Operations Research
Vol. 55, No. 2, March–April 2007, pp. 191–197
issn 0030-364X | eissn 1526-5463 | 07 | 5502 | 0191
doi 10.1287/opre.1070.0384

Thirteen Reasons Why the Vickrey-Clarke-Groves Process Is Not Practical

Michael H. Rothkopf
MSIS Department and RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway, New Jersey 08854-8003, rothkopf@rutcor.rutgers.edu

In theory, the mathematically elegant Vickrey-Clarke-Groves process offers perfect efficiency with dominant truth-revealing strategies. However, it has many serious practical problems. This paper describes these problems and argues that research that aims to maintain the dominant truth-revealing strategies while compromising on the other practical issues is of limited practical value.

Subject classifications: auctions; mechanism design; truth-revealing strategies.
Area of review: OR Forum.
History: Received December 2005; revision received January 2007; accepted January 2007.

Introduction

In 1961, in the first published game-theoretic model of auctions, William Vickrey analyzed an unusual auction form, sealed bidding in which the highest bidder wins but the price is set by the bid of the best losing bid. Vickrey argued that bidders in such auctions should bid their true value.[1] Thus, the bidder with the highest true value would always win, making the auction allocation perfectly efficient. He also proved a "revenue neutrality theorem" that showed that (with independent private values) the expected revenue of the seller was the same for this auction as for standard first-price sealed bidding, English (i.e., progressive) auctions, and Dutch (i.e., descending price) auctions.[2] This work was one of three topics mentioned in Vickrey's Nobel Prize citation. Auctions in which the best losing bid sets the price for the winner are now commonly called Vickrey auctions.

E. H. Clarke (1971) and T. Groves (1973) generalized Vickrey's results to a rather general competitive process. Vickrey auctions are a special case of such Vickrey-Clarke-Groves (VCG) processes. These VCG processes are designed so that it is apparently in the interest of bidders in auctions of multiple items with interdependent values to truthfully bid the value of each combination. This is achieved by having an optimization (in general it is an integer programming problem) determine the best combinations of bids to honor and setting prices by refunding to each successful bidder the increase in the value of the objective function that is due to her bids. This can be done by resolving the optimization problem once for each winning bidder with all of that bidder's bids omitted. Winning bidders pay the amount of their bids but are refunded the difference between the optimal value of the objective function with all bids and its value with its own bids omitted. This refund is sometimes called a "Vickrey payment" or "Vickrey discount." This process assures the bidder that her bids determine whether she wins, but do not affect the price she has to pay.[3] There is no longer a revenue neutrality theorem, but many theorists have been captivated by the idea of dominant truth-revealing strategies and perfect efficiency in combinatorial auctions.[4] Combinatorial auctions are becoming important in commerce, and it is tempting to try to develop theory based upon maintaining the theoretical truth-revealing nature of the process while using approximations to deal with some of the practical problems that arise in trying to implement VCG combinatorial auctions. While incentives are important, they may be an unwise choice when the truth-revealing ideal of VCG process is, in practice, usually an unreachable mirage.

Vickrey auctions were rare in 1990 (see Rothkopf et al. 1990) and remain so, and as far as I know, no one has conducted a *general* VCG process (i.e., not just a simple Vickrey auction or a market-clearing price auction of identical items) in real commerce. I believe that there are important practical reasons for this. It is the purpose of this paper to spell out these reasons. Some of them go back almost three decades (e.g., Groves and Ledyard 1977), and some were first articulated quite recently. Here is a list of the reasons discussed in this paper:

- the fact that the dominant strategy equilibrium is a weak equilibrium and there may exist alternative weak equilibria;
- the nonexistence of dominant strategy equilibria in models that include reasonable bid preparation costs;
- the exponential growth of effort related to bid preparation and bid communication;
- the NP completeness of the winner determination problem;
- problems related to capital limited bidders;
- problems associated with the disclosure of valuable confidential information;

- problems associated with various kinds of cheating including:
  - false bids by the bid taker,
  - conspiracies by competing bidders,
  - conspiracies in two-sided markets between bidders offering to sell and those offering to buy, and
  - the use of false-name bids by single bidders;
- the fact that strategies in sequences of strategy-proof auctions may not be strategy-proof; and
- the fact that the process can be revenue deficient.

