RC22661 (W0205-203) May 28, 2002
Computer Science

IBM Research Report

Multiattribute Reverse Auctions

David C. Parkes
Division of Engineering and Applied Sciences,
Harvard University,
33 Oxford Street,
Cambridge, MA 02138

Jayant Kalagnanam
IBM Research Division,
T.J. Watson Research Center,
P.O. Box 218,
Yorktown Heights, NY 10598

Iterative Multiattribute Vickrey Auctions

David C. Parkes
Division of Engineering and Applied Sciences,
Harvard University,
33 Oxford Street,
Cambridge, MA 02138
parkes@eecs.harvard.edu

Jayant Kalagnanam
IBM Research Division,
T.J. Watson Research Center,
P.O. Box 218,
Yorktown Heights, NY 10598
jayant@us.ibm.com

June 14, 2002

Abstract

Multiattribute auctions extend traditional auction settings. In addition to price, multiattribute auctions allow negotiation over non-price attributes such as weight, color, terms-of-delivery, and promise to improve market efficiency in markets with configurable goods. Multiattribute auctions also provide quite general purpose negotiation mechanisms, for example over the terms of a contracting relationship. We propose a family of iterative primal-dual based multiattribute auction mechanisms, for reverse auction settings with one buyer and many sellers. The auctions support incremental preference revelation from both the buyer and the sellers. The auctions are price-directed, and a straightforward myopic best-response strategy is in equilibrium for sellers, assuming a class of consistent buyer strategies. Moreover, the auctions are efficient with a truthful buyer, and we quantify the maximal possible gain to a buyer from deviating from a truthful strategy.

1 Background

Multiattribute auctions [5] extend the traditional auction setting to allow negotiation over price and attributes, with the final characteristics of the item, as well as the price, determined dynamically through agents' bids. For example, in a procurement problem, a multiattribute auction can allow different suppliers to compete over both attributes values and price. To the extent that other negotiation problems can be formulated as multiattribute allocation problems, multiattribute auction mechanisms also provide mechanisms for automated negotiation outside of e-commerce, such as bargaining over shared resources between distributed computational agents [12, 10], and automated negotiation over the terms of a contracting relationship.

In this paper we apply a linear-programming based methodology to develop a family of iterative multiattribute auctions. Primal-dual analysis is used to construct an iterative auction that terminates with the outcome of a modified Vickrey-Clarke-Groves (VCG) [22, 6, 8] mechanism for the multiattribute allocation problem. A similar approach has yielded successful designs for efficient ascending-price combinatorial auctions, in which bidders demand different combinations of items [18, 2, 20]. The primary difference between the multiattribute allocation problem and the combinatorial allocation problem is that there is private information on both sides of the auction in the multiattribute setting. This complicates the incentive structure of the mechanism design problem, because the winner determination problem in each round of the auction depends on the preferences of the buyer and the sellers.

Iterative mechanisms, that allow participants to provide incremental information about their preferences, are especially important in applications of multiattribute auctions to procurement settings. First, preference elicitation is often *costly* in procurement problems, and bidders would prefer not to have to determine an exact value tradeoff across all different combinations of attribute levels if that can be avoided. Second, it is often important to reveal as little information as possible about costs and preferences in a strategic situation such as procurement, because participants are in a long-term competitive relationship.

One could imagine two reasonable goals for multiattribute auction design: *buyer payoff maximization*, or alternatively *allocative efficiency*, which selects the outcome that maximizes the difference between the value of the buyer and the cost of the seller, across all attribute levels and all sellers. It is well known from classic auction theory that the two goals of payoff maximization and efficiency are typically incompatible [13]. Earlier designs for multiattribute auctions emphasized *optimal* auction design, maximizing the final payoff to the buyer in the auction. Che [5] proposed a buyer payoff-maximizing *one-shot sealed-bid* two-attribute auction protocol with a first price and a second price payoff function. Assuming that buyers announce a scoring function (not necessarily truthful) he shows that the equilibrium behavior for the supplier is to supply at a quality level that maximizes the score minus the true cost to achieve that score. He also derives the optimal scoring functions for the buyer (both for first price and second price protocols) where a quality penalty is applied as a function of the distribution of the cost parameter (which is assumed to be common knowledge). Branco [4] extends this protocol to the case where the seller cost functions are correlated.

Recently, Beil & Wein [1] have proposed an iterative payoff-maximizing auction procedure, for a class of parameterized utility functions (with K parameters) with known functional forms and naive suppliers. The buyers uses K rounds to estimate the seller costs functions deterministically. For the final ($K+1$th) round they design a scoring function so as to maximize buyer payoff by essentially reporting the same score (within ϵ) for the top two suppliers. They assume that the buyer scoring function does not need to be consistent across rounds.

We propose multiattribute auction mechanisms that are efficient under reasonable assumptions about bidder strategies. This goal is similar to the one in the paper by Vulkan & Jennings [23]. We believe that efficiency is a more appropriate goal than utility-maximization for stable long-term market design. We expect that efficient markets will come to dominate the electronic market landscape [**?**]. Buyer payoff maximization is perhaps more appropriate for a one-shot procurement problem, and in a setting in which the buyer has considerable market power [5, 1].

We formulate and solve two variations on the one-buyer many-seller multiattribute allocation problem. In the first, and more general variation, we assume discrete attribute levels and non-linear valuation and cost functions across attributes. In the second variation we assume discrete attribute levels and *preferential-independence*, which is equivalent to an assumption that sellers and the buyer have linear-additive cost and valuation functions across different attribute types. Iterative auctions are proposed for both variations, that maintain prices on different combinations of attribute values and adjust prices across rounds based on bids from sellers and information from the buyer. Prices are non-linear but anonymous in the problem with general agent preferences, and prices cannot be represented as a linear sum over prices on individual levels of each attribute type. The auction for the special case of preferential-independence maintains two separable price components, comprised of linear prices and an individualized price penalty for each seller.

We prove that straightforward myopic best-response is an *ex post* Nash equilibrium for sellers against a class of *consistent* buyer strategies. A *consistent* strategy for the buyer is any strategy that can be represented as a straightforward myopic best-response strategy for some *ex ante* fixed, but perhaps untruthful, cost function. In particular, if the buyer follows a truthful consistent strategy, then the auction terminates with the efficient outcome. We are also able to characterize the maximal benefit to the buyer for a non-truthful strategy, consistent or otherwise.

1.1 Outline

Section 2 formulates the multiattribute allocation problem, and introduces linear program formulations for general non-linear preferences, and for the special case of preferential-independence. A simple example is provided to illustrate each variation. We then establish the integrality of the formulations, and demonstrate that the dual programs compute competitive equilibrium prices. Continuing, in Section 3 we define the VCG mechanism for the efficient multiattribute allocation problem, and demonstrate that it is not budget-balanced. We introduce a budget-balanced approximate VCG mechanism, that remains strategyproof for the sellers but is not strategyproof for the buyer, and bound the effect of manipulation on the efficiency. Finally, we relate the Vickrey payments with the competitive equilibrium prices developed in the LP analysis.

Section 4 introduces the primal-dual based auction for the general multiattribute allocation problem, with non-linear preferences. Theoretical and computational analysis is presented to characterize the performance of the auction. We also present a worked example, and relate the design of the auction to existing designs for iterative combinatorial auctions. Section 5 defines a simpler auction for the case of preferential-independence, with a smaller price space composed of a linear component with a simple additive penalty term.

2 The Multiattribute Allocation Problem

In the multiattribute allocation problem (MAP) there are N sellers, one buyer, and M attributes. Let $\mathcal{I}$ denote the set of sellers, and $\mathcal{J}$ denote the set of attributes. Each attribute, $j \in \mathcal{J}$, has a domain of possible attribute values (or *levels*), denoted with abstract set Θ_j; for example $\Theta_1 = \{\text{red,yellow,green}\}$ if attribute 1 is the color of an item. The joint domain, across all attributes, is denoted $\Theta = \Theta_1 \times \Theta_M$. Each seller, $i \in \mathcal{I}$, has a *cost function*, $c_i(\theta) \geq 0$, for an attribute bundle, $\theta \in \Theta$, and the buyer has a *valuation function*, $v(\theta) \geq 0$. For simplicity, it is useful to assume that Θ contains a null attribute bundle, ϕ, for which $v(\phi) = 0$ and $c_i(\phi) = 0$ for all $i \in \mathcal{I}$.

Definition 1 (multiattribute allocation problem). *Select attribute bundle, θ^*, and seller, i^*, to maximize the difference between buyer value and seller cost:*

$$V(\mathcal{I}) = \max_{\theta \in \Theta, i \in \mathcal{I}} v(\theta) - c_i(\theta) \qquad [\textsc{map}(\mathcal{I})]$$

The solution to the MAP problem, $\textsc{map}(\mathcal{I})$, is the *efficient allocation*, denoted (θ^*, i^*). In later sections, when we introduce incentive-compatible mechanisms for MAP, it will be useful to consider the MAP problem restricted to $(\mathcal{I} \setminus i)$ agents. We write, $[\textsc{map}(\mathcal{I} \setminus i)]$, to denote this restricted problem, and $V(\mathcal{I} \setminus i)$ to denote the value of the solution to $[\textsc{map}(\mathcal{I} \setminus i)]$.

In this paper we assume that agents have *quasilinear utility functions*. The utility to seller, i, for selling an item with attribute bundle, θ, at price p is the difference between the price and its cost, i.e. $u_i(\theta, p) = p - c_i(\theta)$. Similarly, the utility to the buyer for buying an item with attribute bundle, θ, at price p is simply the difference between its value and the price, i.e. $u^B(\theta, p) = v(\theta) - p$.

Notice that we restrict our attention to multiattribute allocation problems in which a single buyer negotiates with multiple sellers in a reverse auction, and will eventually select a *single* seller. This assumption is common in the literature [5, 4, 1], and the single-seller outcome remains important in practical e-commerce applications [1]. Extensions to allow aggregation across multiple sellers, and combinatorial effects, significantly complicate incentive considerations and are left for future work. In particular, the Vickrey payoff to the winning sellers can no longer be supported in competitive equilibrium when the final trade involves multiple sellers.

In general the abstract attribute domains, Θ_j, can allow for both discrete and continuous attribute values. In this paper we focus on the discrete attribute case, and consider two main variations. The first variation, MAP-1, formulates the MAP with general non-linear agent preferences. The second variation, MAP-2, makes a preferential-independence assumption about the valuation and cost functions of agents. In the following sections we introduce each variation, and use primal-dual analysis to characterize the space of *competitive equilibrium* (CE) prices. In competitive equilibrium, the outcome maximizes the surplus for every seller and the buyer, given the prices. This characterization is useful for the design of iterative mechanisms, because we can design auctions that terminate in competitive equilibrium, and support the efficient allocation.

2.1 MAP-1: General Preferences

To apply a primal-dual auction design methodology we need an appropriate linear-programming formulation of the MAP problem. The formulation must be integral, provide enough structure to make an equivalence between the dual solution and competitive equilibrium prices, and also have a useful correspondence to VCG payments. Unlike the combinatorial allocation problem [21, 7], integrality is quite straightforward in the MAP problem. The trick in the formulation is to introduce additional structure, through valid inequalities and lifting, to achieve a useful price structure.

First consider a simple formulation. Introduce variable, $x_i(\theta) \geq 0$, to indicate that attribute bundle, θ, is purchased from seller i at level $x_i(\theta)$. Of course, in an integral solution this will be a 0/1 value.

$$
\begin{aligned}
& \max_{x_i(\theta)} \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} (v(\theta) - c_i(\theta)) x_i(\theta) \\
\text{s.t.} \quad & \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} x_i(\theta) \leq 1 \\
& x_i(\theta) \geq 0
\end{aligned}
$$

The dual formulation to this LP provides no useful structure, simply reducing to $[\min \pi \ : \ \pi \geq v(\theta) - c_i(\theta), \ \forall i, \forall \theta]$, with dual variable, $\pi \geq 0$.

Following the second-order formulation for the combinatorial allocation problem, introduced in Bikchandani & Ostroy [3], we might consider the following

LP formulation, in which additional variables, $x^B(\theta) \geq 0$ are introduced.

$$\begin{aligned}
& \max_{x_i(\theta), x^B(\theta)} \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} (v(\theta) - c_i(\theta)) x_i(\theta) \\
\text{s.t.} \quad & \sum_{\theta \in \Theta} x_i(\theta) \leq 1, \quad \forall i \in \mathcal{I} && (1) \\
& \sum_{i \in \mathcal{I}} x_i(\theta) \leq x^B(\theta), \quad \forall \theta \in \Theta && (2) \\
& \sum_{\theta \in \Theta} x^B(\theta) \leq 1 && (3) \\
& x_i(\theta), x^B(\theta) \geq 0
\end{aligned}$$

The corresponding dual formulation for this LP has some additional structure, but is still not sufficient for primal-dual auction design. Introduce dual variable, π_i, for constraints (1), dual variable, π^B, for constraint (3), and dual variable, $p(\theta)$, for constraints (2).

$$\begin{aligned}
& \min_{\pi_i, \pi^B, p(\theta)} \sum_{i \in \mathcal{I}} \pi_i + \pi^B \\
\text{s.t.} \quad & \pi_i + p(\theta) \geq v(\theta) - c_i(\theta), \quad \forall i \in \mathcal{I}, \forall \theta \in \Theta && (4) \\
& -p(\theta) + \pi^B \geq 0, \quad \forall \theta \in \Theta && (5) \\
& \pi_i, \pi^B, p(\theta) \geq 0
\end{aligned}$$

We would like to interpret dual variables as competitive equilibrium prices, via complementary slackness (CS) conditions. Although variables, $p(\theta)$, can be interpreted as prices, the problem with this primal/dual formulation is that the complementary slackness (CS) conditions do not separate the information about the buyer's valuation function from the information about the sellers' cost functions. This information must be separated across different CS conditions, because we would like to test whether a particular CS condition holds with best-response information from a single agent at the prices.

