Price-Based Information Certificates for Minimal-Revelation Combinatorial Auctions

Citation

Parkes, David C. 2002. Price-based information certificates for minimal-revelation combinatorial auctions. In Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems, ed. J. A Padget, 253-276. Berlin; New York: Springer. Previously published in Lecture Notes in Computer Science 2531: 253-276.

Published Version

doi:10.1007/3-540-36378-5

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4101006>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Price-Based Information Certificates for Minimal-Revelation Combinatorial Auctions

David C. Parkes

Division of Engineering and Applied Sciences, Harvard University,
33 Oxford Street, MA 02138, USA
parkes@eecs.harvard.edu
<http://www.eecs.harvard.edu/~parkes>

Abstract. The equilibrium of the direct-revelation Vickrey-Clarke-Groves (VCG) mechanism for combinatorial auctions requires every agent to provide complete information about its preferences. Not only is this unreasonable in settings with costly preference evaluation, but moreover it is often possible to implement the same outcome with *incomplete* preference information. We formalize the problem of minimal-revelation mechanism design and introduce the concept of minimal information certificates. Linear programming duality theory proves that a class of price-based information certificates are both necessary and sufficient to compute an efficient allocation in a dominant strategy mechanism.

1 Introduction

Mechanism design addresses the problem of decision making in distributed systems in which agents are self-interested, and each agent has private information about its preferences across different outcomes. Mechanism design assumes that agents will behave as individual-rational utility-maximizers, and solves for a game in which the joint actions of self-interested agents implements a desirable system-wide outcome [18].

It is traditional in mechanism design to ignore computational considerations, and focus on *direct revelation mechanisms*, in which agents report all of their private information and the mechanism solves a central optimization problem and implements a particular outcome. However, computational considerations are important when mechanisms are applied to solve hard problems, for example combinatorial allocation problems [33]. Moreover, there is often a subtle interaction between computational and economic properties, with naive approximations changing the economic properties of a mechanism [20, 23]. Computational mechanism design [27, chapter 3] augments the traditional mechanism design focus on incentives with a careful integration of computational concerns.

Combinatorial auctions are important mechanisms in future agent-mediated electronic commerce applications. For example, in a procurement setting within the supply chain a combinatorial auction allows competing suppliers to express volume discounts for service to multiple locations. Focusing on the forward direction, a combinatorial auction has a seller with a set of heterogeneous items,

such as hotel rooms and flights, and multiple buyers that can submit bids on arbitrary bundles of items. Combinatorial auctions present a number of interesting computational problems. The winner-determination problem is NP-hard, equivalent to a weighted set-packing problem. In addition to characterizing tractable special-cases [33, 12], methods have been developed to introduce approximations but maintain strategic properties of combinatorial auctions [23]. The *communication complexity* is another problem, because the number of bundles is exponential in the number of items and an agent can have quite general preferences across bundles. The worst-case communication complexity of a fully-efficient combinatorial auction is exponential, both for direct- and indirect-revelation mechanisms [24], although indirect mechanisms can achieve better average-case performance. One approach to address the communication complexity is to develop structured bidding languages that are compact for particular agent preferences [22, 6]. Another approach is to place explicit restrictions on the expressivity of a language, while being careful to maintain the equilibrium properties of a mechanism [16, 32].

The *valuation problem* is another equally important problem, but one that has received less attention. There are many electronic commerce applications in which it is costly for a bidder to evaluate precise tradeoffs across all possible outcomes. In a combinatorial auction there are an exponential number of bundles, and the valuation problem for any single bundle can be an NP-hard problem [34, 26]. For example, in a shipping logistics problem a bidder might need to solve a local optimization problem to determine its cost to provide a particular schedule of pickup and drop-offs. Moreover, the valuation problem is often the one stage in negotiation that is difficult to automate within electronic markets, often requiring additional information from management. Mechanism design can reduce the complexity of the valuation problem through careful *preference elicitation* [26]. Direct-revelation mechanisms represent one extreme, in which a bidder must compute an exact value for all outcomes to follow an equilibrium strategy. Indirect mechanisms, in which agents respond to dynamic feedback, such as price information, can solve the same problem as a direct-revelation mechanism without agents revealing or computing exact information about their preferences.

Preference elicitation is not a problem that is readily solved by introducing more expressive bidding languages into direct-revelation mechanisms [22]. Although expressive languages can provide a more compact representation of preferences, information transfer from agents to the mechanism in direct-revelation mechanisms remains oblivious to the local valuations of other agents, and bidders (or bidding agents) must still perform enough computation and/or information discovery to allow their bid to completely characterize their preferences. Moreover, bidding *oracles* [22, 32], or programs, within direct-revelation mechanisms, move the problem to the mechanism infrastructure without solving the problem. Let us suppose for the moment that it is cheap to *define* the oracle (while evaluation of the oracle to compute exact preferences is costly), and that problems of trust involved in sending a bidding program to an auctioneer can be overcome. The fundamental problem of designing a minimal-revelation mecha-

nism still remains, because the computational burden of preference elicitation is simply shifted to the mechanism infrastructure. For this reason, we focus here on the core problem of minimizing the amount of preference information that a bidding agent must provide, and we will not care whether this agent is a bidding program submitted by the bidder to the mechanism, or simply a bidder or a bidding agent that interacts with the mechanism.

Indirect mechanisms allow for *adaptive* preference elicitation by a mechanism because agents can compute strategies in response to information provided by the mechanism. The information, for example prices in an ascending-price auction, defines the equilibrium path of a mechanism, and agent strategies along that path provide incremental information about preferences for different outcome. Thus, it is possible in an indirect mechanism, but impossible in a direct mechanism, to elicit just that information that is required to determine the optimal outcome and no more. For example, in an English auction it is sufficient for the two bidders with the highest value to bid up the price until only one bidder, the winner, is left in the auction. The winner reveals a lower-bound on its value for the item, the competing bidder reveals its exact value, while the other losers reveal only upper-bounds on their values.

In this paper, we formalize the problem of minimal-revelation mechanism design and introduce the concept of minimal information certificates. Our main results are for the class of allocatively-efficient combinatorial auctions, that allocate items to maximize the total value across agents. First, we show that the problem of computing an efficient allocation with truthful agents is *informationally equivalent* to the problem of computing *competitive equilibrium* (CE) prices for a large class of query languages. Second, we quantify the informational-cost of incentive-compatibility, or *cost-of-truthfulness*, and show that the problem of computing an efficient allocation with self-interested agents is informationally equivalent to the problem of computing *Universal* CE prices. From this, we show that there is a non-zero cost-of-truthfulness unless a technical condition, *agents-are-substitutes* holds, which requires that items are more like substitutes than complements.

