

# Generalized Value Decomposition and Structured Multiattribute Auctions

Yagil Engel and Michael P. Wellman

University of Michigan, Computer Science & Engineering  
2260 Hayward St, Ann Arbor, MI 48109-2121, USA  
{yagil,wellman}@umich.edu

## ABSTRACT

Multiattribute auction mechanisms generally either remain agnostic about traders' preferences, or presume highly restrictive forms, such as full additivity. Real preferences often exhibit dependencies among attributes, yet may possess some structure that can be usefully exploited to streamline communication and simplify operation of a multiattribute auction. We develop such a structure using the theory of measurable value functions, a cardinal utility representation based on an underlying order over preference differences. A set of local conditional independence relations over such differences supports a generalized additive preference representation, which decomposes utility across overlapping clusters of related attributes. We introduce an iterative auction mechanism that maintains prices on local clusters of attributes rather than the full space of joint configurations. When traders' preferences are consistent with the auction's generalized additive structure, the mechanism produces approximately optimal allocations, at approximate VCG prices.

**Categories and Subject Descriptors:** J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

**General Terms:** Algorithms, Economics

**Keywords:** Auctions, Multiattribute Auctions, Preference Handling

## 1. INTRODUCTION

Multiattribute trading mechanisms extend traditional, price-only mechanisms by facilitating the negotiation over a set of predefined attributes representing various non-price aspects of the deal. Rather than negotiating over a fully defined good or service, a multiattribute mechanism delays commitment to specific configurations until the most promising candidates are identified. For example, a procurement department of a company may use a multiattribute auction to select a supplier of hard drives. Supplier offers may be evaluated not only over the price they offer, but also over various qualitative attributes such as volume, RPM, access time, latency, transfer rate, and so on. In addition, suppliers may offer different contract conditions such as warranty, delivery time, and service.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'07, June 11–15, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-653-0/07/0006 ...\$5.00.

In order to account for traders' preferences, the auction mechanism must extract evaluative information over a complex domain of multidimensional configurations. Constructing and communicating a complete preference specification can be a severe burden for even a moderate number of attributes, therefore practical multiattribute auctions must either accommodate partial specifications, or support compact expression of preferences assuming some simplified form. By far the most popular multiattribute form to adopt is the simplest: an additive representation where overall value is a linear combination of values associated with each attribute. For example, several recent proposals for iterative multiattribute auctions [2, 3, 8, 19] require additive preference representations.

Such additivity reduces the complexity of preference specification exponentially (compared to the general discrete case), but precludes expression of any interdependencies among the attributes. In practice, however, interdependencies among natural attributes are quite common. For example, the buyer may exhibit complementary preferences for size and access time (since the performance effect is more salient if much data is involved), or may view a strong warranty as a good substitute for high reliability ratings. Similarly, the seller's production characteristics (such as "increasing access time is harder for larger hard drives") can easily violate additivity. In such cases an additive value function may not be able to provide even a reasonable approximation of real preferences.

On the other hand, fully general models are intractable, and it is reasonable to expect multiattribute preferences to exhibit *some* structure. Our goal, therefore, is to identify the subtler yet more widely applicable structured representations, and exploit these properties of preferences in trading mechanisms.

We propose an iterative auction mechanism based on just such a flexible preference structure. Our approach is inspired by the design of an iterative multiattribute procurement auction for additive preferences, due to Parkes and Kalagnanam (PK) [19]. PK propose two types of iterative auctions: the first (NLD) makes no assumptions about traders' preferences, and lets sellers bid on the full multidimensional attribute space. Because NLD maintains an exponential price structure, it is suitable only for small domains. The other auction (AD) assumes additive buyer valuation and seller cost functions. It collects sell bids per attribute level and for a single discount term. The price of a configuration is defined as the sum of the prices of the chosen attribute levels minus the discount.

The auction we propose also supports compact price spaces, albeit for levels of *clusters* of attributes rather than singletons. We employ a preference decomposition based on *generalized additive independence* (GAI), a model flexible enough to accommodate interdependencies to the exact degree of accuracy desired, yet providing a compact functional form to the extent that interdependence can be limited. Given its roots in multiattribute utility theory [13],

the GAI condition is defined with respect to the *expected utility* function. To apply it for modeling values for certain outcomes, therefore, requires a reinterpretation for preference under certainty. To this end, we exploit the fact that auction outcomes are associated with continuous prices, which provide a natural scale for assessing magnitude of preference.

We first lay out a representation framework for preferences that captures, in addition to simple orderings among attribute configuration values, the difference in the *willingness to pay* (wtp) for each. That is, we should be able not only to compare outcomes but also decide whether the difference in quality is worth a given difference in price. Next, we build a direct, formally justified link from preference statements over priced outcomes to a generalized additive decomposition of the wtp function. After laying out this infrastructure, we employ this representation tool for the development of a multiattribute iterative auction mechanism that allows traders to express their complex preferences in GAI format. We then study the auction's allocational, computational, and practical properties.

In Section 2 we present essential background on our representation framework, the *measurable value function* (MVF). Section 3 develops new multiattribute structures for MVF, supporting generalized additive decompositions. Next, we show the applicability of the theoretical framework to preferences in trading. The rest of the paper is devoted to the proposed auction mechanism.

## 2. MULTIATTRIBUTE PREFERENCES

As mentioned, most tools facilitating expression of multiattribute value for trading applications assume that agents' preferences can be represented in an additive form. By way of background, we start by introducing the formal prerequisites justifying the additive representation, as provided by multiattribute utility theory. We then present the generalized additive form, and develop the formal underpinnings for measurable value needed to extend this model to the case of choice under certainty.

### 2.1 Preferential Independence

Let  $\Theta$  denote the space of possible outcomes, with  $\preceq$  a preference relation (weak total order) over  $\Theta$ . Let  $A = \{a_0, \dots, a_m\}$  be a set of attributes describing  $\Theta$ . Capital letters denote subsets of variables, small letters (with or without numeric subscripts) denote specific variables, and  $\bar{X}$  denotes the complement of  $X$  with respect to  $A$ . We indicate specific variable assignments with prime signs or superscripts. To represent an instantiation of subsets  $X, Y$  at the same time we use a sequence of instantiation symbols, as in  $X'Y'$ .

**DEFINITION 1.** *A set of attributes  $Y \subset A$  is preferentially independent (PI) of its complement  $Z = A \setminus Y$  if the conditional preference order over  $Y$  given a fixed level  $Z^0$  of  $Z$  is the same regardless of the choice of  $Z^0$ .*

In other words, the preference order over the projection of  $A$  on the attributes in  $Y$  is the same for any instantiation of the attributes in  $Z$ .

**DEFINITION 2.**  *$A = \{a_1, \dots, a_m\}$  is mutually preferentially independent (MPI) if any subset of  $A$  is preferentially independent of its complement.*

The preference relation  $\preceq$  when no uncertainty is modeled is usually represented by a value function  $v$  [17]. The following fundamental result greatly simplifies the value function representation.

**THEOREM 1** ([9]). *A preference order over set of attributes  $A$  has an additive value function representation*

$$v(a_1, \dots, a_m) = \sum_{i=1}^m v_i(a_i)$$

*iff  $A$  is mutually preferentially independent.*

Essentially, the additive forms used in trading mechanisms assume mutual preferential independence over the full set of attributes, including the money attribute. Intuitively that means that willingness to pay for value of an attribute or attributes cannot be affected by the value of other attributes.

A cardinal value function representing an ordering over certain outcomes need not in general coincide with the cardinal utility function that represents preference over lotteries or *expected utility* (EU). Nevertheless, EU functions may possess structural properties analogous to that for value functions, such as additive decomposition. Since the present work does not involve decisions under uncertainty, we do not provide a full exposition of the EU concept. However we do make frequent reference to the following additive independence relations.

**DEFINITION 3.** *Let  $X, Y, Z$  be a partition of the set of attributes  $A$ .  $X$  and  $Y$  are conditionally additive independent given  $Z$ , denoted as  $CAI(X, Y | Z)$ , if preferences over lotteries on  $A$  depend only on their marginal conditional probability distributions over  $X$  and  $Y$ .*

**DEFINITION 4.** *Let  $I_1, \dots, I_g \subseteq A$  such that  $\bigcup_{i=1}^g I_i = A$ .  $I_1, \dots, I_g$  are called generalized additive independent (GAI) if preferences over lotteries on  $A$  depend only on their marginal distributions over  $I_1, \dots, I_g$ .*

An (expected) utility function  $u(\cdot)$  can be decomposed additively according to its (possibly overlapping) GAI sub-configurations.