In the following sections, I will describe these problems and some countermeasures that have been proposed. In a final section, I will sum up the case for VCG processes being impractical mirages, but also indicate that there is some value that might arise from a better understanding of them.

Weak Equilibria

The usual examples of Vickrey auctions involve a priori symmetry. Bidders draw values at random from the same commonly known distribution. However, consider the following situation: Bidder 1 draws a value from the uniform distribution on $[0, 10]$, while bidder 2 draws her value independently from a uniform distribution on $[5, 10]$. It is still a weakly dominant equilibrium strategy for each bidder to bid her value. However, bidder 1 has nothing to lose by bidding 0 (or, equivalently, not bidding) whenever she draws a value of less than 5. With either strategy, bidder 1 wins 25% of the time and makes an expected profit of $5/12 = 0.417$, and bidder 2 wins 75% of the time. However, bidder 2's expected winnings and the expected payment received by the bid taker depend on which of the two alternative equilibrium strategies bidder 1 follows. If bidder 1 bids her true value, the expected payment to the bid taker is $55/12 = 4.583$ and bidder 2's expected profit is $35/12 = 2.917$.[5] If bidder 1 does not bid when her value is less than 5, the bid taker's expected revenue decreases by $5/4$ to $10/3 = 3.333$, and bidder 2's expected profit increases by that amount to $25/6 = 4.167$. All of the extra expected profit of 1.25 comes from the bid taker.

The practical implication of this example is that one cannot expect bidders in VCG processes to make bids that are sure to lose. Because losing bids set the price the winner pays, the bid taker should be very worried about the revenue implications of this and be reluctant to use the VCG process.[6]

Bid Preparation and Bid Communication Costs

Note that the example above assumes that there is no cost of submitting a bid. However, if we add an infinitesimal cost, ε, of submitting a bid, then *not* bidding when her value is less than 5 *must* be part of any equilibrium strategy for bidder 1. The problem is more general than this and does not depend upon asymmetry. Larson and Sandholm (2001) consider a general model of a single-item auction. They assume that the bidders are *deliberative agents*, i.e., that bidders spend costly effort to determine their values (and may spend such effort on estimating the values of other bidders). In one of their models, the more costly the effort that bidders expend, the better their estimates. They show that in this model there are no reasonable dominant strategy equilibria, including in Vickrey auctions. In other words, in the real world in which a bidder must estimate her values and in which the harder she works at it the better she does, the Vickrey auction does not work.

The general VCG process for n items assumes that the bidder will submit information sufficient to determine bids on all $2^n - 1$ nonempty combinations of n items. This exponential growth of the bid preparation effort with the size of the process can present a serious problem for the overall efficiency of the process where the overall efficiency includes bid preparation costs.[7] It also implies the possibility that there is an exponential growth in the amount of information that the bidders must communicate to the bid taker.[8]

Winner Determination Effort

Often, the winner determination problem in VCG processes will be NP-hard. For example, in a combinatorial auction, each of the m bidders must provide information sufficient to determine $2^n - 1$ bids; the bid taker must solve $m + 1$ winner determination problems. One of these can involve as many as $m(2^n - 1)$ nonzero bids, and the other m problems can involve as many as $(m - 1)(2^n - 1)$ nonzero bids. In general, each of these problems is NP complete. This will not be a problem for auctions involving just a few items, but if there are many items, it may well be. Because bids on all possible combinations may be required, the winner determination problem will not, in general, be able to take advantage of the guarantee of computability in any of the special cases described in Rothkopf et al. (1998) unless appropriate bids are guaranteed to be zero. While approximate solution of the winner determination problem in difficult cases may work for some purposes, it will destroy the incentive for truth-revealing strategies that is a key purpose of the VCG process.

Budget Constraints

Capital markets are not perfect and frictionless. Suppose that bidders have budget constraints, i.e., they cannot always arrange the financing to bid their true values. This situation has been studied in Vickrey auctions by Che and Gale (1996, 1998, 2000). More recently, Borgs et al. (2005) showed that budget constraints destroy truthful bidding in VCG processes.[9]

This problem is not restricted to situations in which the bid taker *knows* that one or more of the bidders has a budget constraint. It is also relevant when the bid taker is not sure that none of the bidders has such a constraint. In general, it will be hard for a bid taker to know if bidders are budget constrained. In particular, looking at bids in past auctions will not reveal if a bidder was budget constrained. For example, I was an advisor in an auction for a spectrum license to a bidder who valued the asset at $85 million. He was only able to finance a bid of $65 million. From his bidding, which stopped when the price reached $65 million, the seller would have had no way of knowing that his value was higher.