The competitive equilibrium prices in a combinatorial auction can be both non-linear, such that the price on a bundle of items is not equal to the sum over the constituent items of the bundle, and non-anonymous, such that different agents face different prices for the same bundle. Universal CE prices require that the same condition also holds on the auction problem induced by removing each agent in turn from the auction. In competitive equilibrium every agent prefers the bundle it receives in the efficient allocation and the efficient allocation maximizes the revenue for the seller. Although it is immediate from linear programming duality theory that CE prices provide sufficient information to compute an efficient allocation, the *necessary* direction of the main result is novel. This necessary direction requires some additional assumptions about the structure of the preference-elicitation language, which are in fact met by current proposals for indirect combinatorial auctions. The characterization of the informational cost-of-truthfulness is also novel.

The informational equivalence between the problem of computing CE prices and the problem of computing an efficient allocation sheds new light on the preference-elicitation properties of ascending-price auctions such as iBundle [25, 29] and iBundle Extend & Adjust [31]. More generally, the equivalence provides formal motivation for the continued study of primal-dual based methods for the design of indirect mechanisms because primal-dual algorithms have natural interpretations as ascending-price auctions and terminate with CE prices [29, 4].

2 Preliminaries: Efficient Combinatorial Auctions

Before moving on to introduce a formal framework for minimal-revelation mechanism design it is useful to describe a particular mechanism design problem. This will provide some context to the discussion. In this section, we introduce the *efficient* combinatorial auction problem. This problem is used to illustrate the minimal-revelation framework in Section 3, and is also the problem for which we derive our main results. We provide only a bare description of the combinatorial auction design problem. Many details of practical auction design, such as activity rules [21], winner-determination algorithms [35, 1], bidding languages, and collusion [2] are necessarily left undiscussed.

In a combinatorial auction there is a set of $\mathcal{G}$ discrete items to allocate to $\mathcal{I}$ agents, $i \in \{1, \dots, N\}$. Each agent has a valuation function $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}_+$, that defines its value $v_i(S) \geq 0$ for bundles of items, $S \subseteq \mathcal{G}$. This is the *private values* auction model, in which an agent's value is independent of the values of other agents. We assume *free-disposal*, such that $v_i(S) \leq v_i(S')$ for all $S' \supset S$. An additional assumption is that agents have *quasi-linear* utility functions, such that $u_i(S, p) = v_i(S) - p$, is agent i 's utility for bundle S at price p . This is a common assumption in the auction literature, and tantamount to assuming *risk-neutral* agents.

An allocation, $S = (S_1, \dots, S_N)$ is *feasible*, written $S \in \mathcal{F}$, if no two agents receive the same item, i.e. $S_i \cap S_j = \emptyset$ for all $i \neq j$, and if $S_i \subseteq \mathcal{G}$ for all agents. We focus on the efficient auction design problem, in which the goal as designer is to implement the allocation that maximizes the total value across all agents. The efficient allocation, S^* , solves:

$$S^* = \arg \max_{S \in \mathcal{F}} \sum_{i \in \mathcal{I}} v_i(S_i) \quad [\text{CAP}]$$

An efficient combinatorial auction is a mechanism that solves CAP in equilibrium, such that self-interested agents with private information about their valuations submit bids that implement the efficient allocation when the auction is cleared.

In the classic Vickrey-Clarke-Groves (VCG) [36, 7, 15] mechanism for the combinatorial auction, or *Generalized Vickrey Auction* (GVA), every agent reports a claim about its valuation function to the auctioneer. The GVA implements the allocation that maximizes the total reported value, and computes a payment, $p_{\text{gva}, i}$, for each agent, i , which can be less than the agent's bid price.

Fig. 1. A Price-Based Auction with Proxy-Bidding-Agents

The GVA is strategyproof, with a dominant-strategy equilibrium in which each agent reports its truthful valuation function. In equilibrium the GVA is efficient, and each agent's payment is computed as:

$$p_{\text{gva},i} = v_i(S_i^*) - [V(\mathcal{I}) - V(\mathcal{I} \setminus i)]$$

where $V(\mathcal{I})$ is the value of the efficient allocation, $V(\mathcal{I} \setminus i)$ is the value of the efficient allocation computed without agent i in the auction, and S^* is the efficient allocation. Each agent pays the minimal price it could have bid with hindsight and still received the same bundle, given the bids submitted by other agents. In the special case of a single item the GVA reduces to the Vickrey auction, a second-price sealed-bid auction. The GVA is a direct-revelation mechanism, with complete elicitation of agent preferences.

Many indirect mechanisms have been proposed for the efficient combinatorial auction problem. These are typically price-based auctions, in which agents must bid in response to price information provided by the auction to guide the bidding process [13, 37, 25], although a number of non-price based methods have also been proposed [3, 9]. Indirect mechanisms, such as ascending-price auctions, allow *adaptive* information elicitation and have been demonstrated to outperform direct mechanisms in settings with hard valuation problems and costly preference elicitation [26, 8]. The information feedback provided in an indirect mechanism guides the valuation work of participants, enabling them to make better decisions about which outcomes to refine their value over.

2.1 Proxy Bidding Agents

Proxy bidding agents provide a useful framework within which to understand the goals of minimal-revelation mechanism design. Figure 1 illustrates a proxy-bidder based framework for an iterative VCG-based combinatorial auction [30].

Proxy bidding agents sit between the auction and the bidders, and maintain partial information about agent preferences and compute an equilibrium strategy in response to price information from the auction. In an ascending-price Vickrey auction the equilibrium strategy is simply *myopic best-response* (MBR), or straightforward bidding [30, 31]. In MBR, an agent bids in each round for the bundle(s) that maximize its payoff given the current prices. The role of each proxy agent is to elicit just enough preference information to be able to follow an equilibrium bidding strategy. A MBR strategy can be computed with incomplete information about preferences. For example, it is sufficient to have lower- and upper-bounds on the values of bundles to determine the bundles in the best-response set. This is discussed in more detail in Parkes [27, chapter 8].

Related work proposes an alternative model for partial-revelation VCG mechanisms, in which preference queries are not driven by a search for competitive equilibrium prices [9, 11]. The model consists of a single *elicitor agent*, which acts on behalf of the mechanism and asks a sequence of explicit questions of agents. Example queries include *rank*-queries, e.g. “do you prefer bundle S_1 or S_2 ?”, and *value*-queries, e.g. “what is your value for bundle S_1 ?”, queries. In this collection, Hudson & Sandholm [17] present experimental results that compare the effectiveness of different preference elicitation properties. In addition, Conen & Sandholm [10], propose a *differential elicitation* method to implement VCG mechanisms. Differential elicitation is a price-based approach, although prices need not be ascending. Agents are asked to provide MBR information across a pair of bundles, as the price difference across the bundles is adjusted.

