

# Issues in Computational Vickrey Auctions\*

**Tuomas Sandholm**

sandholm@cs.wustl.edu

Washington University

Department of Computer Science

St. Louis, MO 63130-4899

---

\*This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994. A short early version of this paper appeared at a conference [45].

# Issues in Computational Vickrey Auctions

## Abstract

The Vickrey auction has been widely advocated for multiagent systems. First we review its limitations so as to guide practitioners in their decision of when to use that protocol. These limitations include lower revenue than alternative protocols, lying in non-private-value auctions, bidder collusion, a lying auctioneer, and undesirable revelation of sensitive information. We discuss the special characteristics of Internet auctions: third party auction servers, cryptography, and how proxy agents relate to the revelation principle and fail to promote truth-telling.

Then we uncover limitations of the protocol which stem from computational complexity considerations. These include inefficient allocation and lying in sequential auctions of interrelated items, untruthful bidding under valuation uncertainty, and counterspeculation to make deliberation control (or information gathering) decisions. We also discuss methods for winner and price determination in combinatorial “second-price” auctions, with implications on truth-dominance.

**Keywords:** Auction, multiagent system, proxy bidder, winner determination, electronic commerce.

# 1 Introduction

Auctions provide efficient, distributed and autonomy preserving ways of solving task and resource allocation problems in multiagent systems. While fixed-menu take-it-or-leave-it offers are still quite common in electronic commerce, auctions are the leading vehicle for dynamically priced electronic trades. Electronic auctions are used for consumer-to-consumer, business-to-consumer, and business-to-business electronic commerce, and the success of recent Internet auction companies has been phenomenal. Auctions can be used among cooperative agents, but they also work in open systems consisting of self-interested agents. An auction can be analyzed using noncooperative game theory: what strategies are self-interested agents best off using in the auction (and therefore will use), and will a desirable social outcome—*e.g.* efficient allocation—still follow. The goal is to design the protocols (mechanisms) of the interaction so that desirable social outcomes follow even though each agent acts based on self-interest.

This paper serves three roles:

1. It reviews the known limitations of the Vickrey auction protocol, *i.e.* the second-price sealed-bid auction. Vickrey auctions have been widely advocated and adopted for use in computational multiagent systems [62, 17, 13, 58, 1, 9, 8, 56, 36, 20, 21]. The methods and frequency of adoption suggest that the limitations are not well understood by the multiagent systems and electronic commerce communities. Understanding the shortcomings is important in order not to ascribe desirable characteristics—such as truth-dominance and counterspeculation avoidance—to a protocol when the protocol really does not guarantee them in the setting in question. Misapplications of protocols open the door for manipulation by the agents, which may lead to inefficient outcomes and processes. While this paper focuses on the limitations of the Vickrey auction, there are some auction settings where that protocol is desirable. The explication of the limitations will hopefully help practitioners distinguish when the Vickrey auction is and is not appropriate.
2. It discusses how computational auctions—potentially among computerized agents as bidders

and/or auctioneers—in settings like the Internet differ from traditional auctions among humans, and how these differences affect the choice of auction protocols. Issues of proxy bidders, cryptographic techniques, and trusted third party auction servers are discussed.

3. Perhaps most importantly, it uncovers new limitations of the Vickrey auction that stem from computational complexity and information gathering considerations. While the computational questions apply to auctions among humans as well, they emerge particularly clearly in multiagent systems because the bidders' and/or auctioneer's algorithms need to be constructed, and because the agents' cognitive capabilities can be analytically characterized.

An auction consists of an auctioneer and potential bidders.<sup>1</sup> Auctions are usually discussed regarding situations where the auctioneer wants to sell an item and get the highest possible payment for it while each bidder wants to acquire the item at the lowest possible price. However, settings in which the auctioneer wants to subcontract out a task at the lowest possible price and each bidder wants to handle the task at the highest possible payment, are totally analogous.

The rest of the paper is organized as follows. Section 2 reviews different auction settings and protocols. Section 3 details the known problems regarding the Vickrey auction. These problems have been discovered by auction theorists and practitioners, and they have led to the lack of deployment of Vickrey auctions among humans. The first problem is lower revenue than alternative protocols (Section 3.1). Section 3.2 discusses current Internet auctions and how the simple proxy agents relate to the revelation principle, and how they fail to incent truthful bidding. The other problems discussed in the first part of the paper include lying in non-private-value auctions (Section 3.3), bidder collusion (Section 3.4), an untruthful auctioneer (Section 3.5), and the necessity to reveal sensitive information (Section 3.6). The wide advocacy of Vickrey auctions in multiagent systems suggests that these limitations have not been fully assimilated by the builders of multiagent systems. The first part of the paper serves as a reminder for practitioners.

The second part (Section 4) focuses on limitations that arise mainly from computational con-

---

<sup>1</sup>There are also auctions with multiple bid takers, *i.e.* auctioneers.

siderations. The first of these problems is inefficient allocation and lying in sequential auctions of interrelated items (Section 4.1). That section also discusses combinatorial auctions, algorithms for winner determination, fully expressive bidding languages, and the impacts of approximate winner determination on truth-dominance. The computational complexities of bidding are also elaborated. Section 4.2 shows how untruthful bidding can be beneficial when a bidder's valuation is uncertain. Finally, Section 4.3 shows how counterspeculation is needed for deciding on deliberation actions and information gathering actions.

## 2 Auction settings and protocols

There are three qualitatively different auction settings depending on how an agent's value of the auctioned item is formed.

In *private value* auctions, the value of the good depends only on the agent's own preferences. An example is auctioning off a cake that the winning bidder will eat. The key is that the winning bidder will not resell the item or get utility from showing it off to others, because in such cases the value would depend on other agents' valuations. The agent is often assumed to know its value for the good exactly.

On the other hand, in *common value* auctions, an agent's value of an item depends entirely on other agents' values of it: all agents have the same value for the item. Auctioning treasury bills is an example: nobody inherently prefers having the bills, and the value of the bill comes entirely from reselling possibilities.

In *correlated value* auctions, an agent's value depends partly on its own preferences and partly on others' values. For example, a negotiation within a contracting setting fulfills this criterion. An agent may handle a task itself in which case its local concerns define the cost of handling the task. On the other hand, the agent can recontract out the task in which case the cost depends solely on other agents' valuations.

The rest of this section reviews four simple auction protocols from the literature for allocating

a single item. The following sections compare their properties in the different settings discussed above.

In the *English (first-price open-cry) auction*, each bidder is free to raise his bid. When no bidder is willing to raise anymore, the auction ends, and the highest bidder wins the item at the price of his bid. An agent's strategy is a series of bids as a function of his private value, his prior estimates of other bidder's valuations, and the past bids of others. In private value English auctions, an agent's dominant strategy is to always bid a small amount more than the current highest bid, and stop when his private value price is reached. These strategies lead to the bidding ending when the second highest bidder's valuation (plus epsilon) has been reached. In correlated value auctions the rules are often varied to make the auctioneer increase the price at a constant rate or at a rate he thinks appropriate. Also, sometimes *open-exit* is used where a bidder has to openly declare exiting without a re-entering possibility. The motivation behind these variations is to make more information about the bidders' valuations public. This information will signal to a bidder about the item's value.

In the *first-price sealed-bid auction*, each bidder submits one bid without knowing the others' bids. The highest bidder wins the item and pays the amount of his bid. An agent's strategy is his bid as a function of his private value and his prior beliefs of others' valuations. In general there is no dominant strategy for bidding in this auction. An agent's best strategy is to bid less than his true valuation, but how much less depends on what the others bid. The agent would want to bid the lowest amount that still wins the auction—given that this amount does not exceed his valuation. With common knowledge assumptions regarding the probability distributions of the agents' values, it is possible to determine Nash equilibrium strategies for the agents. For example, in a private value auction where the valuation,  $v_i$ , for each agent,  $i$ , is drawn independently from a uniform distribution between 0 and some positive number,  $\bar{v}$ , there is a Nash equilibrium where every agent,  $i$ , bids  $\frac{a-1}{a}v_i$ , where  $a$  is the number of bidders (see *e.g.* [32]).

