

Preference-Based Characterizations of Truthfulness and the Limited Expressiveness of Order-Based Domains

Florin Constantin and David C. Parkes
Division of Engineering and Applied Sciences
Harvard University, Cambridge, MA 02138
{florin, parkes}@eecs.harvard.edu

Abstract

An important direction in computational mechanism design is to characterize the space of choice functions that can be truthfully implemented. For this, one must carefully describe the class of preferences in a domain. For unrestricted preferences the domain is well-characterized, and small. Moreover, recent work [Lavi *et al.*, 2003] has allowed for “order-based” preferences but found essentially the same (negative) characterization. However, most interesting domains have preferences that are still more structured than those allowed in the order-based model. We highlight this issue in our paper, through many examples, thus demonstrating the limited applicability of this result. We propose extensions to the model of order-based domains to capture new preference structure, and conjecture that more positive characterizations for truthfulness are possible. We also advocate, in proposing a research direction for sufficient conditions for truthfulness, that attention be restricted to natural (critical-value based) payment functions.

1 Introduction

In a *mechanism*, a *center* chooses between a number of outcomes or alternatives. The choice is made based on reports (not necessarily sincere) of preferences (also called types) from selfish *agents*. The center then decides how much the agents should pay given the outcome and their reported preferences. A simple example of a mechanism is an auction where agents submit bids and the auctioneer decides how to allocate goods and fixes payments.

Mechanism design is concerned with designing social choice functions (scfs) that satisfy certain desirable properties. A central property is *truthfulness*: a mechanism is *truthful* if it is a dominant strategy for agents to declare their true types. We say that the choice function f is truthful if there exists some payment scheme p that together with f defines a truthful mechanism. A central question, recently revisited in Lavi *et al.* [2003] is to characterize the truthful choice functions.

There has been a surge of interest in mechanism design, motivated in part by an increasing number of computer sci-

ence applications in which participants are self interested, with distributed systems more often the norm than the exception and with the continued automation of commerce. When characterizing the truthful choice functions it is important to carefully define the preference space because this affects the results. Loosely, the more structured the space of preferences, the more positive the characterization. For completely unconstrained preferences, Roberts [1979] showed that the truthful functions are exactly the *affine-maximizer* (AM) functions. This can be understood as a negative result, because it identifies the class of well-known Groves mechanisms as the only truthful mechanisms for unrestricted preferences.

Lavi *et al.* [2003] recognized an important shortcoming with Robert’s result, noting that many preferences are known to be structured in many interesting domains. In their words, *the assumption of unrestricted domains is not without restriction*. In response, the authors consider the family of *order-based* (OB) domains. OB domains are significant because they allow one to state restrictions on preferences such as *free-disposal*, and *no-externalities* which are the hallmark of many classic preference models. As such, OB domains are able to capture the preference structure of combinatorial auctions. However, despite this new ability to capture preference structure, Lavi *et al.* [2003] recovered essentially the same negative result as Roberts [1979], at least for the case where the scf satisfies the Independence of Irrelevant Alternatives property.

In this paper, we highlight many examples with more preference structure than the OB domains, thus demonstrating the limited applicability of the results in Lavi *et al.* [2003]. The negative result is not as powerful as it first seems. We argue that the assumption of OB domains is not without restriction, and look to minimal and modular extensions to the order-based description language that can capture our counterexamples. This provides motivation to look for alternate characterizations of preference domains, and for positive (i.e. non affine-maximizer) results. In addition, we argue for two additional properties of a useful characterization: that of *natural* payment functions, and *universal applicability*. We return to these requirements in Subsection 3.1.