While ascending-price combinatorial auctions scale to many hundreds of items and many hundreds of bids, these direct elicitation methods currently scale to only tens of items. Moreover, the new results presented in this paper, which characterize the minimal information required to implement an efficient combinatorial auction, also apply to these non price-based elicitation methods.

2.2 Competitive Equilibrium Prices

Competitive equilibrium (CE) prices have an important role in formalizing conditions for minimal preference elicitation in VCG mechanisms, and also in constructing minimal revelation mechanisms.

Let price, $p_i(S) \geq 0$, denote the price for bundle $S \subseteq \mathcal{G}$ to agent i . In general, prices can be non-linear, such that $p_i(S_1) + p_i(S_2) \neq p_i(S_1 \cup S_2) - p_i(S_1 \cap S_2)$ for bundles S_1 and S_2 , and non-anonymous, such that $p_i(S) \neq p_j(S)$, for agents $i \neq j$.

Competitive equilibrium (CE) prices, p_{ce} , satisfy conditions (CS1) and (CS2) with respect to the efficient allocation, S^* :

$$S_i^* = \arg \max_{S \subseteq \mathcal{G}} [v_i(S) - p_i(S)] \quad (\text{CS1})$$

$$S^* = \arg \max_{S \in \mathcal{F}} \sum_{i \in \mathcal{I}} p_i(S_i) \quad (\text{CS2})$$

In competitive equilibrium, every buyer receives a bundle that maximizes its payoff at the prices (CS1), and the allocation maximizes the seller’s revenue at the prices (CS2). Competitive equilibrium prices always exist in the combinatorial auction problem, although they must sometimes be both non-linear and non-anonymous [5]. Moreover, if an allocation is supported in CE then it must be efficient. These results come directly from linear-programming (LP) duality theory, with respect to a hierarchy of LP formulations for CAP. Each formulation introduces additional variables and constraints to strengthen the natural LP relaxation of a mixed-integer formulation of CAP and achieve integrality.

In Section 4 we use CE prices, and a restricted class, *Universal CE prices* to define minimal-information certificates to compute efficient allocations in dominant-strategy mechanisms.

Definition 1. *Universal CE prices, p_{uce} , are prices that are in competitive equilibrium for CAP defined across all agents, $\mathcal{I}$, and also for CAP defined with each agent removed from the auction in turn.*

As an example, consider a single item allocation problem with agents values, $v_1 = 4, v_2 = 8, v_3 = 10$. Non-anonymous CE prices satisfy, $p_3 \geq \max(p_1, p_2)$, $p_3 \leq 10$, $p_2 \geq 8$, and $p_1 \geq 4$. However, prices $p_1 = 6, p_2 = 9, p_3 = 9$ are not Universal CE prices because $p_1 = 6, p_2 = 9$ are not CE prices for the problem with agents $v_1 = 4, v_2 = 8$. Universal CE prices also require $p_2 \leq 8$ and $p_1 \leq p_2$. Combining these constraints, Universal CE prices must satisfy, $4 \leq p_1 \leq 8, p_2 = 8$, and $8 \leq p_3 \leq 10$.

3 Minimal-Revelation Mechanism Design

In this section we introduce a formal framework for minimal-revelation mechanism design. In particular, we formalize minimal preference information with respect to a partial order defined over the space of possible agent valuation functions. We introduce the idea of an information *certificate* for a statement about an aggregate property of agent preferences, such as “allocation $(S_1, \dots, S_N)$ is efficient”. A certificate provides sufficient information about agent preferences to verify the correctness of the claim. A *minimal certificate* is a certificate that contains no additional information about agent preferences beyond that implied by the statement itself.

The objective in minimal-revelation mechanism design is to elicit a minimal certificate for the optimal outcome, such as an efficient allocation, while also satisfying the standard requirements of incentive-compatibility and individual-rationality. We describe the problem in terms of agent valuations, but the discussion can be easily recast in more general terms by thinking of agent *types* instead of valuations.

3.1 Information Sets, Statements, and Certificates

Let V_i define the abstract set of all possible valuation functions for agent i , and let $V = V_1 \times \dots \times V_N$ define the joint product across all agents. Also, let K

denote the set of possible outcomes. For example, in a combinatorial auction the set K corresponds to the set of feasible allocations.

Definition 2. A *preference information set*, $\text{inf}_i \subseteq V_i$ is a space of possible agent valuations that includes the true valuation, $v_i \in \text{inf}_i$.

Similarly, let $\text{inf} \subseteq V$ denote a *joint* preference information set, such that $v = (v_1, \dots, v_N) \in \text{inf}$. Preference elicitation can be formalized as the process of refining beliefs about agent preferences. It is useful to define a partial order, $\preceq$, across information sets, with $\text{inf}_a \preceq \text{inf}_b \Leftrightarrow \text{inf}_a \supseteq \text{inf}_b$. Information set, inf_a , provides less information about preferences than inf_b if $v \in \text{inf}_b \Rightarrow v \in \text{inf}_a$. In a given problem instance the information sets form a complete lattice with respect to $\preceq$ with minimal element, $\text{inf} = V$, and maximal element $\text{inf} = (v_1, \dots, v_N)$.

We now define *statements* and *certificates* for statements.

Definition 3. A *statement*, $X : V \rightarrow \{\text{T}, \text{F}\}$, defines a function from agent values to a truth value.

Examples of statements include: “allocation $(\emptyset, AB)$ is efficient”, “ $v_1(A) = 10$ ”, and “ $v_1(A) > v_2(B)$ ”.

Definition 4. Information set, inf , is a *certificate* for X , written $\text{cert}(\text{inf}, X)$, if the information set satisfies $v \in \text{inf} \Rightarrow X(v)$.

In words, an information set provides a certificate for a statement if the statement holds for all possible preferences consistent with the information set. Only true statements can have certificates because an information set must contain the true preferences of agents.

Definition 5. Information set, inf , is a *minimal certificate* for X , written $\text{cert}_{\min}(\text{inf}, X)$, if the information set is both a certificate, and in addition if $X(v) \Rightarrow v \in \text{inf}$ for all $v \in V$.

The minimal certificate, with respect to partial ordering, $\preceq$, defines the *maximal* space of valuations that are consistent with statement X .