**THEOREM 2** ([13]). *Let  $I_1, \dots, I_g$  be GAI. Then there exist functions  $f_1, \dots, f_g$  such that*

$$u(a_1, \dots, a_m) = \sum_{r=1}^g f_r(I_r). \quad (1)$$

What is now known as the GAI condition was originally introduced by Fishburn [13] for EU, and was named GAI and brought to the attention of AI researchers by Bacchus and Grove [1]. Graphical models and elicitation procedures for GAI decomposable utility were developed for EU [4, 14, 6], for a cardinal representation of the ordinal value function [15], and for an ordinal preference relations corresponding to a TCP-net structure by Brafman et al. [5]. Apart from the work on GAI in the context of preference handling that were discussed above, GAI have been recently used in the context of mechanism design by Hyafil and Boutilier [16], as an aid in direct revelation mechanisms.

As shown by Bacchus and Grove [1], GAI structure can be identified based on a set of CAI conditions, which are much easier to detect and verify. In general, utility functions may exhibit GAI structure not based on CAI. However, to date all proposals for reasoning and eliciting utility in GAI form take advantage of the GAI structure primarily to the extent that it represents a collection of CAI conditions. For example, GAI trees [14] employ triangulation of the CAI map, and Brazunas and Boutilier's [6] conditional set  $C_j$  of a set  $I_j$  corresponds to the CAI separating set of  $I_j$ .

Since the CAI condition is also defined based on preferences over lotteries, we cannot apply Bacchus and Grove's result without first establishing an alternative framework based on priced outcomes. We develop such a framework using the theory of measurable value functions, ultimately producing a GAI decomposition

(Eq. 1) of the wtp function. Readers interested primarily in the multiattribute auction and willing to grant the well-foundedness of the preference structure may skip down to Section 5.

## 2.2 Measurable Value Functions

Trading decisions represent a special case of decisions under certainty, where choices involve multiattribute outcomes and corresponding monetary payments. In such problems, the key decision often hinges on relative valuations of price differences compared to differences in alternative configurations of goods and services. Theoretically, price can be treated as just another attribute, however, such an approach fails to exploit the special character of the money dimension, and can significantly add to complexity due to the inherent continuity and typical wide range of possible monetary outcome values.

We build on the fundamental work of Dyer and Sarin [10, 11] on *measurable value functions* (MVF). As we show below, wtp functions in a quasi-linear setting can be interpreted as MVFs. However we first present the MVF framework in a more generic way, where the measurement is not necessarily monetary. We present the essential definitions and refer to Dyer and Sarin for more detailed background and axiomatic treatment. The key concept is that of *preference difference*. Let  $\theta^1, \theta^2, \vartheta^1, \vartheta^2 \in \Theta$  such that  $\theta^1 \preceq \theta^2$  and  $\vartheta^1 \preceq \vartheta^2$ .  $[\theta^2, \theta^1]$  denotes the preference difference between  $\theta^2$  and  $\theta^1$ , interpreted as the strength, or degree, to which  $\theta^2$  is preferred over  $\theta^1$ . Let  $\preceq^*$  denote a preference order over  $\Theta \times \Theta$ . We interpret the statement

$$[\theta^2, \theta^1] \preceq^* [\vartheta^2, \vartheta^1]$$

as “the preference of  $\vartheta^2$  over  $\vartheta^1$  is at least as strong as the preference of  $\theta^2$  over  $\theta^1$ ”. We use the symbol  $\sim^*$  to represent equality of preference differences.

DEFINITION 5.  $u : D \rightarrow \Re$  is a measurable value function (MVF) wrt  $\preceq^*$  if for any  $\theta^1, \theta^2, \vartheta^1, \vartheta^2 \in D$ ,

$$[\theta^2, \theta^1] \preceq^* [\vartheta^2, \vartheta^1] \Leftrightarrow u(\theta^2) - u(\theta^1) \leq u(\vartheta^2) - u(\vartheta^1).$$

Note that an MVF can also be used as a value function representing  $\preceq$ , since  $[\theta', \theta] \preceq^* [\theta'', \theta]$  iff  $\theta' \preceq \theta''$ .

DEFINITION 6 ([11]). Attribute set  $X \subset A$  is called difference independent of  $\bar{X}$  if for any two assignments  $X^1 \bar{X}' \preceq X^2 \bar{X}'$ ,

$$[X^1 \bar{X}', X^2 \bar{X}'] \sim^* [X^1 \bar{X}'', X^2 \bar{X}']$$

for any assignment  $\bar{X}''$ .

Or, in words, the preference differences on assignments to  $X$  given a fixed level of  $\bar{X}$  do not depend on the particular level chosen for  $\bar{X}$ .

As with additive independence for EU, this condition is stronger than preferential independence of  $X$ . Also analogously to EU, mutual preferential independence combined with other conditions leads to additive decomposition of the MVF. Moreover, Dyer and Sarin [11] have defined analogs of *utility independence* [17] for MVF, and worked out a parallel set of decomposition results.

## 3. ADVANCED MVF STRUCTURES

### 3.1 Conditional Difference Independence

Our first step is to generalize Definition 6 to a conditional version.

DEFINITION 7. Let  $X, Y, Z$  be a partition of the set of attributes  $A$ .  $X$  is conditionally difference independent of  $Y$  given  $Z$ , denoted as  $CDI(X, Y | Z)$ , if

$$\forall \text{ instantiations } \hat{Z}, X^1, X^2, Y^1, Y^2 \\ [X^1 Y^1 \hat{Z}, X^2 Y^1 \hat{Z}] \sim [X^1 Y^2 \hat{Z}, X^2 Y^2 \hat{Z}].$$

Since the conditional set is always the complement, we sometimes leave it implicit, using the abbreviated notation  $CDI(X, Y)$ .

CDI leads to a decomposition similar to that obtained from CAI [17].

LEMMA 3. Let  $u(A)$  be an MVF representing preference differences. Then  $CDI(X, Y | Z)$  iff

$$u(A) = u(X^0, Y, Z) + u(X, Y^0, Z) - u(X^0, Y^0, Z).$$

To complete the analogy with CAI, we generalize Lemma 3 as follows.

PROPOSITION 4.  $CDI(X, Y | Z)$  iff there exist functions  $\psi_1(X, Z)$  and  $\psi_2(Y, Z)$ , such that

$$u(X, Y, Z) = \psi_1(X, Z) + \psi_2(Y, Z). \quad (2)$$

An immediate result of Proposition 4 is that CDI is a symmetric relation.

The conditional independence condition is much more applicable than the unconditional one. For example, if attributes  $a \in X$  and  $b \notin X$  are complements or substitutes,  $X$  cannot be difference independent of  $\bar{X}$ . However,  $X \setminus \{a\}$  may still be CDI of  $\bar{X}$  given  $a$ .

### 3.2 GAI Structure for MVF

A single CDI condition decomposes the value function into two parts. We seek a finer-grain global decomposition of the utility function, similar to that obtained from mutual preferential independence. For this purpose we are now ready to employ the results of Bacchus and Grove [1], who establish that the CAI condition has a *perfect map* [20]; that is, there exists a graph whose nodes correspond to the set  $A$ , and its node separation reflects exactly the complete set of CAI conditions on  $A$ . Moreover, they show that the utility function decomposes over the set of maximal cliques of the perfect map. Their proofs can be easily adapted to CDI, since they only rely on the decomposition property of CAI that is also implied by CDI according to Proposition 4.

THEOREM 5. Let  $G = (A, E)$  be a perfect map for the CDI conditions on  $A$ . Then

$$u(A) = \sum_{r=1}^g f_r(I_r), \quad (3)$$

where  $I_1, \dots, I_g$  are (overlapping) subsets of  $A$ , each corresponding to a maximal clique of  $G$ .

Given Theorem 5, we can now identify an MVF GAI structure from a collection of CDI conditions. The CDI conditions, in turn, are particularly intuitive to detect when the preference differences carry a direct interpretation, as in the case with monetary differences discussed below. Moreover, the assumption or detection of CDI conditions can be performed incrementally, until the MVF is decomposed to a reasonable dimension. This is in contrast with the fully additive decomposition of MVF that requires mutual preferential independence [11].

Theorem 5 defines a decomposition structure, but to represent the actual MVF we need to specify the functions over the cliques.

The next theorem establishes that the functional constituents of MVF are the same as those for GAI decompositions as defined by Fishburn [13] for EU. We adopt the following conventional notation. Let  $(a_1^0, \dots, a_m^0)$  be a predefined vector called the *reference outcome*. For any  $I \subseteq A$ , the function  $u([I])$  stands for the projection of  $u(A)$  to  $I$  where the rest of the attributes are fixed at their reference levels.