2 Preliminaries

We assume that there are n agents, a finite range of alternatives A ($|A| = m$), and a domain of preferences (valuations

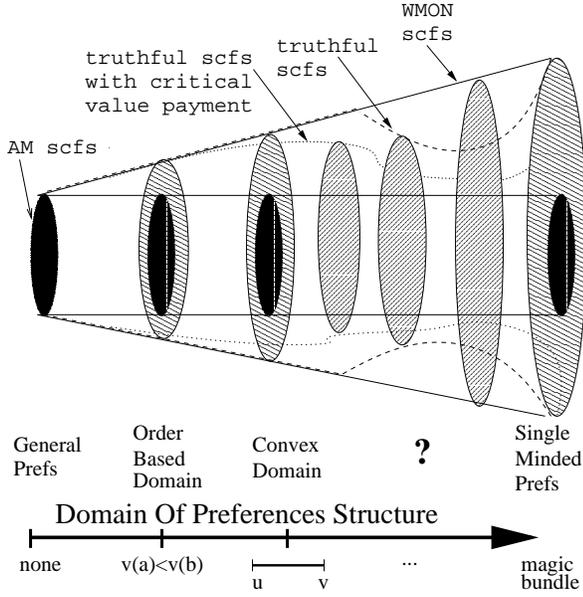


Figure 1: Illustration of the inclusion relations between subclasses of WMON functions on various domains. WMON is equivalent to truthful for order-based or convex domains. First example in Section 7 of the [Saks and Yu, 2005] provides a function that is WMON but not truthful for a non-convex domain. Roberts’ result ([Roberts, 1979]) states that only AM are truthful for the domain of general valuations. For single-minded preferences, [Lehmann *et al.*, 2002] shows that monotonicity is equivalent to truthfulness and thus the class of truthful functions is richer than the class of AM.

over alternatives), $V = V_1 \times \dots \times V_n$, where each agent i has some (private) valuation $v_i \in V_i$ such that $v_i : A \rightarrow \mathbb{R}$. We can regard any $v_i \in V_i$ as a vector in \mathbb{R}^m with a component for each alternative. For $v \in V$ and $1 \leq i \leq n$ we denote $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$. Similarly, we denote $V_{-i} = V_1 \times \dots \times V_{i-1} \times V_{i+1} \times \dots \times V_n$.

A mechanism consists of a choice function $f : V \rightarrow A$ and payments $p = (p_1, \dots, p_n)$ where $p_i = p_i(v, f(v)) : V \times A \rightarrow \mathbb{R}$. Agents are assumed to have *quasi-linear* utilities, i.e. if any agent j has true type v_j and reports type v'_j then, letting $v' = (v'_1, \dots, v'_n)$, the utility for agent i is given by

$$u_i(v_i, f(v')) = v_i(f(v')) - p_i(v', f(v')) \quad (1)$$

The following is a formal definition of truthfulness in a dominant strategy equilibrium:

Definition 1 Mechanism (f, p) is truthful if every agent i maximizes its utility by truthfully reporting its type v_i :

$$u_i(v_i, f(v_i, v'_{-i})) \geq u_i(v_i, f(v'_i, v'_{-i})) \quad \forall v'_i \in V_i \quad (2)$$

for all reports $v'_{-i} \in V_{-i}$ from the other agents.

If the mechanism is viewed as a game with the agents as players, then a mechanism is truthful if and only if it is a dominant strategy for the players to bid truthfully.

3 Preference-Based Characterizations of Truthfulness

In this section, we will review some of the recent papers characterizing truthfulness and draw some connections between

them. Figure 1 serves to illustrate this section. First, it is well-known [Rochet, 1987] that a monotonicity property (known as *weak-monotonicity*) is necessary for truthfulness, whatever the preference domain.

Definition 2 The social choice function f satisfies weak monotonicity WMON if $\forall i, \forall v_{-i} \in V_{-i}, \forall u_i, v_i \in V_i$

$$\left. \begin{array}{l} f(v_i, v_{-i}) = a \\ f(u_i, v_{-i}) = b \end{array} \right\} \Rightarrow u_i(b) - v_i(b) \geq u_i(a) - v_i(a) \quad (3)$$

That is, if the outcome chosen by f changes from a to b and only agent i has changed its report, from v_i to u_i , then it must be that the relative increase in value is weakly higher for the new outcome b than for the old outcome a .