Definition 6. The *information content*, $\text{inf}(X)$, of statement X , is the *minimal certificate* for X .

The information content of a statement, X , is the minimal information implied about agent preferences by the statement. Equivalently, the information content of a statement X captures the weakest set of constraints on agent preferences that still imply X .

Figure 2 illustrates these ideas. The actual preferences, $v \in V$, of the agents are represented by the single point in the domain of preferences. Notice that $v \in \text{inf}_1$ and $v \in \text{inf}_2$, but that information $\text{inf}_1 \preceq \text{inf}_2$ and only inf_2 is a certificate for X . Information set, inf_2 , is not a minimal certificate, because $\text{inf}_2 \subset \text{inf}(X)$.

Fig. 2. Information sets and certificates. The space, $\text{inf}(X)$, inside the inner circle represents the information-content of statement X , or equivalently the minimal certificate for X . Information set inf_2 , but not information set, inf_1 , provides a certificate for statement X .

3.2 Query Languages

A *query language*, $\mathcal{L}$, characterizes the kinds of questions that can be asked in a mechanism about agent preferences. A query language provides a method to refine the preference information set, adding new constraints to represent responses from agents. The definition of the information provided in response to a query assumes *truthful responses*, and will be used in the context of incentive-compatible mechanism design.

Definition 7. A *query language*, $\mathcal{L}$, defines queries, $Q \in \mathcal{L}$, that provide information sets, $Q(v_i)$, about agent preferences when evaluated by agent i , with valuation v_i .

As an example, the query “what is your valuation?”, elicits the singular information set, $\text{inf}_i = \{v_i\}$, from agent i , while the query “is your value for A greater than your value for B ?” elicits the information set that contains all valuations for which $v_i(A) > v_i(B)$ if the response is “yes”, or the information set that contains all valuations for which $v_i(B) \geq v_i(A)$, otherwise.

Let $\text{inf}(Q)$ denote the information set that is implied by the response of an agent to query Q . The information content of a sequence of queries, $Q_0, \dots, Q_t$, is simply:

$$\text{inf}(Q_0, \dots, Q_t) = \text{inf}(Q_0) \cap \dots \cap \text{inf}(Q_t)$$

In Section 4.1 we provide concrete examples of a query language for a combinatorial auction.

Fig. 3. Preference-information sets, for an allocation of a single item and two agents. (a) This information set, $\{v_1 > v_2\}$, is not agent-independent. (b) This information set, $\{v_1 \geq 4, v_2 \leq 4\}$, is agent-independent.

3.3 Information Set Refinements

We introduce two refinements on information sets. The first refinement places restrictions on joint information sets, and is stated in terms of constraints and relationships *across* multiple agents.

Definition 8. *Preference information, inf , is **agent-independent**, if there is an equivalent decomposition, $(\text{inf}_1, \dots, \text{inf}_N)$, for which $\text{inf} = \text{inf}_1 \cap \dots \cap \text{inf}_N$.*

In words, an agent-independent preference information set can be captured exactly as the intersection over individual agent preference information sets. Thinking in terms of constraints on preferences, agent-independence implies that the constraints on joint preferences can be expressed as the union of constraints stated on individual agent preferences. Agent-independence does not restrict the types of agent preferences in the private values model.¹ Instead, it restricts the types of query languages. For example, an information set that follows from information “the value of agent 1 for choice k_1 is greater than the value of agent 2” can not be captured as an agent-independent information set.

As an example, information set, $\text{inf}_c = \{v_1(S) \geq v_2(S)\}$, for agents 1 and 2, is not agent-independent because agent 1’s valuation is stated in relation to agent 2’s valuation. As another example, Figure 3 illustrates agent-independence in a single-item allocation problem, with two agents with values $v_1 = 6$ and $v_2 = 2$. The information is agent-independent in the right-hand plot, but not the left-hand plot.

The second refinement requires *outcome-independent* information, and restricts the space of feasible information sets for a particular agent. As a preliminary, let $\Delta_{\max}(k', k, \text{inf}_i)$ denote the maximal possible difference in agent

¹ However, *correlated-value* and *common-value* preferences, in which agent valuations depend on the valuations of other agents, are not compatible with agent-independent information sets.

i 's value for choices k' and k , given a preference information set. Formally, $\Delta_{\max}(k', k, \text{inf}_i) = \max_{v_i \in \text{inf}_i} [v_i(k') - v_i(k)]$.

Definition 9. Given information set, inf_i , the **worst case preferences**, $\text{WC}(\text{inf}_i, k) \succeq \text{inf}_i$, for outcome $k \in K$, are the set of valuations, $v \in \text{inf}_i$, that satisfy:

$$\begin{aligned} \max_{v \in \text{inf}_i} \sum_{k' \neq k, k' \in K} v_i(k') & \quad (\text{WC}) \\ \text{s.t. } v_i(k) &= \min_{v'_i \in \text{inf}_i} v'_i(k) \end{aligned}$$

In words, worst-case preferences, $\text{WC}(\text{inf}_i, k)$, contain all valuations consistent with information, inf_i , that maximize the total value to agent i for all choices other than k across all valuations that also minimize agent i 's value for choice k .

Definition 10. Information, inf , is **outcome-independent** when

$$v_i \in \text{WC}(\text{inf}_i, k) \Rightarrow v_i(k') - v_i(k) = \Delta_{\max}(k', k, \text{inf}_i), \quad \forall k' \neq k$$

for all choices, $k \in K$, and for all agents, $i \in \mathcal{I}$.

In words, preference information is outcome-independent whenever it is always possible (for all choices, k , and all agents, i) to select a valuation consistent with preference information that simultaneously minimizes the value on choice, k , to agent, i , and maximizes the difference in value to agent i between every other choice and that choice. In other words, the solution to (WC) must simultaneously maximize the difference in value between choice k and all other choices, for all k .

Consider the following examples from combinatorial auctions.

Example 1. Information set, $\text{inf}_d = \{v_1(A) \in [5, 10], v_1(B) \leq v_1(A) + 6, v_1(AB) \leq v_1(A) + v_1(B)\}$, is not outcome-independent. Valuation $v'_1(A) = 5, v'_1(B) = 11, v'_1(AB) = 16$ satisfies $v'_1 \in \text{WC}(\text{inf}_d, A)$, but $\Delta_{\max}(AB, A, \text{inf}_d) = 16$, with valuation $v''_1(A) = 10, v''_1(B) = 16, v''_1(AB) = 26$, which satisfies $v''_1 \in \text{inf}_d$.