Information Revelation

Generally, businessmen do not like to reveal their costs or values. Some auction forms do not require them to do so. A bid in a standard sealed-bid auction is not a statement of true value, but an offer. Similarly, no one knows whether the winner of an English auction would have been willing to bid more than she did. Vickrey auctions, however, if they work as they are supposed to, require bidders to reveal their true costs or values. If it were in fact truly in their interest to do so, it might be argued that businessmen would learn to overcome their reluctance. However, is it *really* in their interest to do so? Suppose that a government agency is taking bids for a major construction project using a Vickrey auction with publicly opened bids, that you submit the winning bid of $100 million (your true cost), and that the next best bid is for $150 million. If this were an isolated transaction, you would expect to have a profit of $50 million. Suppose, however, that in order to complete the project you need to negotiate with others—perhaps local licensing boards, banks, construction companies, and labor unions. Your negotiating position will be weak because all of them will know that you have $50 million you don't "need." You would very much wish that you had bid more than your true cost, because having done so would improve your position in such negotiations. In particular, it would be in your interest to bid more in order to keep a greater proportion of your apparent profit.

This is a general phenomenon. Rothkopf et al. (1990) showed that if third parties capture a fraction of the revealed economic gain in a Vickrey auction, bidders will have incentive to shade their bids so that all of that captured gain is, on average, added to their bids. (See also Englebrecht-Wiggans and Kahn 1991.) Thus, the bid taker can be expected to pay extra if he chooses to hold a Vickrey auction.

There may also be political ramifications of the revelation of a large gap between the best bid and the second-best bid. In a government transaction, the government may be hard-pressed to explain why the large extra payment is warranted.[10] It is proof of the sale's lack of competitiveness. A progressive auction would not disclose this. One way to look at a Vickrey auction is that it induces competitors with market power to behave efficiently by paying them the amount they could extract by using their market power. However, if that payment is large and visible, it may well lead to political pressure to control the market power rather than to buy it off. In turn, this possibility may well lead bidders to change their bids.

In a simple Vickrey auction, it may be possible to avoid this information revelation problem by using cryptography (Nurmi and Salomaa 1993; Franklin and Reiter 1996; Kikuchi et al. 1999; Naor et al. 1999; Jakobsson and Juels 2000; Abe and Suzuki 2002; and Brandt 2002, 2003). However, as far as I know, this has yet to be done. In any case, it may require either bonding of bidders or a cryptographic protocol in which no bidder is required to do anything after she discovers that she has lost the auction (Bradford et al. 2004). Of course, the general VCG process is much more complicated than a simple Vickrey auction. I am not aware of any practical general cryptographic protocols for it, although Yokoo and Suzuki (2004) have considered the topic.

Four Kinds of Cheating

Auction theory is generally developed on the assumption that the auction has rules and that these rules will be followed. In practice, some auctions are more subject to cheating than others, and choice of auction form may depend upon resistance to cheating and the reassurance that this resistance gives to participants. Vickrey auctions are relatively susceptible to two different kinds of cheating, and multi-item VCG processes are susceptible to two more kinds.[11]

First of all, as Robinson (1985) pointed out and on which Graham and Marshall (1987) elaborated, Vickrey auctions (and English auctions as well, but not standard sealed bidding or Dutch auctions) are relatively susceptible to conspiracies by bidders. The problem is that in such auctions collusive conspiracies by bidders are stable. Here is an example to illustrate this. Suppose something is to be auctioned off to two bidders, one of whom values it at $100 and the other of whom values it at $90. In a well-functioning auction, the item should sell for about $90. Suppose, however, that the bidder with the higher value agrees to pay the other bidder a substantial bribe if she fails to top the first bidder's initial bid of $10. In a progressive auction or a Vickrey auction, such a conspiracy is stable. In the Vickrey auction, the bidder who is to win bids $100. If the other bidder bids $90, the winner will win anyway, know that the deal has been broken, and not make the payoff. In a standard sealed-bid auction, the bidder who is supposed to win at $10 has to submit a bid at that amount. If the conspirator who was to lose instead submits a bid of $11, she will win. Because the conspiracy is illegal, there is no way for the intended winner to enforce the conspiracy agreement or avoid losing the auction. Hence, such conspiracies are somewhat less likely.