In the *Dutch (descending) auction*, the seller continuously lowers the price until one of the bidders takes the item at the current price. The Dutch auction is strategically equivalent to the first-price sealed-bid auction because in both games an agent's bid matters only if it is the highest, and no

relevant information is revealed during the auction process. Among humans, Dutch auctions are often used for perishable goods such as flowers and fish. In these auctions, the current price is reduced rapidly so as to make the auction efficient in terms of real time. Note, however, that the other auction protocols discussed in this paper could be executed rapidly as well, see *e.g.* [11]. Dutch auctions are not often used in multiagent systems; the FishMarket project is a notable exception [35]. One reason may be the need to pass price information repeatedly to the bidders, and preferably at the same time. While the network lag can be implicitly removed from the bids, for example, by using timestamping methods, guaranteeing that the price information gets posted to all bidders simultaneously is more difficult. While Dutch auctions are not common in real electronic markets, simulations do exist on the Internet, see *e.g.* [www.mcsr.olemiss.edu/~ccjimmy/auction](http://www.mcsr.olemiss.edu/~ccjimmy/auction).

In the *Vickrey (second-price sealed-bid) auction*, each bidder submits one bid without knowing the others' bids. The highest bidder wins, but at the price of the second highest bid. An agent's strategy is his bid as a function of his private value and prior beliefs of others' valuations. A bidder's dominant strategy in a private value Vickrey auction is to bid his true valuation [60]. If he bids more than his valuation, and the increment made the difference between winning or not, he will end up with a loss if he wins. If he bids less, there is a smaller chance of winning, but the winning price is unaffected.<sup>2</sup> The truth-dominance result means that an agent is best off bidding truthfully no matter what the other bidders are like: what their capabilities, operating environments, bidding plans, *etc.*, are. This has two desirable sides. First, the agents reveal their preferences truthfully which allows globally efficient decisions to be made. Second, the agents need not waste effort in counterspeculating other agents because they do not matter in making the bidding decision.

Vickrey auctions have been widely advocated and adopted for use in computational multiagent systems [58, 1, 9, 8, 21, 56, 36, 59, 13]. For example, versions of the Vickrey auction have been

---

<sup>2</sup>In private value auctions, the Vickrey auction is strategically equivalent to the English auction. They will produce the same allocation at the same prices. On the other hand, in correlated value auctions, the other agents' bids in the English auction signal to the agent about his own valuation. Therefore, English and Vickrey auctions are not strategically equivalent in general, and may lead to different results.

used to allocate computation resources in operating systems [62], to allocate bandwidth in computer networks [20], and to computationally control building heating [17]. On the other hand, Vickrey auctions have not been widely adopted in auctions among humans [37, 39] even though the protocol was invented over 28 years ago [60].

There are severe limitations to the applicability of the Vickrey auction protocol. This paper explores the limitations. Understanding the shortcomings is important in order not to ascribe desirable characteristics—such as truth-dominance and counterspeculation avoidance—to a protocol when the protocol really does not guarantee them in the setting in question. Some of the limitations are independent of computational considerations, but others stem from computational issues. While the computational questions apply to auctions among humans as well, they emerge particularly clearly in multiagent systems because the bidders’ and/or auctioneer’s algorithms need to be constructed, and because the agents’ cognitive capabilities can be analytically characterized. While this paper focuses on the limitations of the Vickrey auction, in some auction settings the Vickrey protocol is desirable. The explication of the limitations will hopefully help practitioners distinguish when the Vickrey auction is and is not appropriate.

### **3 General limitations of the Vickrey auction**

Before discussing the problems that arise in Vickrey auctions due to computational issues, we review some known limitations of the Vickrey auction that do not stem from computational considerations. These limitations alone warrant caution when applying the Vickrey auction, for example to Internet commerce or to multiagent systems in general.

#### **3.1 Lower revenue than with the English auction**

In isolated private value or common value auctions, each one of the four auction protocols (English, Dutch, first-price sealed-bid, and Vickrey) allocates the auctioned item Pareto efficiently to the



bidder who values it the most.<sup>3</sup> Although all four are Pareto efficient in the allocation, the ones with dominant strategies (Vickrey and English auction) are more efficient in the sense that no effort is wasted in counterspeculating the other bidders.

One might expect the first-price auctions (first-price sealed-bid and Dutch) to give higher expected revenue to the auctioneer because in second-price auctions (the Vickrey auction is second-price by definition, and the English auction is second-price in effect because the winning bidder only has to bid as high as the second highest bidder is willing to raise plus  $\epsilon$ ) the auctioneer only gets the second price. This is not the case, however, because in first-price auctions the bidders are motivated to lie by biasing their bids downward. The *revenue-equivalence theorem* [28, 29, 24, 22] states that all four auction protocols produce the same expected revenue to the auctioneer in private value auctions where the values that bidders place on the item are independently drawn from an identical distribution, and bidders are risk-neutral.

Among risk-averse bidders, the Dutch and the first-price sealed-bid protocols give higher expected revenue to the auctioneer than the Vickrey or English auction protocols [23, 27].<sup>4</sup> This is because in the former two protocols, a risk-averse agent can insure himself by bidding more than what is optimal for a risk-neutral agent. So, since agents take on the preferences of the real world parties that they represent, and most real world parties are risk-averse in practice, it may make sense for the auctioneer to choose one of the former two protocols. On the other hand, a risk-averse auctioneer achieves higher expected utility via the Vickrey or English auction protocols than via the Dutch or the first-price sealed-bid protocol (given that the bidders are risk-neutral). This is because the four protocols lead to the same expected revenue, and the former two have less variance.

The fact that revenue equivalence holds in private value auctions—with the assumptions men-

---

<sup>3</sup>This holds as long as the auctioneer always sells the item. On the other hand, the optimal auction protocol for private value auctions where the bidder's valuations are drawn independently from the same distribution is a modified second-price auction where the seller also submits a bid. Unfortunately the seller's best strategy is not to bid his true valuation: he should overbid [31]. This may cause the seller to inefficiently keep the item.

<sup>4</sup>Interestingly, among risk-seeking agents a *third-price auction* can lead to expected revenue that is higher than the expected value of the highest valuation among the bidders [30].

tioned above—does not mean that it usually holds in practice: most auctions are not pure private value auctions. In non-private value auctions with at least three bidders, the English auction (especially the open-exit variant) leads to higher revenue than the Vickrey auction. The reason is that the willingness of other bidders to state high prices causes a bidder to increase his own valuation of the auctioned item. In these types of auctions, both the English and the Vickrey protocols produce greater expected revenue to the auctioneer than the first-price sealed-bid auction or its equivalent, the Dutch auction.

The fact that the bidders can update their valuations based on the others' bids also has interesting repercussions on proxy bidding in Internet auctions, as the next subsection will discuss.

### 3.2 Proxy bidder agents and the revelation principle in Internet auctions

Most existing Internet auction houses use the English auction protocol or some variant of it. However, many of these auction sites allow the bidder to tell his maximum bid to a *proxy bidder* “agent” that will bid in the English auction by always bidding a minimum increment over the current high bid, and exiting when the user’s maximum bid is reached, see for example [www.ebay.com](http://www.ebay.com) or [www.webauction.com](http://www.webauction.com).