Example 2. Information set, $\text{inf}_e = \{v_1(A) \in [5, 10], v_1(B) \leq v_1(A) + 6, v_1(AB) \leq v_1(A) + 5, v_1(AB) \leq v_1(B) + 5\}$, is outcome-independent. Valuation, $v'_1(A) = 5, v'_1(B) = 11, v'_1(AB) = 10$, satisfies $v'_1 \in \text{WC}(\text{inf}_e, A)$, and $\Delta_{\max}(AB, A, \text{inf}_e) = 5$, $\Delta_{\max}(B, A, \text{inf}_e) = 6$. Valuation, $\tilde{v}_1(A) = 10, \tilde{v}_1(B) = 0, \tilde{v}_1(AB) = 5$ satisfies $\tilde{v}_1 \in \text{WC}(\text{inf}_e, B)$, and $\Delta_{\max}(AB, B, \text{inf}_e) = 5$, $\Delta_{\max}(A, B, \text{inf}_e) = 10$. Valuation, $\hat{v}_1(A) = 10, \hat{v}_1(B) = 16, \hat{v}_1(AB) = 0$ satisfies $\hat{v}_1 \in \text{WC}(\text{inf}_e, AB)$, and $\Delta_{\max}(B, AB, \text{inf}_e) = 16$, $\Delta_{\max}(A, AB, \text{inf}_e) = 10$.

3.4 Minimal Information Revelation

We can now formally define the minimal revelation mechanism design problem. Note that these definitions are presented *ex post*, and such provide a strong (perhaps too strong) requirement for an indirect mechanism that must elicit preferences *online* without access to an oracle.

Introduce an *information cost function*, $c(\text{inf}) \geq 0$, to capture the total cost to agents to provide an information set. It is reasonable to assume monotonic cost functions, such that $\text{inf}_a \preceq \text{inf}_b \Rightarrow c(\text{inf}_a) \leq c(\text{inf}_b)$. In other words, it cannot be cheaper to provide more information than less.² In this paper, we assume that this cost information, or some useful heuristic, is known to the mechanism designer.

Definition 11. *The **ex post minimal-revelation mechanism design problem**, for statement X , is to elicit the minimal cost certificate:*

$$\begin{aligned} \underline{\text{inf}}(X, v) &= \min_{\text{inf} \subseteq V} c(\text{inf}) && [\text{MRMDP}(X, c)] \\ \text{s.t. } \text{inf}(X) &\preceq \text{inf} \\ &(\text{IC}), (\text{IR}), \dots \end{aligned}$$

where $c(\cdot)$ is the information-cost function and (IC) and (IR) represent the incentive-compatibility and individual-rationality constraints of mechanism design.

Incentive-compatibility (IC) constraints in a direct-revelation mechanism, make truth-revelation a game-theoretic equilibrium, such that an agent cannot improve its expected payoff by misstating its valuation information. Informally, the idea in an indirect revelation mechanism, in which the strategy space allows an agent to provide partial information about its preferences, is to make truthful, but partial, revelation a game-theoretic equilibrium. Similarly, individual-rationality (IR) constraints ensure that an agent's expected payoff in equilibrium is non-negative. The VCG mechanism is an example of an (IC) and (IR) direct-revelation mechanism. The next section discusses indirect variations on the VCG mechanism that retain (IC) and (IR) properties. We will find that (IC) can impose additional requirements on the information that is required to solve a problem. For example, in the Vickrey auction we need exact information about the second-highest value in order to implement the Vickrey price.

We also consider a couple of additional constraints on the minimal-revelation mechanism design problem. The first variation introduces a language-based constraint, that seeks to elicit minimal information relative to a particular language.

² We allow an agent to provide more information for the *same* cost, for example when the best method for a bidder to value a particular bundle is to value a number of smaller bundles, in which case the additional information about the value of those smaller bundles is available for free.

Definition 12. *The language-constrained minimal-revelation mechanism design problem, for statement X , is to select the minimal cost queries from language $\mathcal{L}$ that elicit a certificate for statement X .*

$$\begin{aligned} \inf_{\mathcal{L}}(X, v) &= \min_{Q \in \mathcal{L}} c(Q(v)) & [\text{MRMDP}(\mathcal{L}, X, c)] \\ \text{s.t. } \quad \inf(X) &\preceq Q(v) \\ &(\text{IC}), (\text{IR}), \dots \end{aligned} \tag{1}$$

where $Q \in \mathcal{L}$ denotes a set of queries in the language, $c(\cdot)$ is the information cost function, and $Q(v)$ is the information implied by the response to the query.

This language-constrained variation is useful when the language is *coarse* and does not allow precise minimal preference elicitation. The second variation introduces an information-set based constraint, and requires agent- and outcome-independent information sets.

Definition 13. *The information-constrained minimal-revelation mechanism design problem, for statement X , is to select the minimal cost agent- and outcome-independent information sets that solve:*

$$\begin{aligned} \min_{\inf_1 \subseteq V_1, \dots, \inf_N \subseteq V_N} & c(\inf_1 \cap \dots \cap \inf_N) \\ \text{s.t. } \quad \inf(X) &\preceq \inf_1 \cap \dots \cap \inf_N \\ &(\text{IC}), (\text{IR}), \dots \end{aligned}$$

where the information sets, $\inf_i$, are also restricted to outcome-independent information sets.

This information-constrained variation is useful because it simplifies the information equivalence analysis in the next section.

4 Minimal-Revelation Combinatorial Auctions

In this section, we prove an *information equivalence* result between the problem of computing an efficient allocation and the problem of computing competitive equilibrium prices. The problem of computing an efficient allocation with minimal preference elicitation can be reduced to the problem of computing a set of CE prices with minimal preference elicitation. The *necessary* direction, that a price certificate is necessary for an efficient allocation, is novel.

In addition to providing some theoretical support for the use of ascending combinatorial auctions, this reduction is useful because simple minimal-information queries exist to verify that prices are CE. Before presenting the main result we introduce a query language that guarantees that preference information sets are both agent- and outcome-independent.

In the following, $\text{EFF}(S)$, is a statement that S is efficient, and $\text{CE}(p, S)$ is a statement that prices p are in competitive equilibrium with allocation S . The corresponding functions, that evaluate to *true* if and only if the statement holds for preferences, $v \in V$, are written $\text{EFF}(S, v)$ and $\text{CE}(p, S, v)$.

4.1 Best-Response Query Language

A canonical query language, the *best-response* query language, $\mathcal{L}_{\text{br}}$, can be constructed from simple price-based queries.