A particular concern in Vickrey auctions is cheating by the bid taker. If the bid taker makes up an artificial bid that is for an amount between the two best bids, it will capture some of the profit due to the winning bidder. For example, if the high bid is $15 and the second highest bid is $10, an artificial bid of $14 will raise the winning bidder's payment by \$4. Having a trusted third party handle the bids is not a complete cure for this problem, because if the bid taker can approximately anticipate the bids by the two highest sincere bidders, it can procure an insincere bid by an ally to cut the winning bidder's profit.[12] This kind of cheating has been modeled in Rothkopf and Harstad (1995), and an instance of cheating by a bid taker in a Vickrey auction is documented in Lucking-Reiley (2000).[13]

Sakurai et al. (1999) pointed out a new vulnerability of the VCG process—"false-name" bids. Consider an auction in which there are two bidders and three items, a, b, and c, for sale. Bidder 1 values $\{a, b, c\}$ at \$2 and bidder 2 values $\{a, b, c\}$ at \$3. Neither values any smaller aggregation of items. In the honestly conducted Vickrey auction, bidder 2 will win all three items and pay \$2. However, suppose that bidder 2 decides instead to submit bids under three different names—2a, 2b, and 2c. Suppose that 2a bids \$1 for $\{a\}$; 2b bids \$1 for $\{b\}$, and 2c bids \$1 for $\{c\}$. Now, the three bidders 2a, 2b, and 2c each win the item they bid on, but none of them have to make any net payment. For example, without bidder 2a, the revenue would be \$2; with him, the apparent value is \$3, so while he pays \$1 for $\{a\}$, he gets a Vickrey payment, i.e., a refund, of his apparent contribution to the surplus, of \$1. A bid taker is potentially vulnerable to such false-name bids unless he has complete control of the identities and contractual arrangements between all of the bidders.

Finally, Hobbs et al. (2000) examine the potential use of the VCG process in electricity markets and pointed out the possibility of conspiracies between suppliers and buyers at the expense of a market maker operating the transmission grid.[14] If an electricity supplier and an electricity user both connect with the transmission grid at the same point, they can increase both of their Vickrey payments from the system operator by artificially increasing their quantities—supply in one case and demand in the other. To detect such cheating, the grid operator would need to monitor each bidder individually and not just the net flow at the point that the two connect to the grid. With a commodity that can be supplied over time (unlike electricity, where supply must meet demand instantaneously), a private deal between the conspirators to recycle the supply would serve the same end and would not be detectable by measuring within-system deliveries.

Sequences of Auctions

All auction forms can have difficulties with sequences of auctions. In theory, a significant advantage of strategy-proof auctions is that a bidder has a dominant strategy and does not have to worry about the economics or strategies of competitors. However, Juda (2005) and Juda and Parkes (2006) have recently pointed out that this does not apply to sequences of otherwise strategy-proof auctions. The following example is from Juda's paper:

> Concrete mixer Alice values acquiring one ton of sand before Wednesday at \$1,000. Bob will hold a Vickrey auction for one ton of sand on Monday and another such auction on Tuesday.

As a practical mater, Alice does not have a dominant bidding strategy in Monday's auction because her value in that auction depends on her assessment of the level of competition in Tuesday's auction. Because economically important auctions tend not to be isolated events, this will often be a serious concern. In making practical auction design trade-offs, the potential for being strategy-proof of VCGs will be lost when the auctions are part of a larger sequence of commercial transactions. Juda and Parkes (2006) suggest the use of options to deal with this.