**THEOREM 6.** *Let  $G = (A, E)$  be a perfect map for the CDI condition on  $A$ , and  $\{I_1, \dots, I_g\}$  a set of maximal cliques as defined in Theorem 5. Then the functional decomposition from that theorem can be defined as*

$$\begin{aligned} f_1 &= u([I_1]), \text{ and for } r = 2, \dots, g \\ f_r &= u([I_r]) + \sum_{k=1}^{r-1} (-1)^k \sum_{1 \leq i_1 < \dots < i_k < r} u\left(\left[\bigcap_{s=1}^k I_{i_s} \cap I_r\right]\right) \end{aligned} \quad (4)$$

The proof directly shows that if graph  $G = (A, E)$  is a perfect map of CDI,  $u(A)$  decomposes to a sum over the functions defined in (4).<sup>1</sup> Thus this proof does not rely on the decomposition result of Theorem 5, only on the existence of the perfect map.

To summarize, the results of this section generalize additive MVF theory. In particular it justifies the application of methods recently developed under the EU framework [1, 4, 14, 6] to representation of value under certainty.

## 4. WILLINGNESS-TO-PAY AS AN MVF

### 4.1 Construction

In this section we apply measurable value to represent differences of willingness to pay for outcomes. We assume that the agent has a preference order over outcome space, represented by a set of attributes  $A$ , and an attribute  $p$  representing monetary consequence. Note that in evaluating a purchase decision,  $p$  would correspond to the agent's money holdings net of the transaction (i.e., wealth after purchase), not the purchase price. An outcome in this space is represented for example by  $(\theta', p')$ , where  $\theta'$  is an instantiation of  $A$  and  $p'$  is a value of  $p$ . We further assume that preferences are *quasi-linear* in  $p$ , that is there exists a value function of the form  $v(A, p) = u(A) + L(p)$ , where  $L$  is a positive linear function.<sup>2</sup>

The quasi-linear form immediately qualifies money as a measure of preference differences, and establishes a monetary scale for  $u(A)$ .

**DEFINITION 8.** *Let  $v(A, p) = u(A) + L(p)$  represent  $\preceq$ , where  $p$  is the attribute representing money. We call  $u(A)$  a willingness-to-pay (wtp) function.*

Note that wtp may also refer to the seller's "willingness to accept" function. The wtp  $u(A)$  is a *cardinal* function, unique up to a positive linear transformation. Since

$$(\theta_1, p') \preceq (\theta_2, p'') \Leftrightarrow u(\theta_1) - u(\theta_2) \leq L(p'' - p'),$$

(where  $\theta_1, \theta_2 \in \Theta$ , the domain of  $A$ ) the wtp function can be used to choose among priced outcomes.

<sup>1</sup>This proof and most other proofs in this paper are omitted for space consideration, and are available in an online appendix.

<sup>2</sup>In many procurement applications, the deals in question are small relative to the enterprises involved, so the quasi-linearity assumption is warranted. This assumption can be relaxed to a condition called *corresponding tradeoffs* [17], which does not require the value over money to be linear. To simplify the presentation, however, we maintain the stronger assumption.

Naturally, elicitation of wtp function is most intuitive when using direct monetary values. In other words, we elicit a function in which  $L(p) = p$ , so  $v(A, p) = u(A) + p$ . We define a reference outcome  $(\theta^0, p^0)$ , and assuming continuity of  $p$ , for any assignment  $\hat{\theta}$  there exists a  $\hat{p}$  such that  $(\hat{\theta}, \hat{p}) \sim (\theta^0, p^0)$ . As  $v$  is normalized such that  $v(\theta^0, p^0) = 0$ ,  $\hat{p}$  is interpreted as the wtp for  $\hat{\theta}$ , or the *reserve price* of  $\hat{\theta}$ .

**PROPOSITION 7.** *The wtp function is an MVF over differences in the reserve prices.*

We note that the wtp function is used extensively in economics, and that all the development in Section 3 could be performed directly in terms of wtp, relying on quasi-linearity for preference measurement, and without formalization using MVFs. This formalization however aligns this work with the fundamental difference independence theory by Dyer and Sarin.

In addition to facilitating the detection of GAI structure, the CDI condition supports elicitation using local queries, similar to how CAI is used by Braziunas and Boutilier [6]. We adopt their definition of *conditional set* of  $I_r$ , noted here  $S_r$ , as the set of neighbors of attributes in  $I_r$  not including the attributes of  $I_r$ . Clearly,  $S_r$  is the separating set of  $I_r$  in the CDI map, hence  $\text{CDI}(I_r, V_r)$ , where  $V_r = A \setminus (I_r \cup S_r)$ . From the definition of CDI, for any  $V_r^1, V_r^2$  we have:

$$u(I_r^1 S_r^0 V_r^1) - u(I_r^2 S_r^0 V_r^1) = u(I_r^1 S_r^0 V_r^2) - u(I_r^2 S_r^0 V_r^2).$$

Eliciting the wtp function therefore amounts to eliciting the utility (wtp) of one full outcome (the reference outcome  $\theta^0$ ), and then obtaining the function over each maximal clique using monetary differences between its possible assignments (technique known as *pricing out* [17]), keeping the variables in the conditional set fixed. These *ceteris paribus* elicitation queries are local in the sense that the agent does not need to consider the values of the rest of the attributes. Furthermore, in eliciting MVFs we can avoid the global scaling step that is required for EU functions. Since the preference differences are extracted with respect to specific amounts of the attribute  $p$ , the utility is already scaled according to that external measure. Hence, once the conditional utility functions  $u([I_j])$  are obtained, we can calculate  $u(A)$  according to (4).

This last step may require (in the worst case) computation of a number of terms that is exponential in the number of max cliques. In practice however we do not expect the intersection of the cliques to go that deep; intersection of more than just a few max cliques would normally be empty. To take advantage of that we can use the search algorithm suggested by Braziunas and Boutilier [6], which efficiently finds all the nonempty intersections for each clique.

### 4.2 Optimization

As shown, the wtp function can be used directly for pairwise comparisons of priced outcomes. Another preference query often treated in the literature is optimization, or choice of best outcome, possibly under constraints.

Typical decisions about exchange of a good or service exhibit what we call *first-order preferential independence (FOPI)*, under which most or all single attributes have a natural ordering of quality, independent of the values of the rest.<sup>3</sup> For example, when choosing a PC we always prefer more memory, faster CPU, longer warranty, and so on. Under FOPI, the unconstrained optimization of

<sup>3</sup>This should not be mistaken with the highly demanding condition of *mutual preferential independence*, that requires all tradeoffs between attributes to be independent.

unpriced outcomes is trivial, hence we consider choice among attribute points with prices. Since any outcome can be best given enough monetary compensation, this problem is not well-defined unless the combinations are constrained somehow.

A particularly interesting optimization problem arises in the context of negotiation, where we consider the utility of both buyers and sellers. The *multiattribute matching problem* (MMP) [12] is concerned with finding an attribute point that maximizes the surplus of a trade, or the difference between the utilities of the buyer and the seller,  $u_b(A) - u_s(A)$ . GAI, as an additive decomposition, has the property that if  $u_b$  and  $u_s$  are in GAI form then  $u_b(A) - u_s(A)$  is in GAI form as well. We can therefore use combinatorial optimization procedures for GAI decomposition, based on the well studied variable elimination schemes (e.g., [15]) to find the best trading point. Similarly, this optimization can be done to maximize surplus between a trader's utility function and a pricing system that assigns a price to each level of each GAI element, and this way guide traders to their optimal bidding points. In the rest of the paper we develop a multiattribute procurement auction that builds on this idea.