In Figure 1, we draw the outer-cone to illustrate the space of WMON functions. For unrestricted domains, WMON is also sufficient for truthfulness.¹ Seemingly positive, this result which says that *all* WMON functions are truthful for unrestricted preferences, turns negative when one realizes that the space of WMON functions is *exactly that of affine maximizers for unrestricted preferences*. This is not so for other preference domains.

Definition 3 f is an affine maximizer if $\forall a \in A \exists \gamma_a \in \mathbb{R}$ and $\forall i \exists w_i \in \mathbb{R}$ such that

$$f(v) \in \operatorname{argmax}_{a \in A} \left\{ \sum_{i=1}^n w_i v_i(a) + \gamma_a \right\} \quad (4)$$

The affine maximizers are illustrated in Figure 1 through the central (cylindrical) volume. So, for general preferences (to the left) we have this volume taking all of the size of the outer WMON cone. Moreover, notice that the class of Groves mechanisms (e.g. via the Vickrey-Clarke-Groves (VCG) payment scheme), provide truthful payments for all affine maximizers. On the other hand, Lehmann *et al.* [2002] had demonstrated that while WMON was also sufficient for single-minded bidders (see Section 4), the class of WMON functions included non affine-maximizers for that preference domain. Thus, the WMON outer-cone is larger than the AM cylinder for single-minded preferences (to the right).

Lavi *et al.* [2003] considers a model that allows for some structure to preferences. This is the class of *order-based* (OB) domains:

Definition 4 A domain V_i is order-based (OB) if it is defined by a set \mathcal{C} of (in)equalities of the form $v_i(a) \{<, \leq, =\} v_i(b)$ or $v_i(a) = 0$ where $a, b \in A$. That is, all $v_i \in V_i$ and only them satisfy the (in)equalities in \mathcal{C} . A domain $V = V_1 \times \dots \times V_n$ is order-based if each V_i is order-based.

We provide many examples of using the OB model in the later sections in the paper. For now, just note that the domain of unrestricted domains is trivially OB by using no constraints, i.e. $\mathcal{C} = \emptyset$. Lavi *et al.* [2003] establish that WMON remains sufficient for truthfulness in OB domains. Recently, Saks and Yu [2005] extend this result and establish that WMON also remains sufficient for truthfulness in the more expressive class of *convex* domains. A domain is

¹This is a corollary of the sufficiency of WMON in order-based domains [Lavi *et al.*, 2003].

convex if for any $v_i, v'_i \in V_i$ then $\lambda v_i + (1 - \lambda)v'_i \in V_i$ for $\lambda \in [0, 1]$. The following is immediate:

Proposition 1 Any order-based domain V_i is convex.

The converse implication is not true: it can be shown that the domain of linear threshold preferences (Subsection 5.2) is convex but not order-based.

Again, we should not view this sufficiency of WMON for order-based as a positive result because (in their main result) Lavi et al. [2003] establish that the WMON functions remain essentially equivalent to the affine maximizer functions for OB domains. We also note that Gui et al. [2004] adopt a graph-theoretic approach to truthfulness and establish a number of preference domains for which WMON is sufficient for truthfulness, some of which are not order-based. However, neither them nor Saks and Yu [2005] consider how this class relates to the affine-maximizers.

Some of the preference domains that we consider lie outside of order-based and convex domains (see Figures 1 and 2). It is in these (important) domains that new characterization results are required.

3.1 Properties Beyond Truthfulness

In proposing a research direction to characterize sufficient conditions for truthfulness in more structured domains, we also advocate looking for conditions that provide the following two additional properties:

- **Natural payments.** We would prefer a characterization of truthful functions that are supported with a *natural* payment function. We favor the “critical value” payment function that generalizes the VCG payment for non affine-maximizer functions. In particular, for outcome $a^* = f(v)$, denote payment function:

$$p_i(v, a^*) = \min_{v'_i \in V_i} v'_i(a^*) \quad (5)$$

s.t. $f(v'_i, v_{-i}) \in E_i(a^*)$

where $E_i(a^*) \subseteq A$ is the set of alternatives for which agent i is indifferent between that alternative and a^* , given its valuation. It is not hard to show that this is precisely the VCG payment rule for the case that f is an affine maximizer. Moreover, this is precisely the payment rule in the auction family of Lehmann et al. [2002] for single-minded preferences. Yokoo [2003] also states the importance of critical-value payment functions.