As an example, the information about the valuation, $v(\theta)$, and the cost, $c_i(\theta)$, of seller i is required to evaluate the following CS condition.

$$x_i(\theta) > 0 \;\Rightarrow\; \pi_i + p(\theta) = v(\theta) - c_i(\theta)$$

Combined with dual constraint (4), this implies that if attribute bundle, θ', from seller i is the efficient outcome then it is necessary that $v(\theta') - c_i(\theta') - p(\theta') = \max_{\theta \in \Theta}(v(\theta) - c_i(\theta) - p(\theta))$. In otherwords, it is necessary that θ' maximizes the price-adjusted surplus over all outcomes involving seller i. But the valuation function is *private* to the buyer, and the cost function is *private* to the seller, so this property cannot be tested by announcing price $p(\theta)$ and asking the buyer and seller for their best-response at the price.

Linear program formulation, MAP-1, resolves this problem. The formulation separates the valuation of the buyer and the costs of the sellers in the objective function, and leads to a useful set of CS conditions.

$$\max_{x_i(\theta), x^B(\theta)} \sum_{\theta \in \Theta} v(\theta) x^B(\theta) - \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} c_i(\theta) x_i(\theta) \qquad \text{[MAP-1]}$$

$$\text{s.t.} \quad \sum_{\theta \in \Theta} x_i(\theta) \leq 1, \quad \forall i \in \mathcal{I} \qquad (6)$$

$$\sum_{i \in \mathcal{I}} x_i(\theta) \geq x^B(\theta), \quad \forall \theta \in \Theta \qquad (7)$$

$$\sum_{\theta \in \Theta} x^B(\theta) \leq 1 \qquad (8)$$

$$x_i(\theta), x^B(\theta) \geq 0$$

Notice that constraint (7) is reversed from (2) in this formulation, although it is straightforward to demonstrate that constraints (7) and (2) hold with equality for all $\theta \in \Theta$ at the optimal solution to both formulations.

Before continuing, it is useful to verify that MAP-1 is indeed integral. First notice that the choice of $x^B(\theta)$ is extremal with respect to the choice of $x_i(\theta)$, which is to say that $x^B(\theta) = \sum_i x_i(\theta)$ for all θ in an optimal solution. The problem reduces to $[\max_{x_i(\theta)} \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} (v(\theta) - c_i(\theta)) x_i(\theta) \ : \ \sum_{i \in \mathcal{I}} \sum_{\theta \in \Theta} x_i(\theta) \leq 1]$, and integrality of $x_i(\theta)$ and $x^B(\theta)$ follows.

The dual LP formulation becomes:

$$\min_{\pi_i, \pi^B, p(\theta)} \sum_{i \in \mathcal{I}} \pi_i + \pi^B \qquad \text{[DMAP-1]}$$

$$\text{s.t.} \quad p(\theta) + \pi^B \geq v(\theta), \quad \forall \theta \in \Theta \qquad (9)$$

$$\pi_i - p(\theta) \geq -c_i(\theta), \quad \forall i \in \mathcal{I}, \forall \theta \in \Theta \qquad (10)$$

$$\pi_i, \pi^B, p(\theta) \geq 0$$

Looking now at the CS conditions for MAP-1 and DMAP-1, we can interpret the values of variables, $p(\theta)$, in an optimal dual solution as non-linear competitive equilibrium prices. First, notice that once values, $p(\theta)$, are defined, call them prices, the optimal values for π_i and π^B with respect to those prices are easy to state.

$$\pi_i = \max_{\theta \in \Theta} [p(\theta) - c_i(\theta), 0] \qquad (11)$$

$$\pi^B = \max_{\theta \in \Theta} [v(\theta) - p(\theta), 0] \qquad (12)$$

In words, π_i, is the maximal utility to seller i at prices $p(\theta)$ and π^B is the maximal utility to the buyer, since we assume quasilinear utility functions.

Then, the first pair of interesting CS conditions, that relate to the preferences of seller i, are:

$$\pi_i > 0 \Rightarrow \sum_{\theta \in \Theta} x_i(\theta) = 1 \qquad \text{(CS-1)}$$

$$x_i(\theta) > 0 \Rightarrow \pi_i - p(\theta) = -c_i(\theta) \qquad \text{(CS-2)}$$

Given (11), condition (CS-1) states that if a seller has positive utility for some outcome at the prices, then the seller must win the auction. In particular, condition (CS-2) states that the attribute bundle must maximize the seller's utility at the prices.

The second pair of interesting CS conditions, that relate to the preferences of the buyer, are:

$$\pi^B > 0 \Rightarrow \sum_{\theta \in \Theta} x^B(\theta) = 1 \tag{CS-3}$$

$$x^B(\theta) > 0 \Rightarrow p(\theta) + \pi^B = v(\theta) \tag{CS-4}$$

Given (12), condition (CS-3) states that if the buyer has positive utility for some outcome at the prices, then some trade must occur. In particular, condition (CS-4) states that the attribute bundle must maximize the buyer's utility at the prices.

This analysis demonstrates that prices, $p(\theta)$, in an optimal dual solution correspond to competitive equilibrium prices. The efficiency of the corresponding allocation follows immediately from the CS conditions.

Definition 2 (competitive equilibrium). *Prices, $p(\theta)$, are competitive equilibrium prices, if there is an allocation (θ^*, i^*) that maximizes the utility of the buyer, and the utility of every seller, at the prices.*

Proposition 1 (efficiency). *The allocation, (θ^*, i^*), that is supported in competitive equilibrium is efficient.*

Formulation MAP-1 succeeds in separating the preference information of the buyer from that of the sellers; individual agents, namely the buyer and the sellers, can provide sufficient information in best-response bids to prices to verify CS conditions.

In general, there are many possible CE prices for a given instance of the MAP problem. It is useful to characterize the range of CE prices, because the incentive properties of the auction depend on the actual prices selected when the auction terminates.

First, feasibility constraint (10), with $\pi_i = 0$ for $i \neq i^*$, reduces to:

$$p(\theta) \leq \min \left[\min_{i \neq i^*} c_i(\theta), \pi_{i^*} + c_{i^*}(\theta) \right], \quad \forall \theta \in \Theta$$

Combining this with constraint (9), we need:

$$\pi^B \geq \max_{\theta \in \Theta} \left\{ v(\theta) - \min \left[\min_{i \neq i^*} c_i(\theta), \pi_{i^*} + c_{i^*}(\theta) \right] \right\} \tag{*}$$

We first show that this condition trivially holds whenever the expression is maximized for some $\theta' \in \Theta$ for which $\pi_{i^*} + c_{i^*}(\theta') < \min_{i \neq i^*} c_i(\theta')$. In this case, we require $\pi^B \geq v(\theta') - \pi_{i^*} - c_{i^*}(\theta')$, which is satisfied for any θ' because $\pi^B + \pi_{i^*} = v(\theta^*) - c_{i^*}(\theta^*)$ in equilibrium, and (i^*, θ^*) is the efficient solution.

Now, writing $(\tilde{\theta}, \tilde{i})$ to denote the *second-best* allocation, i.e. the solution to $[\text{MAP}(\mathcal{I} \setminus i^*)]$, condition (*) simplifies to:

$$\pi^B \geq \max_{\theta \in \Theta} \left[v(\theta) - \min_{i \neq i^*} c_i(\theta) \right] = v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}) \tag{13}$$

Competitive equilibrium also requires, $\pi_{i^*} \geq 0$, that together with $\pi^B + \pi_{i^*} = v(\theta^*) - c_{i^*}(\theta^*)$, becomes:

$$\pi^B \leq v(\theta^*) - c_{i^*}(\theta^*) \tag{14}$$

Finally, substituting $p(\theta^*) = v(\theta^*) - \pi^B$, and constructing the prices on attribute bundles, $\theta \neq \theta^*$, to give $\pi_i \leq 0$ for all $i \neq i^*$, and also to ensure that the buyer prefers attributes θ^* to any other attributes, then the space of optimal dual solutions corresponds to:

$$c_{i^*}(\theta^*) \quad \leq p(\theta^*) \quad \leq v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta})) \tag{15}$$

$$v(\theta) - v(\theta^*) + p(\theta^*) \quad \leq p(\theta) \quad \leq \min_{i \neq i^*} c_i(\theta), \qquad \forall \theta \tag{16}$$

It is trivial to check that (16) holds for any values, $p(\theta^*)$, that satisfy (15).

We can now define the *maximal* CE prices, which are prices that maximize the payment to the efficient seller across all competitive equilibrium prices. In Section 3 we relate the maximal CE prices to the payments in the VCG mechanism for the MAP problem.

Lemma 1 (maximal CE prices). *The maximal CE prices, given efficient solution (θ^*, i^*) and second-best solution, $(\tilde{\theta}, \tilde{i})$, set price $p(\theta^*) = v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$ and $p(\theta)$ for $\theta \neq \theta^*$ to satisfy (16).*

The maximal CE prices are characterized by the property that the second-best seller, $\tilde{i}$, is *pivotal*, in that its maximal utility across all attributes is exactly zero. We will see this characteristic again in the next section, when we derive maximal CE prices for the preferential-independence MAP problem.

Lemma 2 (second-best utility). *The utility of the second-best seller, $\tilde{i}$, is exactly zero for the second-best attribute bundle, $\tilde{\theta}$, at the maximal CE prices.*

Proof. Substitute, $p(\theta^*) = v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$, into (16). □

Later, it is useful to write, $\overline{p}_{\text{ce}}(\theta)$, to denote the particular maximal CE prices in which the price on attributes, $\theta \neq \theta^*$, is maximized, i.e. $\overline{p}_{\text{ce}}(\theta) = \min_{i \neq i^*} c_i(\theta)$ for all $\theta \neq \theta^*$, and $\overline{p}_{\text{ce}}(\theta^*) = v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$.

2.2 MAP-2: Preferential-Independence

A classic assumption in multiattribute utility theory [11] is that agents' utilities satisfy *preferential-independence*. An attribute $j \in \mathcal{J}$ is said to be preferentially independent of $j' \neq j$ if preferences for specific levels of θ_j do not depend on the level of attribute $\theta_{j'}$. Let $k \in \Theta_j$ denote the k^{th} level of attribute j, and let $x_{jk} \in \{0,1\}$ equal 1 if level k of attribute j is selected in attribute bundle θ. It is required that $\sum_{k \in \Theta_j} x_{jk} \leq 1$ for all attributes, $j \in \mathcal{J}$.

A cost function for seller i that satisfies preferential independence can be expressed as:

$$c_i(\theta) = \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} c_{ijk} x_{jk}$$

where c_{ijk} is the marginal cost to seller i if level k of attribute j is selected. Similarly, the buyer's valuation function can be expressed as

$$v(\theta) = \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} v_{jk} x_{jk}$$

where v_{jk} is the marginal value to seller i if level k of attribute j is selected.

Introduce variables x_{ijk}, x^B_{jk}, and y_i. Perhaps the most natural formulation of the preferential-independence special case, given MAP-1, is as follows:

$$\max_{x_{ijk}, x^B_{jk}, y_i} \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} v_{jk} x^B_{jk} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} c_{ijk} x_{ijk}$$

$$\text{s.t.} \quad \sum_{k \in \Theta_j} x_{ijk} \leq y_i, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J} \tag{17}$$

$$\sum_{i \in \mathcal{I}} y_i \leq 1 \tag{18}$$

$$\sum_{i \in \mathcal{I}} x_{ijk} \geq x^B_{jk}, \quad \forall j \in \mathcal{J}, \forall k \in \mathcal{K} \tag{19}$$

$$x_{ijk}, x^B_{jk}, y_i \geq 0$$

Constraints (17) correspond to constraints (6) in MAP-1, and constraints (19) correspond to constraints (7). The additional constraint, (18), is introduced to ensure that at most one seller is selected in the outcome, which makes constraints that correspond to constraints (8) in MAP-1 redundant. Although the dual formulation for this LP has some useful price structure, it is not quite suitable for primal-dual auction design. To see this, introduce dual variables, π_{ij}, for constraints (17), dual variable, π, for constraint (18), and dual variables, p_{jk}, for constraints (19). Variables, p_{jk} can be interpreted as prices, with π_{ij} representing the maximal utility of seller i on attribute j at prices p_{jk}, and the dual problem is to set prices that minimize the maximal total utility, across all sellers:

$$
\begin{aligned}
& \min \pi \\
& \text{s.t.} \quad p_{jk} \geq v_{jk}, \quad \forall j \in \mathcal{J}, \forall k \in \Theta_j \\
& \pi_{ij} - p_{jk} \geq -c_{ijk}, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}, \forall k \in \Theta_j \\
& -\sum_{j \in \mathcal{J}} \pi_{ij} + \pi \geq 0, \quad \forall i \in \mathcal{I} \\
& \pi_{ij}, \pi, p_{jk} \geq 0
\end{aligned}
$$

There are two problems with this formulation. First, the prices do not achieve the goal of solving the MAP with incomplete preference revelation from the buyer. Notice that the complementary slackness conditions for primal variable, x^B_{jk}, state that $x^B_{jk} > 0 \Rightarrow p_{jk} = v_{jk}$; in other words, these attribute levels must be priced at the value of the buyer. Second, there are no optimal dual prices that support the VCG payments, which we will need to define an iterative primal-dual based auction with useful incentive properties.