Such “agents” can be implemented by the trusted auction server via one global thread of execution. As the high bid changes, all the bidders’ stored maximum bids are checked, in first-come-first-serve order, to see if any “agent” would like to raise its bid, and the first one that would raise gets to do so. This check is repeated until no agent is willing to raise its bid.

Another alternative is to give each agent its own thread of execution. This is implemented in *eMediator*, a next generation electronic commerce server that we have built [49]. In addition to having their own threads of execution, the agents in *eMediator* are mobile: they can move to other sites on the net [16]. Users can program their agents for *eMediator* using Java. In addition, an HTML page is provided for non-programmers for preference elicitation, and the user’s mobile agent is automatically programmed based on the preferences. In many auction settings, these automatically programmed agents will bid optimally on the user’s behalf based on game-theoretically

predetermined strategies. *eMediator* also incorporates other novel features such as combinatorial bidding [48], bidding with price-quantity graphs, leveled commitment contracts [55, 53, 46], and a safe exchange planner [46, 51].

From a game theoretic point of view, the “agents” in traditional Internet auctions convert the auction protocol from an English auction to a Vickrey auction: the participant with the highest willingness to pay gets the item at the price of the willingness to pay of the second highest participant. This is an interesting real world manifestation of the *revelation principle* [22]. It states that any outcome that can be supported in equilibrium via a complex protocol can be supported in an equilibrium via a protocol where the agents reveal their types truthfully in a single step. The proof is based on having the new protocol incorporate a virtual player for each real world participant such that the virtual player will find and play the best strategy for the original complex protocol on behalf of the real world participant—given that the participant reveals his preferences to the virtual player. Because the virtual player will play optimally for the participant, the participant is motivated to reveal his preferences truthfully. Each “agent” in current Internet auctions is a materialization of such a theoretical virtual player.

Each “agent” in current Internet auctions plays optimally on behalf of the real world participant in private value auctions: it keeps increasing the bid by a minimal amount and stops when the participant’s revealed maximum is reached. However, this is not the best strategy in non-private value auctions. Instead, the “agent” should take into account in its strategy how the others’ bids affect the participant’s valuation. Therefore, the current “agents” do not play optimally for the participants in general. It follows that a participant is not necessarily best off by revealing her maximum willingness to pay truthfully to her “agent”—unlike the Internet auction sites suggest. Since English auctions have higher expected revenue than Vickrey auctions in non-private settings, the current over-simplification of the “agents” tends to hurt the sellers.

### 3.3 Bidders lying in non-private-value auctions

Most auctions are not pure private value auctions: an agent's valuation of a good depends at least in part on the other agents' valuations of that good. For example in contracting settings, a bidder's evaluation of a task is affected by the prices at which the agent can subcontract the task or parts of it out to other agents. This type of recontracting is commonly allowed in automated versions of the contract net protocol also [44, 47, 2, 57].

Common value auctions (and correlated value auctions to a certain extent) suffer from the *winner's curse*. If an agent bids its valuation and wins the auction, it will know that its valuation was too high because the other agents bid less. Therefore winning the auction amounts to a loss in utility. Knowing this in advance, agents should bid less than their valuations [28, 32]. This is the best strategy in common value Vickrey auctions also. So, even though the Vickrey auction promotes truthful bidding in private-value auctions, it fails to induce truthful bidding in most auction settings.

### 3.4 Vulnerability to bidder collusion

One problem with all four of the auction protocols is that they are not collusion proof [34]. The bidders could coordinate their bid prices so that the bids stay artificially low. In this manner, the bidders get the item at a lower price than they normally would.

The English auction and the Vickrey auction actually self-enforce some of the most likely collusion agreements. Therefore, from the perspective of deterring collusion, first-price sealed-bid or Dutch auctions are preferable. The following example from [32] shows this. Let bidder Smith have value 20, and every other bidder have value 18 for the auctioned item. Say that the bidders collude by deciding that Smith will bid 6, and everyone else will bid 5. In an English auction this is self-enforcing, because if one of the other agents exceeds 5, Smith will observe this, and will be willing to go all the way up to 20, and the cheater will not gain anything from breaking the coalition agreement. In the Vickrey auction, the collusion agreement can just as well be that Smith bids 20,

because Smith will get the item for 5 anyway. Bidding 20 removes the incentive from any bidder to break the coalition agreement by bidding between 5 and 18, because no such bid would win the auction. On the other hand, in a first-price sealed-bid auction, if Smith bids anything below 18, the other agents have an incentive to bid higher than Smith's bid because that would cause them to win the auction at a profit. The same holds for the Dutch auction.

However, for collusion to occur under the Vickrey auction, the first-price sealed-bid auction, or the Dutch auction, the bidders need to identify each other before the submission of bids. Otherwise a non-member of the coalition could win the auction. On the other hand, in the English auction this is not necessary, because the bidders identify themselves by shouting bids. To prevent this, the auctioneer can organize a computerized English auction where the bidding process does not reveal the identities of the bidders.

### **3.5 Vulnerability to a lying auctioneer**

Insincerity of the auctioneer may be a problem in the Vickrey auction, and this problem is exacerbated in Internet auctions where the sellers and buyers often do not know or trust each other. The auctioneer may overstate the second highest bid to the highest bidder unless that bidder can verify it. An overstated second offer would give the highest bidder a higher bill than she would receive if the auctioneer were truthful.

Cheating by the auctioneer has been suggested as one of the main reasons why the Vickrey auction protocol has not been widely adopted in auctions among humans [39]. In another paper, two formal models of cheating by the auctioneer are discussed [37]. The first model is game theoretic. It analyses the situation where the auctioneer can choose to use a first-price sealed-bid protocol or a Vickrey protocol. The bidders' equilibrium behavior creates positive incentives for all auctioneers, except the type most prone to cheat, to choose standard first-price sealed-bid auctions. The second model assumes simple, not rational, bidders. They bid honestly as long as the auctioneer has not been caught cheating, but after catching a cheating auctioneer, the bidders will bid as if the auctioneer always cheats. The result is that a seller with probabilistic opportunities to cheat, and finite abilities

to resist cheating, will cheat and be caught in finite time and thereafter have no reason to conduct Vickrey auctions.

To address the problem of a lying auctioneer, cryptographic electronic signatures could be used by the bidders so that the auctioneer could actually present the second best bid to the winning bidder—and would not be able to alter it. However, this would not preclude the auctioneer from having some phony bidder bid at the last moment just below the winning bidder's bid [60]. This problem can be tackled by using an (automated) trusted third party auction server that reveals the bids to the seller only once the auction has closed instead of having the seller be the auctioneer directly. The other three auction protocols (English, Dutch, and first-price sealed-bid) do not suffer from lying by the auctioneer because the highest bidder gets the item at the price that she stated in the bid.

The auctioneer may also have other tools at his disposal. For example, he may place a bid himself to guarantee that the item will not be sold below a certain price. This can also be achieved by having a reservation price which might not be public to the bidders. However, for example in the Vickrey auction, the auctioneer is motivated to bid more than his true reservation price. This is because there is a chance that his bid will be second highest in which case it determines the item's price. Such overbidding leads to the possibility that the auctioneer ends up inefficiently keeping the item even though some bidders' valuations exceed the auctioneer's valuation. Unfortunately, no other auction protocol can lead to higher allocative efficiency either [31]

Some of the other protocols are vulnerable to different kinds of further manipulations. For example, in non-private value auctions with the English auction protocol, the auctioneer can use *shills* that bid in the auction in order to make the real bidders increase their valuations of the item. This is not possible in the sealed-bid protocols or the Dutch protocol because the bidders do not observe the others' bids.