Revenue Deficiency

As mentioned above, when used for deciding on transactions involving multiple buyers and multiple sellers, the VCG process obtains efficiency by paying both sellers and buyers with market power to behave efficiently. An important problem is that the VCG process provides no source for this revenue. Presumably, it must be raised from some sort of tax or overhead. However, the efficiency effects of that tax or overhead have not been taken into account in the claim that the VCG process is efficient. It has been shown that the VCG process is the least revenue deficient of any process that produces perfect (theoretical) efficiency (Krishna and Perry 1997). However, as far as I know it has not been shown that a process with lower taxes or overheads that was less than perfectly efficient in its allocations would not be more efficient when the taxes or overheads are taken into account.[15] Even leaving aside the other problems discussed above, it seems reasonable to assume that the overall optimum will involve a trade-off of efficiency in allocation against the efficiency in keeping the taxes/overheads down.

There is another sense in which the VCG process can be revenue deficient. The payments can be too low for the process to be stable. To see this, let's return to a variant of the example we used to discuss false-name bids. This time, however, let's assume that the false-name bidders are the real bidders. Thus, bidder 1 values $\{a, b, c\}$ at \$2 and bidder 2a values $\{a\}$ at \$1; bidder 2b values $\{b\}$ at \$1; and bidder 2c values $\{c\}$ at \$1. The result, as we observed above, is that bidders 2a, 2b, and 2c win and each pay nothing. This payment is too low for the result to be stable. It is not in the core of a game involving the bidders and the bid taker. Bidder 1 can go to the bid taker and truthfully say, "Ignore those bidders. I will give you \$1 for $\{a, b, c\}$, and we will both be better off." See Ausubel and Milgrom (2006).

Which of These Reasons Are Most Important?

I have been asked by reviewers to comment on which of these reasons that VCG mechanisms are not practical are most important. In some ways, it is as difficult to answer this question as it would be to say which of 13 bullets that hit a shooting victim contributed most to his demise. In some situations, it would be clear that there are several shots that alone would have been fatal. However, the real answer is "it depends." Some concerns are critical in some contexts and unimportant in others. For example, in government sales of extremely valuable assets, the political repercussions of revealing the gap between large offers and small revenue could be a dominant concern. In other contexts, perhaps an industrial procurement auction with lots of competition and smaller amounts at stake, this could be a minor issue. To quote Paul Klemperer (2002, p. 187), "auction design is *not* 'one size fits all.' "

Conclusions

VCG processes have great theoretical appeal. They are a dominant strategy mechanism. This means that, in theory, a bidder's decision to use the strategy they call for does not depend on what the bidder thinks her competitors' strategies are, and she need spend no effort in trying to find them out or to keep her competitors from learning her strategy. In theory, they promise perfect efficiency. In some circumstances, they produce, in theory, expected revenue equivalent to other common auction forms. They are also mathematically elegant and allow proof of many interesting theoretical results that would be hard to derive for other auction forms. However, they are just not practical. They do not work the way the (simple) theory says they should. We have discussed a baker's dozen of different reasons why in practice they fail to live up to their promise: unfavorable alternative weak equilibria, the nonexistence of dominant strategy equilibria in the face of reasonably modeled bid preparation costs, exponentially growing transaction costs both for bidders and in winner determination, unreliability in the face of possible budget constraints, revelation of information better kept private, susceptibility to serious revenue deficiencies, the failure of dominant strategies in sequences of auctions, and susceptibility to four different kinds of cheating. Note that all of these reasons assume rational behavior. We leave to others the discussion of concerns based on behavioral or bounded rationality factors.

This paper is a comment on the *practicality* of VCG auctions. It is not a comment on the *theoretical* value of knowledge about VCG processes. Because finding equilibrium strategies in combinatorial auctions is extraordinarily difficult except in VCG processes, there may well be useful insights to be had from such knowledge. For example, Mishra and Parkes (2007) analyze an iterative version of the VCG process. Ausubel's analysis of his recent proposal for a dynamic auction for heterogeneous goods (Ausubel 2006) makes use of a comparison with the VCG process, as does Jehiel et al. (2007); so does work on approximations to the VCG process such as Cavallo (2006), Faltings (2004), and Nisan and Ronen (2000). However, real combinatorial auctions are complex. Ultimately, to be useful, research on them needs to deal with the complexities. Research that tries to maintain theoretical dominant strategies at the expense of compromises on other important aspects of the process seems to have limited direct practical value. I certainly do not wish to be absolutist about even this. There will be situations in which a VCG process may be a useful approximation of a complicated situation. For example, Harstad and Rothkopf (2000) showed that the Vickrey auction can, on occasion, be a better approximation of certain realistic progressive auctions than the standard "Japanese" variant of the English auction analyzed by Milgrom and Weber (1982). Furthermore, if no information about competitive bids in an English auction except the current best bid is revealed to bidders—this is often the case in industrial procurement auctions—then the Vickrey auction will often be an excellent model of the English auction. In addition, computerized bidding agents may be able to be programmed to avoid some of the 13 problems discussed here. Nonetheless, there is great need for analysis focused on the phenomenon at issue, not just on the mirage of perfection that the VCG process generally proves to be.