We introduce linear program MAP-2 to resolve these problems. MAP-2 introduces an additional variable, $x_i \geq 0$, and replaces constraints (17) with constraints (20) and (21). Secondly, MAP-2 introduces valid inequalities, (22) and (23); note that (22) is redundant given (18) and (21), and (23) is redundant given (18), (19), (20) and (21).

$$
\max_{x_{ijk}, x^B_{jk}, x_i, y_i} \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} v_{jk} x^B_{jk} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{k \in \Theta_j} c_{ijk} x_{ijk} \qquad \text{[MAP-2]}
$$

$$
\text{s.t.} \quad \sum_{k \in \Theta_j} x_{ijk} \leq x_i, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J} \tag{20}
$$

$$
x_i \leq y_i, \quad \forall i \in \mathcal{I} \tag{21}
$$

$$
\sum_{i \in \mathcal{I}} y_i \leq 1 \tag{18}
$$

$$
\sum_{i \in \mathcal{I}} x_{ijk} \geq x^B_{jk}, \quad \forall j \in \mathcal{J}, \forall k \in \mathcal{K} \tag{19}
$$

$$
x_i \leq 1, \quad \forall i \in \mathcal{I} \tag{22}
$$

$$
\sum_{k \in \Theta_j} x^B_{jk} \leq \sum_{i \in \mathcal{I}} y_i, \quad \forall j \in \mathcal{J} \tag{23}
$$

$$
x_{ijk}, x^B_{jk}, x_i, y_i \geq 0
$$

Proposition 2 (integrality). *Linear program, MAP-2, is integral.*

Proof. Ignoring redundant constraints (22) and (23), we first note that $x^B_{jk} = \sum_{i \in \mathcal{I}} x_{ijk}$ in the optimal solution. Let $u_{ijk} = v_{jk} - c_{ijk}$ and suppose, w.o.l.g., an ordering over k s.t. $u_{ij1} \geq u_{ij2} \geq \ldots$, for all i, j. Then, the optimal setting is $x_{ij1} = x_i = y_i$ when $u_{ij1} \geq 0$, for all i, j, with $x_{ij1} = 0$ otherwise, and $x_{ijk} = 0$ for all $k \neq 1$, all i, j. Taking the interesting case, that $x_{ij1} = y_i$ for all y_i, the problem now reduces to $\max_{y_i} \sum_{i\in\mathcal{I}} \left(\sum_{j\in\mathcal{J}} u_{ij1}\right) y_i = \sum_{i\in\mathcal{I}} V_i y_i$, where $V_i = \sum_{j\in\mathcal{J}} u_{ij1}$. Integrality of y_i, x_{ijk}, and x^B_{jk}, follows. $\square$

To construct the dual, introduce variables π_{ij}, Δ_i, π^B, p_{jk}, π_i, and π^B_j, to correspond to primal constraints (20), (21), (18), (19), (22), and (23) respectively.

$$\min_{\pi_{ij},\Delta_i,\pi^B,p_{jk},\pi_i,\pi^B_j} \pi^B + \sum_{i\in\mathcal{I}} \pi^i \qquad \text{[DMAP-2]}$$

$$\text{s.t.} \quad \pi^B \geq \sum_{j\in\mathcal{J}} \pi^B_j + \Delta_i, \quad \forall i \in \mathcal{I} \qquad (24)$$

$$\pi^B_j \geq v_{jk} - p_{jk}, \quad \forall j \in \mathcal{J}, \forall k \in \Theta_j \qquad (25)$$

$$\pi_i \geq \sum_{j\in\mathcal{J}} \pi_{ij} - \Delta_i, \quad \forall i \in \mathcal{I} \qquad (26)$$

$$\pi_{ij} \geq p_{jk} - c_{ijk}, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}, \forall k \in \Theta_j \qquad (27)$$

$$\pi_{ij}, \Delta_i, \pi^B, p_{jk}, \pi_i, \pi^B_j \geq 0$$

Taken together, variables p_{jk} and Δ_i can be interpreted as providing prices for attribute bundle $\theta \in \Theta$. The prices have an additive linear component, defined across p_{jk}, and a non-linear and non-anonymous component, defined as penalty, Δ_i. The effective price for attribute bundle, θ, from seller i is

$$p_i(\theta) = \sum_{j\in\mathcal{J}} \sum_{k\in\Theta_j} \theta_{jk} p_{jk} - \Delta_i$$

where $\theta_{jk} = 1$ if and only if the value of attribute j is assigned to the k^{th} level in the domain for that attribute. The term, $\Delta_i \geq 0$, represents a price penalty from the base prices for seller i.

As before, once values on prices p_{jk} and Δ_i are defined, then the optimal values for π_i, π_{ij}, π^B and π^B_j are easy to state.

$$\pi_i = \max[\sum_{j\in\mathcal{J}} \pi_{ij} - \Delta_i, 0] \qquad (28)$$

$$\pi_{ij} = \max_{k\in\Theta_j} [p_{jk} - c_{ijk}, 0] \qquad (29)$$

$$\pi^B = \sum_{j\in\mathcal{J}} \pi^B_j + \max_i \Delta_i \qquad (30)$$

$$\pi^B_j = \max_{k\in\Theta_j} [v_{jk} - p_{jk}, 0] \qquad (31)$$

In words, π_i is the maximal utility to seller i at prices p_{jk}, Δ_i, which is evaluated as the sum over the seller's maximal utility for each attribute, π_{ij}, and π^B is the maximal utility to the buyer, which is evaluated as the sum over the buyer's maximal utility for each attribute, π^B_j, plus the *maximum* over all penalties. Notice that the buyer gets to select both the best linear price terms and the best penalty terms.

As desired, the CS conditions demonstrate that an allocation is efficient if and only if there is a dual solution that corresponds to prices that support the allocation in competitive equilibrium. The interesting CS conditions that relate to the preferences of seller i are:

$$\pi_{ij} > 0 \Rightarrow \sum_{k \in \Theta_j} x_{ijk} = x_i \tag{CS-1}$$

$$\pi_i > 0 \Rightarrow x_i = 1 \tag{CS-2}$$

$$x_i > 0 \Rightarrow \pi_i = \sum_{j \in \mathcal{J}} \pi_{ij} - \Delta_i \tag{CS-3}$$

$$x_{ijk} > 0 \Rightarrow \pi_{ij} = p_{jk} - c_{ijk} \tag{CS-4}$$

Given (28) and (29), these conditions state that if seller i has positive overall utility for some trade, given prices (p_{jk}, Δ_i), then the seller should win the auction, and the attribute bundle should maximize its utility.

The interesting CS conditions that relate to the preferences of the buyer are:

$$\pi_j^B > 0 \Rightarrow \sum_{k \in \Theta_j} x_{jk}^B = \sum_{i \in \mathcal{I}} y_i \tag{CS-5}$$

$$\pi^B > 0 \Rightarrow \sum_{i \in \mathcal{I}} y_i = 1 \tag{CS-6}$$

$$y_i > 0 \Rightarrow \pi^B = \sum_{j \in \mathcal{J}} \pi_j^B + \Delta_i \tag{CS-7}$$

$$x_{jk}^B > 0 \Rightarrow \pi_j^B = v_{jk} - p_{jk} \tag{CS-8}$$

Given (30) and (31), if a buyer has positive utility for some attribute bundle at prices p_{jk}, and taking the maximal penalty term across all sellers, then the efficient trade is the attribute bundle that maximizes its utility.

Then, since $x_i = y_i$ and $x_{jk}^B = \sum_i x_{ijk}$ in the optimal solution, the CS conditions provide the following interpretation of competitive equilibrium prices.

Definition 3 (MAP-2 competitive equilibrium). *Prices, (p_{jk}, Δ_i), are competitive equilibrium prices, if and only if there is an allocation, (θ^*, i^*), such that attribute bundle, θ^*, maximizes both the utility of the seller, i^*, and the utility of the buyer at the prices, and no other seller has positive utility for any attribute bundle.*

An immediate implication is that it is necessary that the penalty that corresponds to the efficient seller is maximal across all seller penalties.

Proposition 3 (efficiency). *The allocation that is supported in competitive equilibrium is efficient.*

There are many possible CE prices for a given instance of the preferential independence MAP problem. In general, CE prices for the preferential-independence special case can be non-anonymous, with penalty $\Delta_i \neq \Delta_j$ for some sellers $i \neq j$, although it is always possible to construct anonymous CE prices by increasing Δ_i to all losing sellers, to equal the penalty associated with the winning seller.

We formulate a restricted dual problem to better understand the role of price penalties, and also to characterize the maximal CE prices. We show that penalty terms are necessary when the winning seller does not dominate the other sellers, in terms of cost, across all attribute types. Completely linear CE prices only exist when this dominance condition holds.

Let i^* denote the efficient seller, and k_j^* denote the efficient level of attribute j, denoted k^* when the context of attribute j is clear. Let Γ enumerate the space of feasible attribute bundles.

The restricted dual, [RDMAP-2], computes the maximal CE prices, that maximize the payment to the seller across all CE prices. The constraints are constructed to make prices satisfy the CS conditions with the efficient attribute bundle and efficient seller.

$$\max_{p_{jk},\Delta_i} \quad \sum_{j\in\mathcal{J}} p_{jk^*} \ - \ \Delta_{i^*} \tag{RDMAP-2}$$

$$\text{s.t.} \quad v_{jk^*} \geq p_{jk^*} \geq c_{i^*jk^*}, \quad \forall j \in \mathcal{J} \tag{32}$$

$$v_{jk_j^*} - p_{jk_j^*} \geq v_{jk} - p_{jk}, \quad \forall j \in \mathcal{J}, \forall k \neq k_j^* \tag{33}$$

$$p_{jk_j^*} - c_{i^*jk_j^*} \geq p_{jk} - c_{i^*jk}, \quad \forall j \in \mathcal{J}, \forall k \neq k_j^* \tag{34}$$

$$\sum_{j\in\mathcal{J}} p_{jk^*} - c_{i^*jk^*} \geq \Delta_{i^*} \tag{35}$$

$$\Delta_{i^*} \geq \Delta_i, \quad \forall i \neq i^* \tag{36}$$

$$\Delta_i \geq \sum_{(j,k)\in\gamma} p_{jk} - c_{ijk}, \quad \forall \gamma \in \Gamma, \forall i \neq i^* \tag{37}$$

$$p_{jk}, \Delta_i \geq 0$$

There are always CE prices in which the penalty is the same across all sellers. As a simple example, prices $p_{jk} = v_{jk}$ and penalty $\Delta = V(\mathcal{I} \setminus i^*)$ define anonymous CE prices for the preferential-independence MAP problem; constraints (32) to (35) in [RDMAP-2] hold with these prices. However, non-zero penalty terms are required whenever the efficient seller does not dominate the other sellers, in terms of *value−cost*, across all attribute types.

Definition 4 (cost-dominates). *Seller i is said to cost-dominate seller i' if*

$$\max[0, \max_{k\in\Theta_j} v_{jk} - c_{ijk}] \geq \max[0, \max_{k\in\Theta_j} v_{jk} - c_{i'jk}], \quad \forall j \in \mathcal{J}$$

Proposition 4 (linear prices). *Linear CE prices, with zero penalty terms, exist in the preferential-independence MAP problem if and only if the efficient seller cost-dominates every other seller.*

Proof. Assume penalty, $\Delta_i = 0$, for all sellers i. Conditions (36) and (35) hold, but (37) requires $0 \geq \sum_{(j,k)\in\gamma} p_{jk} - c_{ijk}$, for all $\gamma \in \Gamma$, and all $i \neq i^*$, which implies $p_{jk} \leq \min_{i\neq i^*} c_{ijk}$ for all j, k. To construct minimal prices, p_{jk}, and satisfy (32) and (33), set $p_{jk_j^*} = c_{i^*jk_j^*}$ and $p_{jk} = v_{jk} - v_{jk_j^*} + c_{i^*jk_j^*}$, for all j, all $k \neq k_j^*$. Putting this together, we require, $v_{jk^*} - c_{i^*jk^*} \geq v_{jk} - \min_{i\neq i^*} c_{ijk}$, for all j, k, which is equivalent to $v_{jk^*} - c_{i^*jk^*} \geq \max_{i\neq i^*, k\in\Theta_j} v_{jk} - c_{ijk}$, for all j; in other words seller i^* must dominate all other sellers. $\square$

To characterize the set of maximal CE prices it is useful to reformulate [RDMAP-2] as:

$$\max_{p_{jk^*}} \left(\sum_{j\in\mathcal{J}} p_{jk^*} - \max_{i\neq i^*} \sum_{j\in\mathcal{J}} \max\left[0, \max_{k\in\Theta_j}(v_{jk} - v_{jk^*} + p_{jk^*} - c_{ijk})\right]\right) \tag{RD}$$

$$\text{s.t.} \quad c_{i^*jk^*} \leq p_{jk^*} \leq v_{jk^*}, \quad \forall j \in \mathcal{J}$$

Constraints (32) in [RDMAP-2] are explicitly maintained in this formulation. Then, we fix values for p_{jk}, and assign values $\Delta_i = \max_{\gamma\in\Gamma} \sum_{(j,k)\in\gamma} p_{jk} - c_{ijk}$ and $\Delta_{i^*} = \max_{i\neq i^*} \Delta_i$. This provides penalties, Δ_i, that minimize Δ_{i^*} and satisfy (36) and (37). Assume for now that (35) holds. Now, fix values for p_{jk^*}, and assign prices $p_{jk} = \max(0, v_{jk} - v_{jk^*} + p_{jk^*})$, for $k \neq k^*$. These are the minimal prices that satisfy (33), and selected to minimize Δ_{i^*}. Constraint (34) also holds with these prices. To see this, we must show $\max(0, v_{jk} - v_{jk^*} + p_{jk^*}) \leq p_{jk^*} - c_{i^*jk^*} + c_{i^*jk}$, for all $k \neq k^*$ and all j. Case (i), with $v_{jk} - v_{jk^*} + p_{jk^*} \geq 0$ holds because $v_{jk^*} - c_{i^*jk^*} \geq v_{jk} - c_{i^*jk}$. Case (ii), with $v_{jk} - v_{jk^*} + p_{jk^*} < 0$ holds because $p_{jk^*} \geq c_{i^*jk^*} \Rightarrow p_{jk^*} - c_{i^*jk^*} + c_{i^*jk} \geq 0$.