### 3.6 Undesirable private information revelation

Because the Vickrey auction has truthful bidding as the dominant strategy in private value auctions, agents often bid truthfully. This leads to bidders revealing their true valuations. Sometimes this information is sensitive, and the bidders would prefer not to reveal it. For example, after winning a contract with a low bid, a company's subcontractors figure out that the company's production cost is low, and therefore the company is making larger profits than the subcontractors thought. It has been observed that when such auction results are revealed, the subcontractors will want to renegotiate their deals to get higher payoffs [39]. This has been suggested—along with the problem of a lying auctioneer—as one of the main reasons why the Vickrey auction protocol is not widely used in auctions among humans [39].

First-price auction protocols do not expose a bidder's valuation as clearly because the bid is subject to strategic lying: the bid is based on the agent's model of other bidders, and this possibly inaccurate model is not known by the subcontractors. Therefore, these auction types may be more desirable than the Vickrey auction when valuations are sensitive.

## 4 Problems arising from computational limitations

Most of auction theory has studied auctions without reference to deliberation. However, deliberation considerations are of key importance in many auction settings. In this era when auctioneers and bidders are being computerized, these computational questions are becoming increasingly apparent. Fortunately the deliberative characteristics of computational agents can be analytically modeled and studied. This section discusses computational issues in auctions and how they can be viewed as limitations in the applicability of the Vickrey auction protocol which has been widely advocated for use in computational auction settings.

## 4.1 Inefficient allocation and lying in interrelated auctions

In addition to single-item auctions, Vickrey auctions have been widely studied in the allocation of multiple units of a single good [28]. However, the case of auctioning heterogeneous interrelated goods has received little attention. On the other hand this is the setting of many real world problems where computational agents are used [52, 54, 53, 44, 36].

This section discusses cases where heterogeneous items are auctioned one at a time, and an agent's valuations of these items are interdependent, *i.e.* not additive. Such valuations prevail for example in delivery task allocation in transportation problems [44], in bandwidth allocation [25, 26, 20, 21], and in airport landing slot allocation [33]. They also occur in auctioning collectibles on the Internet, *e.g.* when the bidders are interested in acquiring a complete set of Star Wars figures, Beanie babies, *etc.*

We first demonstrate that the optimal allocation is not reached if the bidders treat the auctions independently and bid truthfully.

**Example 4.1** *Figure 1 presents a simple example of a transportation problem with two delivery tasks:  $t_1$  and  $t_2$ . Task  $t_1$  is auctioned before  $t_2$ . The auctioneer wants to get the tasks handled while paying agents 1 and 2 as little as possible for handling them. The initial locations of the two agents are presented in the figure. To handle a task, an agent needs to move to the beginning of the delivery task (arrow), and take a parcel from there to the end of the arrow. An agent's movement incurs the same cost irrespective of whether it is carrying a parcel or not. The agents need not return to their initial locations. The costs for handling tasks (subscripted by the name of the agent) can be measured from the figure:  $c_1(\{t_1\}) = 2$ ,  $c_1(\{t_2\}) = 1$ ,  $c_1(\{t_1, t_2\}) = 2$ ,  $c_2(\{t_1\}) = 1.5$ ,  $c_2(\{t_2\}) = 1.5$ , and  $c_2(\{t_1, t_2\}) = 2.5$ . Say that these costs are common knowledge. Clearly the globally optimal allocation is the one where agent 1 handles both tasks.*

*We now show that this allocation is not reached if agents treat the auctions independently and bid truthfully. In the first auction of the example, task  $t_1$  is allocated. Agent 1 bids  $c_1(\{t_1\}) = 2$ , and agent 2 bids  $c_2(\{t_1\}) = 1.5$ . The task is allocated to agent 2. In the second auction, task  $t_2$*



is allocated. Agent 1 bids  $c_1(\{t_2\}) = 1$ , and agent 2 bids  $c_2(\{t_2\}) = 1.5$ , so  $t_2$  is allocated to agent 1. The resulting allocation of the two tasks is suboptimal. If agent 2 takes the ownership of  $t_1$  into account when bidding for  $t_2$ , then it will bid  $c_2(\{t_1, t_2\}) \Leftrightarrow c_2(\{t_1\}) = 2.5 \Leftrightarrow 1.5 = 1$ . In this case  $t_2$  may be allocated to either agent. In both cases the resulting allocation of the two tasks is still inefficient.

Alternatively, the agents could incorporate full lookahead into their auction strategies. As the next example shows, this way the optimal allocation is reached, but agents do not bid their true myopic per-item costs.

**Example 4.2** *In the last auction of Example 4.1—i.e. the auction for task  $t_2$ —each bidder is best off bidding its own costs that takes into account the tasks that the bidder already has (because truth-telling is a dominant strategy for a risk-neutral bidder in a single item Vickrey auction). Let us look at the auction of  $t_2$ . If agent 1 has  $t_1$ , it will bid  $c_1(\{t_1, t_2\}) \Leftrightarrow c_1(\{t_1\}) = 2 \Leftrightarrow 2 = 0$ , and  $c_1(\{t_2\}) = 1$  otherwise. If agent 2 has  $t_1$ , it will bid  $c_2(\{t_1, t_2\}) \Leftrightarrow c_2(\{t_1\}) = 2.5 \Leftrightarrow 1.5 = 1$ , and  $c_2(\{t_2\}) = 1.5$  otherwise. So, if agent 1 has  $t_1$ , it will win  $t_2$  at the price 1.5, and get a payoff of  $1.5 \Leftrightarrow 0 = 1.5$  in the second auction, while agent 2 gets zero. On the other hand, if agent 2 has  $t_1$ , the bids for  $t_2$  are equal, and both agents get a zero payoff in the second auction irrespective of which bidder gets  $t_2$ . Therefore it is known that getting  $t_1$  in the first auction is worth an extra 1.5 to agent 1 while nothing extra to agent 2. So, in the auction for  $t_1$ , agent 1's dominant strategy is to bid  $c_1(\{t_1\}) \Leftrightarrow 1.5 = 2 \Leftrightarrow 1.5 = 0.5$ . This is lower than agent 2's bid  $c_2(\{t_1\}) \Leftrightarrow 0 = 1.5 \Leftrightarrow 0 = 1.5$ , so agent one gets  $t_1$ . In the second auction agent 1 gets  $t_2$  as discussed above. So the globally optimal allocation is reached. However, agent 1 bids 0.5 for  $t_1$  instead of 2, which would be the truthful bid if the auctions were treated independently without lookahead.*

Put together, lookahead is a key feature when sequentially auctioning interrelated items. To date it has not been adequately addressed in computational multiagent systems that use Vickrey auctions, and it is a common misunderstanding that Vickrey auctions promote single-shot truth-telling even in interrelated auctions.

### 4.1.1 Comparison of computation in sequential vs. combinatorial auctions

An alternative to sequential auctioning of the interdependent items would be to open them all for auction in parallel. However, some of the same problems prevail. For example, when bidding for an item, the bidder does not know its valuation because it depends on which other items the bidder wins, which in turn depends on how others will bid (in sealed-bid auctions this is not known to the bidder, and in open-cry auctions it may become known only later).

One solution to this problem is to allow the bidders to place bids for combinations of items [33, 43, 44, 26]. The rest of this section will compare the computational aspects of protocols that allow this against protocols that do not. The next subsection discusses the auctioneer's complexity, and the subsection after that addresses the complexity on the bidders' side.

#### Complexity of winner and price determination

The determination of winners—*i.e.*, determining what items each bidder gets—is easy in non-combinatorial auctions. It can be done by picking the highest bidder for each item separately. This takes  $O(an)$  time where  $a$  is the number of bidders, and  $n$  is the number of items. In such auctions, determining the Vickrey price of each item is equally easy: it can be done in  $O(an)$  time by simply finding the second highest bid for each item.