## 5. GAI IN MULTIATTRIBUTE AUCTIONS

### 5.1 The Multiattribute Procurement Problem

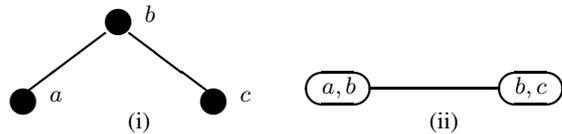
In the procurement setting a single buyer wishes to procure a single good, in some configuration  $\theta \in \Theta$  from one of the candidate sellers  $s_1, \dots, s_n$ . The buyer has some private valuation function (wtp)  $u_b : \Theta \rightarrow \mathcal{R}$ , and similarly each seller  $s_i$  has a private valuation function (willingness-to-accept). For compliance with the procurement literature we refer to seller  $s_i$ 's valuation as a cost function, denoted by  $c_i$ . The *multiattribute allocation problem* (MAP) [19] is the welfare optimization problem in procurement over a discrete domain, and it is defined as:

$$i^*, \theta^* = \arg \max_{i, \theta} (u_b(\theta) - c_i(\theta)). \quad (5)$$

To illustrate the need for a GAI price space we consider the case of traders with non-additive preferences bidding in an additive price space such as in PK's auction AD. If the buyer's preferences are not additive, choosing preferred levels per attribute (as in auction AD) admits undesired combinations and fails to guide the sellers to the efficient configurations. Non-additive sellers face an *exposure problem*, somewhat analogous to traders with complementary preferences that participate in simultaneous auctions. A value  $a^1$  for attribute  $a$  may be optimal given that the value of another attribute  $b$  is  $b^1$ , and arbitrarily suboptimal given other values of  $b$ . Therefore bidding  $a^1$  and  $b^1$  may result in a poor allocation if the seller is "outbid" on  $b^1$  but left "holding"  $a^1$ .<sup>4</sup>

Instead of assuming full additivity, the auction designer can come up with a GAI preference structure that captures the set of common interdependencies between attributes. If traders could bid on clusters of interdependent attributes, it would solve the problems discussed above. For example, if  $a$  and  $b$  are interdependent (meaning  $\text{CDI}(a, b)$  does not hold), we should be able to bid on the cluster  $ab$ . If  $b$  in turn depends on  $c$ , we need another cluster  $bc$ . This is still better than a general pricing structure that solicits bids for the cluster  $abc$ . We stress that each trader may have a different set of interdependencies, and therefore to be completely general the

<sup>4</sup>If only the sellers are non-additive, the auction design could potentially alleviate this problem by collecting a new set of bids each round and "forgetting" bids from previous rounds, and also guiding non-additive sellers to bid on only one level per attribute in order to avoid undesired combinations.



**Figure 1: (i) CDI map for  $\{a, b, c\}$ , reflecting the single condition  $\text{CDI}(a, c)$ . (ii) The corresponding GAI network.**

GAI structure needs to account for all.<sup>5</sup> However, in practice many domains have natural dependencies that are mutual to traders.

### 5.2 GAI Trees

Assume that preferences of all traders are reflected in a GAI structure  $I_1, \dots, I_g$ . We call each  $I_r$  a GAI element, and any assignment to  $I_r$  a *sub-configuration*. We use  $\theta_r$  to denote the sub-configuration formed by projecting configuration  $\theta$  to element  $I_r$ .

**DEFINITION 9.** Let  $\alpha$  be an assignment to  $I_r$  and  $\beta$  an assignment to  $I_{r'}$ . The sub-configurations  $\alpha$  and  $\beta$  are consistent if for any attribute  $a_j \in I_r \cap I_{r'}$ ,  $\alpha$  and  $\beta$  agree on the value of  $a_j$ . A collection  $\nu$  of sub-configurations is consistent if all pairs  $\alpha, \beta \in \nu$  are consistent. The collection is called a cover if it contains exactly one sub-configuration  $\alpha_r$  corresponding to each element  $I_r$ ,  $r \in \{1, \dots, g\}$ .

Note that a consistent cover  $\{\alpha_1, \dots, \alpha_g\}$  represents a full configuration, which we denote by  $(\alpha_1, \dots, \alpha_g)$ .

A *GAI network* is a graph  $G$  whose nodes correspond to the GAI elements  $I_1, \dots, I_g$ , with an edge between  $I_r, I_{r'}$  iff  $I_r \cap I_{r'} \neq \emptyset$ . Equivalently, a GAI network is the clique graph of a CDI-map. In order to justify the compact pricing structure we require that for any set of optimal configurations (wrt a given utility function), with a corresponding collection of sub-configurations  $\gamma$ , all consistent covers in  $\gamma$  must be optimal configurations as well. To ensure this (see Lemmas 8 and 10), we assume a GAI decomposition in the form of a tree or a forest (the *GAI tree*). A tree structure can be achieved for any set of CDI conditions by triangulation of the CDI-map prior to construction of the clique graph (GAI networks and GAI trees are defined by Gonzales and Perny [14], who also provide a triangulation algorithm). Under GAI, the buyer's value function  $u_b$  and sellers' cost functions  $c_i$  can be decomposed as in (1). We use  $f_{b,r}$  and  $f_{i,r}$  to denote the local functions of buyer and sellers (respectively), according to (4).

For example, consider the procurement of a good with three attributes,  $a, b, c$ . Each attribute's domain has two values (e.g.,  $\{a^1, a^2\}$  is the domain of  $A$ ). Let the GAI structure be  $I_1 = \{a, b\}, I_2 = \{b, c\}$ . Figure 1 shows the simple CDI map and the corresponding GAI network, which is a GAI tree. Here, sub-configurations are assignments of the form  $a^1 b^1, a^1 b^2, b^1 c^1$ , and so on. The set of sub-configurations  $\{a^1 b^1, b^1 c^1\}$  is a consistent cover, corresponding to the configuration  $a^1 b^1 c^1$ . In contrast, the set  $\{a^1 b^1, b^2 c^1\}$  is inconsistent.

### 5.3 The GAI Auction

We define an iterative multiattribute auction that maintains a GAI pricing structure: that is, a price  $p^t(\cdot)$  corresponding to each sub-configuration of each GAI-tree element. The price of a configuration  $\theta$  at time  $t$  is defined as

$$p^t(\theta) = \sum_{r=1}^g p^t(\theta_r) - \Delta.$$

<sup>5</sup>We relax this requirement in Section 6.

Bidders submit *sub-bids* on sub-configurations and on an additional global discount term  $\Delta$ .<sup>6</sup> Sub-bids are always submitted for current prices, and need to be resubmitted at each round, therefore they do not need to explicitly carry the price. The set of *full bids* of a seller contains all consistent covers that can be generated from that seller's current set of sub-bids. The existence of a full bid over a configuration  $\theta$  represents the seller's willingness to accept the price  $p^t(\theta)$  for supplying  $\theta$ .

At the start of the auction, the buyer reports (to the auction, not to sellers) her complete valuation in GAI form. The initial prices of sub-configurations are set at some level above the buyer's valuations, that is,  $p^1(\theta_r) > f_{b,r}(\theta_r)$  for all  $\theta_r$ . The discount  $\Delta$  is initialized to zero. The auction has the dynamics of a *descending clock auction*: at each round  $t$ , bids are collected for current prices and then prices are reduced according to price rules. A seller is considered *active* in a round if she submits at least one full bid. In round  $t > 1$ , only sellers who were active in round  $t - 1$  are allowed to participate, and the auction terminates when no more than a single seller is active. We denote the set of sub-bids submitted by  $s_i$  by  $B_i^t$ , and the corresponding set of full bids is

$$B_i^t = \{\theta = (\theta_1, \dots, \theta_g) \in \Theta \mid \forall r. \theta_r \in B_i^t\}.$$

In our example, a seller could submit sub-bids on a set of sub-configurations such as  $a^1 b^1$  and  $b^1 c^1$ , and that combines to a full bid on  $a^1 b^1 c^1$ .

The auction proceeds in two phases. In the first phase (A), at each round  $t$  the auction computes a set of preferred sub-configurations  $\mathcal{M}^t$ . Section 5.4 shows how to define  $\mathcal{M}^t$  to ensure convergence, and Section 5.5 shows how to efficiently compute it.

In phase A, the auction adjusts prices after each round, reducing the price of every sub-configuration that has received a bid but is not in the preferred set. Let  $\epsilon$  be the prespecified price increment parameter. Specifically, the phase A price change rule is applied to all  $\theta_r \in \bigcup_{i=1}^n B_i^t \setminus \mathcal{M}^t$ :

$$p^{t+1}(\theta_r) \leftarrow \max(p^t(\theta_r) - \frac{\epsilon}{g}, f_{b,r}(\theta_r)). \quad [A]$$

The RHS maximum ensures that prices do not get reduced below the buyer's valuation in phase A.

Let  $M^t$  denote the set of configurations that are consistent covers in  $\mathcal{M}^t$ :

$$M^t = \{\theta = (\theta_1, \dots, \theta_g) \in \Theta \mid \forall r. \theta_r \in \mathcal{M}^t\}$$

The auction switches to phase B when all active sellers have at least one full bid in the buyer's preferred set:

$$\forall i. B_i^t = \emptyset \vee B_i^t \cap M^t \neq \emptyset. \quad [\text{SWITCH}]$$

Let  $T$  be the round at which [SWITCH] becomes true. At this point, the auction selects the buyer-optimal full bid  $\eta_i$  for each seller  $s_i$ .

$$\eta_i = \arg \max_{\theta \in B_i^T} (u_b(\theta) - p^T(\theta)). \quad (6)$$

In phase B,  $s_i$  may bid only on  $\eta_i$ . The prices of sub-configurations are fixed at  $p^T(\cdot)$  during this phase. The only adjustment in phase B is to  $\Delta$ , which is increased in every round by  $\epsilon$ . The auction terminates when at most one seller (if exactly one, designate it  $s_i$ ) is active. There are four distinct cases:

1. All sellers drop out in phase A (i.e., before rule [SWITCH] holds). The auction returns with no allocation.

<sup>6</sup>The discount term could be replaced with a uniform price reduction across all sub-configurations.

