

3.2 Costs in On-Line Auctions

The costs to human participants in agent-mediated on-line auctions are *cognitive* (information gathering and processing) and *communicative*. The cost characteristics of auctions vary, of course, depending on the mechanism, the goods being auctioned, and the complexity of the local problem of the bidders. We shall see that on-line agent-mediated auctions have different cost structures from more traditional off-line auctions. We claim that the cognitive costs to human participants are now much more important than communication costs.

Traditional considerations of the cognitive costs of auctions include the cost of strategic bidding in a non-incentive compatible auction. Clearly, incentive-compatible auctions require less computation by bidders than auctions which are not incentive-compatible, and thus are preferable (IC_B, Table 1). Estimating how much other bidders are willing to pay (in order to bid slightly more than what one thinks the second highest bidder will bid) is not easy, although in agent-mediated auctions this cost can be transferred from a human participant to its bidding agent, and is less important than in traditional auctions.

Another key cost is widely neglected: the cognitive cost of *deciding the value* of an auctioned object. In economic theory the value of an object is assumed to be *known* information (private-value assumption) or set exogenously (common-value assumption), but is not, to the best of our knowledge, assumed to be *computable* at some cost. This is a cognitive cost that cannot be easily transferred to an electronic bidding agent. That there is a non-trivial cost for a buyer to decide what her 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.

Different auction structures require vastly different expected effort by buyers to determine how much they value a good being auctioned. Bidding agents that participate in a Vickrey auction must always know the true reservation price for an auctioned good, in advance of the bidding, or risk letting the good go at a price that the buyer would have been willing to pay. However, if the bidding stops (with the agent winning) before the price gets close to the reservation price in an English auction, then the buyer saves the cognitive cost of deciding exactly how much she is willing to pay. Similarly, if the bidding starts well above the buyer's reservation price, there is no need for the bidding-agent to know the true value for the good. A bidding agent only needs to know the true value when other buyers have similar reservation prices. 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, and therefore prefer auctions that are "reservation-price-not-necessary" (nRP, Table 1).

Communication costs in on-line auctions are much lower than in off-line auctions (Beam *et al.* 1996). Buyers can now use agents (either their own or ones provided by the auction house) to bid repeatedly, from distributed locations, and auctioneers can cheaply handle the electronic bids. Mobile agents bidding over a slow or unreliable network (Lukose and Huberman 1997) can also migrate to the auction server node and execute at that node. The only important costs are due to human cognition: deciding what to buy, and how much one is willing to pay. As in off-line auctions, non-incentive-compatible auctions have high cognitive cost and should be avoided in on-line auctions, although this cognitive cost can be bore by the bidding agent. However, in on-line auctions it is worth having an agent submit many bids if one can potentially avoid having to compute one's reservation price for an object. Thus progressive auctions are preferred to sealed-bid auctions because they enable *progressive reasoning*.

3.3 Trust in On-Line Auctions

Another important, but secondary, consideration that drives auction design is trust. A sealed-bid second-price auction is only incentive-compatible with a trusted auctioneer. If the auctioneer is not trusted then second price auctions have little meaning. A non-trusted auctioneer can receive all the bids, and then place a bid just below the highest outside bid received, and charge the highest bidder a price just below its winning bid (Rothkopf *et al.* 1990). Assuming that this can be done without risk of detection, the action is always revenue increasing and the good will be sold to the same (outside) bidder, but at a higher price. A bidding agent will treat a non-trusted Vickrey auction as a first-price sealed-bid auction, and behave strategically (nT, Table 1).

When might auctioneers not be trustworthy? It is not unusual for firms to auction off surplus items which they own. In this case the auctioneer is often the seller, and thus has a conflict of interest. A self-interested auctioneer should not be trusted by rational bidding agents. An institutional solution to untrusted auctioneers in on-line Vickrey auctions is to provide *trust-certified* third parties that operate the auctions.

The English auction is also susceptible to seller manipulation, but strategic action is not risk-free, and requires the help of a third-party, known as a *shill* (IC_S, Table 1). Sellers can drive up the price that they receive by having third parties bid for them. The goal is to compete with an interested bidder, and push up the price close to that bidder's reservation price for the item. There is clearly a risk involved for the seller in bidding above her reservation price, as she may be left with the item. In off-line auctions where the bidders are in the same room, the shill (acting for the seller) can often use visual feedback to learn when the bidder is close to her true reservation price and reduce its risk. An on-line auction does not permit such feedback, but price-manipulation is still possible given prior information on the likely reservation prices of bidders in the auction.

We might also compare the auction protocols with regard to buyer collusion. Buyer collusion that occurs prior to an auction is hard to prevent. 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 was extracted from the seller. This type of collusion is possible in any type of auction, and generally hard to detect. 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. One key property that is necessary for dynamic collusion is the identification of bidders. An on-line English auction can offer protection from bidder collusion by posting all bids anonymously.