Most of the recent characterizations that we reviewed fail to satisfy this desideratum.

- **Universal applicability.** We would prefer a characterization that is *universal* for some large class of structured preferences. Here, we draw an analogy with the Groves mechanisms. Groves mechanisms are “universal” in that they apply whatever the preference domain. This universal applicability is surely one reason for their appeal. While recognizing that new characterizations require new structural assumptions, we wish to identify requirements that are both minimal and easy to interpret and apply.

4 Combinatorial Auctions with Structure

In the following two sections we investigate whether several examples of preference domains are order-based and thus il-

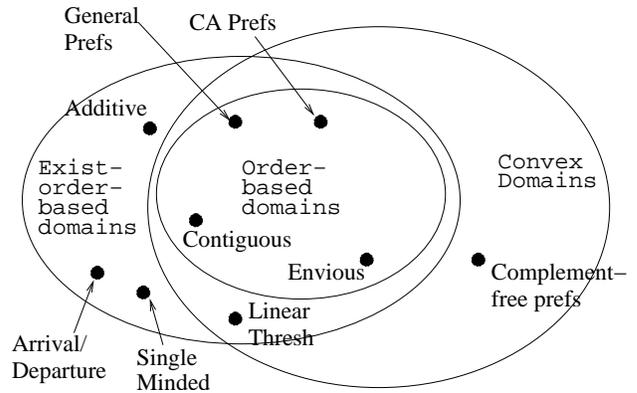


Figure 2: Domain classes and example domains.

lustrate where the negative affine-maximizer result does not hold. In Section 6 we propose extensions to the notion of order-based domains that capture our negative examples. Figure 2 provides an overview of most of our results.

We first consider domains of preferences in combinatorial auctions (CAs) but with additional preference structure. In a CA, the auctioneer has k different items to sell to the n agents. An alternative $a = (a_1, \dots, a_n)$ in a CA is an allocation of the goods auctioned with agent i getting bundle a_i such that the a_i 's are disjoint. The domain of CA preferences is order-based, being defined by the set of inequalities \mathcal{C}_{CA} described below:

- Free disposal. (“if I am allocated extra items then my value does not go down”): $v_i(a) \leq v_i(b), \forall a, b \in A$ s.t. $a_i \subseteq b_i$
- No externalities. (“I only care about my own bundle”): $v_i(a) = v_i(b), \forall a, b \in A$ s.t. $a_i = b_i$
- Normalization. (“getting nothing has value 0 for me”): $v_i(a) = 0, \forall a \in A$ s.t. $a_i = \emptyset$

In the CA setting, a very restricted preference domain is the one of single-minded bidders, where agents only care if they get a certain set of items, the others being irrelevant.

Definition 5 A bidder is single-minded if there exist $w \in \mathbb{R}$ and a “magic bundle” S^* such that

$$v_i(a) = \begin{cases} w & \text{if } S^* \subseteq a_i \\ 0 & \text{otherwise} \end{cases} \quad \forall a \in A$$

If $|S^*| = 1$ then we call this preference the Unknown-Item-Unknown-Value (UIUV).

Note that all single-minded preferences certainly satisfy the constraints in \mathcal{C}_{CA} . However, the domain of single-minded preferences is not order-based.²

In fact, we prove the following stronger result, which provides a powerful tool for many of our negative examples:

Proposition 2 Let V_i be a preference domain that satisfies the constraints in \mathcal{C}_{CA} , contains all UIUV preferences but

²It is also simple to see that the domain of single-minded preferences is not convex. Note that the domain of *known* single-minded preferences (where the magic bundle is known beforehand) is order-based.

does not contain all possible CA preferences. Then V_i can not be order-based.