To show (35), we need $\max_{i\neq i^*}\left[\max_{\gamma\in\Gamma} \sum_{(j,k)\in\gamma} p_{jk} - c_{ijk}\right] \leq \sum_{j\in\mathcal{J}} p_{jk^*} - c_{i^*jk^*}$ Consider seller $i' \neq i^*$, and the worst-case scenario in which $p_{jk} = \max(0, v_{jk} - v_{jk^*} + p_{jk^*}) = v_{jk} - v_{jk^*} + p_{jk^*}$ for all $k \neq k^*$. But, $\max_{\gamma\in\Gamma} \sum_{(j,k)\in\gamma}(v_{jk} - v_{jk^*} + p_{jk^*} - c_{i'jk}) \leq \sum_{j\in\mathcal{J}} p_{jk^*} - c_{i^*jk^*}$ because $\max_{\gamma\in\Gamma} \sum_{(j,k)\in\gamma} v_{jk} - c_{i'jk} \leq \sum_{j\in\mathcal{J}} v_{jk^*} - c_{i^*jk^*}$ for all $i' \neq i^*$. Turning to the objective function, the penalty to the winning seller, i^*, $\Delta_{i^*} = \max_{i\neq i^*, \gamma\in\Gamma} \sum_{(j,k)\in\gamma}(p_{jk} - c_{ijk})$ is equivalent to $\max_{i\neq i^*} \sum_{j\in\mathcal{J}} \max\left[0, \max_{k\in\Theta_j}(\max(0, v_{jk} - v_{jk^*} + p_{jk^*}) - c_{ijk})\right]$. The nested max in this expression is redundant.

As before, let $\tilde{i}$ denote the second-best seller, and let $(\tilde{k}_1, \ldots, \tilde{k}_m)$ denote the second-best attribute bundle. Often we write $\tilde{k}$ as shorthand for $\tilde{k}_j$, when the context of attribute j is clear.

Lemma 3 (maximal CE prices). *Given the efficient outcome, $(i^*, k_1^*, \ldots, k_m^*)$, and the second-best outcome, $(\tilde{i}, \tilde{k}_1, \ldots, \tilde{k}_m)$, prices that satisfy:*

$$\max[c_{i^*jk^*}, v_{jk^*} - (v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}})] \leq p_{jk^*} \leq v_{jk^*}, \quad \forall j \in \mathcal{J} \tag{38}$$

$$\Delta_{i^*} = \Delta_{\tilde{i}} = \sum_{j\in\mathcal{J}} p_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}} \tag{39}$$

$$p_{j\tilde{k}_j} = v_{j\tilde{k}_j} - v_{jk_j^*} + p_{jk_j^*}, \quad \forall j \in \mathcal{J}, \forall \tilde{k}_j \neq k_j^* \tag{40}$$

define maximal CE prices for the preferential-independence MAP problem.

Proof. First, if seller i' maximizes the second term in the objective in [RD], then maximal CE prices, p_{jk^*}, must set $\max_{k\in\Theta_j}(v_{jk} - v_{jk^*} + p_{jk^*} - c_{i'jk^*}) \geq 0$ for all j. Otherwise, if this fails from some j', then a small, $\epsilon > 0$, increase to $p_{j'k^*}$ increase the objective value by ϵ. Let $\mathcal{P}$ denote the set of prices that satisfy this condition. With this, the problem reduces to

$$\max_{p_{jk^*}\in\mathcal{P}} \sum_{j\in\mathcal{J}} p_{jk^*} - \max_{i\neq i^*} \sum_{j\in\mathcal{J}} \max_{k\in\Theta_j}(v_{jk} - v_{jk^*} + p_{jk^*} - c_{ijk})$$

Prices, $\sum_j p_{jk^*}$, cancel, and we find that all solutions that satisfy $\mathcal{P}$ have the same value. To characterize prices that satisfy $\mathcal{P}$, note that we need $p_{jk^*} \geq v_{jk^*} - (v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}})$, so that seller the second-best seller, $\tilde{i}$, has positive surplus at prices $p_{j\tilde{k}}$. □

At the maximal CE prices, the maximal utility to the second-best seller is exactly zero.

Lemma 4. *The second-best seller has zero utility for the second-best attributes,* $(\tilde{k}_1, \ldots, \tilde{k}_m)$*, at the maximal CE prices.*

Proof. First, substitute $p_{jk^*} = \max[c_{i^*jk^*}, v_{jk^*} - (v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}})]$ into $p_{j\tilde{k}_j} = v_{j\tilde{k}_j} - v_{jk^*_j} + p_{jk^*_j}$ to verify that $p_{j\tilde{k}_j} \geq c_{\tilde{i}j\tilde{k}_j}$ for all j. Then, $\sum_j \pi_{ij} = \sum_j p_{j\tilde{k}_j} - c_{\tilde{i}j\tilde{k}} = \Delta_{\tilde{i}}$. □

To complete the description of maximal CE prices, we construct the prices, p_{jk}, on attribute levels other than k^* and $\tilde{k}$, and the penalties, Δ_i, for sellers other than i^* and $\tilde{i}$, as:

$$\begin{aligned} \Delta_{\tilde{i}} \quad \geq \Delta_i \quad &\geq \sum_{j\in\mathcal{J}} \max[0, \max_{k\in\Theta_j}(p_{jk} - c_{ijk})] \\ &, \qquad \forall i \notin \{i^*, \tilde{i}\} \qquad (41) \\ \max[0, v_{jk} - v_{jk^*} + p_{jk^*}] \quad \leq p_{jk} \quad &\leq p_{jk^*} - c_{i^*jk^*} + c_{i^*jk} \\ &, \qquad \forall j \in \mathcal{J}, \forall k \notin \{\tilde{k}_j, k^*_j\} \qquad (42) \end{aligned}$$

There is some flexibility in setting the remaining values, Δ_i, for $i \neq i^*, \tilde{i}$, and p_{jk}, for $k \neq k^*, \tilde{k}$. Constraints (41) make the penalty, Δ_i, large enough to satisfy (37), so that seller $i \neq i^*$ does not want to trade at the adjusted prices, but no larger than the critical penalty value, $\Delta_{\tilde{i}}$, defined by the second-best seller. Constraints (42) fix the prices to maintain (33) and (34), and CS conditions for the buyer and the winning seller.

As a special case, we can set:

$$\Delta_i = \sum_{j\in\mathcal{J}} p_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}}, \quad \forall i \in \mathcal{I} \tag{43}$$

$$p_{jk} = \max[0, v_{jk} - v_{jk^*} + p_{jk^*}], \quad \forall j \in \mathcal{J}, \forall k \notin \{\tilde{k}_j, k^*_j\} \tag{44}$$

It follows that the simple CE prices introduced earlier, $p_{jk} = v_{jk}$ for all j, k and $\Delta_i = V(\mathcal{I} \setminus i^*)$, are themselves maximal CE prices.

We can check that these prices (43) and (44) satisfy conditions (41) and (42). First, constraint (42) holds because $p_{jk^*} \geq c_{i^*jk^*}$ and $v_{jk^*} - c_{i^*jk^*} \geq v_{jk} - c_{i^*jk}$. For (41), assume otherwise, that $\sum_{j \in \mathcal{J}} \max[0, \max_{k \in \Theta_j}(v_{jk} - v_{jk^*} + p_{jk^*} - c_{i'jk})] > \Delta_{\tilde{i}}$ for some $i' \notin \{i^*, \tilde{i}\}$, and prove a contradiction. Let k'_j (or k' where the context is clear) denote the efficient level of attribute j for seller i'. Consider new prices, $p'_{jk^*} = \max(c_{i'jk'} + v_{jk^*} - v_{jk'}, p_{jk^*})$. Then, substituting for $\Delta_{\tilde{i}}$ and replacing p_{jk^*} with p'_{jk^*}, we have $\sum_{j \in \mathcal{J}}(v_{jk'} - v_{jk^*} + p'_{jk^*} - c_{i'jk'}) > \sum_{j \in \mathcal{J}}(v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}} + p'_{jk^*} - v_{jk^*})$. This is a contradiction, because $\sum_{j \in \mathcal{J}}(v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}}) \geq \sum_{j \in \mathcal{J}}(\max_{k \in \Theta_j}(v_{jk} - c_{ijk}))$ for all $i \notin \{i^*, \tilde{i}\}$.

The following lemma provides conditions for the existence of linear maximal CE prices, in which the penalty terms are all zero.

Lemma 5 (maximal and linear CE prices). *It is necessary and sufficient for the efficient seller to cost-dominate the second-best seller, and for the second-best seller to cost-dominate every other seller, for the existence of linear and maximal CE prices, i.e. with the penalty set to zero for all sellers.*

Proof. We already know that it is necessary for the efficient seller to dominate the other sellers for anonymous CE prices with zero penalty. So, assume this, and assume $\Delta_i = 0$, for all i, which implies that $p_{j\tilde{k}_j} = c_{\tilde{i}j\tilde{k}_j}$, for all $j \in \mathcal{J}$ and $\tilde{k}_j \neq k^*_j$. Then, to minimize prices p_{jk_j} on $k_j \neq k^*_j$, set $p_{jk^*} = v_{jk^*} - (v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}})$, for all j. This satisfies (38) because seller i^* dominates seller $\tilde{i}$. Finally, this implies that $p_{jk} = \max(0, v_{jk} - (v_{jk^*_j} - p_{jk^*_j}))$, for all $j \in \mathcal{J}$, and all $k \notin \{\tilde{k}_j, k^*_j\}$. We need $p_{jk} = \max(0, v_{jk} - (v_{jk^*_j} - p_{jk^*_j}))$, satisfy $p_{jk} \leq c_{ijk}$, for all $i \neq i^*, \tilde{i}$ and all j, k, so that (41) holds. Substituting for $p_{jk^*_j}$, and taking the interesting case that $v_{jk} - (v_{j\tilde{k}_j} - c_{\tilde{i}j\tilde{k}_j}) > 0$, we need $v_{jk} - (v_{j\tilde{k}_j} - c_{\tilde{i}j\tilde{k}_j}) \leq c_{ijk}$, for all $i \notin \{i^*, \tilde{i}\}$. This holds if and only if the second-best seller $\tilde{i}$ dominates every other seller on all attributes. ◻

Notice that a stronger dominance requirement than that for the existence of linear CE prices is required for the existence of linear and maximal CE prices.

In the next section we establish an equivalence between maximal CE prices and VCG payments, both for the non-linear MAP problem and this preferential-independence variation.

3 The MAP Mechanism Design Problem

Mechanism design addresses the problem of implementing an outcome in a distributed problem to maximize an objective function, where the optimal solution depends on private information held by individual participants and individual participants are self-interested and willing to misrepresent their private information if that is to their own personal advantage. Jackson [9] provides a useful review of mechanism design theory, describing important characterization results, both positive and negative, for the classes of social choice functions that can be implemented in equilibrium.

Auctions are mechanisms for resource allocation problems, in which the private information of participants represents their preferences over different allocations and prices. Desirable equilibrium properties of a multiattribute auction mechanism include:

- *Efficiency.* The outcome solves the MAP problem, selecting the seller, and attribute levels, that maximize the difference between the value of the buyer and the cost of the seller, over all possible outcomes.

- *Buyer-optimality.* The outcome maximizes the expected utility of the buyer, the difference between its value for the item and the price that it pays given sellers with cost functions drawn from a particular distribution.

- *Individual-rationality.* The utility of the buyer and every seller is non-negative.

- *Budget-balance.* The total payment made by the buyer equals the payments received by the sellers.

Taking efficiency as the primary goal, one positive result, for the general MAP problem, is provided by the VCG [22, 6, 8] mechanism. The VCG mechanism is efficient and individual-rational, and implements the outcome in a *dominant-strategy* equilibrium. This is a robust solution concept because truthful bidding is a dominant strategy for each agent irrespective of the preferences and strategies of other agents. However, the VCG mechanism has two main problems as a mechanism for multiattribute auctions. First, the VCG mechanism is not budget-balanced; the payment by the buyer is typically less than the payment received by the winning seller, and the VCG requires a subsidy from the auctioneer. Second, the VCG mechanism is a direct revelation mechanism, which requires that agents reveal, and compute, complete information about their preferences. This is often undesirable in practice, both because participants are in long-term strategic relationships and also because preference evaluation can be quite costly in industrial settings.

Indeed, a classic result in mechanism design, the Myerson-Satterthwaite impossibility result [14], states that we cannot expect an efficient and *budget-balanced* mechanism, even in Bayesian-Nash equilibrium, for problems with private information on two sides of a market. Thus, unlike the combinatorial allocation problem, in which it is a common assumption that the seller has no reservation price for items [7], we must expect budget-balance to fail for any efficient mechanism for the MAP problem.

Taking buyer-optimality as the primal goal, and holding budget-balance as a constant, Che [5] describes an optimal multiattribute auction mechanism for a special case, with two attributes and continuous attribute levels. Following the optimal auction design approach introduced in Myerson's [13] classic paper, Che formulates an optimal reservation price to incorporate into the VCG mechanism. Truth-revelation remains a dominant strategy for sellers in Che's auction, although the buyer selects an optimal reservation price in a Bayes-Nash equilibrium, given a belief about the costs of suppliers. The effect of the reservation price is to decrease the average payment received by the winning seller, at the cost of an occasional forfeiting of an efficient trade.

3.1 A Modified VCG Mechanism

We propose a modified VCG mechanism, in which the buyer's payment is defined to equal the payment received by the seller. In general, this payment is greater than its VCG payment, this difference opens up an opportunity to the buyer for non-truthful preference revelation. In the following we use $\hat{v}$, $\hat{c}_1$, etc. to indicate that the sellers and the buyer are not assumed to follow truthful strategies; the intention is to allow $\hat{v} \neq v$ and $\hat{c}_i \neq c_i$.

Definition 5 (MODIFIED-VCG). *Given bids* $(\hat{c}_1, \ldots, \hat{c}_n)$ *from the sellers, and ask* $\hat{v}$ *from the buyer, the modified VCG implements the outcome that solves*

$$\max_{\theta \in \Theta, i \in \mathcal{I}} \hat{v}(\theta) - \hat{c}_i(\theta)$$

and the buyer makes payment

$$\hat{p}_{\text{mvick}}(\hat{v}, \hat{c}) = \hat{c}_{\hat{i}}(\hat{\theta}) + (\hat{V}(\mathcal{I}) - \hat{V}(\mathcal{I} \setminus \hat{i})$$

to the winning seller, where $(\hat{\theta}, \hat{i})$ *denotes the outcome,* $\hat{V}(\mathcal{I})$ *the reported value minus cost of this outcome, and* $\hat{V}(\mathcal{I} \setminus i)$ *the reported value minus cost of the best outcome without seller* i.