3.4 A Suitable On-line Auction Type

The most important properties for an on-line auction with a trusted auctioneer are motivated by cognitive cost considerations. The auction should be buyer-incentive-compatible (IC_B) to make strategic bidding unnecessary, and reservation-price-not-necessary (nRP) to enable progressive reasoning. We can see from Table 1 that the two auction formats with these properties (ignoring the redundant ascending second-price auction) are the ascending first-price and descending second-price auctions. Even with a trusted auctioneer, the Vickrey auction is not suitable because we wish to avoid the high cognitive cost of computing the reservation price for every auction. When we also consider the trust properties of auctions (nT), we see that the descending second-price auction is susceptible to manipulation by a trusted auctioneer in the same way as the Vickrey auction, while the English auction is secure to a non-trusted auctioneer because the winning bidder pays the price that she bid. With our claim that human cognition is much more important than the low cost of agent-mediated communication and monitoring, the most suitable auction mechanism for agent-mediated on-line auctions is the English auction. The Vickrey auction is less useful on-line (even with a trusted auctioneer) because communication is cheap and cognitive costs matter. Sandholm (1996) has noted some other limitations of Vickrey auctions in computational multiagent systems.

4 Commercial On-Line Auctions

There are many on-line auctions that are currently in use (Yahoo! 1998). Some sites are run by discounters selling multiple goods, often from odd lots (Major vendor sites), while others are open to the public and anyone can sell (Open public sites), see Table 2. Research platforms for computational auctions are also on the Internet (Wurman *et al.* 1997). It is illuminating to look at the design of real on-line auctions in the light of the above discussion.

Table 2. Commercial On-Line Auctions. **Key**: M – Major Vendor Auction, P – Open Public Auction

Company	**URL**	**Bidding Agent**	**Reservation Price Auction**	**Mechanism**	**M / P**
Auctionworks	<http://www.auctionworks.com>	No	Yes, agent	English	P
BiddersParadise	<http://www.biddersparadise.com>	Yes	No	English	P
AuctionUniverse	<http://www.auctionuniverse.com>	Yes	Yes	English	P
eBay's AuctionWeb	<http://www.ebay.com>	Yes	Yes	English	P
Haggle	<http://www.haggle.com>	Yes	Yes	English	P
onSale	<http://www.onsale.com>	No	No	English	M
zAuction	<http://www.zauction.com>	No	No	English	M
Surplus Direct	<http://www.surplusauction.com>	No	No	First-Price Sealed-Bid	M

Many commercial on-line auctions use the English auction as the underlying mechanism to match buyers and sellers of a single good, although a sealed-bid auction is sometimes used in the final stages of an auction to avoid end-game strategic bids that might seek to take advantage of Internet communication latencies. The auctions often provide agents that will automatically compete in an auction up to the "reservation price" of a buyer (bidding-agent, Table 2). The agents monitor a selection of active auctions and place new bids (increased by a minimum increment from the current bid) when the current highest bid is below the stated reservation price of the buyer, and held by another bidding-agent. Buyers typically receive an e-mail message when their current reservation price is reached. This provides a mechanism for buyers to avoid the cognitive costs of determining exactly how much they value an object, they can initially provide a lower-bound on their (unknown) true reservation price, and improve the bound only as needed, when the current price in the auction exceeds it.

The bidding agents are really implementing a second-price auction, but doing so in a way that allows the current second-highest bid to be publically posted so that other potential buyers can judge whether it is worth their time to enter the bidding. The bidding agents *reveal* the mechanism behind the incentive-compatibility of second-price sealed-bids, and "help screens" that accompany auction sites encourage users to give the bidding agents their true reservation price (e.g. eBay's AuctionWeb, Table 2). A buyer in a non-trusted auction might prefer to implement and execute her own bidding agent, given that the bidding agents hold information on her reservation price that could be used to the advantage of a non-trusted auctioneer.

In another complication of the simple auction, sellers are often permitted to start the auction at a price below their reservation price to generate interest. An auction is sometimes marked as a "reservation price auction" if this is the case. The seller does not have to sell the good if the highest bid falls below her reservation price, and a buyer that holds the best bid when the auction closes will not always receive the good. A better way to generate interest without allowing a bidding agent to hold an invalid highest bid, is to provide a "seller-agent" that simply bids for the good on behalf of the seller if an outside bid is received that is lower than a predetermined reservation price (e.g. Auctionworks, Table 2).

5 Conclusions and Recommendations

Real auctions differ from those typically analyzed by economists and computer scientists. Issues of cost and trust drive these differences. On-line auctions are characterized by low communication costs, so sealed-bid auctions have little advantage. The largest "overhead" cost in participating in an on-line auction is deciding what to buy and how much one is willing to pay for it. This is a much greater problem than in off-line auctions where goods are typically sold in selected lots to a group of interested expert buyers. Agent that search the Internet for goods that satisfy the constraints of a human buyer may find many items that seem suitable (Guttman *et al.* 1998), but that may in fact be unsuitable or of poor quality. English auctions allow human buyers to commit semi-autonomous agents to act on their behalf in on-line auctions, enabling progressive reasoning about the exact value of a good, with better estimates only required as necessary. This frees the buyer of the need to value every single good exactly. English auctions are also secure to price-manipulation by a strategic auctioneer, except at some risk through a third-party shill.