Proof. Suppose for a contradiction that V_i was order-based and let \mathcal{C} be the set of constraints that defines it. Then \mathcal{C} must contain all constraints for CAs \mathcal{C}_{CA} . However, only these are not enough as V_i is not the domain of all possible CA valuations. Thus \mathcal{C} must contain at least one inequality of the form $v_i(a) \{<, \leq, =\} v_i(b)$ where $a_i \not\subseteq b_i$. Assume wlog that this relation is $v_i(a) \leq v_i(b)$. This inequality is not satisfied by any UIUV preference v for item x with $x \in a_i \setminus b_i$ that has value 0 for bundle b_i and some positive value for bundle a_i , contradiction. \square

Note that we can replace the domain of UIUV preferences in the proposition with any domain D with the following property: for any alternatives a and b there exist valuations $u, v \in D$ such that $u(a) > u(b)$ but $v(a) < v(b)$. The key property in the proof is preserved: no ordering inequality that holds for all $v \in D$ can be required for any two alternatives a and b .

We will be using Proposition 2 extensively. For now, note:

Corollary 1 *Single-minded preferences are not order-based.*

Also, note that if some preference domain D_1 is not order-based and $UIUV \subseteq D_2 \subset D_1$, then D_2 is not order-based.

4.1 Bounded XOR and Unit-Demand Preferences

Let the class of “bounded XOR preferences” denote valuations $v_i \in V_i$ for which there is some set of bundles \mathcal{X} , $|\mathcal{X}| < 2^n$, for which $v_i(S) = \max_{S' \in \mathcal{X}} v_i(S')$ for all bundles S . Note that this class assumes knowledge that no bidder needs to express a value for every possible bundle. Thus, some valuations, e.g. any valuation that associates a different value for each bundle, can not be expressed in this domain. We also have the following immediate corollary:

Corollary 2 *Bounded XOR preferences are not order-based.*

It also follows as an immediate corollary that *unit-demand preferences are not order-based.*

4.2 Complement-Free Preferences

Again in a CA setting, consider the domain of complement-free preferences, i.e. preferences that do not have any complementarity relation between bundles (see [Lehmann *et al.*, 2001] for a characterization of this domain):

$$v_i(a_i) + v_i(b_i) \geq v_i(a_i \cup b_i) \quad \forall a_i \cap b_i = \emptyset \quad (6)$$

The domain of complement-free preferences is not order-based as it contains all UIUV preferences and Proposition 2 applies.

4.3 Additive preferences

A valuation v_i is *additive (OR)* if there exist a number k and k “atoms” (bundles) $B^{(1)}, \dots, B^{(k)}$ with values $s^{(1)}, \dots, s^{(k)}$ such that the value of any bundle a_i is the value of the maximum packing of the atoms in a_i (or 0 if no atom is included in a_i).³ The domain of additive preferences is not order-based

³Formally, the evaluation function for an agent i and an allocation $a \in A$ is: $v_i(a)$ equals 0 if $B^{(h)} \not\subseteq a_i, \forall 1 \leq h \leq k$, or $\max \sum_{j=1}^r s^{(h_j)}$ over all decompositions $B^{(h_1)} \cup \dots \cup B^{(h_r)} \subseteq a_i$ of disjoint atoms $B^{(h_j)}$'s otherwise.

because it contains all UIUV preferences but is not fully expressive, and thus Proposition 2 applies.

Corollary 3 *Additive preferences are not order-based.*

4.4 Contiguous preferences

Consider a setting where ordered items are auctioned (for example CPU time shares on a single-processor machine). Let contiguous preferences denote preferences of agents that only have positive value for a contiguous allocation of resources on the machine. Free disposal holds, so that if an agent receives two sequences then its value is the *max* of its value for each individual sequence. However, we will not assume knowledge of any additional structure (such as each agent needs a particular length of allocation).