Proposition 5. *Truthful bidding is a weakly dominant strategy for a seller in the modified VCG mechanism.*

Proof. Consider seller 1, with bid, $\hat{c}_1(\theta)$, and fix the bids from the other sellers and the ask from the buyer. First, suppose that seller 1 is selected in the outcome. Then the seller's utility is $u_1(\hat{\theta}, p) = -c_1(\hat{\theta}) + \hat{p}_{\text{mvick}} = -c_1(\hat{\theta}) + \hat{c}_1(\hat{\theta}) + \hat{V}(\mathcal{I}) - \hat{V}(\mathcal{I} \setminus 1)$, which reduces to $u_1(\hat{\theta}, p) = -c_1(\hat{\theta}) + \hat{v}(\hat{\theta}) - \hat{V}(\mathcal{I} \setminus 1)$. The seller's bid has no effect on the value of the last term, but determines the value of the first two terms indirectly via the choice of $\hat{\theta}$. Because this attribute value is selected in the mechanism to maximize the difference between reported value and reported cost, the seller's dominant strategy is to report a truthful cost function. The utility from this outcome is non-negative, and the seller should submit this bid rather some bid that leaves it out of the outcome. $\square$

As a special case, notice that if the buyer is truthful and the efficient and second-best sellers compete on the same attribute bundle, i.e. $\theta^* = \tilde{\theta}$, then $p_{\text{mvick}} = c_{\tilde{i}}(\theta^*)$, and the lowest cost seller receives a payment that is equal to the minimal price that the second-best seller was willing to accept. As another special case, notice that if $V(\mathcal{I} \setminus i^*) = 0$, but $V(\mathcal{I}) > 0$, that $p_{\text{mvick}} = v(\theta^*)$, and the winning seller extracts all the surplus from the buyer.

Theorem 1. *The modified-VCG is budget-balanced and efficient with a truthful buyer.*

Proof. Immediate, because the dominant strategy for sellers is truthful revelation of costs, and the outcome is selected to maximize the difference between reported value and reported cost. Budget-balance holds by construction. □

Lemma 6 (seller individual-rational). *The modified VCG mechanism is ex post individual-rational for a rational seller, whatever the strategy of the other agents.*

Proof. The winning seller, $\hat{i}$, has utility $\pi_{\hat{i}} = p_{\text{mvick}} - c_{\hat{i}}(\hat{\theta}) = \hat{V}(\mathcal{I}) - \hat{V}(\mathcal{I} \setminus \hat{i})$, which is non-negative for any strategies, and preferences, of the other sellers and the buyer. □

Lemma 7 (buyer individual-rational). *The modified VCG mechanism is ex post individual-rational for any buyer strategy that is truthful, or under-reports its value.*

Proof. If the buyer is truthful, then its utility, $\pi^B = v(\hat{\theta}) - p_{\text{mvick}} = \hat{V}(\mathcal{I} \setminus \hat{i})$, is non-negative for any seller strategies. Similarly, if $\hat{v}(\theta) \leq v(\theta)$ for all θ, then the buyer's utility for the outcome, $(v(\hat{\theta}) - \hat{v}(\hat{\theta})) + \hat{V}(\mathcal{I} \setminus \hat{i})$, remains non-negative. □

Following the direction of Che [5], having shown that truthful bidding is a weakly dominant strategy for sellers, we could at this stage compute a Bayesian-Nash equilibrium, in which the buyer determines a reported value, $\hat{v}$, to maximize her expected payoff with respect to beliefs about the distribution over the cost functions of the sellers.

However, our main goal in this paper is to derive iterative variations on these modified VCG mechanisms, and we leave the full analysis of a Bayesian-Nash equilibrium for future work. Instead, we make a reasonable assumption about the class of buyer strategies in the iterative auctions, and then characterize seller equilibrium strategies with respect to that assumption. For now, we choose to bound the maximal ex post benefit to a perfectly informed buyer for deviation from truthful bidding.[1]

Definition 6 (value of manipulation). *The value of manipulation is the maximal ex post utility gain available to an agent in a mechanism in comparison to a truthful strategy.*

The payoff to the winning seller for a truthful bid is $V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$, and the payoff to the buyer in this equilibrium is $V(\mathcal{I} \setminus i^*)$. The winning seller and the buyer share the overall surplus in the system, $V(\mathcal{I})$.

[1] Parkes et al. [17, 16] introduce a similar analysis of VCG-based budget-balanced mechanisms to clear combinatorial exchanges, referring to the *degree of manipulation freedom* of a mechanism.

Proposition 6. *The value of manipulation to the buyer in the modified VCG is $V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$.*

Proof. Fix the bids, $(\hat{c}_1, \ldots, \hat{c}_n)$, from the sellers, and construct the *ex post* decision problem facing the buyer. The buyer should report a valuation to solve

$$\max_{\hat{v} \in V} [v(\theta_{\text{mvick}}(\hat{v}, \hat{c})) - p_{\text{mvick}}(\hat{v}, \hat{c})]$$

where V is the space of valuation functions. In the best case, there exists a bid that equates $\hat{V}(\mathcal{I})$ and $V(\mathcal{I} \setminus i^*)$, where $\hat{V}(\mathcal{I})$ denotes the reported surplus of the optimal solution. In this case, the optimal strategy, $\hat{v}$, selects the efficient outcome, to maximize $v(\hat{\theta}) - c_{\tilde{i}}(\hat{\theta})$, and the maximal gain in surplus is $V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$. □

As a simple special case, the following buyer strategy is *ex post* rational for a buyer whenever the second-best seller, $\tilde{i}$, has a lower cost than the efficient seller, i^*, on the second-best attribute bundle, $\tilde{\theta}$.

$$v_{\text{manip}}(\theta) = \begin{cases} v(\theta) & , \forall \theta \neq \theta* \\ v(\theta^*) + \delta, & \text{,otherwise} \end{cases}$$

for $\delta = (V(\mathcal{I}) - V(\mathcal{I} \setminus i)) - \max(0, c_{\tilde{i}}(\tilde{\theta}) - c_{i^*}(\tilde{\theta}))$. To verify this, show that $\hat{V}(\mathcal{I}) = V(\mathcal{I}) \geq \hat{V}(\mathcal{I} \setminus i^*) = V(\mathcal{I} \setminus i^*) + \delta$ after the adjustment. We have $\delta = V(\mathcal{I}) - V(\mathcal{I} \setminus i)$, and $\hat{V}(\mathcal{I}) = V(\mathcal{I}) = \hat{V}(\mathcal{I} \setminus i^*)$. Thus, $p_{\text{mvick}} = c_{i^*}(\theta^*)$ and $\pi^B = v(\theta^*) - c_{i^*}(\theta^*) = V(\mathcal{I})$.

Notice that the value of manipulation to the buyer can be expected to be considerably less than the buyer's surplus from truthful bidding, $V(\mathcal{I})$ in many problems. As the auction becomes more competitive, and $V(\mathcal{I}) - V(\mathcal{I} \setminus i^*) \to 0$, then the gains from manipulation available to the seller tend to zero. Together with the reality of risk-averse and poorly-informed buyers, this provides some support for buyer truth-revelation in the iterative multiattribute reverse auctions that we propose in Sections 4 and 5.

3.2 Dual Prices and Vickrey Payments

In this section we show an equivalence between VCG payments and maximal CE prices. This equivalence is critical in the theoretical analysis of the iterative multiattribute auctions that we propose in the following sections.

First, consider the dual prices in DMAP-1, that are non-linear prices, $p(\theta)$, for attribute levels $\theta = (\theta_1, \ldots, \theta_m)$.

Theorem 2. *The Vickrey payment in problem MAP-1 is supported in the maximal competitive equilibrium.*

Proof. From Lemma 1, we have $p(\theta^*) = v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$, which is exactly $c_i(\theta^*) + V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$, since $(\tilde{i}, \tilde{\theta})$ is the second-best solution. □

Now, consider the dual prices in DMAP-2, in which agent values and costs satisfy the preferential-independence conditions.

Theorem 3. *The Vickrey payment in problem MAP-2 is supported in the maximal competitive equilibrium.*

Proof. The payment made to seller i^* at maximal CE prices is $\sum_{j \in \mathcal{J}} v_{jk^*} - \sum_{j \in \mathcal{J}} v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}_j}$, where $(\tilde{i}, \tilde{k}_1, \ldots, \tilde{k}_m)$ is the second-best solution. Again, this is exactly $c_i(\theta^*) + V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$. □

4 Auction 1: Non-Linear Preferences

In this section we propose auction, NonLinear&Discrete, which is a descending-price multiattribute auction. The auction implements a primal-dual algorithm for the MAP-1/DMAP-1 formulation of the MAP problem, and terminates with the outcome of the modified VCG mechanism when agents follow straightforward bidding strategies. We show that straightforward bidding is an *ex post* best-response strategy for sellers, for a reasonable class of buyer strategies.

We initially assume that $V(\mathcal{I} \setminus i^*) > 0$, i.e. that there is an efficient trade without the optimal seller. This assumption will be relaxed in Section 4.4, in which we provide a slight variation of the auction to handle this special case.

Figure 1 provides the top-level structure of the auction. The auction proceeds in rounds $t \geq 1$, and maintains *ask prices*, $p^t(\theta) \geq 0$, on every attribute bundle $\theta \in \Theta$. These prices are anonymous, so that every seller faces the same prices, but can be completely non-linear, with $p(\theta) \neq p(\theta_1) + p(\theta_2)$, for some partition of $\theta \in \Theta$ into θ_1 and θ_2. The auction also maintains a provisional allocation, `alloc`, which indicates the current winning seller, attribute bundle, and price. Auction NonLinear&Discrete is parameterized with a minimal bid increment, ϵ, which determines the rate at which prices are decreased across rounds. Prices are initialized to a set of high ask prices, denoted $p^1(\theta) \leftarrow p_\infty(\theta)$. It is sufficient that the initial ask price, $p^1(\theta)$, on attributes θ is at least as large as $\min_{i \neq i^*} c_i(\theta)$ for the analysis in Section 4.1 to hold.

```
AUCTION NonLinear&Discrete:
t ← 0; p^1(θ) ← p_∞(θ); alloc ← ∅;
while (¬quiescence) {
   t ← t + 1;
   bid_i ← BID_i(p^t);
   alloc ← ASK(bid_1,...,bid_n);
   p^{t+1} ← update_prices(p^t,bid_1,...,bid_n,alloc)
}
return(alloc);
```

Figure 1: Auction Non-linear&Discrete.

In each round a seller can submit bids, represented $\texttt{bid}_i \leftarrow \text{BID}_i(p^t)$ in Figure 1, on multiple attribute bundles. Each attribute bundle, θ', within a bid is associated with a *bid price*, which indicates that the seller is prepared to provide attributes, θ', for any price greater than or equal to that price.

The bid price must be less than or equal to the ask price for the attribute bundle, except in two special cases: (1) a seller can repeat a bid that is successful in the current provisional allocation at the same price, even if the ask price has decreased across rounds; (2) a seller can take an "ϵ-discount", and bid at ϵ above the ask price, on any attribute bundle and in any round, but can never bid a lower price on that attribute bundle in any future round.

The bids ($\texttt{bid}_1, \ldots, \texttt{bid}_n$) are collected from the sellers, and then passed to the buyer. In this step, denoted $\texttt{alloc} \leftarrow \texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$, the buyer is asked for its preferred bid across all bid sets. We discuss a general method in Section 4.5 to introduce a *proxy agent*, to reduce the extent to which the buyer must be active in this winner-determination step. The new provisional allocation, `alloc`, stores the winning bid (the attribute bundle and the bid price), and the index of the winning bidder.

The ask prices are updated at the end of each round, in step $p^{t+1} \leftarrow \texttt{update\_prices}(p^t, \texttt{bid}_1, \ldots, \texttt{bid}_n, \texttt{alloc})$. The price changes are based on bids from sellers that are not in the current provisional allocation, call these the *unsuccessful agents*. The new ask prices, p^{t+1}, are reduced to ϵ, the minimal bid increment, below the lowest unsuccessful bid price, as follows:

$$p^{t+1}(\theta) = \min\left(p^t(\theta), \min_{i \notin \texttt{alloc}} p^t_{\text{bid},i}(\theta) - \epsilon\right), \quad \forall \theta \in \Theta$$

where $i \notin \texttt{alloc}$ indicates that seller i is not in the current allocation, and $p^t_{\text{bid},i}(\theta)$ is the bid price from seller i in the current round, for any attribute bundle, θ, for which the seller submitted a bid in the current round, and ∞ otherwise.

The auction terminates whenever it is in `quiescence`, which is defined to hold whenever the ask prices have not changed for two consecutive rounds. At termination, the provisional allocation becomes the final allocation. For now, because we assume, $V(\mathcal{I} \setminus i^*)$, this final allocation will always be individual-rational for a truthful buyer with equilibrium seller strategies.

4.1 Primal-Dual Analysis

To continue, we assume that every seller, and the buyer, follow a straightforward *myopic best-response* bidding strategy. Just as in *i*Bundle [15, 18], we use duality theory to establish the efficiency of the auction for this straightforward bidding strategy. Later, in Section 4.2, we will establish that myopic best-response is an *ex post* best-response for sellers against a reasonable class of buyer strategies. This set, the *consistent strategies*, includes a myopic best-response strategy for the buyer.

We will assume that the buyer follows a *truthful* and consistent strategy, i.e. myopic best-response. The analysis trivially extends to the case in which the buyer follows some non-truthful consistent strategy, for a valuation $\hat{v} \neq v$. In this case the auction implements a primal-dual algorithm for the MAP problem instance defined on valuation $\hat{v}$ and seller cost functions. At this stage we also assume that there is a second-best solution, i.e. $V(\mathcal{I} \setminus i^*) > 0$. We discuss a simple extension to the auction to handle this special case, when there is only one seller that can provide services for less cost than the buyer's value in Section 4.4.