Endnotes

1. The argument is simple. Consider the effect of deviations from a policy of bidding one's true value. If a bidder bids more than her true value, doing so either has no effect or causes her to win. If it causes her to win, she will regret it because she will pay more than her value. Similarly, if she bids less than her true value and this causes her to lose, she will regret it. Nothing about this argument depends on the strategies of other bidders. This makes bidding one's true value a dominant strategy, i.e., one that is best regardless of the strategies of other bidders.
2. Myerson (1981) generalized this proof.
3. This process will work even if the bid taker has restrictions on winning bids (e.g., no bidder may win more than seven licenses) that can be included in the optimization problems.
4. There are important exceptions. For example, Pekec and Rothkopf (2003), Milgrom (2004), and Aususbel and Milgrom (2006) make many of the points made here.
5. There are two equally likely cases: Bidder 1 draws a value on $[0, 5)$ and bidder 1 draws a value on $[5, 10]$. In the first case, bidder 2 always wins and pays an average of $5/2$ if bidder 1 has bid, and 0 otherwise. The second case is symmetric; each bidder wins one half of the time. The average price is $5 + 5/3$ and the average value of the winner is $5 + 10/3$.
6. English auctions without reserve prices also suffer from this problem. Standard sealed bidding and Dutch auctions do not. A bid taker who knows the situation precisely can set a reserve price, five in this example, to protect himself, but assuming that the *bid taker* has such knowledge is problematical. Often, bidders know much more about their business than the bid taker does.

7. Just as use of an English single-item auction may reduce bid preparation costs relative to a Vickrey auction by narrowing the range of relevant values for an item that a bidder needs to consider, the use of an iterative variant of the VCG process can lower the costs of bid preparation. It also may change the incentives of bidders by allowing communication among them through their bids.
8. Often, this burden may be reduced or alleviated by recent work on preference elicitation (e.g., Lahaie and Parkes 2004 and Sandholm and Boutilier 2006) and on concise bidding languages (e.g., Nisan 2006).
9. Budget constraints are an important particular case of a general problem. The desirable properties of the VCG depend on an assumption of quasi linearity. In other words, it assumes that the net value to a bidder of a group of items won with a bid of b is that value less b.
10. This was the case when New Zealand used Vickrey auctions to sell spectrum. In one auction, a firm that bid NZ$100,000 paid only NZ$6. In another, a firm that offered NZ$7,000,000 paid only NZ$5,000. See McMillan (1994).
11. All auctions are subject to some kinds of cheating. Here we discuss the extra vulnerabilities of Vickrey and VCG auctions. Also note that many kinds of cheating are explicitly illegal in some jurisdictions; however, we will not attempt to be attorneys here.
12. Even with shaky information, this can sometimes be a low-risk stratagem. If, through miscalculation, the insincere bid turns out to be the best bid, it may be able to be disqualified.
13. Bidders in standard sealed-bid auctions and Dutch auctions do not face this potential problem. Bidders in English auctions may face shills but have some chance of observing them and reacting during the auction.
14. Use of VCG processes for making markets where both sides bid is sometimes proposed. See, for example, MackieMason and Varian (1995) and McGuire (1997). In general, as discussed below, there are extra issues associated with such use with both sides of the market bidding. In particular, a source of funding for the Vickrey payments needs to be found, and raising these funds may have strategy implications and will normally have efficiency implications.
15. Faltings (2004) and Cavallo (2006) have begun wrestling with this problem.