Winner determination in combinatorial auctions is more difficult. Let  $X$  be the set of items to be auctioned. Then any agent,  $i$ , could place any bid  $b_i(S)$  for any combination  $S \subseteq X$ . If agent  $i$  does not place a bid for combination  $S$ , we set  $b_i(S) = 0$ . Let

$$\bar{b}(S) = \max_i b_i(S) \tag{1}$$

Now, winner determination in a combinatorial auction is the following problem:

$$\max_{\mathcal{S}} \sum_{S \in \mathcal{S}} \bar{b}(S) \tag{2}$$

where  $\mathcal{S}$  is a valid outcome, *i.e.* an outcome where each item is allocated to only one bidder:  $[S \subseteq X, T \subseteq X, S \in \mathcal{S}, T \in \mathcal{S}] \Rightarrow S \cap T = \emptyset$ . The problem can be solved in  $O(3^n)$  time using

dynamic programming [38]. However, no known algorithm can solve the problem in polynomial time in the size of the input if only some combinations have received bids: that problem is the same as *weighted set packing*, which is  $\mathcal{NP}$ -complete [19]. The problem can be made tractable by placing severe restrictions on what combinations can be bid on [38], but such restrictions can lead to inefficient outcomes because the bidders are faced with similar uncertainties as in bidding for interrelated items sequentially.

Another approach for optimal winner determination is to allow all combinations to be bid on, and to capitalize on the fact that in practice the space of bids is usually extremely sparsely populated. For example, even if there are only 100 items to be auctioned, there are  $2^{100} \approx 1$  combinations, and it would take longer than the life of the universe to bid on all of them even if every person in the world submitted a bid per second. Sparseness of bids implies sparseness of the allocations,  $\mathcal{S}$ , that actually need to be checked. We recently devised an algorithm that only checks those allocations [48]. The details of the algorithm are beyond the scope of this paper. The algorithm scaled to hundreds of items and thousands of bids in minutes on a general-purpose uniprocessor workstation.

Naive methods for winner determination are based on the implicit assumption that each agent's bids are locally superadditive:  $b_i(S \cup S') \geq b_i(S) + b_i(S')$ . But what would happen if agent 1 bid  $b_1(\{1\}) = 5$ ,  $b_1(\{2\}) = 4$ , and  $b_1(\{1, 2\}) = 7$ ? The auctioneer could allocate items 1 and 2 to agent 1 separately, and that agent's bid for the combination would value at  $5 + 4 = 9$  instead of 7. This could be fixed by using an auction protocol that only allocates combinations according to bids on combinations, *i.e.*, restricts to making at most one bid win per agent. Alternatively this could be handled by a protocol that allows the allocation of multiple bids to an agent but only if the agent did not submit a lesser bid for the corresponding combination. We have also developed a protocol where the bidders themselves can submit mutually non-exclusive (OR) bids as well as mutually exclusive (XOR) bids. This allows the bidders to express general preferences, and our winner determination algorithm works in this case as well [48].

With such bidding languages that allow bidders to express general preferences, it is possible to generalize the Vickrey pricing rule to the combinatorial case to obtain truthful bidding in dominant

strategy equilibrium. Following the Groves–Clarke mechanism [14, 7], the amount that an agent needs to pay should be computed as the sum of the others’ winning bids had the agent not submitted any bids, minus the sum of the others’ winning bids in the actual optimal allocation. Therefore, the winner determination problem would have to be solved once overall, and once per winning agent without any of that agent’s bids. Just removing one winning bid at a time would not lead to an incentive compatible mechanism, *i.e.*, one where the agents are motivated to bid truthfully in dominant strategy equilibrium.

The algorithms used by the auctioneer also have repercussions on the bidders’ strategies. If either winner determination or price determination is done only approximately, incentive compatibility can be lost.

### Complexity of bidding

Computational complexity is also present in each bidder’s strategy, for example in the following ways:

1. Determining the valuation of an isolated item may be intractable. For example, when a dispatch center evaluates the cost of taking on a transportation task, it would need to compute the solution to a vehicle routing problem with its old tasks and the new task. Then it would need to compute the solution to a vehicle routing problem with its old tasks only. The cost of accepting the new task is the cost of the former solution minus the cost of the latter. However, solving the two vehicle routing problems would mean solving two  $\mathcal{NP}$ -complete problems. Therefore, in practice, the cost of the task has to be approximated [44, 46].
2. The setting becomes more complex if the valuation of the item depends on what other items the agent has. In a combinatorial auction this would mean calculating the valuation on potentially all combinations of items. In a sequential auction, determining the valuation would require finding out what other items the agent will get. This requires lookahead in the game tree. Solving the whole game tree is often intractable, but acting myopically is not desirable either:

partial search in the game tree would be more appropriate. However, several open questions remain. How does one evaluate the node when further search could be done from it but it is decided that further search is not worthwhile? How does one decide how deep to search (this may vary along different paths of the same game tree)?

3. Even if one assumes that the other bidders bid nonstrategically, at some points of the game tree of a sequential auction one needs to know (at least probabilistically [5]) what the others will bid because that affects what items one will get. This may require solving their optimization problems, which again may be intractable. While avoidance of counterspeculation was one of the original reasons suggested for adopting the Vickrey auction, lookahead requires speculation in the sense of trying to guess how high others are going to bid. Other speculative issues in sequential Vickrey auctions have been discussed for example in [18].
4. Dominant strategy mechanisms, such as the Vickrey auction, do not require common knowledge of the priors, but as discussed, many of those mechanisms lose their dominant strategy properties when one takes their limitations into account. Moving then to the Nash equilibrium solution concept, or one of its refinements, introduces the problem that those solution concepts assume common knowledge, which is usually unobtainable [10]. Instead of common knowledge, in reality there is a recursive modeling tree: what does 1 think that 2 thinks that 3 thinks that 1 thinks that 3 thinks that 2 thinks... [6, 12]. It is often intractable to deduce in that tree—even if information is limited. Again, partial search in this tree would be appropriate [61].

Search is a central paradigm in AI, and several efficient engineering answers to similar questions have been provided over the years [40]. Recently, more normative methods for deliberation control have been developed, but they still do not reach the goal of provably optimal reasoning in general because they make simplifying assumptions such as myopic search control, costless meta-reasoning, and conditioning the search control on engineer-chosen features [42, 3], or they guarantee only optimal composition of underlying components which themselves may not be optimal [41, 4, 63, 50, 15].

This is an invaluable important topic of continuing research especially for game theory: in order to devise normative theories of how agents should act in non-cooperative equilibrium, one also needs to prescribe how the agents should deliberate. Therefore we envision that each agent’s deliberation actions need to be modeled as part of the agent’s strategy—just like physical actions. This approach would allow one to make the best computational tradeoffs, *e.g.*, in the four complexity-generating items above simultaneously.

Once such computation-incorporating equilibrium analyses are done, one can seriously start to compare auction protocols to each other. One can start to ask the key questions such as what combinational bids should be allowed, and how parallel vs. sequential the protocol should be.

## 4.2 Untruthful bidding with local uncertainty

Bidders often have uncertainty about their valuation of the auctioned item. This valuation may be inherently uncertain. Computational agents may have additional uncertainty regarding the valuation because computing it can be complex, and the computation may not have finished by the time of the auction. Such computational complexities arise for example in task allocation auctions where evaluating a task set requires solving  $\mathcal{NP}$ -complete problems [52, 54, 53, 44, 36]. Also, an agent may be better off by carrying out an approximate valuation calculation before bidding, and investing the remaining detailed computation only if he wins the item (computation is still important at this time because the agent may need to decide how to act with the item, *e.g.*, how to incorporate a delivery task into the weekly vehicle routing solution).