Definition 14. A *price-based query*, $Q(p)$, elicits from an agent the set of bundles that maximize its utility at prices p , or the emptyset if no bundle has positive utility.

Example 3. Consider prices, $p(A) = 5, p(B) = 5, p(AB) = 7$. Response, $\{A\}$ to query, $Q(p)$, provides information set, $\text{inf} = \{v_i(A) - 5 \geq v_i(B) - 5; v_i(A) - 5 \geq v_i(AB) - 7; v_i(A) - 5 \geq 0\}$.

We continue to assume truth-revelation from agents in response to queries, an assumption that will be justified very soon, with incentive-compatible auction design.

Definition 15. A *best-response query language*, $\mathcal{L}_{\text{br}}$, contains all queries, Q , with information content equal to the information content of some sequence of price queries.

Formally, $Q \in \mathcal{L}_{\text{br}}$ if and only if $\text{inf}(Q) = \text{inf}(Q'_0, \dots, Q'_t)$, for some sequence, $Q'_0, \dots, Q'_t$, of price queries, for any agent preferences.

Our intention in introducing $\mathcal{L}_{\text{br}}$ is not that queries be *executed* as a sequence of price-based queries. Rather, we will analyze necessary and sufficient properties of information certificates in efficient and incentive compatible mechanisms that contain all queries that can be expressed within the best-response query language. As an example, $\mathcal{L}_{\text{br}}$ is rich enough to contain all query types previously proposed for iterative combinatorial auctions in Parkes [28] and Conen & Sandholm [9, 11, 10]. This includes the following queries: *value queries*, what is your value for bundle S ?; *rank queries*, what is your preference ordering across bundles S_1 and S_2 ?; *bounded-value queries*, is your value for bundle S less (or greater) than x ?; and *weighted-rank*, or *differential* queries, is your difference in value between bundles S_1 and S_2 less (or greater) than x ?

For all these queries, there exists a reduction to a sequence of price-based queries. As an example, the rank query, “is $v_i(S_1) > v_i(S_2)$ ” can be implemented with a single price-query with prices, $p_i(S_1) = 0, p_i(S_2) = 0$, and $p_i(S) = \infty$, for all $S \neq \{S_1, S_2\}$. Similarly, the value query, “what is your value for S ” can be implemented with a sequence of price-queries, with prices selected according to a standard binary-search method. With a slight variation, language $\mathcal{L}_{\text{br}}$ is also rich enough to represent approximate preference revelation via partially-specified constraint-network formulations of valuations [9].³ The variation allows an agent to respond with, $\perp$, to indicate that it does not have enough information to respond to price-based query.

³ In a constraint-network model, bundles are represented as vertices with lower- and upper-bounds on values, and ordering information between pairs of bundles is represented with weighted edges, with weights that define upper- and lower-bounds on the difference in value between a pair of bundles.

Example 4. Consider a problem with 2 goods, A and B , and 2 agents. Agent 1 has values, $v_1(A) = 4, v_1(B) = 3, v_1(AB) = 8$, and agent 2 has values $v_2(A) = 1, v_2(B) = 6, v_2(AB) = 9$ [9]. Let $\text{EFF}(S^*, v)$ denote a function that evaluates to *true* if and only if S^* is an efficient allocation, given agent values v . Complete information, $\text{inf}_a = (v_1, v_2)$, provides a trivial certificate of efficiency. Rank information, that specifies a value-decreasing order over agent preferences, i.e. $(AB, A, B, \emptyset)$ from agent 1 and $(AB, B, A, \emptyset)$ from agent 2, is not a certificate; allocations $(AB, \emptyset)$ and $(\emptyset, AB)$ can both be efficient with respect to this information. A price-based query, with prices $p(A) = 3, p(B) = 4$, and $p(AB) = 7$, elicits information set, $\text{inf}_b = \{v_1(A) - 3 \geq v_1(B) - 4; v_1(A) - 3 = v_1(AB) - 7; v_1(A) - 3 \geq 0; v_2(B) - 4 \geq v_2(A) - 3; v_2(B) - 4 = v_2(AB) - 7\}$, which provides a certificate that prices are in competitive equilibrium with allocation $S^* = (A, B)$, and immediately also provides a certificate that allocation, S^* , is efficient.

Crucially, language $\mathcal{L}_{\text{br}}$, with or without $\perp$, generates preference information that is both agent- and outcome-independent. We omit the proof in the interest of space.

Proposition 1. *The best-response query language, $\mathcal{L}_{\text{br}}$, constructs information sets that are both agent- and outcome- independent.*

The main results are stated for agent- and outcome-independent information sets, and hold in particular for mechanisms with queries restricted to language $\mathcal{L}_{\text{br}}$.

4.2 Minimal-Revelation: Truthful Agents

We first show that the problem of computing an efficient allocation with truthful agents is informationally equivalent to the problem of computing competitive equilibrium prices.

Proposition 2 (sufficient). *A preference information set that provides a certificate of competitive equilibrium, for some prices $\hat{p}$ and allocation $\hat{S}$, also provides a certificate of efficiency for allocation $\hat{S}$.*

Proof. This direction is easy, following immediately from standard duality arguments. We have $\text{cert}(\text{inf}, \text{CE}(\hat{p}, \hat{S})) \Rightarrow (\text{CE}(\hat{p}, \hat{S}, v), \forall v \in \text{inf}) \Rightarrow (\text{EFF}(\hat{S}, v), \forall v \in \text{inf}) \Rightarrow \text{cert}(\text{inf}, \text{EFF}(\hat{S}))$ $\square$

Let $v_{\text{WC}}(\text{inf}_i, S)$ denote a worst-case valuation for agent i and bundle S , given information set, inf_i . Also, let $\text{BR}_i(p, v_i)$ denote the set of utility-maximizing bundles for agent i with valuation v_i at prices p (solving CS1), and let $\pi_s(p)$ denote the set of revenue-maximizing allocations for the seller at prices p (solving CS2).