Definition 7 (seller myopic best-response). *Bids,* $\texttt{BID}_i(p^t)$*, are myopic best-response from seller* i*, given minimal bid increment,* ϵ*, when*

$$\texttt{BID}_i(p^t) = \{(\theta, p^t(\theta)) \ : \ p^t(\theta) - c_i(\theta) + \epsilon \geq \max[0, \max_{\theta' \in \Theta}(p^t(\theta') - c_i(\theta'))]\}$$

Definition 8 (buyer myopic best-response). *Strategy,* $\texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$*, is a myopic best-response for the buyer, when the provisional outcome,* $(\texttt{seller}^t, \texttt{attr}^t, \texttt{price}^t)$*, satisfies*

$$v(\texttt{attr}^t) - \texttt{price}^t + \epsilon \geq \max_{i \in \mathcal{I}} \max_{(\theta,p) \in \texttt{bid}_i} (v(\theta) - p)$$

and $(\texttt{attr}^t, \texttt{price}^t) \in \texttt{bid}_{\texttt{seller}^t}$.

In other words, in myopic best-response (MBR), seller i takes prices as given and bids for the attribute bundles that ϵ-maximize its surplus, and the buyer selects the utility-maximizing bid.

Recall that $\overline{p}_{\text{ce}}(\theta)$ denotes the set of maximal CE prices, given an instance of the MAP problem. Ask prices, $p^t(\theta)$, in round t of the auction, are said to *ϵ-dominate* the maximal CE prices, $\overline{p}_{\text{ce}}(\theta)$, if $p^t(\theta) + \epsilon \geq \overline{p}_{\text{ce}}(\theta), \quad \forall \theta \in \Theta$.

Lemma 8. *Auction* NON-LINEAR&DISCRETE *maintains ask prices that ϵ-dominate the maximal CE prices if agents follow MBR strategies.*

Proof. By induction, with induction hypothesis (i.h.) that $p^t(\theta) \geq \overline{p}_{\text{ce}}(\theta)$ for all θ in round $t \geq 1$. The base case is trivial, as long as $p_\infty(\theta) \geq c_i(\theta)$ for at least two sellers, for every $\theta \geq \Theta$. Prove the inductive case by case analysis on conditions for prices to decrease from round t to $t+1$. First, the ask price is reduced on vectors that receive unsuccessful bids from sellers $i \neq \hat{i}^t$. Consider bid (θ', i'), some $i' \neq \hat{i}^t$, and suppose $p^t(\theta') < \overline{p}_{\text{ce}}(\theta')$. Now, by the i.h. the seller will continue to ϵ-prefer vector θ' at all CE prices with $p^t(\theta')$; similarly, the buyer will continue to ϵ-prefer vector $\hat{\theta}^t$. Thus, this price cannot be part of a set of CE prices, a contradiction, and the i.h. is established for this case. A similar argument can be made for the second case, in which the ask price is reduced on vectors offered by the one remaining successful seller. □

The allocation selected at each round in the auction is a feasible primal solution and the ask prices are a feasible dual solution. Notice that CS conditions (CS-2) and (CS-3) always hold, in every round, and that (CS-1) holds when the auction terminates because there is then only one seller with non-negative utility for some outcome and that seller is in the provisional allocation.

Let (θ^*, i^*) denote the efficient outcome, and $(\tilde{\theta}, \tilde{i})$ denote the second-best outcome.

Lemma 9. *At termination the ask prices on bundles, θ^* and $\tilde{\theta}$, equal the maximal CE prices on those bundles.*

Proof. Given Lemma 8, we show that termination implies $p(\theta^*) \leq \overline{p}_{\rm ce}(\theta^*)$ and $p(\tilde{\theta}) \leq \overline{p}_{\rm ce}(\tilde{\theta})$. First, assume $p(\theta^*) > \overline{p}_{\rm ce}(\theta^*)$, and prove a contradiction. This implies that $p(\theta^*) > v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$, that agent i^* is still bidding. Given termination, we must have $p(\tilde{\theta}) \leq c_{\tilde{i}}(\tilde{\theta})$ so that the second-best agent is not still bidding. In addition, we need $v(\theta^*) - p(\theta^*) \geq v(\tilde{i}) - p(\tilde{\theta})$, so that the buyer selects θ^* over $\tilde{i}$. This is not trivially satisfied because we assume that $V(\mathcal{I} \setminus i^*)$, which implies that $v_{\tilde{i}}(\tilde{\theta}) - p(\tilde{\theta}) > 0$. This gives a contradiction with $p(\theta^*) > v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$, and $p(\tilde{\theta}) \leq c_{\tilde{i}}(\tilde{\theta})$. Second, assume that $p(\tilde{\theta}) > \overline{p}_{\rm ce}(\tilde{\theta})$, and prove a contradiction. Termination requires, $p(\theta^*) \leq c_{i^*}(\theta^*)$, else seller i^* still bids, and we have $p(\tilde{\theta}) > c_{\tilde{i}}(\tilde{\theta})$ and $v(\theta^*) - p(\theta^*) \leq v(\tilde{\theta}) - p(\tilde{\theta})$. This implies $v(\theta^*) - c_{i^*}(\theta^*) \leq v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta})$, a contradiction. □

Finally, we show (CS-4), which requires that the buyer maximizes its utility with attribute θ^* across all possible attributes. It is useful to express (CS-4) as two conditions:

$$x^B(\theta) > 0 \;\Rightarrow\; v(\theta) - p(\theta) = \max_{\theta' \in \Theta} [v(\theta') - p(\theta')] \tag{CS-4a}$$

$$x^B(\theta) > 0 \;\Rightarrow\; v(\theta) - p(\theta) \geq 0 \tag{CS-4b}$$

Clearly, (CS-4a) and (CS-4b) imply (CS-4). We have (CS-4b) at the end of the auction, by Lemma 9. Let $\Theta^c \subseteq \Theta$ denote the set of attribute bundles that have changed in price during the auction, and define a relaxed (CS-4a):

$$x^B(\theta) > 0 \;\Rightarrow\; v(\theta) - p(\theta) = \max_{\theta \in \Theta^c} [v(\theta) - p(\theta)] \tag{CS-4a'}$$

Lemma 10. *Condition (CS-4a) holds at the end of the auction.*

Proof. First, prove by induction on rounds, t, that (CS-4a') holds. The base case, in round $t = 1$, is trivial. Then, to prove the inductive case, note that the seller in the current provisional allocation continues to bid for the same attribute bundle across rounds, and that the ask price is only decreased on a bundle with less utility than the attribute bundle in the provisional allocation. Finally, at the end of the auction, we show that (CS-4a') implies (CS-4a). As long as the initial prices, $p^1(\theta) \geq \overline{p}_{\rm ce}(\theta)$, then the buyer has more utility from attributes θ^* at $\overline{p}_{\rm ce}(\theta^*)$ than on any bundles that have not changed price. This follows immediately from CE, because the buyer has more utility for θ^* than other attribute vectors at the CE prices, and therefore also at higher prices. □

Lemma 11 (termination). *Auction* NON-LINEAR&DISCRETE *terminates with consistent agent strategies.*

Proof. While the auction is open the price, $p^t(\theta)$, is reduced on at least one attribute bundle, θ, for which at least one seller, i, has $c_i(\theta) \leq p^t(\theta)$ by MBR. Termination follows, because $\mathcal{I}$ and Θ are finite, and $c_i(\theta) \geq 0$, $\forall \theta, \forall i$. □

Putting this all together, we have the main result.

Theorem 4 (Vickrey outcome). *Auction* NON-LINEAR&DISCRETE *terminates with the outcome of the modified VCG mechanism for agents that follow MBR strategies, as* $\epsilon \to 0$.

Proof. The auction terminates with prices and an allocation that satisfy CS conditions (CS1)-(CS4), and with the price on the efficient attribute bundle equal to the Vickrey payment. □

A more careful analysis can also consider the ϵ approximation within the myopic best-response agent strategies, and derive corresponding bounds on the quality of the approximation to the efficient allocation.

4.2 Game-Theoretic Analysis

In this section we justify the assumption of seller MBR. We show that the set of *ex ante* consistent strategies form a space of reasonable strategies for a buyer in Auction NONLINEAR&DISCRETE. In these strategies the buyer follows a myopic best-response strategy for some, perhaps untruthful valuation function, and chooses this valuation before the auction begins and not during the auction as information is revealed about seller cost functions via their bids. Let V denote the space of valuation functions.

Definition 9 (ex ante consistent strategy). *The buyer follows an ex ante consistent strategy if she chooses a misrepresentation function,* $f : V \to V$, *at the start of the auction, and then follows a MBR strategy with respect to valuation* $f(v)$ *during the auction.*

Given a consistent buyer strategy, then myopic best-response is an *ex post* Nash equilibrium for sellers in Auction NONLINEAR&DISCRETE, without placing any restrictions on the space of seller strategies. This robustness to manipulation is inherited from the seller strategyproofness of the modified VCG mechanism.

Theorem 5 (ex post Nash equilibrium). *Myopic best-response is an* ex post *Nash equilibrium for sellers in Auction* NONLINEAR&DISCRETE *against a buyer that follows an ex ante consistent strategy.*

Proof. Assume a consistent buyer strategy, for some $\hat{v}$, and with out loss of generality assume that sellers follow MBR for some cost functions, $c_2, \ldots, c_n$. We show that for any strategy, $\hat{s}_1$ selected by seller 1, we can construct an equivalent MBR strategy for some cost function, $\hat{c}_1 \neq c_1$, and that implements the same outcome. This implies that strategy, $\hat{s}_1$, selects the VCG outcome for a reported cost function, $\hat{c}_1$, and because truthful bidding is an *ex post* Nash equilibrium for sellers in the modified VCG the seller can do no better than following a truthful MBR strategy. To construct the equivalent cost function, $\hat{c}_1$, given strategy, $\hat{s}_1$, suppose that the auction terminations with seller 1 providing attributes $\hat{\theta}$ at price $\hat{p}$. Cost function $\hat{c}_1(\hat{\theta}) = \hat{p} - \epsilon$, $\hat{c}_1(\theta') = \infty$ for all $\theta' \neq \hat{\theta}$, and a small $\epsilon > 0$ selects the same outcome. The maximum CE price is defined by the other agents, and this cost function will satisfy (CS-1) and (CS-2) with the same prices and allocation. □

This analysis shows that myopic best-response is a sequential best-response of a seller, against any consistent buyer strategy, and against the MBR strategies of other sellers, whatever their cost functions. Although not a dominant strategy for sellers, as it require MBR by other sellers and also a consistent buyer strategy, this is quite a strong solution concept because MBR form a Nash equilibrium *whatever* the costs of other sellers. MBR is not a dominant strategy for sellers, even against a consistent buyer strategy, because other sellers can condition their strategies on information revealed during the auction. Similarly, MBR strategy is not an *ex post* best-response for a seller against any buyer strategy, because the buyer can condition her strategy on information revealed during the auction.

Proposition 7 (value of manipulation). *The maximal value of manipulation to the buyer in the modified VCG, with sellers that follow myopic best-response strategies, is* $V(\mathcal{I}) - V(\mathcal{I} \setminus i^*)$.

Proof. This follows from the analysis of the modified-VCG mechanism. □

4.3 Computational Analysis

Consider a problem with $m = |\mathcal{J}|$ attributes and $l = \max_j |\Theta_j|$ attribute-levels, such that $|\Theta| = O(l^m)$. Let $V_{\max} = \max_{i \in \mathcal{I}}[\max_\theta p_\infty(\theta) - c_i(\theta)]$.

Theorem 6 (complexity). *Auction* NON-LINEAR&DISCRETE *converges in* $O\left(\frac{l^m V_{\max}}{\epsilon}\right)$ *rounds, with minimal bid increment* ϵ.

Proof. The maximal number of rounds that seller i can be unsuccessful in the auction and still have non-negative surplus at the prices is $N_i = [l^m \max_\theta (p_\infty(\theta) - c_i(\theta))/\epsilon]$. After $N_{\max} = \max_i N_i$ rounds, at most $N_{\max}$ valid price decreases remain, one for each provisional winner in each round. Running for a further $N_{\max}$ rounds must take care of this. □

In otherwords, the number of rounds to convergence of NON-LINEAR&DISCRETE is worst-case exponential in m, the number of attributes.

4.4 Special Case: Limited Competition

In this section we briefly consider an auction variation that is required to handle the case in which there is no positive surplus outcome without the efficient seller, i.e. $V(\mathcal{I} \setminus i^*) = 0$. It is also possible that $V(\mathcal{I}) = 0$, and that "no trade" is the efficient outcome. In both of these cases, the bid selected at the end of auction NONLINEAR&DISCRETE will be priced above the value of the buyer.

We propose a simple extension to the auction NONLINEAR&DISCRETE to handle this case. Once the regular auction terminates, the buyer is asked to submit a reported valuation function, $\hat{v}$. The auction then continues, implementing a myopic best response strategy for the buyer, with valuation $\hat{v}$, until: (1) either the seller walks away because the ask price drops to low; or, (2) until the bid price first drops below the reported value on some attribute bundle. In case (1) the auction terminates with no trade. In case (2) the auction terminates with the payment in the modified VCG mechanism.

We implement the final stage of this extended auction with a sealed-bid from the buyer because at this stage the negotiation is 1:1, between the buyer and the final seller, and we wish to prevent adaptive buyer strategies that seek to extract more surplus from the seller and drive the price down below that in the modified VCG mechanism. This sealed-bid stage forces the buyer to select a consistent strategy for this final phase. Instead, we could simply continue to assume a buyer-consistent strategy for this final phase of the auction.

4.5 Acceleration: Proxy Buyer Agents

The basic auction, described in Figure 1, includes the step $\texttt{alloc} \leftarrow \texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$, in which the buyer is asked to select a bid from the bids submitted by sellers at the current ask prices. The buyer is involved in this winner determination step because the best outcome depends not only on the bid price, but also on the buyer's value for different outcomes. At one extreme, if the auctioneer had complete information about the valuation function of the buyer, then this winner determination step could be completed automated. However, this would require that the buyer provides complete information about her valuations for different outcomes at the start of the auction, and we would like to avoid this preference elicitation cost.