In settings where a bidder has uncertainty about his valuation, a risk-neutral bidder is best off bidding his *expected* valuation in a single-shot private value Vickrey auction. This is a dominant strategy. However, in practice, most agents are risk-averse. Computational agents take on the preferences of the real world parties that they represent. Therefore most computational agents will be risk-averse as well. We now show that risk-averse agents are not best off bidding truthfully in the Vickrey auction. Therefore it is nonobvious that the Vickrey auction protocol can really be used in computational systems to avoid lying in practice.

**Example 4.3** Let the bidder's utility function be  $U(x) = \begin{cases} 2x & \text{if } x \leq 0 \\ x & \text{if } x > 0. \end{cases}$  The concavity of this function represents risk aversion of the agent. Let the bidder's valuation,  $v$ , be uniformly distributed between 0 and 1. We now show that the bidder can increase his expected utility by bidding  $E[v] \Leftrightarrow \epsilon$  instead of  $E[v]$ . We analyze the situation based on what the highest bid,  $b$ , coming from the other agents might be.

**Case 1:**  $b \leq E[v] \Leftrightarrow \epsilon$ . In this case, the agent wins the auction at price  $b$  when bidding  $E[v]$  or  $E[v] \Leftrightarrow \epsilon$ . Therefore the expected utility is unaffected by bidding  $E[v] \Leftrightarrow \epsilon$  instead of  $E[v]$ .

**Case 2:**  $b \geq E[v]$ . In this case, the agent loses the auction when bidding  $E[v]$  or  $E[v] \Leftrightarrow \epsilon$ . Therefore the expected utility  $U(0) = 0$  is unaffected by bidding  $E[v] \Leftrightarrow \epsilon$  instead of  $E[v]$ .

**Case 3:**  $E[v] \Leftrightarrow \epsilon < b < E[v]$ . In this case, the bidder loses the auction when bidding  $E[v] \Leftrightarrow \epsilon$ , but wins it when bidding  $E[v]$ . Therefore, the utility from bidding  $E[v] \Leftrightarrow \epsilon$  is  $U(0) = 0$ . The expected utility from bidding  $E[v]$  is

$$\begin{aligned} \int_{-\infty}^{\infty} U(v \Leftrightarrow b) dv &= \int_0^b 2(v \Leftrightarrow b) dv + \int_b^1 v \Leftrightarrow b dv \\ &= \Leftrightarrow \frac{1}{2} b^2 \Leftrightarrow b + \frac{1}{2} \end{aligned}$$

which is less than zero when  $b > \frac{-2+\sqrt{8}}{2} \approx 0.41$  (i.e. in the range of case 3). So, bidding  $E[v]$  has smaller expected utility than bidding  $E[v] \Leftrightarrow \epsilon$ .

### 4.3 Counterspeculation for deciding on information gathering

One of the main original motivations for using the Vickrey auction was that each bidder has a dominant strategy (of telling the truth), i.e. a bidder's best strategy does not depend on other bidders. Therefore the bidders will not waste effort in counterspeculating each other. This would lead to global savings.

This section shows that there are cases where the Vickrey auction fails to have this desirable property. Let us look at a setting where a bidder has uncertainty regarding his valuation, but can pay to remove this uncertainty. This situation often occurs among computational agents, for

example because the only way to evaluate a good (or cost of taking on a task) might be to carry out a costly computation—*e.g.*, solving combinatorial problems as discussed earlier. Alternatively the payment can be viewed as the cost of solving a prediction problem, or as the cost of performing an information gathering action (this is a common decision that needs to be made by automated agents in Internet commerce), or as the cost paid to an expert oracle. The following theorem states that in such a setting, the Vickrey auction protocol does not avoid counterspeculation. In particular, the decision of whether or not to pay to resolve the uncertainty depends on the other bidders.

**Proposition 4.1** *In a private value Vickrey auction with uncertainty about a bidder's valuation, the bidder's best (deliberation or information gathering) strategy can depend on the other bidders.*

**Proof.** Let there be two risk-neutral bidders: 1 and 2. Let agent 1's valuation,  $v_1$ , be uniformly distributed between 0 and 1, *i.e.* agent 1 does not know his valuation exactly. Let agent 2's exact valuation,  $v_2$ , be common knowledge. Let  $0 \leq v_2 < \frac{1}{2}$ , which implies  $E[v_1] > v_2$ .

Let agent 1 have the choice of finding out his exact valuation,  $v_1$ , before the auction by paying a cost,  $c$ . Now, should agent 1 take this informative but costly action?

No matter what agent 1 chooses here, agent 2 will bid  $v_2$  because bidding one's valuation is a dominant strategy in a single-shot private value Vickrey auction for a risk-neutral bidder.

If agent 1 chooses to not pay  $c$ , agent 1 should bid  $E[v_1] = \frac{1}{2}$  because bidding one's expected valuation is a risk-neutral agent's dominant strategy in a single-shot private value Vickrey auction. Now agent 1 gets the item at price  $v_2$ . If agent 1's valuation  $v_1$  turns out to be less than  $v_2$ , agent 1 will suffer a loss. Agent 1's expected payoff is

$$E[\pi_{noinfo}] = \int_0^1 v_1 \Leftrightarrow v_2 dv_1 = \frac{1}{2} \Leftrightarrow v_2$$

On the other hand, if agent 1 chooses to pay  $c$  for the exact information, he should bid  $v_1$  because bidding one's valuation is a risk-neutral bidder's dominant strategy in a single-shot private value Vickrey auction. In this case agent 1 gets the item if *and only if*  $v_1 \geq v_2$ . Note that now the agent has no chance of suffering a loss, but on the other hand he has invested  $c$  in the information. So,



agent 1's expected payoff is

$$\begin{aligned} E[\pi_{info}] &= \int_0^{v_2} c dv_1 + \int_{v_2}^1 v_1 \Leftrightarrow v_2 \Leftrightarrow c dv_1 \\ &= \frac{1}{2} v_2^2 \Leftrightarrow v_2 + \frac{1}{2} \Leftrightarrow c \end{aligned}$$

Agent 1 should choose to buy the information iff

$$\begin{aligned} E[\pi_{info}] &\geq E[\pi_{noinfo}] \\ \Leftrightarrow \frac{1}{2} v_2^2 \Leftrightarrow v_2 + \frac{1}{2} \Leftrightarrow c &\geq \frac{1}{2} \Leftrightarrow v_2 \\ \Leftrightarrow \frac{1}{2} v_2^2 &\geq c \\ \Leftrightarrow v_2 &\geq \sqrt{2c} \quad (\text{because } v_2 \geq 0) \end{aligned}$$

So, agent 1's best choice of action depends on agent 2's valuation,  $v_2$ . Therefore, agent 1 benefits from counterspeculating agent 2.  $\square$

## 5 Conclusions

Vickrey auctions have been widely advocated and adopted for use in computational multiagent systems. This auction protocol has certain desirable properties—such as truth-promotion and counterspeculation avoidance—in limited settings. It is important to clearly understand these limitations in order not to use the protocol when inappropriate, and in order not to trust the protocol to have certain desirable properties when it really does not have them in the particular setting.

The first part of the paper reviewed known problems regarding the Vickrey auction. These include lower revenue than alternative protocols, promotion of lying in non-private-value auctions, bidder collusion, vulnerability to a lying auctioneer, and the necessity to reveal sensitive information. Bidder collusion can be reduced by electronic auctions because the bidders may be unable to identify each other. The proposed mechanisms for avoiding lying by the auctioneer include cryptographic signatures on bids, and the use of trusted third party auction servers.