Proposition 3 *Contiguous preferences are order-based.*

Proof. The proof is immediate once one recognizes that we can redefine the alternative space so that each alternative $a \in A$ gives each agent i either no allocation or a contiguous allocation. Then this is a CA preference domain in this new alternative space. \square

Note that this result is very brittle. If we now say that it is also known in the domain that each bidder has an (unknown beforehand) minimal length of allocation for which it has value then we lose the order-based property.

4.5 Attribute-Based Preferences

In some domains there is an attribute to goods and this attribute is important in defining a bidder’s valuation. For instance, in online auctions [Lavi and Nisan, 2000], this attribute is *time* and we can model preferences such as “I have positive value for outcome t iff $t_1 \leq t \leq t_2$,” (where t_1 and t_2 are interpreted as arrival and departure times not known beforehand). In general, an attribute-based preference is one where there is some condition $g_i(a) \in \{0, 1\}$ on an alternative $a \in A$ that determines whether or not the agent has value for the alternative (with $v_i(a) > 0$ if and only if $g_i(a)$). Condition g_i is private to each agent (or perhaps it is a parameterized condition for which the parameters, e.g. t_1 and t_2 , are unknown to the mechanism).

Proposition 4 *Attribute-based preferences are not order-based.*

Proof. Just note that for all $a_1, a_2 \in A$ we can find some preferences for which $g_i(a_1) = 1$ and $g_i(a_2) = 0$ and vice versa. Thus, no order-based inequalities can be put forth for this domain. As this domain is not the domain of unrestricted preferences we conclude that it is not order-based. \square

5 Multi-Unit Auctions with Structure

Consider a multi-unit auction (MUA) setting, where k identical goods are for sale to the n agents. An allocation a here is defined by a vector of values a_1, \dots, a_n where a_i represents the number of items allocated to agent i and $a_1 + \dots + a_n \leq k$. Out of the three axioms for CAs we will just assume normalization (“If I get no items then I have value 0 for this allocation”).

5.1 Envious Preferences

Consider a preference such as: “I have positive value for an allocation only if my share is the biggest”, i.e. $v_i(a) > 0$ iff $a_i \geq a_j \forall j \neq i$. We call these *envious preferences*.

Proposition 5 *Envious preferences are order-based.*

Proof. The set \mathcal{C} that defines this domain contains all inequalities of the form:

- $v_i(a) = 0$ iff $\exists j \neq i$ s.t. $a_i < a_j$
- $v_i(a) > v_i(\emptyset) = 0$ iff $\forall j \neq i, a_i \geq a_j$

□

Free-disposal does not hold for envious preferences: if there are 3 agents and 20 items agent 1 has positive value for the allocation vector (8, 7, 5) but zero value for (9, 11, 0). However, one can add “limited” free-disposal to envious preferences, in the sense that $a'_i \geq a_i \geq a_j \forall j \Rightarrow v_i(a'_i) \geq v_i(a_i) > 0$. Envious preferences with limited free-disposal remain order-based.

5.2 Linear Threshold preferences

Again, in a multi-unit auction, consider preferences like “I have positive value for an allocation iff I get at least r items”, where r is not known beforehand. A more formal definition of the domain THR of linear threshold preferences is: $v_i \in \text{THR}$ iff $\exists 1 \leq r = r(v_i) \leq k$ s.t. $v_i(a) = 0$ if $a_i < r$ and $v_i(a) > 0$ if $a_i \geq r$.

The following result shows that linear threshold preferences are not order-based. The proof uses the fact that all inequalities for normalization, free-disposal and no externalities should be used in defining this domain and shows that no other constraint can be added.

Proposition 6 *The domain THR is not order-based.*

Proof. Suppose for a contradiction that THR was order-based, defined by the set of inequalities \mathcal{C} . Note first that any linear threshold valuation satisfies normalization, free disposal and no externalities so all the inequalities characterizing those properties must be in \mathcal{C} . However, those are not sufficient as there are many MUA preferences that are not linear threshold. Let us analyze what kind of constraints we can add to the ones already in \mathcal{C} .