2. All active sellers drop out in the same round in phase B. The auction selects the best seller ( $s_i$ ) from the preceding round, and applies the applicable case below.
3. The auction terminates in phase B with a final price above buyer's valuation,  $p^T(\eta_i) - \Delta > u_b(\eta_i)$ . The auction offers the winner  $s_i$  an opportunity to supply  $\eta_i$  at price  $u_b(\eta_i)$ .
4. The auction terminates in phase B with a final price  $p^T(\eta_i) - \Delta \leq u_b(\eta_i)$ . This is the ideal situation, where the auction allocates the chosen configuration and seller at this resulting price.

The overall auction is described by high-level pseudocode in Algorithm 1. As explained in Section 5.4, the role of phase A is to guide the traders to their efficient configurations. Phase B is a one-dimensional competition over the surplus that remaining seller candidates can provide to the buyer. In Section 5.5 we discuss the computational tasks associated with the auction, and Section 5.6 provides a detailed example.

---

#### Algorithm 1 GAI-based multiattribute auction

---

```

collect a reported valuation,  $\hat{v}$  from the buyer
set high initial prices,  $p^1(\theta_r)$  on each level  $\theta_r$ , and set  $\Delta = 0$ 
while not [SWITCH] do
  collect sub-bids from sellers
  compute  $\mathcal{M}^t$ 
  apply price change by [A]
end while
compute  $\eta_i$ 
while more than one active seller do
  increase  $\Delta$  by  $\epsilon$ 
  collect bids on  $(\eta_i, \Delta)$  from sellers
end while
implement allocation and payment to winning seller

```

---

## 5.4 Economic Analysis

When the optimal solution to MAP (5) provides negative welfare and sellers do not bid below their cost, the auction terminates in phase A, no trade occurs and the auction is trivially efficient. We therefore assume throughout the analysis that the optimal (seller, configuration) pair provides non-negative welfare.

The buyer profit from a configuration  $\theta$  is defined as<sup>7</sup>

$$\pi_b(\theta) = u_b(\theta) - p(\theta)$$

and similarly  $\pi_i(\theta) = p(\theta) - c_i(\theta)$  is the profit of  $s_i$ . In addition, for  $\mu \subseteq \{1, \dots, g\}$  we denote the corresponding set of sub-configurations by  $\theta_\mu$ , and define the profit from a configuration  $\theta$  over the subset  $\mu$  as

$$\pi_b(\theta_\mu) = \sum_{r \in \mu} (f_{b,r}(\theta_r) - p(\theta_r)).$$

$\pi_i(\theta_\mu)$  is defined similarly for  $s_i$ . Crucially, for any  $\mu$  and its complement  $\bar{\mu}$  and for any trader  $\tau$ ,

$$\pi_\tau(\theta) = \pi_\tau(\theta_\mu) + \pi_\tau(\theta_{\bar{\mu}}).$$

The function  $\sigma_i : \Theta \rightarrow \mathcal{R}$  represents the welfare, or surplus function  $u_b(\cdot) - c_i(\cdot)$ . For any price system  $p$ ,

$$\sigma_i(\theta) = \pi_b(\theta) + \pi_i(\theta).$$

<sup>7</sup>We drop the  $t$  superscript in generic statements involving price and profit functions, understanding that all usage is with respect to the (currently) applicable prices.

Since we do not assume anything about the buyer's strategy, the analysis refers to profit and surplus with respect to the face value of the buyer's report. The functions  $\pi_i$  and  $\sigma_i$  refer to the true cost functions of  $s_i$ .

**DEFINITION 10.** *A seller is called Straightforward Bidder (SB) if at each round  $t$  she bids on  $\mathcal{B}_i^t$  as follows: if  $\max_{\theta \in \Theta} \pi_i^t(\theta) < 0$ , then  $\mathcal{B}_i^t = \emptyset$ . Otherwise let*

$$\begin{aligned}\Omega_i^t &\subseteq \arg \max_{\theta \in \Theta} \pi_i^t(\theta) \\ \mathcal{B}_i^t &= \{\theta_r \mid \theta \in \Omega_i^t, r \in \{1, \dots, g\}\}.\end{aligned}$$

Intuitively, an SB seller follows a *myopic best response* strategy (MBR), meaning they bid myopically rather than strategically by optimizing their profit with respect to current prices. To calculate  $\mathcal{B}_i^t$  sellers need to optimize their current profit function, as discussed in Section 4.2.

The following lemma bridges the apparent gap between the compact pricing and bid structure and the global optimization performed by the traders.

**LEMMA 8.** *Let  $\Psi$  be a set of configurations, all maximizing profit for a trader  $\tau$  (seller or buyer) at the relevant prices. Let  $\Phi = \{\theta_r \mid \theta \in \Psi, r \in \{1, \dots, g\}\}$ . Then any consistent cover in  $\Phi$  is also a profit-maximizing configuration for  $\tau$ .*

*Proof sketch (full proof in the online appendix):* A source of an element  $\theta_r$  is a configuration  $\hat{\theta} \in \Psi$  from which it originated (meaning,  $\hat{\theta}_r = \theta_r$ ). Starting from the supposedly suboptimal cover  $\theta^1$ , we build a series of covers  $\theta^1, \dots, \theta^L$ . At each  $\theta^j$  we flip the value of a set of sub-configurations  $\mu_j$  corresponding to a subtree, with the sub-configurations of the configuration  $\hat{\theta}^j \in \Psi$  which is the source of the parent  $\gamma_j$  of  $\mu_j$ . That ensures that all elements in  $\mu_j \cup \{\gamma_j\}$  have a mutual source  $\hat{\theta}^j$ . We show that all  $\theta^j$  are consistent and that they must all be suboptimal as well, and since all elements of  $\theta^L$  have a mutual source, meaning  $\theta^L = \hat{\theta}^L \in \Psi$ , it contradicts optimality of  $\Psi$ .

**COROLLARY 9.** *For SB seller  $s_i$ ,*

$$\forall t, \forall \theta' \in \mathcal{B}_i^t, \pi_i^t(\theta') = \max_{\theta \in \Theta} \pi_i^t(\theta).$$

Next we consider combinations of configurations that are only within some  $\delta$  of optimality.

**LEMMA 10.** *Let  $\Psi$  be a set of configurations, all are within  $\delta$  of maximizing profit for a trader  $\tau$  at the prices, and  $\Phi$  defined as in Lemma 8. Then any consistent cover in  $\Phi$  is within  $\delta g$  of maximizing utility for  $\tau$ .*

This bound is tight, that is for any GAI tree and a non-trivial domain we can construct a set  $\Psi$  as above in which there exists a consistent cover whose utility is exactly  $\delta g$  below the maximal.

Next we formally define  $\mathcal{M}^t$ . For connected GAI trees,  $\mathcal{M}^t$  is the set of sub-configurations that are part of a configuration within  $\epsilon$  of optimal. When the GAI tree is in fact a forest, we apportion the error proportionally across the disconnected trees. Let  $G$  be comprised of trees  $G_1, \dots, G_h$ . We use  $\theta_j$  to denote the projection of a configuration  $\theta$  on the tree  $G_j$ , and  $g_j$  denotes the number of GAI elements in  $G_j$ .

$$\mathcal{M}_j^t = \{\theta_r \mid \pi_b^t(\theta_j) \geq \max_{\theta'_j \in \Theta_j} \pi_b^t(\theta'_j) - g_j \frac{\epsilon}{g}, r \in G_j\}$$

Then define  $\mathcal{M}^t = \bigcup_{j=1}^h \mathcal{M}_j^t$ .

Let  $e_j = g_j - 1$  denote the number of edges in  $G_j$ . We define the connectivity parameter,  $e = \max_{j=1, \dots, h} e_j$ . As shown below, this connectivity parameter is an important factor in the performance of the auction.

**COROLLARY 11.**

$$\forall \theta' \in M^t, \pi_b^t(\theta') \geq \max_{\theta \in \Theta} \pi_b^t(\theta) - (e + 1)\epsilon$$

In the fully additive case this loss of efficiency reduces to  $\epsilon$ . On the other extreme, if the GAI network is connected then  $e + 1 = g$ . We also note that without assuming any preference structure, meaning that the CDI map is fully connected,  $g = 1$  and the efficiency loss is again  $\epsilon$ .