The proxy bidder agents of current Internet auction houses are a real world materialization of

the revelation principle. In private value auctions they convert the English auction into a Vickrey auction. However, in non-private value auctions they do not bid optimally on the user's behalf unlike the auction sites claim. It follows that users are not motivated to reveal their valuations truthfully to the proxy bidder agents.

The second part of the paper presented our results regarding limitations of the protocol that mainly stem from computational considerations. While these problems prevail in auctions among humans, they become particularly apparent when the algorithms for the auctioneer and bidder agents need to be designed. They also become increasingly analyzable since the deliberative capabilities of computational agents can be precisely modeled.

The first problem is inefficient allocation and lying in sequential auctions of interrelated items. The problem was demonstrated via two simple example games. Combinatorial auctions were discussed as a candidate remedy. Winner determination is easy in non-combinatorial auctions, but  $\mathcal{NP}$ -complete in combinatorial auctions. Restricting the allowable combinations to bid on is one way of making winner determination tractable, but that introduces some of the same inefficiencies that non-combinatorial auctions have. A new algorithm for optimal winner determination in the non-restricted setting was briefly mentioned. It capitalizes on the sparseness of bids. Bidding languages that allow users to express general preferences were also presented. In combinatorial auctions, the Vickrey-Groves-Clarke pricing mechanism introduces added computational complexity. Also, truth-dominance ceases to hold if approximate algorithms are used for winner or price determination. We also discussed the computational complexities involved in bidding, and how those complexities cause the Vickrey auction to no longer promote truth-telling and avoid counterspeculation.

We showed how insincere bidding can be beneficial when a risk-averse bidder is uncertain about his valuation. Finally, we showed the need for counterspeculation to make deliberation control—or information gathering—decisions when an agent has local uncertainty. So, here too the Vickrey auction loses the dominant strategy property.

In the future, systems will increasingly be designed, built, and operated in a distributed manner. A larger number of systems will be used by multiple real-world parties. The problem of coordinating

these parties and avoiding manipulation cannot be tackled by technological or economic methods alone. Instead, the successful solutions are likely to emerge from a deep understanding and careful hybridization of both. In particular, when taking into account the agents' computational constraints by modeling each agent's computational steps as actions that the agent intentionally takes, many of the classic results from game theory cease to apply.

## References

- [1] Agorics, Inc. Going, going, gone! A survey of auction types. <http://www.agorics.com/agorics/auctions/bibliography.html> (1996).
- [2] Andersson, M. R., and Sandholm, T. W. Time-quality tradeoffs in reallocative negotiation with combinatorial contract types. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)* (Orlando, FL, 1999), pp. 3–10.
- [3] Baum, E. B., and Smith, W. D. A Bayesian approach to relevance in game playing. *Artificial Intelligence* 97, 1–2 (1997), 195–242.
- [4] Boddy, M., and Dean, T. Deliberation scheduling for problem solving in time-constrained environments. *Artificial Intelligence* 67 (1994), 245–285.
- [5] Boutilier, C., Goldszmidt, M., and Sabata, B. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)* (Stockholm, Sweden, Aug. 1999), pp. 527–534.
- [6] Brainov, S., and Sandholm, T. W. Auctions without common knowledge. In *International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Agent-Mediated Electronic Commerce (AMEC)* (Stockholm, Sweden, Aug. 1999), pp. 1–12.
- [7] Clarke, E. H. Multipart pricing of public goods. *Public Choice* 11 (1971), 17–33.

- [8] Drexler, K. E., and Miller, M. S. Incentive engineering for computational resource management. In *The Ecology of Computation*, B. A. Huberman, Ed. North-Holland, 1988, pp. 231–266.
- [9] Edelman, J. Web page mention of the use of Vickrey auction in AgentTCL. *http://www.cs.dartmouth.edu/~jxe/resume.html* (1996).
- [10] Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. *Reasoning About Knowledge*. MIT Press, 1995.
- [11] Fujishima, Y., McAdams, D., and Shoham, Y. Speeding up ascending-bid auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)* (Stockholm, Sweden, Aug. 1999), pp. 554–559.
- [12] Gmytrasiewicz, P. J., and Durfee, E. H. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)* (San Francisco, CA, June 1995), pp. 125–132.
- [13] Gomber, P., Schmidt, C., and Weinhardt, C. Efficiency and incentives in MAS-coordination. In *Proceedings of the 5th European Conference on Information Systems (ECIS)* (Cork, Ireland, 1997), vol. 2, pp. 1040–1051.
- [14] Groves, T. Incentives in teams. *Econometrica* 41 (1973), 617–631.
- [15] Horvitz, E. Reasoning about beliefs and actions under computational resource constraints. In *Proceedings of Third Workshop on Uncertainty in Artificial Intelligence* (Seattle, Washington, July 1987), American Association for Artificial Intelligence, pp. 429–444. Also in L. Kanal, T. Levitt, and J. Lemmer, ed., *Uncertainty in Artificial Intelligence 3*, Elsevier, 1989, pps. 301–324.
- [16] Huai, Q., and Sandholm, T. W. Mobile agents in an electronic auction house. In *Mobile Agents in the Context of Competition and Cooperation (MAC3) Workshop at the Third International Conference on Autonomous Agents (AGENTS)* (Seattle, WA, 1999), pp. 24–33.

- [17] Huberman, B., and Clearwater, S. H. A multi-agent system for controlling building environments. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)* (San Francisco, CA, June 1995), pp. 171–176.
- [18] Jackson, M., and Peck, J. Speculation and price fluctuations with private, extrinsic signals. *Journal of Economic Theory* 55 (1991), 274–295.
- [19] Karp, R. M. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, Eds. Plenum Press, NY, 1972, pp. 85–103.
- [20] MacKie-Mason, J. K., and Varian, H. R. Pricing the Internet. In *Proceedings of the Public Access to the Internet Conference* (May 1993), JFK School of Government.
- [21] MacKie-Mason, J. K., and Varian, H. R. Some FAQs about usage-based pricing, 1994. Available from <ftp://alfred.sims.berkeley.edu/pub/Papers/useFAQs.html>. Also in Proceedings of WWW '94, Chicago, Illinois, and in Proceedings of the Association of Research Librarians 1994. To appear in Computers and ISDN.
- [22] Mas-Colell, A., Whinston, M., and Green, J. R. *Microeconomic Theory*. Oxford University Press, 1995.
- [23] Maskin, E. S., and Riley, J. Optimal auctions with risk-averse buyers. *Econometrica* 52, 6 (1984), 1473–1518.
- [24] McAfee, R. P., and McMillan, J. Auctions and bidding. *Journal of Economic Literature* 25 (1987), 699–738.
- [25] McAfee, R. P., and McMillan, J. Analyzing the airwaves auction. *Journal of Economic Perspectives* 10, 1 (1996), 159–175.
- [26] McMillan, J. Selling spectrum rights. *Journal of Economic Perspectives* 8, 3 (1994), 145–162.
- [27] Milgrom, P., and Weber, R. A theory of auctions and competitive bidding. *Econometrica* 50 (1982), 1089–1122.