- Can we add any inequality of the form $v_i(a) < v_i(b)$ for some $a, b \in A$? Because of free disposal and no externalities we need $a_i < b_i$. The linear threshold valuation that has value 1 for any allocation that assigns at least a_i items to agent i does not satisfy this inequality, so we can not add such an inequality.
- Can we add any constraint of the form $v_i(a) = v_i(b)$ for $a_i < b_i$? No, as any linear threshold valuation with b_i as threshold does not satisfy this constraint.
- ...

In a similar fashion, we get that we can not add any other constraints to \mathcal{C} . Therefore we can not construct a set of constraints that defines THR, i.e. THR is not order-based. □

Thus, the negative result does not apply to resource-allocation settings with this “linear-threshold” structure.

6 Extensions to the Order-Based Model

In this section, we propose some minimal and modular extensions to the OB model that can capture our counterexamples. We believe that the extensions should form the basis of future research in truthful characterizations.

6.1 Exist-order-based domains

Some of the examples have illustrated the fact that order-based domains can not handle the quantifier “ \exists ” but are appropriate for enforcing universal linear constraints via the quantifier “ \forall ”. We can extend the notion of order-based domains to include these counterexamples: we allow “there exists” statements and inequalities that use the outcomes quantified by the “ \exists ” statements. We aim the following notion as a minimal extension of order-based:

Definition 6 *A domain V_i is exist-order-based if it is defined by the a set \mathcal{C} of constraints of the following form:*

Unconstrained: $v_i(a) \{<, \leq, =\} v_i(b)$ or $v_i(a) = 0$ where $a, b \in A$ are fixed.

Constrained Fix $k \geq 1$, an arbitrary relation $\mathcal{R}(\overset{k+1}{\cdot}, \dots, \cdot)$ and an ordering relation \preceq in $\{<, \leq, >, \geq, =\}$. Then $\forall v_i \in V_i$ there exist reference outcomes $a^{(1)}, \dots, a^{(k)} \in A$ (that may depend on v_i) such that:

if $\mathcal{R}(a, a^{(1)}, \dots, a^{(k)})$ then

$$v_i(a) \preceq \{v_i(a^{(h)}), v_i(b)\} \text{ or} \\ v_i(a) = 0$$

where $a, b \in A$ are fixed and h is some index, $1 \leq h \leq k$. That is, all $v_i \in V_i$ and only them satisfy the constraints in \mathcal{C} .

The unconstrained (in)equalities hold unconditionally for any $v_i \in V_i$. The constrained (in)equalities allow the choice of the reference outcomes for each $v_i \in V_i$, allowing thus richer structure than order-based domains: in the definition of order-based domains only unconstrained (in)equalities are allowed. The relations and the outcomes a, b in the constrained inequalities are the same across all $v_i \in V_i$: it is only the reference outcomes that change from one v_i to another. We can now revisit some of our earlier example domains:

• **Linear threshold preferences.** We need the constraints for free-disposal, normalization and no externalities. We have two sets of reference outcomes: the $a^{(h)}$ ’s and $b^{(h)}$ ’s are the allocations where agent i gets $r - 1$ and r items, respectively. Then the constrained inequalities are:

- any allocation a that gives i exactly $r - 1$ items has value 0: if $\exists h$ s.t. $a = a^{(h)}$ then $v_i(a) = 0$.
- any allocation a that gives i exactly r items has value greater than 0: if $\exists h$ s.t. $a = b^{(h)}$ then $v_i(a) > v_i(\emptyset) = 0$
- free disposal extends the properties above to all allocations where agent i gets less than $r - 1$ or more than r items

• **Attribute-Based Preferences.** As an example, consider the class of arrival-departure preferences, in which an agent has positive value for an item received between time t_1 and time t_2 . Given the exist-order-based model, we can define two

reference outcomes: $t^{(1)}$ and $t^{(2)}$, representing the agent’s arrival and departure time. Recall that we identify an outcome with the time t the agent is served. The constrained inequalities are:

- If $t^{(1)} \leq t \leq t^{(2)}$ then $v_i(t) > 0$.
- If $t < t^{(1)}$ or $t > t^{(2)}$ then $v_i(t) = 0$.