Lemmas 12 through 15 show that through the price system, the choice of buyer preferred configurations, and price change rules, Phase A leads the buyer and each of the sellers to their mutually efficient configuration.

**LEMMA 12.**  *$\max_{\theta \in \Theta} \pi_b^t(\theta)$  does not change in any round  $t$  of phase A.*

*PROOF.* We prove the lemma per each tree  $G_j$ . The optimal values for disconnected components are independent of each other hence if the maximal profit for each component does not change the combined maximal profit does not change as well. If the price of  $\theta'_j$  was reduced during phase A, that is  $p^{t+1}(\theta'_j) = p^t(\theta'_j) - \delta$ , it must be the case that some  $w \leq g_j$  sub-configurations of  $\theta'_j$  are not in  $\mathcal{M}_j^t$ , and  $\delta = \frac{w\epsilon}{g}$ . The definition of  $\mathcal{M}_j^t$  ensures

$$\pi_b^t(\theta'_j) < \max_{\theta \in \Theta} \pi_b^t(\theta_j) - g_j \frac{\epsilon}{g}.$$

Therefore,

$$\pi_b^{t+1}(\theta') = \pi_b^t(\theta') + \delta = \pi_b^t(\theta') + \frac{w\epsilon}{g} \leq \max_{\theta \in \Theta} \pi_b^t(\theta_j).$$

This is true for any configuration whose profit improves, therefore the maximal buyer profit does not change during phase A.  $\square$

**LEMMA 13.** *The price of at least one sub-configuration must be reduced at every round in phase A.*

*PROOF.* In each round  $t < T$  of phase A there exists an active seller  $i$  for whom  $\mathcal{B}_i^t \cap M^t = \emptyset$ . However to be active in round  $t$ ,  $\mathcal{B}_i^t \neq \emptyset$ . Let  $\hat{\theta} \in \mathcal{B}_i^t$ . If  $\forall r, \hat{\theta}_r \in M^t$ , then  $\hat{\theta} \in M^t$  by definition of  $M^t$ . Therefore there must be  $\hat{\theta}_r \notin M^t$ . We need to prove that for at least one of these sub-configurations,  $\pi_b^t(\hat{\theta}_r) < 0$  to ensure activation of rule [A].

Assume for contradiction that for any  $\hat{\theta}_r \in \hat{\mathcal{M}}^t$ ,  $\pi_b^t(\hat{\theta}_r) \geq 0$ . For simplicity we assume that for any  $\theta_r$ ,  $\pi_b^1(\theta_r)$  is some product of  $\frac{\epsilon}{g}$  (that can be easily done), and that ensures that  $\pi_b^t(\hat{\theta}_r) = 0$  because once profit hits 0 it cannot increase by rule [A].

If  $\hat{\theta}_r \in \hat{\mathcal{M}}^t, \forall r = 1, \dots, g$  then  $\pi_b^t(\hat{\theta}) = 0$ . This contradicts Lemma 12 since we set high initial prices. Therefore some of the sub-configurations of  $\hat{\theta}$  are in  $\mathcal{M}^t$ , and WLOG we assume it is  $\hat{\theta}_1, \dots, \hat{\theta}_k$ . To be in  $\mathcal{M}^t$  these  $k$  sub-configurations must have been in some preferred full configuration, meaning there exists  $\theta' \in M^t$  such that

$$\theta' = (\hat{\theta}_1, \dots, \hat{\theta}_k, \theta'_{k+1}, \dots, \theta'_g)$$

Since  $\hat{\theta} \notin M^t$  It must be that case that  $\pi_b^t(\hat{\theta}) < \pi_b^t(\theta')$ . Therefore

$$\pi_b^t(\theta'_{k+1}, \dots, \theta'_g) > \pi_b^t(\hat{\theta}_{k+1}, \dots, \hat{\theta}_g) = 0$$

Hence for at least one  $r \in \{k + 1, \dots, g\}$ ,  $\pi_b^t(\theta'_r) > 0$  contradicting rule [A].  $\square$

LEMMA 14. *When the solution to MAP provides positive surplus, and at least the best seller is SB, the auction must reach phase B.*

PROOF. By Lemma 13 prices must go down in every round of phase A. Rule [A] sets a lower bound on all prices therefore the auction either terminates in phase A or must reach condition [SWITCH].

We set the initial prices are high such that  $\max_{\theta \in \Theta} \pi_b^1(\theta) < 0$ , and by Lemma 12  $\max_{\theta \in \Theta} \pi_b^t(\theta) < 0$  during phase A. We assume that the efficient allocation  $(\theta^*, i^*)$  provides positive welfare, that is  $\sigma_{i^*}(\theta^*) = \pi_b^t(\theta^*) + \pi_{i^*}^t(\theta^*) > 0$ .  $s_{i^*}$  is SB therefore she will leave the auction only when  $\pi_{i^*}^t(\theta^*) < 0$ . This can happen only when  $\pi_b^t(\theta^*) > 0$ , therefore  $s_{i^*}$  does not drop in phase A hence the auction cannot terminate before reaching condition [SWITCH].  $\square$

LEMMA 15. *For SB seller  $s_i$ ,  $\eta_i$  is  $(e + 1)\epsilon$ -efficient.*

PROOF.  $\eta_i$  is chosen to maximize the buyer's surplus out of  $B_i^t$  at the end of phase A. Since  $B_i^t \cap M^t \neq \emptyset$ , clearly  $\eta_i \in M^t$ . From Corollary 11 and Corollary 9, for any  $\theta$ ,

$$\begin{aligned} \pi_b^T(\eta_i) &\geq \pi_b^T(\tilde{\theta}) - (e + 1)\epsilon \\ \pi_i^T(\eta_i) &\geq \pi_i^T(\tilde{\theta}) \\ \Rightarrow \sigma_i(\eta_i) &\geq \sigma_i(\tilde{\theta}) - (e + 1)\epsilon \end{aligned}$$

$\square$

This establishes the approximate bilateral efficiency of the results of Phase A (at this point under the assumption of SB). Based on Phase B's simple role as a single-dimensional bidding competition over the discount, we next assert that the overall result is efficient under SB, which in turn proves to be an approximately ex-post equilibrium strategy in the two phases.

LEMMA 16. *If sellers  $s_i$  and  $s_j$  are SB, and  $s_i$  is active at least as long as  $s_j$  is active in phase B, then*

$$\sigma_i(\eta_i) \geq \max_{\theta \in \Theta} \sigma_j(\theta) - (e + 2)\epsilon$$

THEOREM 17. *Given a truthful buyer and SB sellers, the auction is  $(e + 2)\epsilon$ -efficient: the surplus of the final allocation is within  $(e + 2)\epsilon$  of the maximal surplus.*

Following PK, we rely on an equivalence to the one-sided VCG auction to establish incentive properties for the sellers. In the one-sided multiattribute VCG auction, buyer and sellers report valuation and cost functions  $\hat{u}_b, \hat{c}_i$ , and the buyer pays the sell-side VCG payment to the winning seller.

DEFINITION 11. *Let  $(\theta^*, i^*)$  be the optimal solution to MAP. Let  $(\tilde{\theta}, \tilde{i})$  be the best solution to MAP when  $i^*$  does not participate. The sell-side VCG payment is*

$$VCG(\hat{u}_b, \hat{c}_i) = \hat{u}_b(\theta^*) - (\hat{u}_b(\tilde{\theta}) - \hat{c}_i(\tilde{\theta})).$$

It is well-known that truthful bidding is a dominant strategy for sellers in the one-sided VCG auction. It is also shown by PK that the maximal regret for buyers from bidding truthfully in this mechanism is  $u_b(\theta^*) - c_{i^*}(\theta^*) - (u_b(\tilde{\theta}) - \hat{c}_i(\tilde{\theta}))$ , that is, the *marginal product* of the efficient seller.

Usually in iterative auctions the VCG outcome is only nearly achieved, but the deviation is bounded by the minimal price change. We show a similar result, and therefore define  $\delta$ -VCG payments.

DEFINITION 12. *Sell-side  $\delta$ -VCG payment for MAP is a payment  $p$  such that*

$$VCG(\hat{u}_b, \hat{c}_i) - \delta \leq p \leq VCG(\hat{u}_b, \hat{c}_i) + \delta.$$

When payment is guaranteed to be  $\delta$ -VCG sellers can only affect their payment within that range, therefore their gain by falsely reporting their cost is bounded by  $2\delta$ .