Prices in progressive auctions provide feedback that allows interested human buyers to reduce the cognitive cost of reasoning about their value for a good and avoid difficult decisions. For example, if the posted current price is high (relative to what the buyer is willing to pay) in an ascending-price auction then there is no need to bid. If the posted price is low, then the buyer can have her agent start bidding, without going through the effort of determining the exact price that she is willing to pay. Only if the posted price in the auction is comparable to the buyer's true reservation price does she need to determine her true reservation price. We do not recommend progressive auctions because of the affiliated-value information effect that causes bidders to adjust their valuation upwards on the basis of high valuations from other bidders (McAfee and McMillan 1987), but rather to allow progressive reasoning. We assume that every buyer has the information necessary to reason about her true value for a good, but that this cognitive process is non-trivial.

This suggests a new design for bidding agents that participate in on-line English auctions: an autonomous *reservation-price-agent* for buyers that have an exact valuation for the good being auctioned, and a semi-autonomous *progressive-price-agent* for buyers that only have an approximate valuation. The reservation-price-agent will compete in relevant auctions until either the current asking price is above the reservation price, or the agent wins. The progressive-price-agent requires a *lower-bound* on the true reservation price of the buyer, and (optionally) an *upper-bound*. The agent will behave autonomously while the current highest bid in a relevant auction is either less than the lower-bound (in which case the agent will hold that bid), or the current highest bid is above the upper-bound (in which case the agent will no longer monitor that auction). The agent will request more information from the buyer, in the form of refinements on the bounds whenever the current highest bid is in the range of uncertain value between the two bounds. The agent will stop monitoring the auction when the upper-bound is lowered to below the current asking price or the current asking price increases to above the upper-bound. In future work we intend to further investigate the design of useful semi-autonomous market-aware agents for electronic commerce.

6 Acknowledgments

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

References

Beam, C., Segev, A., Shanthikumar, J. G.: Electronic Negotiation through Internet-based Auctions. CITM Working Paper 96-WP-1019, December 1996. <http://haas.berkeley.edu/ citm/wp-1019-summary.html>

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>

Guttman, R., Moukas, A., Maes, P.: Agent-mediated Electronic Commerce: A Survey. To appear, Knowledge Engineering Review (June 1998) <http://ecommerce.media.mit.edu/papers/ker98.ps>

Lukose, R. M., Huberman, B. A.: A Methodology for Managing Risk in Electronic Transactions over the Internet. Third International Conference on Computing in Economics and Finance (CEFES'97). <http://www.parc.xerox.com/spl/groups/dynamics/www/banking.html>

McAfee, R. P., McMillan, J.: Auctions and Bidding. Journal of Economic Literature **25** (1987) 699–738

Milgrom, P., Weber, R. J.: A Theory of Auctions and Competitive Bidding. Econometrica **50(5)** (1982) 1089–1122

Milgrom, P.: Auctions and Bidding: A primer. Journal of Economic Perspectives **3(3)** (1989) 3–32

Rothkopf, M. H., Teisberg, T. J., Kahn, E. P.: Why are Vickrey auctions rare? Journal of Political Economy **98(1)** (1990) 94–109

Sandholm, T., Lesser, V.: Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework. First International Conference on Multiagent Systems (ICMAS-95) 328–335
<ftp://ftp.cs.umass.edu/pub/lesser/sandholm-icmas95-issues.ps>

Sandholm, T.: Limitations in the Vickrey Auction in Computational Multiagent Systems. Second International Conference on Multiagent Systems (ICMAS-96) 299–306 <http://www.cs.wustl.edu/ sandholm/vickrey.ICMAS96.ps>

Varian, H.: Economic Mechanism Design for Computerized Agents. First USENIX Workshop on Electronic Commerce, July 1995.
<ftp://alfred.sims.berkeley.edu/pub/Papers/mechanism-design.ps.Z>

Vickrey, W.: Counter Speculation, Auctions, and Competitive Sealed Tenders. Journal of Finance **16** (1961) 8–37

Wellman, M. P., Wurman, P. R.: Market-aware Agents for a Multiagent World. Forthcoming, Robotics and Autonomous Systems.
<ftp://ftp.eecs.umich.edu/people/wellman/maamaw.ps>

Wurman, P. R., Wellman, M. P., Walsh, W. E.: The Michigan internet AuctionBot: A Configurable Auction Server for Human and Software Agents. Forthcoming, Autonomous Agents'98.
<ftp://ftp.eecs.umich.edu/people/wellman/agents98wurman.ps.Z>

Yahoo! Online Auction URL: <http://www.yahoo.com/Business_and_Economy/Companies/Auctions/Online_Auctions>