• **Single Minded Bidders.** We need all CA unconstrained inequalities. We only have one reference outcome: some allocation a^{S^*} where agent i gets its magic bundle S^* . The constrained inequalities are:

- any allocation a that gives agent i a superset of S^* has value equal to s , the value of S^* to the agent: if $a_i^{S^*} = S^* \subseteq a_i$ then $v_i(a) = s$
- any allocation a that does not give agent i a superset of S^* has value 0: if $a_i^{S^*} = S^* \not\subseteq a_i$ then $v_i(a) = 0$

In a similar fashion it can be shown that attribute-based preferences or bounded-XOR preferences are exist-order-based.

6.2 General Bidding Languages

The exist-order-based extension does not help to express “value” related features of structured domains, such as those in *complement-free* domains and with *OR* preferences.

For complement-free, consider the following extension of order-based:

Definition 7 A domain is multi-order-based if it can be defined by a set \mathcal{C} of constraints of the form:

- $v_i(a) = 0$
- $v_i(a^{(1)}) + \dots + v_i(a^{(r)}) \{<, \leq, =\} v_i(b^{(1)}) + \dots + v_i(b^{(s)})$ for some integers r, s

where the a ’s and the b ’s are alternatives in A that do not depend on v_i . That is, all $v_i \in V_i$ and only them satisfy the constraints in \mathcal{C} .

In the definition of order-based domains, we always have $r = s = 1$. The domain of complement-free preferences can be expressed by introducing inequalities of the form Eq. (6) for all bundles a_i and b_i , and using the fact that in CAs we have no externalities: $v_i(a) = v_i(a_i)$. It can be shown that attribute-based preferences (Subsection 4.5) are not multi-order-based.

In general, this pushes us in the direction of introducing features from bidding languages into the descriptive language for preference domains. For instance, in order to include OR preferences in the definition of order-based, one can add a *max_packing* operator to the syntax of order-based domains in a similar fashion to the method we used for the “there exists” quantifier.

This can be done for any other bidding language \mathcal{L} : if agents are restricted to having preferences expressible in \mathcal{L} , then we can extend order-based to contain that domain as well by adding \mathcal{L} semantics. This is what we did for complement-free preferences and preferences that rely on the “there exists” quantifier.

7 Conclusions

We have considered the role of structured preferences in the characterization of truthful mechanisms. Through many examples, we have suggested that the current (negative) characterizations may not carry forward to realistic domains in which there can be considerable structure on preferences. This motivates the need for continued research in the characterization of truthful mechanisms. In this vein, we have suggested some minimal extensions to the order-based model due to Lavi et al. [2003], introducing new “there exists” semantics in addition to more expressiveness to constrain the relative values on alternatives. In closing, we reemphasize our view that it will be important to seek sufficient conditions on truthfulness that also provide for the existence of natural payments and for methods with universal applicability.

Acknowledgments

This work is supported in part by NSF grant IIS-0238147.

References

- [Gui et al., 2004] Hongwei Gui, Rudolf Muller, and Rakesh Vohra. Dominant strategy mechanisms with multidimensional types. 2004.
- [Lavi and Nisan, 2000] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, 2000.
- [Lavi et al., 2003] Ron Lavi, Ahuva Mu’alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of FOCS ’03*, 2003.
- [Lehmann et al., 2001] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *EC ’01: Proceedings of the 3rd ACM conference on Electronic Commerce*, 2001.
- [Lehmann et al., 2002] Daniel Lehmann, Liadan Ita O’Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, September 2002.
- [Roberts, 1979] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences*, pages 321–348. North-Holland, Amsterdam, 1979.
- [Rochet, 1987] Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16:191–200, 1987. North-Holland.
- [Saks and Yu, 2005] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *ACM EC*, 2005.
- [Yokoo, 2003] Makoto Yokoo. The characterization of strategy/false-name proof combinatorial auctions protocols: price-oriented, rationing-free protocol. In *Proceedings of 18th IJCAI*, 2003.