As an intermediate method, we propose to introduce a proxy agent, between the buyer and the auction. A similar idea was proposed in Parkes & Ungar [19] in the context of an iterative combinatorial auction. The role of the proxy buyer agent is to maintain approximate and incomplete information about the valuation function of the buyer, and automate the winner-determination step as much as possible. Simply observing the choices of the buyer in each round, and assuming a consistent strategy, the proxy agent could build up a *constraint network* to represent the buyer's valuation. For example, if the buyer chooses attributes θ_1 over θ_2 in some round, at prices p_1 and p_2, then this implies that $v(\theta_1) - p_1 \geq v(\theta_2) - p_2$, and provides value information. The constraint network could then be used to prune bids that cannot be in the MBR set for the buyer, and send only the undominated set of bids to the buyer.

In addition to reducing the preference elicitation cost in the auction, for example by carefully structuring elicitation queries to the buyer to collect enough preference information to follow a MBR strategy, proxy agents may also be useful in reducing the strategy space available to the buyer, and also in speeding-up the auction. First, a proxy agent can enforce *ex post* consistency across rounds, so that the buyer at least follows a strategy that is consistent with some fixed valuation function– even if this valuation is not selected by the buyer before the start of the auction. The extent to which enforcing *ex post* consistency is reduces agents' ability to manipulate the outcome of the auction is an interesting empirical question.

Second, even incomplete information about the preferences of the buyer can be used to speed-up the progress of the auction by the automatic propagation of price changes on one attribute bundle to price changes on other attribute bundles in anticipation of the valuation function of the buyer. In the earlier example, if θ_1 is selected over θ_2, at prices p_1 and p_2, and the price on θ_1 then drops in a future round, to p'_1, then any CE prices must set the price on θ_2, to no greater than $p'_1 - (p_1 - p_2)$. Thus, the price decrease on θ_1 propagates to θ_2, and price changes are accelerated.

5 Auction 2: Preferential-Independence

In this section we first propose a simple auction for a special case of the preferential-independence MAP problem, in which there are *linear* maximal CE prices, and no penalty terms are required. Then, we propose a primal-dual based auction design for the general preferential-independence setting, and present a partial theoretical analysis of its properties.

5.1 Special Case: Cost-Dominance and Linear Prices

In problems in which the efficient seller cost-dominates the second-best seller, and the second-best seller cost-dominates all other sellers, the MAP problem decouples across attribute types and can be solved with m independent auctions, one for each attribute. In addition, the auction problem for each attribute type can be solved with the auction for the general non-linear MAP problem, applied to the special case of a single attribute.

AUCTION LINEAR&DISCRETE:
$t \leftarrow 0$; $p^1(j,k) \leftarrow p_\infty(j,k)$; $\Delta_i \leftarrow 0$; $\texttt{alloc}_j \leftarrow \emptyset$;
while (¬quiescence) {
    $t \leftarrow t+1$;
    for each $j \in \mathcal{J}$ {
        $\texttt{bid}_i(j) \leftarrow \texttt{BID}_i(p^t(j), j)$;
        $\texttt{alloc}_j \leftarrow \texttt{ASK}(\texttt{bid}_1(j), \ldots, \texttt{bid}_n(j))$;
        $p^{t+1}(j) \leftarrow \texttt{update\_prices}(p^t(j), \texttt{bid}_1(j), \ldots, \texttt{bid}_n(j), \texttt{alloc}_j)$;
    }
}
return(alloc);

Figure 2: Auction LINEAR&DISCRETE.

Figure 2 illustrates the structure of Auction LINEAR&DISCRETE, which is a simultaneous descending price auction, with one auction for each attribute type.

The auction maintains separate prices, $(p^t(j,1),\ldots,p^t(j,m))$, for each attribute, and holds the penalty terms at zero. The price adjustment and winner-determination process is fully decoupled across attribute types, although no auction for a single attribute terminates until all auctions are in `quiescence`. Consider the auction for attribute j. In each round, bids $\texttt{bid}_i(j)$, are received from sellers on attribute levels, and the buyer is asked to select a bid, which is assigned to $\texttt{alloc}_j$. The ask price for level k' of attribute j is decreased at the end of a round whenever some seller submits an unsuccessful bid on k'. The new ask price is set to the minimal bid increment, ϵ, below the minimal price across all such unsuccessful bids.

From the perspective of the buyer, she can mix-and-match bids in any particular round from across multiple sellers. No explicit coordination is required across the attributes, because at the end of the auction the efficient seller wins the auction for each attribute, and the same second-best seller sets the price that the winning seller receives. This follows from the seller cost-dominance property.

Theorem 7. *Auction* LINEAR&DISCRETE *terminates with an efficient attribute bundle and CE prices in the preferential-independence MAP problem, when the efficient seller cost-dominates all other sellers.*

Furthermore, the auction terminates with maximal CE prices. As discussed earlier, this requires a slightly stronger dominance requirement.

Theorem 8. *Auction* LINEAR&DISCRETE *terminates with the outcome of the modified VCG mechanism in the preferential-independence MAP problem, when the efficient seller cost-dominates the second-best seller, and the second-best seller cost-dominates all other sellers.*

Proof. Maximal CE prices for the non-linear MAP problem set the price, $p(\theta^*) = v(\theta^*) - (v(\tilde{\theta}) - c_{\tilde{i}}(\tilde{\theta}))$, on the efficient attribute bundle, θ^*. This is the price that just leaves the buyer indifferent between that attribute bundle, and the efficient outcome without seller, i^*, i.e. outcome $(\tilde{\theta},\tilde{i})$. For a single attribute, j, this reduces to the expression, $p_{jk^*} = v_{jk^*} - (v_{j\tilde{k}} - c_{\tilde{i}j\tilde{k}})$, which defines maximal CE prices for the preferential-independence MAP problem. $\square$

5.2 General Case

Figure 3 provides the top-level structure of an auction design, ADDITIVE&DISCRETE, for the general preferential-independence MAP problem.

The auction proceeds in rounds, $t \geq 1$, and maintains linear price terms, $p^t(j,k)$, on level $k \in \Theta_j$, and a single price penalty term that applies to all sellers. We know that this is sufficient for the existence of CE prices. The overall ask price to seller i in round t, on attribute bundle, $(k_1,\ldots,k_m) \in \Theta$, is determined as:

$$p_i^t(k_1,\ldots,k_m) = \sum_{j\in\mathcal{J}} p^t(j,k_j) - \Delta^t$$

Auction ADDITIVE&DISCRETE maintains a provisional allocation, `alloc`, which indicates the current winning seller, attribute bundle, and price. Prices are initialized to a set of high ask prices, $p_\infty(j,k)$, which must be at least as large as maximal CE prices, for example greater than the value, $v_{j,k}$, of the buyer. The initial penalty term, Δ^1, is set to zero.

AUCTION ADDITIVE&DISCRETE:
$t \leftarrow 0$; $p^1(j,k) \leftarrow p_\infty(j,k)$; $\Delta^1 \leftarrow 0$; $\texttt{alloc} \leftarrow \emptyset$;
`while (¬quiescence) {`
    $t \leftarrow t+1$;
    $\texttt{bid}_i \leftarrow \texttt{BID}_i(p^t, \Delta^t)$;
    $(\texttt{alloc}, \texttt{best}_1, \ldots, \texttt{best}_m) \leftarrow \texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$;
    $(p^{t+1}, \Delta^{t+1}) \leftarrow \texttt{update\_prices}(p^t, \Delta^t, \texttt{bid}, \texttt{alloc}, \texttt{best})$
`}`
`return(alloc);`

Figure 3: Auction ADDITIVE&DISCRETE.

In each round a seller can submit bids, represented $\texttt{bid}_i \leftarrow \texttt{BID}_i(p^t, \Delta^t)$. Each bid, $\texttt{bid}_i = (\texttt{bid}_{i1}, \ldots, \texttt{bid}_{im})$, can include multiple attribute-levels and bid prices for each attribute, with $\texttt{bid}_{ij} = \{(k_1, p_{\text{bid},i}(j,1)), \ldots, (k_l, p_{\text{bid},i}(j,l))\}$, with the levels $k_l \in \Theta_j$. A seller must also state an overall bid penalty, $\Delta_{\text{bid},i}$.

Taken together, a bid indicates that a seller will provide an attribute bundle, $(k_1, \ldots, k_m)$, composed of attributes included in its bid, at price:

$$p_{\text{bid},i}(k_1, \ldots, k_m) = \sum_{j \in \mathcal{J}} p_{\text{bid},i}(j, k_j) - \Delta_{\text{bid},i} \tag{45}$$

The bid price, $p_{\text{bid},i}(j,k)$, on level $k \in \Theta_j$ of attribute j, must be less than or equal to the ask price, $p^t_{j,k}$, except in two special cases: (1) a seller can repeat a bid for an attribute level that is successful in the current provisional allocation at the same price, even if the ask price has decreased across rounds; (2) a seller can take an ϵ-discount, and bid at ϵ-above the ask price, on any attribute level in any round, but can never bid a lower price on that attribute level in any future round.

The bid penalty, $\Delta_{\text{bid},i}$, must be greater than or equal to the *ask* penalty, Δ^t, except that a seller can take an ϵ-discount, and bid a penalty at ϵ below Δ^t. This can be used in any round, but once exercised a seller can never submit a higher bid penalty in any future round, and will be excluded from the auction if the ask penalty increases.

The bids, $(\texttt{bid}_1, \ldots, \texttt{bid}_n)$, are collected from the sellers, and then passed to the buyer. In this step, denoted $(\texttt{alloc}, \texttt{best}_1, \ldots, \texttt{best}_m) \leftarrow \texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$, the buyer is asked for the following information:

- the provisional allocation, `alloc`, which must consist of a combination of attribute levels from a *single seller*, and is priced at the penalty-adjusted overall bid price for that seller.

- the set of utility-maximizing bid components, $\texttt{best}_j$, for each attribute type, $j \in \mathcal{J}$, drawn across the bids from all sellers, and selected at bid prices *before* penalties. Multiple components indicates indifference across the components at the current prices.

While, `alloc`, is used to update the provisional allocation, and to determine the attribute bundle when the auction terminates, the auction also uses the additional information to determine price updates. At the end of each round, in step $(p^{t+1}, \Delta^{t+1}) \leftarrow \texttt{update\_prices}(p^t, \Delta^t, \texttt{bid}, \texttt{alloc}, \texttt{best})$, prices are decreased based on bids from sellers that are not in the current provisional allocation, `alloc`. Call these sellers the *unsuccessful* sellers.

The price-update rule is defined as:

- for any unsuccessful seller, i', and for any attribute, j, for which the current winning seller, $i \in \texttt{alloc}$, is also submitting a bid within the preferred set, $\texttt{best}_j$, decrease the price on all levels that receive a bid from i' (except the level that receives a bid from the current winning seller), to the minimum of the current bid price and ϵ below the bid price of the unsuccessful seller.

- for any seller that only submits bids at or below the ask price on an attribute for which the current winning seller is not best, then increase the penalty to the maximum of the current penalty and ϵ above the bid penalty of the seller.

Notice that the first part of this price-update rule is different than that presented for the *dominance* special case. In this variation, the price $p^t(j,k)$ on the level of an attribute is only decreased for attributes in which the current best overall seller, $i \in \texttt{alloc}$, is also in the utility-maximizing set, $\texttt{best}_j$, for each attribute j. The price penalty is used to ensure progress is made toward CE prices when a seller only submits competitive bids on attributes for which the seller in `alloc` is not in the utility-maximizing set.

5.2.1 Theoretical Analysis

We first *assume* that bidders and the seller follow a straightforward MBR bidding strategy in the auction. We prove that the auction terminates with the efficient attribute bundle, maximal CE prices, and therefore the outcome of the modified VCG mechanism. As before, this then provides incentives to make MBR an *ex post* equilibrium for the sellers.

Let $p^t(\theta)$ denote short-hand for the sum of the linear ask prices on attribute bundle θ, similarly for $c_i(\theta)$, $v(\theta)$, and $p_{\mathrm{bid},i}(\theta)$ (which denotes the bid price before the penalty is applied). Also, let $\texttt{bid}_i(j)$ denote the bids from seller i on attribute j. We first define the MBR strategies in this auction.

Definition 10 (seller myopic best-response). *Bids,* $\texttt{BID}_i(p^t, \Delta^t)$*, are myopic best-response from seller* i*, given minimal bid increment,* ϵ*, when*

$$\texttt{BID}_i(p^t, \Delta^t) = \{(\theta, p^t(\theta), \Delta^t) \ : \ p^t(\theta) - \Delta^t - c_i(\theta) + \epsilon \geq \max[0, \max_{\theta' \in \Theta}(p^t(\theta') - c_i(\theta')) - \Delta^t]\}$$

Definition 11 (buyer myopic best-response). *Strategy,* $\texttt{ASK}(\texttt{bid}_1, \ldots, \texttt{bid}_n)$, *is a myopic best-response for the buyer, when the provisional allocation* $\texttt{alloc}$, *specifies attribute bundle* $\hat{\theta}$ *and seller* $\hat{i}$ *that solves:*

$$v(\hat{\theta}) - p_{\mathrm{bid},\hat{i}}(\hat{\theta}) - \Delta_{\mathrm{bid},\hat{i}} + \epsilon \geq \max[0, \max_{i \in \mathcal{I}, \theta \in \Theta} v(\theta) - p_{\mathrm{bid},i}(\theta)]$$

and for every attribute j, *the set,* $\texttt{best}_j$, *solves:*

$$(k'_j, i') \in \texttt{best}_j \Rightarrow v_{jk'_j} - p_{\mathrm{bid},i'}(j,k') + \epsilon \geq \max[0, \max_{i \in \mathcal{I}, k \in \texttt{bid}_i(j)} (v_{jk} - p_{\mathrm{bid},i}(j,k))]$$

It is easy to see that conditions (CS-1), (CS-3), (CS-4), (CS-5), (CS-6), and (CS-7), hold in each round. Notice that (CS-7) holds because all penalty terms are equal in each round. On termination it is immediate that (CS-2) also holds.

