Characterizing truthful mechanisms with convex type spaces

Aaron Archer
AT&T Labs – Research, 180 Park Avenue, Florham Park, NJ 07932.
and
Robert Kleinberg
Computer Science Department, Cornell University, Ithaca, NY 14853.

Mechanism design studies the implementation of *allocation functions* (sometimes called *social choice functions*) when relevant information resides with self-interested agents who may misreport their information if it is rational to do so. A central question in the field is to characterize which allocation functions are *truthful*, meaning that they can be combined with a payment function that induces players to report their preferences truthfully. Characterization theorems for truthful mechanisms, when they exist, are a boon to the mechanism designer since they allow us to reduce problems of optimal mechanism design to algorithm design problems in which the algorithm that computes the allocation function is required to satisfy additional constraints.

This article may be viewed as an annotated discussion of the results in [Archer and Kleinberg 2008].

1. PRELIMINARIES

In this article we focus on mechanism design problems with a single player, also known as principal-agent problems. In some sense, this restriction is without loss of generality: fixing all other agents' bids, agent i faces a single-player mechanism, and a multi-player mechanism is truthful if and only if each of the single-player mechanisms in this ensemble is truthful. This reduction from multi-player mechanisms to single-player mechanisms applies across a range of different solution concepts — including dominant-strategy, Nash, and Bayes-Nash equilibrium. See [Archer and Kleinberg 2008] for details.[1]

Abstractly, the environment of a single-player mechanism design problem can be de-

[1] The statement that multi-player mechanism design reduces to single-player mechanism design overlooks the fact that designing an n-player social choice function whose 1-player restrictions are simultaneously truthful is much more difficult than designing an n-tuple of truthful single-player mechanisms. Also, the reduction for Bayes-Nash equilibrium is somewhat subtle, and works only if the players' *types* (defined shortly) are independently distributed.

Authors' addresses: aarcher@research.att.com, rdk@cs.cornell.edu. Robert Kleinberg is supported in part by NSF CAREER award CCF-0643934 and NSF award CCF-0729102.

scribed by specifying two things: a set $\mathcal{O}$ of *outcomes*, and a set $\mathcal{T}$ of *player types*. Each type $\mathbf{x} \in \mathcal{T}$ is a mapping from outcomes to real numbers; thus $\mathcal{T}$ is a subset of the function space $\mathbb{R}^{\mathcal{O}}$. If $\mathbf{x}$ is a player's type and $\mathbf{a}$ is an outcome, then $\mathbf{x}(\mathbf{a})$ is interpreted as the player's valuation for outcome $\mathbf{a}$. We do not assume that $\mathcal{O}$ is finite, nor do we assume that $\mathcal{T}$ is contained in a finite-dimensional subspace of the vector space $\mathbb{R}^{\mathcal{O}}$. However, we *do* assume that $\mathcal{T}$ is a convex subset of $\mathbb{R}^{\mathcal{O}}$. In other words, for any two types $\mathbf{x}, \mathbf{y} \in \mathcal{T}$ and any parameter $\lambda \in [0, 1]$, the function $\mathbf{a} \mapsto \lambda\mathbf{x}(\mathbf{a}) + (1 - \lambda)\mathbf{y}(\mathbf{a})$ is also an element of $\mathcal{T}$.

Example 1 (Finite outcome set)**.** For mechanism design problems with a finite number of outcomes, n, the type space $\mathcal{T}$ is simply a subset of $\mathbb{R}^n$ and a type $\mathbf{x} \in \mathcal{T}$ is simply a vector in $\mathbb{R}^n$ whose components represent the player's valuation for each of the n outcomes. The definition of convex type spaces given above reduces to the standard definition of convex subsets of $\mathbb{R}^n$.

Example 2 (Dot-product valuations)**.** Suppose that $\mathcal{O}$ is a subset of $\mathbb{R}^n$ and that for every type $\mathbf{x} \in \mathcal{T}$ there is a corresponding vector $\mathbf{v}_{\mathbf{x}} \in \mathbb{R}^n$ such that $\mathbf{x}(\mathbf{a}) = \mathbf{v}_{\mathbf{x}} \cdot \mathbf{a}$ for all $\mathbf{a} \in \mathcal{O}$. We express this situation by saying that the player has "dot-product valuations" and we will ignore the distinction between $\mathbf{x}$ and $\mathbf{v}_{\mathbf{x}}$. Thus, instead of considering $\mathcal{T}$ to be a subset of $\mathbb{R}^{\mathcal{O}}$, we will consider it to be a subset of $\mathbb{R}^n$. Once again, the definition of convex type spaces given above reduces to the standard definition of convex subsets of $\mathbb{R}^n$. Note that Example 1 is a special case of a dot-product valution, in which $\mathcal{O}$ is represented by the set of standard basis vectors $\{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n\}$.

Here are two settings where dot-product valuations are natural. In the first setting, the components of the outcome vector $\mathbf{a}$ represent amounts of n different goods allocated to the player, the components of the type vector $\mathbf{x}$ represent the player's value per unit of each of these goods, and these values are additive. In the second setting, there are n "pure" outcomes (which may or may not be in $\mathcal{O}$), each outcome vector $\mathbf{a} \in \mathcal{O}$ represents a lottery over the pure outcomes, and the player's value for a lottery is its expected value for the random pure outcome chosen.