LEMMA 18. *When sellers are SB, the payment in the end of GAI auction is sell-side  $(e + 2)\epsilon$ -VCG.*

THEOREM 19. *SB is an  $(3e + 5)\epsilon$  ex-post Nash Equilibrium for sellers in GAI auction. That is, sellers cannot gain more than  $(3e + 5)\epsilon$  by deviating.*

In practice, however, sellers are unlikely to have the information that would let them exploit that potential gain. They are much more likely to lose from bidding on their less attractive configurations.

## 5.5 Computation and Complexity

The size of the price space maintained in the auction is equal to the total number of sub-configurations, meaning it is exponential in  $\max_r |I_r|$ . This is also equivalent to the *tree-width* (plus one) of the original CDI-map. For the purpose of the computational analysis let  $d_j$  denote the domain of attribute  $a_j$ , and  $I = \bigcup_{r=1}^g \prod_{j \in I_r} d_j$ , the collection of all sub-configurations. The first purpose of this sub-section is to show that the complexity of all the computations required for the auction depends only on  $|I|$ , i.e., no computation depends on the size of the full exponential domain.

We are first concerned with the computation of  $\mathcal{M}^t$ . Since  $M^t$  grows monotonically with  $t$ , a naive application of optimization algorithm to generate the best outcomes sequentially might end up enumerating significant portions of the fully exponential domain. However as shown below this plain enumeration can be avoided.

PROPOSITION 20. *The computation of  $\mathcal{M}^t$  can be done in time  $O(|I|^2)$ . Moreover, the total time spent on this task throughout the auction is  $O(|I|(|I| + T))$ .*

The bounds are in practice significantly lower, based on results on similar problems from the probabilistic reasoning literature [18].

One of the benefits of the compact pricing structure is the compact representation it lends for bids: sellers submit only sub-bids, and therefore the number of them submitted and stored per seller is bounded by  $|I|$ . Since the computation tasks:  $B_i^t \neq \emptyset$ , rule [SWITCH] and choice of  $\eta_i$  are all involving the set  $B_i^t$ , it is important to note that their performance only depend on the size of the set  $\mathcal{B}_i^t$ , since they are all subsumed by the combinatorial optimization task over  $\mathcal{B}_i^t$  or  $\mathcal{B}_i^t \cap \mathcal{M}^t$ .

Next, we analyze the number of rounds it takes for the auction to terminate. Phase B requires  $\max_{i=1, \dots, n} \pi_i^T(\eta_i) \frac{1}{\epsilon}$ . Since this is equivalent to price-only auctions, the concern is only with the time complexity of phase A. Since prices cannot go below  $f_{b,r}(\theta_r)$ , an upper bound on the number of rounds required is

$$T \leq \sum_{\theta_r \in I} (p^1(\theta_r) - f_{b,r}(\theta_r)) \frac{g}{\epsilon}$$

However phase A may converge faster. Let the initial negative profit chosen by the auctioneer be  $m = \max_{\theta \in \Theta} \pi_b^1(\theta)$ . In the worst case phase A needs to run until  $\forall \theta \in \Theta. \pi_b(\theta) = m$ . This happens for example when  $\forall \theta_r \in I. p^t(\theta_r) = f_{b,r}(\theta_r) + \frac{m}{g}$ . In general, the closer the initial prices reflect buyer valuation, the faster phase A converges. One extreme is to choose  $p^1(\theta_r) =$

	$I_1$				$I_2$			
	$a^1b^1$	$a^2b^1$	$a^1b^2$	$a^2b^2$	$b^1c^1$	$b^2c^1$	$b^1c^2$	$b^2c^2$
$f_b$	65	50	55	70	50	85	60	75
$f_1$	35	20	30	70	65	65	70	61
$f_2$	35	20	25	25	55	110	70	95

**Table 1: GAI utility functions for the example domain.**  $f_b$  represents the buyer’s valuation, and  $f_1$  and  $f_2$  costs of the sellers  $s_1$  and  $s_2$ .

$f_{b,r}(\theta_r) + \frac{m}{g}$ . That would make phase A redundant, at the cost of full initial revelation of buyer’s valuation as done in other mechanisms discussed below. Between this option and the other extreme, which is  $\forall \alpha, \hat{\alpha} \in I, p^1(\alpha) = p^1(\hat{\alpha})$  the auctioneer has a range of choices to determine the right tradeoff between convergence time and information revelation. In the example below the choice of a lower initial price for the domain of  $I_1$  provides some speedup by revealing a harmless amount of information.

Another potential concern is the communication cost associated with the Japanese auction style. The sellers need to send their bids over and over again at each round. A simple change can be made to avoid much of the redundant communication: the auction can retain sub-bids from previous rounds on sub-configurations whose price did not change. Since combinations of sub-bids from different rounds can yield sub-optimal configurations, each sub-bid should be tagged with the number of the latest round in which it was submitted, and only consistent combinations *from the same round* are considered to be full bids. With this implementation sellers need not resubmit their bid until a price of at least one sub-configuration has changed.

## 5.6 Example

We use the example settings introduced in Section 5.2. Recall that the GAI structure is  $I_1 = \{a, b\}$ ,  $I_2 = \{b, c\}$  (note that  $e = 1$ ). Table 1 shows the GAI utilities for the buyer and the two sellers  $s_1, s_2$ . The efficient allocation is  $(s_1, a^1b^2c^1)$  with a surplus of 45. The maximal surplus of the second best seller,  $s_2$ , is 25, achieved by  $a^1b^1c^1$ ,  $a^2b^1c^1$ , and  $a^2b^2c^2$ . We set all initial prices over  $I_1$  to 75, and all initial prices over  $I_2$  to 90. We set  $\epsilon = 8$ , meaning that price reduction for sub-configurations is 4. Though with these numbers it is not guaranteed by Theorem 17, we expect  $s_1$  to win on either the efficient allocation or on  $a^1b^2c^2$  which provides a surplus of 39. The reason is that these are the only two configurations which are within  $(e + 1)\epsilon = 16$  of being efficient for  $s_1$  (therefore one of them must be chosen by Phase A), and both provide more than  $\epsilon$  surplus over  $s_2$ ’s most efficient configuration (and this is sufficient in order to win in Phase B).

Table 2 shows the progress of phase A. Initially all configuration have the same cost (165), so sellers bid on their lowest cost configuration which is  $a^2b^1c^1$  for both (with profit 80 to  $s_1$  and 90 to  $s_2$ ), and that translates to sub-bids on  $a^2b^1$  and  $b^1c^1$ .  $\mathcal{M}^1$  contains the sub-configurations  $a^2b^2$  and  $b^2c^1$  of the highest value configuration  $a^2b^2c^1$ . Price is therefore decreased on  $a^2b^1$  and  $b^1c^1$ . After the price change,  $s_1$  has higher profit (74) on  $a^1b^2c^2$  and she therefore bids on  $a^1b^2$  and  $b^2c^2$ . Now (round 2) their prices go down, reducing the profit on  $a^1b^2c^2$  to 66 and therefore in round 3  $s_1$  prefers  $a^2b^1c^2$  (profit 67). After the next price change the configurations  $a^1b^2c^1$  and  $a^1b^2c^2$  both become optimal (profit 66), and the sub-bids  $a^1b^2$ ,  $b^2c^1$  and  $b^2c^2$  capture the two. These configurations stay optimal for another round (5), with profit 62. At this point  $s_1$  has a full bid (in fact two full bids:  $a^1b^2c^2$  and  $a^1b^2c^1$ ) in  $M^5$ , and

t	$I_1$				$I_2$			
	$a^1b^1$	$a^2b^1$	$a^1b^2$	$a^2b^2$	$b^1c^1$	$b^2c^1$	$b^1c^2$	$b^2c^2$
1	75	75	75	75	90	90	90	90
2	75	$s_1, s_2$	75	75	$s_1, s_2$	*	90	90
3	75	$s_2$	$s_1$	75	$s_2$	*	*	$s_1$
4	75	$s_1, s_2$	71	75	$s_2$	*	$s_1$	*
5	75	$s_2$	$s_1$	75	$s_2$	*, $s_1$	90	86
6	71	59	67	75	74	90	86	86
7	71	55	67	75	70	90	82	86
8	67	55	67	75	66	90	82	86
9	*	$s_2$	*, $s_1$	*	*	*, $s_1$	$s_2$	*, $s_1$
	67	51	67	75	66	90	78	86
	*, $s_2$		*, $s_1$	*	*, $s_2$	*, $s_1$		*, $s_1$

**Table 2: Auction progression in phase A. Sell bids and designation of  $\mathcal{M}^t$  (using \*) are shown below the price of each sub-configuration.**

therefore she no longer changes her bids since the price of her optimal configurations does not decrease.  $s_2$  sticks to  $a^2b^1c^1$  during the first four rounds, switching to  $a^1b^1c^1$  in round 5. It takes four more rounds for  $s_2$  and  $\mathcal{M}^t$  to converge ( $\mathcal{M}^{10} \cap B_2^{10} = \{a^1b^1c^1\}$ ).