We must establish (CS-8) to prove that the auction terminates in competitive equilibrium. This condition requires that the final level selected in the bid of the last remaining seller in the auction maximizes the buyer's utility across *all* attribute levels, and not just across levels restricted to seller bids.

First, we will *assume* that (CS-8) holds on termination, and prove that it follows that the auction will terminate with *maximal* CE prices, out of the space of all possible CE prices.

Lemma 12. *At the end of the auction the prices,* p_{jk}, *on the levels* $k \notin \{k^*_j, \tilde{k}_j\}$, *satisfy conditions (41) and (42), that are required conditions for maximal CE prices, if agents follow MBR strategies and preferential-independence holds.*

Proof. Condition (CS-8) implies that the best seller is also best on each attribute, and from the price-update rule this implies that the other prices are bid down until the buyer is indifferent between the efficient level and the alternative levels. □

Lemma 13. *At least two sellers have a non-negative utility at the penalty-adjusted prices in every round, if agents follow MBR strategies and preferential-independence holds.*

Proof. The price penalty is only increased when there is an unsuccessful seller that faces the same linear price components in the next round. The best seller in each round also faces the same linear price components. Furthermore, both these sellers have non-negative surplus at the current prices, and continue to bid in the next round. □

Lemma 14. *At termination, the prices on the efficient attribute levels, and the penalty term, satisfy the conditions (38), (39) and (40), that are required for maximal CE prices, if agents follow MBR strategies and preferential independence holds.*

Proof. From Lemma 13, at termination the penalty is just enough to leave the second-best seller with zero utility at the adjusted prices. This is the requirement on the penalty term for maximal CE prices. Secondly, because the final prices are CE, then the linear price terms on the efficient attribute levels and the second-best attribute levels are both greater than cost, for the efficient and second-best sellers respectively. Combined with the conditions on the non-efficient price levels, (41) and (42) this is sufficient to show that these prices are within the required range. □

We have an outline of a proof that condition (CS-8) holds on termination. First, define a relaxed definition of (CS-8):

$$x^B_{jk} > 0 \Rightarrow v_{jk} - p_{jk} = \max_{k' \in \Theta_j} v_{jk'} - p_{jk'}, \quad \forall j \in C \tag{CS-8a}$$

$$x^B_{jk} > 0 \Rightarrow v_{jk} - p_{jk} \geq 0, \quad \forall j \tag{CS-8b}$$

where $C \subseteq \mathcal{J}$. Taken together, (CS-8a), for $C = \mathcal{J}$, and (CS-8b) imply condition (CS-8).

Claim. Condition (CS-8) holds at the end of the auction, if agents follow MBR strategies and preferential-independence holds.

The proof structure will first demonstrate that (CS-8a) holds for a monotonically increasing set, C, over the course of the auction. This step uses the fact that linear price terms are only increased on attribute levels for which the winning seller is submitting the best attribute level. Then we show that (CS-8b) and $C = \mathcal{J}$ hold at the end of the auction.

Finally, we present the main result for this auction. At present we present this result as a claim, as a complete proof depends on establishing the previous claim.

Claim. Auction Additive&Discrete terminates, and computes the efficient outcome and Vickrey payoff to the winner seller, if the agents follow MBR, as $\epsilon \to 0$, in the preferential independence MAP problem.

Proof. (outline) The auction maintains a feasible primal (the offers of $\hat{i}^t$) and a feasible dual solution in each round. The auction terminates with solutions that satisfy CS conditions. Combining the earlier lemmas, at termination the prices satisfy the conditions for maximal CE prices, and therefore support the payment in the modified-VCG mechanism. □

As before, we can justify the MBR strategy of sellers for any *consistent* buyer strategy during the auction (Definition 9). We choose to omit this proof, which is very similar to the proof for ex post truth-revelation in Auction Non-Linear&Discrete. The method of the proof is to establish that for MBR strategies of other sellers, and a consistent buyer strategy, any seller strategy selects some outcome of the modified VCG mechanism.

Theorem 9. *Seller MBR is an ex post Nash equilibrium of auction* Additive&Discrete *for any consistent strategy from the buyer, in the preferential-independence MAP problem.*

The price space in this auction is smaller than in NONLINEAR&DISCRETE, with a polynomial number of prices instead of an exponential number of prices. The number of rounds to convergence is also asymptotically better than for auction NONLINEAR&DISCRETE, in this case the number of rounds is *polynomial* in m, n, $W_{\max}$, and $1/\epsilon$.

For $m = |\mathcal{J}|$ attributes and $l = \max_j |\Theta_j|$ attribute-levels, and $W_{\max} = \max_i[\max_j \max_k p_\infty(j,k) - c_i(j,k)]$, auction ADDITIVE&DISCRETE converges in rounds polynomial in m, n, $W_{\max}$, and $1/\epsilon$.

Theorem 10 (complexity). *Auction* ADDITIVE&DISCRETE *converges in* $O\left(\frac{lmW_{\max}}{\epsilon}\right)$ *rounds, with minimal bid increment* ϵ*, for the preferential-independence MAP problem.*

Proof. The proof is similar to that for NONLINEAR&DISCRETE. The key observation is that any unsuccessful agent still bidding in the auction faces a smaller *effective* price on its current bid in the next round, either through a lower $p^{t+1}(j,k)$ or a higher penalty Δ^{t+1}. □

6 Computational Results

A fundamental aspect of using an iterative scheme is that the informational complexity associated with eliciting the complete cost and value functions from the suppliers and buyers is mitigated. A standard modified-VCG (direct revelation mechanism) would require a complete specification of agent cost and value functions. An additional consideration is that in long-term relationships, suppliers are circumspect about completely revealing their true cost functions or participating in designs where this can be easily inferred. The iterative mechanism presented in this paper addresses both these concerns: (i) The iterative mechanism elicits cost information from the suppliers on a pure "need to know" basis, and (ii) a completely rational buyer who uses the bid information to infer the cost functions of the suppliers still cannot precisely infer the cost function. In this section we provide some computational results (based on a simple simulation) to illustrate these points.

For the purpose of measuring the residual uncertainty about agent preferences, let us assume preferential-independence, with cost function $c_i(\theta) = \sum_j w_{i,j} c_{i,j}(\theta_j)$ to seller i and value function $v(\theta) = \sum_j w_j^B v_j(\theta_j)$ for the buyer. Furthermore, assume that the marginal costs, $c_{i,j}(\theta_j)$, and marginal values, $v_j(\theta_j)$, are *known* to the auctioneer and the only unknowns are the weights, $w_{i,j}$ and w_j^B.

The simulation is set up as follows. Consider the auction NONLINEAR&DISCRETE, with a commodity with 3 attributes ($j = 1,2,3$) with levels $k_1 = 3, k_2 = 2$, and $k_3 = 4$ for the attributes. The auction is assumed to have one buyer and $N = 5$ suppliers. Each agent is randomly assigned a (monotonically increasing) cost for each attribute level and a random weight $w_{i,j}$ for each attribute. Similarly, a cost function for the buyer is generated by randomly assigning a value for each attribute level and a weight for each attribute. All attribute bundles are

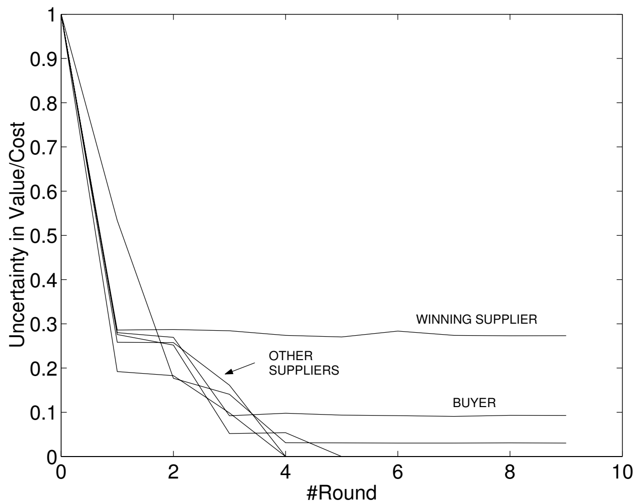

Figure 4: Incremental Information Revelation: Fractional residual uncertainty vs. Auction round.

initially priced at some high price $p_{j,k} = p_{\text{hi}}$. In each round the myopic best response (MBR)is computed by the agent and reported to the auctioneer. Similarly, the buyer indicates the best bid (among those received from suppliers) at the current price level. The prices are updated according to the rules specified in the auction.

Notice that for each buyer, MBR suggests that the chosen attribute bundle dominates all other bundles at the current price. This can be characterized by a set of linear constraints (in $w_{i,j}$) that indicate the utility for the MBR bundle is greater than all other bundles in each round. In each round, myopic best-response from each supplier reveals additional information about its cost structure, in terms of a set of linear constraints in weight space, and the volume of the feasible polytope decreases. We use the *volume* of the polytope described by the (linear) constraints on weights $w_{i,j}$ and w_j^B as a measure to quantify the residual preference-uncertainty. The preference-uncertainty volume is computed with a simple *Monte Carlo* approach.

Figure 4 shows how this volume decreases (averaged over 5 instances) for the suppliers and buyers in the auction instance specified above, with simple uniformly generated costs, values, and weights. Auction NonLinear&Discrete converges within around 10 rounds, with a residual uncertainty of around 30% for the winning supplier, 10% for the buyer, and between 0 and 5% for the other suppliers. Notice that the residual uncertainties (1-5%) are important since this still leaves the question of the exact cost structure of the agents unspecified in future interactions. Also notice that the information reveled by each buyer is restricted to indicating (in each round) a most preferred bundle, which is linear in the number of rounds.

7 Conclusions

Multiattribute auctions are central to procurement activity where the buyers and suppliers are engaged in long-terms relations. In such settings it is important that negotiation protocols provide allocative efficiency to sustain the relationship rather than pure utility maximization for the buyer. In addition, due to the asymmetry of the relations (big buyers and small suppliers), and due to the cost of preference elicitation, it is important that these protocols solve problems with minimal information revelation. The iterative techniques described in paper provide mechanisms for multiattribute negotiation that are allocatively efficient while preserving (to the extent possible) the cost and value information of the participants.

The analysis presented in this paper has interesting connections with other results in the literature. For example the myopic best response equilibrium strategy for suppliers against an ex ante consistent buyer strategy is similar to the result in Che [5]. The need for a consistent strategy arises in our context simply because of the iterative nature of the design which provides the buyer an opportunity to change the scoring rule in each round. In a direct revelation mechanism (such as in [5]) the buyer is forced to provide a single consistent scoring rule. However, the equilibrium strategy derived here is more general since it allows for more than *attributes* and requires less information revelation from the supplier. Although the analysis in this paper has focused on allocative efficiency, it might be equally interesting to construct an optimal scoring function (that is consistent across rounds). Of course this would not exploit the information that is revealed in each round as is done in [1]. However, once the buyer begins to tune the scoring function in each round, myopic best response is no longer an equilibrium strategy for the supplier and a naive supplier has to be assumed to keep this behavior. The question that remains open for future research is whether one can design an optimal scoring strategy that is consistent across rounds with respect to preferences across bids but at the same time changes the scoring function from round to round (within this constraint) based on bid information. This would be the middle ground between a naive supplier and a strategic supplier who finds it in her best interest to the a myopic best responder while still allowing for buyer utility maximization.

Acknowledgments

We wish to thank William Walsh, and anonymous reviewers, for helpful comments on an earlier draft of this paper.

References

[1] Damian R Beil and Lawrence M Wein. An inverse-optimization-based auction mechanism to support a multi-attribute RFQ process. Technical report, Operations Research Center, MIT, 2001.

[2] Sushil Bikchandani, Sven de Vries, James Schummer, and Rakesh V Vohra. Linear programming and Vickrey auctions. Technical report, Anderson Graduate School of Management, U.C.L.A., 2001.

[3] Sushil Bikchandani and Joseph M Ostroy. The package assignment model. Technical report, Anderson School of Management and Department of Economics, UCLA, 1999.

[4] F Branco. The design of multidimensional auctions. *RAND Journal of Economics*, 28:63–81, 1997.

[5] Y K Che. Design competition through multidimensional auctions. *RAND Journal of Economics*, 24:668–680, 1993.

[6] E H Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[7] Sven de Vries and Rakesh Vohra. Combinatorial auctions: A brief survey. Technical report, MEDS Department, Kellogg Graduate School of Management, Northwestern University, 2000.

[8] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[9] Matthew Jackson. *Mechanism Theory*. 2000. Forthcoming in Encyclopedia of Life Support Stystems.

[10] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: Prospects, methods and challenges. *Int. Journal of Group Decision and Negotiation*, 10, 2001.

[11] R L Keeney and H Raiffa. *Decision Making with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, 1993.

[12] Sarit Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94:79–97, 1997.

[13] Robert B Myerson. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.

[14] Robert B Myerson. Mechanism design by an informed principal. *Econometrica*, 51:1767–1797, 1983.

[15] David C Parkes. *i*Bundle: An efficient ascending price bundle auction. In *Proc. 1st ACM Conf. on Electronic Commerce (EC-99)*, pages 148–157, 1999.

[16] David C. Parkes, Jayant R. Kalagnanam, and Marta Eso. Achieving budget-balance with vickrey-based payment schemes in combinatorial exchanges. Technical report, IBM Research Report RC 22218, 2001.

[17] David C Parkes, Jayant R Kalagnanam, and Marta Eso. Vickrey-based surplus distribution in combinatorial exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001. To appear.

[18] David C Parkes and Lyle H Ungar. Iterative combinatorial auctions: Theory and practice. In *Proc. 17th National Conference on Artificial Intelligence (AAAI-00)*, pages 74–81, 2000.

[19] David C Parkes and Lyle 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.

[20] David C Parkes and Lyle H Ungar. An ascending-price generalized Vickrey auction. Technical report, Division of Engineering and Applied Sciences, Harvard University, 2002.

[21] Michael H Rothkopf, Aleksandar Pekeč, and Ronald M Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[22] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[23] N Vulkan and N R Jennings. Efficient mechanisms for the supply of services in multi-agent environments. *Decision Support Systems*, 28:5–19, 2000.