Acknowledgments

The author is grateful to Ronald Harstad, Aleksandar Pekec, and three attentive anonymous referees for helpful comments. Any remaining errors or heresies are the responsibility of the author.

References

Abe, M., K. Suzuki. 2002. $M+1$st price auction using homorphic encryption. *Proc. 5th Internat. Workshop on the Practice and Theory of Public Key Encryption* (*PKC 2002*), Paris, France. *Lecture Notes in Computer Science*, Vol. 2274. Springer Verlag, 115–224.

Ausubel, L. 2006. An efficient dynamic auction for heterogeneous commodities. *Amer. Econom. Rev.* **96** 602–629.

Ausubel, L. M., P. Milgrom. 2006. The lovely but lonely Vickrey auction. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*. MIT Press, Cambridge, MA, 17–40.

Borgs, C., J. Chayes, N. Immorlica, M. Mahdian, A. Saberi. 2005. Multiunit auctions with budget-constrained bidders. *Proc. Sixth ACM Conf. Electronic Commerce* (*EC'05*), 44–51.

Bradford, P. G., S. Park, M. H. Rothkopf. 2004. Protocol completion incentive problems in cryptographic Vickrey auctions. *Proc. Seventh Internat. Conf. Electronic Commerce Res.* (*ICECR-7*), Dallas, TX, 55–64.

Brandt, F. 2002. A verifiable bidder-resolved auction protocol. R. Falcone, S. Barber, L. Korba, M. Singh, eds. *Proc. First Internat. Joint Conf. Autonomous Agents and Multi-Agent Systems* (*AAMAS 2002*), Bologna, Italy, 18–25.

Brandt, F. 2003. Fully private auctions in a constant number of rounds. *Proc. Financial Cryptography* (*FC 2003*), *Lecture Notes in Computer Science*, Vol. 2742. Springer Verlag, 223–238.

Cavallo, R. 2006. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. *Proc. Fifth Internat. Joint Conf. Autonomous Agents and Multi-Agent Systems* (*AAMAS'06*), Hakodate, Japan.

Che, Y., J. Gale. 1996. Expected revenue of all-pay auctions and first-price sealed bid auctions with budget constraints. *Econom. Lett.* **50** 373–380.

Che, Y., J. Gale. 1998. Standard auctions with financially constrained bidders. *Rev. Econom. Stud.* **65** 1–21.

Che, Y., J. Gale. 2000. The optimal mechanism for selling to a budget-constrained buyer. *J. Econom. Theory* **92** 198–233.

Clarke, E. 1971. Multipart pricing of public goods. *Public Choice* **8** 19–33.

Englebrecht-Wiggans, R., C. M. Kahn. 1991. Protecting the winner: Second-price versus oral auctions. *Econom. Lett.* **35** 243–248.

Faltings, B. 2004. A budget-balanced, incentive-compatible scheme for social choice. *Agent-Mediated E-Commerce* (*AMEC*) *VI*, *Lecture Notes in Computer Science*, Vol. 3435. Springer Verlag, 59–72.

Franklin, M. K., M. K. Reiter. 1996. The design and implementation of a secure auction service. *IEEE Trans. Software Engrg.* **32**(5) 302–312.

Graham, D., R. Marshall. 1987. Collusive bidder behavior at second-price and English auctions. *J. Political Econom.* **95** 1217–1239.

Groves, T. 1973. Incentives in teams. *Econometrica* **41** 617–631.

Groves, T., J. Ledyard. 1977. Some limitations on demand revealing processes. *Public Choice* **29**(2) 107–124.

Harstad, R. M., M. H. Rothkopf. 2000. An "alternating recognition" model of English auctions. *Management Sci.* **46** 1–12.

Hobbs, B. F., M. H. Rothkopf, L. C. Hyde, R. P. O'Neill. 2000. Evaluation of a truthful revelation auction for energy markets with nonconcave benefits. *J. Regulatory Econom.* **18**(1) 5–32.

Jakobsson, M., A. Juels. 2000. Mix and match: Secure function evaluation via ciphertexts. T. Okamoto, ed. *Advances in Cryptography* (*ASIACRYPT'00*)*, Lecture Notes in Computer Science*, Vol. 1976. Springer Verlag, 162–177.