- [28] Milgrom, P. R. The economics of competitive bidding: a selective survey. In *Social goals and social organization: Essays in memory of Elisha Pazner*, L. Hurwicz, D. Schmeidler, and H. Sonnenschein, Eds. Cambridge University Press, 1985, ch. 9, pp. 261–292.
- [29] Milgrom, P. R. Auction theory. In *Advances in Economic Theory: Fifth World Conference* (1987), T. Bewley, Ed., Cambridge University Press.
- [30] Monderer, D., and Tennenholtz, M. Optimal auctions revisited. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)* (Madison, WI, July 1998), pp. 32–37.
- [31] Myerson, R. B. Optimal auction design. *Mathematics of Operation Research* 6 (1981), 58–73.
- [32] Rasmusen, E. *Games and Information*. Basil Blackwell, 1989.
- [33] Rassenti, S. J., Smith, V. L., and Bulfin, R. L. A combinatorial auction mechanism for airport time slot allocation. *Bell J. of Economics* 13 (1982), 402–417.
- [34] Robinson, M. S. Collusion and the choice of auction. *Rand Journal of Economics* 16, 1 (1985), 141–145.
- [35] Rodriguez-Aguilar, J. A., Noriega, P., Sierra, C., and Padget, J. FM96.5: A Java-based electronic auction house. In *In Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97)* (1997).
- [36] Rosenschein, J. S., and Zlotkin, G. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [37] Rothkopf, M. H., and Harstad, R. M. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business* 68, 2 (1995), 257–267.
- [38] Rothkopf, M. H., Pekeč, A., and Harstad, R. M. Computationally manageable combinatorial auctions. *Management Science* 44, 8 (1998), 1131–1147.
- [39] Rothkopf, M. H., Teisberg, T. J., and Kahn, E. P. Why are Vickrey auctions rare? *Journal of Political Economy* 98, 1 (1990), 94–109.

- [40] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [41] Russell, S., and Subramanian, D. Provably bounded-optimal agents. *Journal of Artificial Intelligence Research 1* (1995), 1–36.
- [42] Russell, S., and Wefald, E. *Do the right thing: Studies in Limited Rationality*. The MIT Press, 1991.
- [43] Sandholm, T. W. A strategy for decreasing the total transportation costs among area-distributed transportation centers. In *Nordic Operations Analysis in Cooperation (NOAS): OR in Business* (Turku School of Economics, Finland, 1991).
- [44] Sandholm, T. W. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)* (Washington, D.C., July 1993), pp. 256–262.
- [45] Sandholm, T. W. Limitations of the Vickrey auction in computational multiagent systems. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)* (Keihanna Plaza, Kyoto, Japan, Dec. 1996), pp. 299–306.
- [46] Sandholm, T. W. *Negotiation among Self-Interested Computationally Limited Agents*. PhD thesis, University of Massachusetts, Amherst, 1996. Available at <http://www.cs.wustl.edu/~sandholm/dissertation.ps>.
- [47] Sandholm, T. W. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium Series: Satisficing Models* (Stanford University, CA, Mar. 1998), pp. 68–75.
- [48] Sandholm, T. W. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)* (Stockholm, Sweden, 1999), pp. 542–547. Extended version: Washington University, Department of Computer Science technical report WUCS-99-01, January.

- [49] Sandholm, T. W. *eMediator*: A next generation electronic commerce server. In *AAAI-99 Workshop on AI in Electronic Commerce* (Orlando, FL, July, 1999), pp. 46–55. Extended version: Washington University, Department of Computer Science technical report WUCS-99-02, January.
- [50] Sandholm, T. W., and Lesser, V. R. Utility-based termination of anytime algorithms. In *ECAI Workshop on Decision Theory for DAI Applications* (Amsterdam, The Netherlands, 1994), pp. 88–99. Extended version: Univ. of Mass. at Amherst, Comp. Sci. Tech. Report 94-54.
- [51] Sandholm, T. W., and Lesser, V. R. Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)* (Montreal, Canada, Aug. 1995), pp. 694–701.
- [52] Sandholm, T. W., and Lesser, V. R. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)* (San Francisco, CA, June 1995), pp. 328–335. Reprinted in *Readings in Agents*, Huhns and Singh, eds., pp. 66–73, 1997.
- [53] Sandholm, T. W., and Lesser, V. R. Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)* (Portland, OR, Aug. 1996), pp. 126–133.
- [54] Sandholm, T. W., and Lesser, V. R. Coalitions among computationally bounded agents. *Artificial Intelligence* 94, 1 (1997), 99–137. Special issue on Economic Principles of Multiagent Systems. Early version appeared at the International Joint Conference on Artificial Intelligence, pages 662–669, 1995.
- [55] Sandholm, T. W., Sikka, S., and Norden, S. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)* (Stockholm, Sweden, 1999), pp. 535–540. Extended version: Washington University, Department of Computer Science technical report WUCS-99-04.



- [56] Smart Market. Web page mention of the use of Vickrey auction in Smart Market. [http://mortadello.wu-wien.ac.at/~iwork/pricing/smart\\_market.html](http://mortadello.wu-wien.ac.at/~iwork/pricing/smart_market.html) (1996).
- [57] Smith, R. G. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers C-29*, 12 (Dec. 1980), 1104–1113.
- [58] Sun Microsystems. Web page mention of the use of Vickrey auction in Webmart. <http://www.sun.com:80/960201/cover/webmart.html> (1996).
- [59] Varian, H. Mechanism design for computerized agents. In *Usenix Workshop on Electronic Commerce* (New York, July 1995).
- [60] Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1961), 8–37.
- [61] Vidal, J. M., and Durfee, E. H. Recursive agent modeling using limited rationality. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)* (San Francisco, CA, 1995), pp. 376–383.
- [62] Waldspurger, C. A., Hogg, T., Huberman, B., Kephart, J. O., and Stornetta, W. S. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering* 18, 2 (1992), 103–117.
- [63] Zilberstein, S., and Russell, S. Optimal composition of real-time systems. *Artificial Intelligence* 82, 1–2 (1996), 181–213.

TUOMAS SANDHOLM is assistant professor of computer science at Washington University in St. Louis. He received the Ph.D. and M.S. degrees in computer science from the University of Massachusetts at Amherst in 1996 and 1994. He earned an M.S. with distinction in Industrial Engineering and Management Science from the Helsinki University of Technology, Finland in 1991. His research interests include electronic commerce, artificial intelligence, game theory, multiagent

systems, rational resource-bounded reasoning, auctions, automated negotiation and contracting, coalition formation, machine learning, constraint satisfaction, and combinatorial optimization. He has nine years of experience building multiagent systems for electronic commerce, and is chief scientist of a startup company. He has published over 95 academic papers, and received the NSF CAREER award.

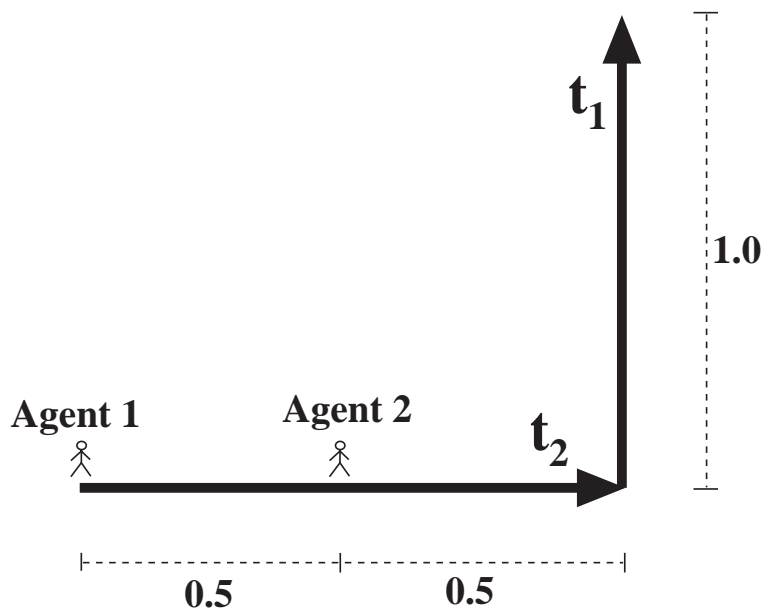


Figure 1: *Small example problem with two agents and two delivery tasks (bold arrows).*