Proposition 3 (necessary). *An agent- and outcome-independent preference information set that provides a certificate of efficiency for some allocation, $\hat{S}$, must also provide a certificate of competitive equilibrium, for some prices, $\hat{p}$.*

Proof. First, $\text{cert}(\text{inf}, \text{EFF}(\hat{S})) \Rightarrow (\text{EFF}(\hat{S}, v), \forall v \in \text{inf})$, which implies $(\exists p \cdot \text{CE}(p, \hat{S}, v), \forall v \in \text{inf})$, by standard duality arguments. Expanding the definition of CE prices, and because preferences are agent-independent, we have $(\exists p \cdot \forall i \cdot (\hat{S}_i \in \text{BR}_i(p, v_i) \wedge \hat{S} \in \pi_s(p)), \forall v_1 \in \text{inf}_1, \dots, \forall v_N \in \text{inf}_N)$, which as a special case implies $\exists p \cdot \forall i \cdot (\hat{S}_i \in \text{BR}_i(p, v_{\text{WC}}(\text{inf}_i, \hat{S}_i)) \wedge \hat{S} \in \pi_s(p))$. Let $\hat{p}$ denote the equilibrium prices for this worst-case valuation special-case, so that $(\hat{S}_i \in \text{BR}_i(\hat{p}, v_{\text{WC}}(\text{inf}_i, \hat{S}_i)) \wedge \hat{S} \in \pi_s(\hat{p})), \forall i$. Now, with outcome-independence, every agent will continue to prefer bundle $\hat{S}_i$ at any valuation other than the worst-case valuation for its efficient bundle, and we have $(\forall i \cdot \hat{S}_i \in \text{BR}_i(\hat{p}, v_i) \wedge \hat{S} \in \pi_s(\hat{p})), \forall v \in \text{inf}$. Finally, this implies $\text{CE}(\hat{p}, \hat{S}, v), \forall v \in \text{inf}$, by basic duality, and we have a certificate, $\text{cert}(\text{inf}, \text{CE}(\hat{p}, \hat{S}))$. $\square$

Combining Proposition 1, and Propositions 2 and 3, we immediately have the following result, which establishes the information equivalence between the minimal-revelation efficient allocation problem and the minimal-revelation competitive equilibrium prices problem.

Theorem 1 (information equivalence). *In a distributed algorithm restricted to queries in the best-response query language it is necessary and sufficient to elicit an information certificate for CE prices to compute an efficient allocation.*

This result is stated for a *distributed algorithm*, and not a *mechanism*, because for now we continue to assume that agents respond truthfully to queries. Minimal CE price certificates are easy to construct *ex post* from queries, once the efficient allocation is known. The *ex post* minimal query problem reduces to finding the minimal-cost price-based query from the set of CE prices.

Proposition 4. *The ex post minimal-revelation efficient allocation problem, given queries restricted to the best-response query language, $\mathcal{L}_{\text{br}}$, is equivalent to:*

$$\begin{aligned} \underline{\text{inf}}(\mathcal{L}_{\text{br}}, \text{EFF}, v) &= \min_{(p, S) \in \mathcal{CE}} c(Q(p)) & [\text{MP}(\mathcal{L}_{\text{br}}, \text{EFF}, c)] \\ \text{s.t. } \text{inf}(\text{CE}(p, S)) &\preceq Q(p) \end{aligned}$$

where $\mathcal{CE}$ is the set of competitive equilibrium prices, and $c(\cdot)$ is the cost metric.

This demonstrates the power of the information equivalence between the efficient-allocation and competitive-equilibrium price problems. The query to generate a minimal certificate for CE prices is easy, just announce the prices as a price-query and ask agents for their best-response sets. We are not aware of a corresponding simple construction for the *ex post* minimal query to elicit an information certificate to verify directly that a particular allocation is efficient.

Of course, we are really interested in constructing indirect mechanisms that achieve this *ex post* performance. A competitive analysis, to compare the performance of an indirect mechanism with this *ex post* ideal, is left for future work. We already notice that an asynchronous English auction explicitly achieves CE prices by increasing the price until only one agent is left. Similarly, an ascending-price

combinatorial auction such as iBundle [29], terminates with competitive equilibrium prices. The precise performance of ascending-price auctions such as iBundle depends on details such as whether bids and price-updates are synchronous or asynchronous across agents, and on the order with which agents interact with the auction.

4.3 Minimal-Revelation: Self-Interested Agents

We can now quantify the informational cost of incentive-compatibility, or the *cost-of-truthfulness*, and show that the problem of computing an efficient allocation with self-interested agents is informationally equivalent to the problem of computing Universal CE (UCE) prices. From this we show that there is a non-zero cost-of-truthfulness unless items are more like substitutes than complements.

The equivalence result is stated for an *ex post* Nash equilibrium, in which no agent can gain from a unilateral deviation at the end of an auction, even with complete information about the bids placed by other agents.

Theorem 2. *An information certificate for Universal CE prices is necessary and sufficient to compute the efficient allocation in a revenue-maximizing ex post Nash equilibrium, at least for agent- and outcome-independent information sets.*

Proof. (sketch) The *sufficient* direction is easy to show. Given UCE prices, p_{uce} , the VCG payments are computed as $p_{\text{gva},i} = p_{\text{uce},i}(S_i^*) - [P(\mathcal{I}) - P(\mathcal{I} \setminus i)]$, where S^* is the efficient allocation and $P(\mathcal{I})$ is the maximal revenue to the seller at prices, p_{uce} , over all allocations, and $P(\mathcal{I} \setminus i)$ is the maximal revenue to the seller at prices, p_{uce} , over all allocations that exclude agent i [31]. The VCG payments provide an *ex post* Nash equilibrium in an indirect mechanism, and the VCG mechanism maximizes the expected revenue across all efficient auctions [19]. The proof of the *necessary* direction has two main steps. First, the uniqueness of the VCG mechanism amongst revenue-maximizing, efficient, and Bayes-Nash incentive-compatible mechanisms [14, 19], implies that an efficient indirect mechanism in *ex post* Nash, that is also revenue-maximizing across efficient mechanisms, must implement the outcome of the VCG mechanism. Second, a similar argument to that in Proposition 3, demonstrates that as long as information is agent- and outcome-independent then $\text{cert}(\text{inf}, \text{EFF}(S))$ and $\text{cert}(\text{inf}, \text{VCG}(p))$, taken together, imply a certificate for UCE prices. $\square$

This extension to UCE prices provides a clean characterization of the additional information elicitation costs imposed by agent incentives. UCE prices are a subset of CE prices, and can therefore be expected to require more preference information in some cases. Take as a hypothesis that *minimal CE prices are the prices that minimize information revelation across all CE prices*. The following condition is necessary and sufficient for minimal CE prices to be Universal CE prices, and under the hypothesis necessary and sufficient for a zero informational cost-of-truthfulness.

Definition 16 (agents are substitutes).

$$V(\mathcal{I}) - V(\mathcal{I} \setminus L) \geq \sum_{i \in L} [V(\mathcal{I}) - V(\mathcal{I} \setminus i)], \quad \forall L \subset \mathcal{I}$$

In words, the marginal product of any set of agents must be greater than the total marginal product of each agent individually. This condition holds in problems in which items are substitutes, but quickly breaks down in problems with complements between items [2].