Jehiel, P., M. Meyer-Ter-Vehn, B. Moldovanu. 2007. Mixed bundling auctions. *J. Econom. Theory*. Forthcoming.

Juda, A. I. 2005. The sequential auction problem: An analysis and solution. Mimeo, Harvard Business School, Boston, MA.

Juda, A. I., D. Parkes. 2006. The sequential auction problem on eBay: An empirical analysis and a solution. *ACM Conf. Electronic Commerce EC-06*,

Kikuchi, H., M. Harkavy, J. D. Tygar. 1999. Multi-round anonymous auction protocols. *TIEICE: Trans. Comm./Electronics/Inform. and Systems*, 62–69.

Klemperer, P. 2002. What really matters in auction design. *J. Econom. Perspect.* **16** 169–189.

Krishna, V., M. Perry. 1997. Efficient mechanism design. Working paper, Pennsylvania State University, University Park, PA.

Lahaie, S., D. C. Parkes. 2004. Applying learning algorithms to preference elicitation. *Proc. Fifth ACM Conf. Electronic Commerce*, 180–188.

Larson, K., T. Sandholm. 2001. Costly valuation calculation in auctions. *Proc. Theoret. Aspects Rationality and Knowledge* (*TARK VIII*), Siena, Italy, 169–182.

Lucking-Reiley, D. 2000. Vickrey auctions in practice: From nineteenth-century philately to twenty-first century e-commerce. *J. Econom. Perspect.* **14**(2) 183–192.

Mackie-Mason, J. K., H. Varian. 1995. A spatial "smart market" for electric power and transmission. Working paper, Department of Economics, University of Michigan, Ann Arbor, MI.

McGuire, B. 1997. Power auctions and intertemporal production costs: Dealing with unit commitment. Mimeo, Energy Institute, University of California at Berkeley, Berkeley, CA.

McMillan, J. 1994. Selling spectrum rights. *J. Econom. Perspect.* **8**(3) 145–162.

Milgrom, P. 2004, *Putting Auction Theory to Work*. Cambridge University Press, Cambridge, UK.

Milgrom, P., R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50** 1089–1122.

Mishra, D., D. Parkes. 2007. Ascending price Vickrey auctions for general valuations. *J. Econom. Theory*, **132** 335–366.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6** 58–73.

Naor, M., B. Pinkas, R. Sumner. 1999. Privacy preserving auctions and mechanism design. *Proc. First ACM Conf. Electronic Commerce*, ACM, 129–139.

Nisan, N. 2006. *Bidding Languages for Combinatorial Auctions*. P. Cramton, Y. Shoham, R. Steinberg, eds. MIT Press, Cambridge, MA, 215–232.

Nisan, N., A. Ronen. 2000. Computationally feasible VCG mechanisms. *ACM Conf. Electronic Commerce 2000 EC-2000*.

Nurmi, H., A. Salomaa. 1993. Cryptographic protocols for Vickrey auctions. *Group Decision and Negotiation* **4** 363–373.

Pekec, A., M. H. Rothkopf. 2003. Designing combinatorial auctions. *Management Sci.* **49** 1485–1503.

Robinson, M. S. 1985. Collusion and choice of auction. *RAND J. Econom.* **16** 141–145.

Rothkopf, M. H., R. M. Harstad. 1995. Two models of bid-taker cheating in Vickrey auctions. *J. Bus.* **68** 257–267.

Rothkopf, M. H., A. Pekec, R. M. Harstad. 1998. Computationally manageable combinational auctions. *Management Sci.* **44** 1131–1147.

Rothkopf, M. H., T. J. Teisberg, E. P. Kahn. 1990. Why are Vickrey auctions rare? *J. Political Econom.* **98** 94–109.

Sakurai, Y., M. Yokoo, S. Matsubara. 1999. An efficient approximate algorithm for winner determination in combinatorial auctions. *Proc. Second ACM Conf. Electronic Commerce* (*EC-00*), ACM Press, New York, 8–37.

Sandholm, T., C. Boutilier. Preference elicitation in combinatorial auctions. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*. MIT Press, Cambridge, MA, 233–264.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16** 8–37.

Yokoo, M., K. Suzuki. 2004. Secure generalized Vickrey auctions without third-party servers. *Eighth Internat. Financial Cryptography Conf.* (*FC-2004*).