After round 9 the auction sets  $\eta_1 = a^1b^2c^1$  (which yields more buyer profit than  $a^1b^2c^2$ ) and  $\eta_2 = a^1b^1c^1$ . For the next round (10)  $\Delta = 8$ , increased by 8 for each subsequent round. Note that  $p^9(a^1b^1c^1) = 133$ , and  $c_2(a^1b^1c^1) = 90$ , therefore  $\pi_2^T(\eta_2) = 43$ . In round 15,  $\Delta = 48$  meaning  $p^{15}(a^1b^1c^1) = 85$  and that causes  $s_2$  to drop out, setting the final allocation to  $(s_1, a^1b^2c^1)$  and  $p^{15}(a^1b^2c^1) = 157 - 48 = 109$ . That leaves the buyer with a profit of 31 and  $s_1$  with a profit of 14, less than  $\epsilon$  below the VCG profit 20.

The welfare achieved in this case is optimal. To illustrate how some efficiency loss could occur consider the case that  $c_1(b^2c^2) = 60$ . In that case, in round 3 the configuration  $a^1b^2c^2$  provides the same profit (67) as  $a^2b^1c^2$ , and  $s_1$  bids on both. While  $a^2b^1c^2$  is no longer optimal after the price change,  $a^1b^2c^2$  remains optimal on subsequent rounds because  $b^2c^2 \in \mathcal{M}^t$ , and the price change of  $a^1b^2$  affects both  $a^1b^2c^2$  and the efficient configuration  $a^1b^2c^1$ . When phase A ends  $B_1^{10} \cap M^{10} = \{a^1b^2c^2\}$  so the auction terminates with the slightly suboptimal configuration and surplus 40.

## 6. DISCUSSION

### 6.1 Preferential Assumptions

A key aspect in implementing GAI based auctions is the choice of the preference structure, that is, the elements  $\{I_1, \dots, I_g\}$ . In some domains the structure can be more or less robust over time and over different decision makers. When this is not the case, extracting reliable structure from sellers (in the form of CDI conditions) is a serious challenge. This could have been a deal breaker for such domains, but in fact it can be overcome. It turns out that we can run this auction without any assumptions on sellers’ preference structure. The only place where this assumption is used in our analysis is for Lemma 8. If sellers whose preference structure does not agree with the one used by the auction are guided to submit only one full bid at each round, or a set of bids that does not yield undesired consistent combinations, all the properties of the auction

still hold. Locally, the sellers can optimize their profit functions using the union of their GAI structure with the auction's structure. It is therefore essential only that the buyer's preference structure is accurately modeled. Of course, capturing sellers' structures as well is still preferred since it can speed up the execution and let sellers take advantage of the compact bid representation.

In both cases the choice of clusters may significantly affect the complexity of the price structure and the runtime of the auction. It is sometimes better to ignore some weaker interdependencies in order to reduce dimensionality. The complexity of the structure also affects the efficiency of the auction through the value of  $e$ .

## 6.2 Information Revelation Properties

In considering information properties of this mechanism we compare to the standard approach for iterative multiattribute auctions, which is based on the theoretical foundations of Che [7]. In most of these mechanisms the buyer reveals a scoring function and then the mechanism solicits bids from the sellers [3, 22, 8, 21] (the mechanisms suggested by Beil and Wein [2] is different since buyers can modify their scoring function each round, but the goal there is to maximize the buyer's profit). Whereas these iterative procurement mechanisms tend to relieve the burden of information revelation from the sellers, a major drawback is that the buyer's utility function must be revealed to the sellers before receiving any commitment. In the mechanisms suggested by PK and in our GAI auction above, buyer information is revealed only in exchange for sell commitments. In particular, sellers learn nothing (beyond the initial price upper bound, which can be arbitrarily loose) about the utility of configurations for which no bid was submitted. When bids are submitted for a configuration  $\theta$ , sellers would be able to infer its utility relative to the current preferred configurations only after the price of  $\theta$  is driven down sufficiently to make it a preferred configuration as well.

## 6.3 Conclusions

We propose a novel exploitation of preference structure in multiattribute auctions. Rather than assuming full additivity, or no structure at all, we model preferences using the GAI decomposition. We developed an iterative auction mechanism directly relying on the decomposition, and also provided direct means of constructing the representation from relatively simple statements of willingness-to-pay. Our auction mechanism generalizes PK's preference modeling, while in essence retaining their information revelation properties. It allows for a range of tradeoffs between accuracy of preference representation and both the complexity of the pricing structure and efficiency of the auction, as well as tradeoffs between buyer's information revelation and the time required for convergence.

## 7. ACKNOWLEDGMENTS

This work was supported in part by NSF grants IIS-0205435 and IIS-0414710, and the STIET program under NSF IGERT grant 0114368. We are grateful to comments from anonymous reviewers.

## 8. REFERENCES

- [1] F. Bacchus and A. Grove. Graphical models for preference and utility. In *Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 3–10, Montreal, 1995.
- [2] D. R. Beil and L. M. Wein. An inverse-optimization-based auction for multiattribute RFQs. *Management Science*, 49:1529–1545, 2003.
- [3] M. Bichler. *The Future of e-Markets: Multi-Dimensional Market Mechanisms*. Cambridge University Press, 2001.
- [4] C. Boutilier, F. Bacchus, and R. I. Brafman. UCP-networks: A directed graphical representation of conditional utilities. In *Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 56–64, Seattle, 2001.
- [5] R. I. Brafman, C. Domshlak, and T. Kogan. Compact value-function representations for qualitative preferences. In *Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 51–59, Banff, 2004.
- [6] D. Braziunas and C. Boutilier. Local utility elicitation in GAI models. In *Twenty-first Conference on Uncertainty in Artificial Intelligence*, pages 42–49, Edinburgh, 2005.
- [7] Y.-K. Che. Design competition through multidimensional auctions. *RAND Journal of Economics*, 24(4):668–680, 1993.
- [8] E. David, R. Azoulay-Schwartz, and S. Kraus. An English auction protocol for multi-attribute items. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Artificial Intelligence*, pages 52–68. Springer, 2002.
- [9] G. Debreu. Topological methods in cardinal utility theory. In K. Arrow, S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*. Stanford Univ. Press, 1959.
- [10] J. S. Dyer and R. K. Sarin. An axiomatization of cardinal additive conjoint measurement theory. Working Paper 265, WMSI, UCLA, February 1977.
- [11] J. S. Dyer and R. K. Sarin. Measurable multiattribute value functions. *Operations Research*, 27:810–822, 1979.
- [12] Y. Engel, M. P. Wellman, and K. M. Lochner. Bid expressiveness and clearing algorithms in multiattribute double auctions. In *Seventh ACM Conference on Electronic Commerce*, pages 110–119, Ann Arbor, MI, 2006.
- [13] P. C. Fishburn. Interdependence and additivity in multivariate, unidimensional expected utility theory. *Intl. Economic Review*, 8:335–342, 1967.
- [14] C. Gonzales and P. Perny. GAI networks for utility elicitation. In *Ninth Intl. Conf. on the Principles of Knowledge Representation and Reasoning*, pages 224–234, Whistler, BC, 2004.
- [15] C. Gonzales and P. Perny. GAI networks for decision making under certainty. In *IJCAI-05 Workshop on Advances in Preference Handling*, Edinburgh, 2005.
- [16] N. Hyafil and C. Boutilier. Regret-based incremental partial revelation mechanisms. In *Twenty-first National Conference on Artificial Intelligence*, pages 672–678, Boston, MA, 2006.
- [17] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, 1976.
- [18] D. Nilsson. An efficient algorithm for finding the M most probable configurations in probabilistic expert systems. *Statistics and Computing*, 8(2):159–173, 1998.
- [19] D. C. Parkes and J. Kalagnanam. Models for iterative multiattribute procurement auctions. *Management Science*, 51:435–451, 2005.
- [20] J. Pearl and A. Paz. Graphoids: A graph based logic for reasoning about relevance relations. In B. Du Boulay, editor, *Advances in Artificial Intelligence II*. 1989.
- [21] J. Shachat and J. T. Swarthout. Procurement auctions for differentiated goods. IBM Research Report RC22587, IBM T.J. Watson Research Laboratory, 2002.
- [22] N. Vulkan and N. R. Jennings. Efficient mechanisms for the supply of services in multi-agent environments. *Decision Support Systems*, 28:5–19, 2000.