In related work, Parkes & Ungar [31] propose an extension to *i*Bundle, in which enough information about agent preferences is determined to implement the VCG outcome. The auction, *i*BEA, collects additional information beyond that collected by *i*Bundle precisely when minimal CE prices are not Universal CE prices.

5 Closing Remarks

Preference elicitation is an important problem in mechanism design for agent-mediated electronic commerce. Preference elicitation also presents an important and interesting theoretical challenge in computational mechanism design. This paper makes some initial progress towards a formal theory for the design of minimal-revelation VCG mechanisms. We demonstrate the central role that price-based methods play in minimal-revelation mechanism design. The equivalence results provide some justification for price-based designs for iterative combinatorial auctions, such as *i*Bundle [26, 29] and AkBA [37].

Current price-based auctions structure an ascending search through price space, but this is of course just one possible search method. It would be interesting to understand the preference-revelation tradeoffs in other price-based search heuristics. One useful property of ascending-price auctions, that would be lost in other price-based methods, is the *monotonicity* provided for local agent valuation problems. For example, once a particular bundle in an ascending-price auction is too expensive it will always be too expensive in future rounds and an agent knows that it will not need to refine its value in the future.

The centralized elicitor-agent methods of Conen & Sandholm [9], provide immediate and full information about agent preference information to the mechanism, as preference elicitation queries are executed. In comparison, in the proxy-bidding agents/MBR setting described in this paper, preference information is only provided to the mechanism in the form of best-responses to prices. We plan to investigate the role of asynchronous price-updates and accelerated price increases via extended “virtual auction rounds” to leverage the additional preference information that proxy agents may have in our setting. Another topic of current work is to measure the effect of price-discrimination on the preference-elicitation costs of ascending-price auctions, and to measure the cost-of-truthfulness in ascending-price combinatorial auctions.

Acknowledgments

I would like to thank the anonymous reviewers, and the participants of the AMEC'IV workshop, for their helpful comments.

References

1. A. Andersson, M. Tenhunen, and F. Ygge. Integer programming for auctions with bids for combinations. In *Proc. 4th International Conference on Multi-Agent Systems (ICMAS-00)*, 2000.
2. L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1:1–42, 2002.
3. J. S. Banks, J. O. Ledyard, and D. Porter. Allocating uncertain and unresponsive resources: An experimental approach. *The Rand Journal of Economics*, 20:1–25, 1989.
4. S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra. Linear programming and Vickrey auctions. Technical report, Anderson Graduate School of Management, U.C.L.A., 2001.
5. S. Bikhchandani and J. M. Ostroy. The package assignment model. *Journal of Economic Theory*, 2002. forthcoming.
6. C. Boutilier and H. Hoos. Bidding languages for combinatorial auctions. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001.
7. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
8. O. Compte and P. Jehiel. On the virtues of the ascending price auction: New insights in the private value setting. Technical report, UCL, 2000.
9. W. Conen and T. Sandholm. Preference elicitation in combinatorial auctions. In *Proc. 3rd ACM Conf. on Electronic Commerce (EC-01)*. ACM Press, New York, 2001.
10. W. Conen and T. Sandholm. Differential-revelation VCG mechanisms for combinatorial auctions. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Artificial Intelligence*. 2002. This volume.
11. W. Conen and T. Sandholm. Partial-revelation VCG mechanism for combinatorial auctions. In *Proc. 18th Nat. Conf. on Artificial Intelligence (AAAI'02)*, 2002. to appear.
12. S. de Vries and R. V. Vohra. Combinatorial auctions: A survey. *Inform Journal on Computing*, 2002. forthcoming.
13. C. DeMartini, A. M. Kwasnica, J. O. Ledyard, and D. Porter. A new and improved design for multi-object iterative auctions. Technical Report SSWP 1054, California Institute of Technology, 1998. Revised March 1999.
14. J. Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
15. T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
16. R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. Technical report, Technion and Stanford, 2001.
17. B. Hudson and T. Sandholm. Effectiveness of preference elicitation in combinatorial auctions. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Artificial Intelligence*. 2002. This volume.

18. M. O. Jackson. Mechanism theory. In *The Encyclopedia of Life Support Systems*. EOLSS Publishers, 2000.
19. V. Krishna and M. Perry. Efficient mechanism design. Technical report, Pennsylvania State University, 1998. Available at: <http://econ.la.psu.edu/~vkrishna/vcg18.ps>.
20. D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce (EC-99)*, pages 96–102, 1999.
21. P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 2000.
22. N. Nisan. Bidding and allocation in combinatorial auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, pages 1–12, 2000.
23. N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, pages 242–252, 2000.
24. N. Nisan and I. Segal. The communication complexity of efficient allocation problems. Technical report, Hebrew University and Stanford University, 2002.
25. D. C. Parkes. iBundle: An efficient ascending price bundle auction. In *Proc. 1st ACM Conf. on Electronic Commerce (EC-99)*, pages 148–157, 1999.
26. D. C. Parkes. Optimal auction design for agents with hard valuation problems. In *Proc. IJCAI-99 Workshop on Agent Mediated Electronic Commerce*, pages 206–219, July 1999. Stockholm.
27. D. C. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, May 2001. <http://www.cis.upenn.edu/~dparkes/diss.html>.
28. D. C. Parkes. An iterative generalized Vickrey auction: Strategy-proofness without complete revelation. In *Proc. AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*. AAAI Press, March 2001.
29. D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)*, pages 74–81, 2000.
30. D. C. Parkes and L. H. Ungar. Preventing strategic manipulation in iterative auctions: Proxy agents and price-adjustment. In *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)*, pages 82–89, 2000.
31. D. C. Parkes and L. H. Ungar. An ascending-price generalized Vickrey auction. Technical report, Harvard University, 2002.
32. A. Ronen. Mechanism design with incomplete languages. In *Proc. 3rd ACM Conf. on Electronic Commerce (EC'01)*, 2001.
33. M. H. Rothkopf, A. Pekeć, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
34. T. W. Sandholm. An implementation of the Contract Net Protocol based on marginal-cost calculations. In *Proc. 11th National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, July 1993.
35. T. W. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *Proc. 17th Int. Joint Conf. on Artificial Intelligence*, 2001.
36. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
37. P. R. Wurman and M. P. Wellman. A_kBA: A progressive, anonymous-price combinatorial auction. In *Second ACM Conference on Electronic Commerce*, pages 21–29, 2000.