Actually, dot-product valuations are more general than at first they seem. Whenever $\mathcal{T}$ is contained inside a finite-dimensional linear subspace of $\mathbb{R}^{\mathcal{O}}$, we may choose a set of types $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ that form a basis for this subspace, and embed $\mathcal{O}$ into $\mathbb{R}^n$ by mapping outcome $\mathbf{a}$ to the vector $(\mathbf{x}_1(\mathbf{a}), \mathbf{x}_2(\mathbf{a}), \ldots, \mathbf{x}_n(\mathbf{a}))$. If we treat $\mathcal{O}$ as a subset of $\mathbb{R}^n$ in this way, the reader can now check that every $\mathbf{x} \in \mathcal{T}$ defines a linear function $\mathcal{O} \xrightarrow{\mathbf{x}} \mathbb{R}$, so it is represented by a vector $\mathbf{v}_{\mathbf{x}}$ as in the first paragraph of this example. Thus, every mechanism design problem with a finite-dimensional type space can be expressed in the form of dot product valuations.

Example 3. In many natural mechanism design problems, the type space cannot be modeled as a subset of a finite-dimensional vector space. For example, suppose we are designing a mechanism that assigns to the player some quantity of a divisible good. The set of outcomes $\mathcal{O}$ is thus $\mathbb{R}_+$, and the set of types $\mathcal{T}$ might be the set of all monotonically increasing functions from $\mathbb{R}_+$ to $\mathbb{R}_+$. This type space is convex, because a convex combination of increasing functions is increasing. Figure 1 illustrates two types in $\mathcal{T}$, and a convex combination of them.

We focus on *direct revelation* mechanisms, which means that the action available to the player is to announce a *bid* to the mechanism, which is some type in $\mathcal{T}$ but may or may

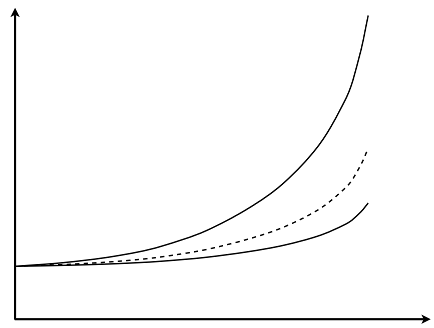

Fig. 1. Two types, represented as functions from positive reals to positive reals, and a convex combination of the two.

not be its true type.[2] Continuing with definitions, let us define an *allocation function* to be a function $f : \mathcal{T} \to \mathcal{O}$ and a *payment function* to be a function $p : \mathcal{T} \to \mathbb{R}$. We interpret $f(\mathbf{x})$ to be the outcome chosen when the player bids $\mathbf{x}$, while $p(\mathbf{x})$ is the payment made to the player in this case. We assume the player has a quasi-linear utility function, meaning that a player with type $\mathbf{x}$ values the pair $(\mathbf{a}, q)$ at $\mathbf{x}(\mathbf{a}) + q$.

A *mechanism* is an ordered pair (f, p) consisting of an allocation function f and payment function p. We say that (f, p) is *truthful* if the player always maximizes utility by bidding its true type; in other words, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{T}$ we have

$$\mathbf{x}(f(\mathbf{x})) + p(\mathbf{x}) \geq \mathbf{x}(f(\mathbf{x}')) + p(\mathbf{x}'). \tag{1}$$

We say that f is a truthful allocation if there exists a payment function p such that (f, p) is a truthful mechanism.

Characterization theorems for truthful mechanisms supply criteria for determining whether f is truthful without having to explicitly search for the payment function p that makes (f, p) into a truthful mechanism. The most general such characterization theorem, due to Rochet [1987], proves that f is truthful if and only if it satisfies the following *cycle monotonicity* property (CMON) for all $k > 0$:

$$\forall \mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathcal{T} \quad \sum_{i=1}^{k} \mathbf{x}_i(f(\mathbf{x}_i)) \geq \sum_{i=1}^{k} \mathbf{x}_{i-1}(f(\mathbf{x}_i)), \tag{2}$$

where the subscripts are interpreted mod k. With dot product valuations, this may be rewritten as:

$$\forall \mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathcal{T} \quad \sum_{i=1}^{k} f(\mathbf{x}_i) \cdot (\mathbf{x}_i - \mathbf{x}_{i-1}) \geq 0. \tag{3}$$

CMON is a simple and general criterion that is necessary and sufficient for truthfulness even when $\mathcal{T}$ is infinite-dimensional and non-convex. However, it is surprisingly hard to

[2]The famous *revelation principle* tells us that this restriction is without loss of generality.

apply this criterion to design truthful mechanisms. In fact, in the literature on algorithmic mechanism design we are aware of only one example [Lavi and Swamy 2007] in which a mechanism was successfully designed by directly enforcing the CMON iequalities.

The special case of CMON when $k = 2$ is called *weak monotonicity* (WMON) and is expressed by the following inequality for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{T}$:

$$\mathbf{x}_1(f(\mathbf{x}_1)) + \mathbf{x}_2(f(\mathbf{x}_2)) \geq \mathbf{x}_1(f(\mathbf{x}_2)) + \mathbf{x}_2(f(\mathbf{x}_1)). \quad (4)$$

With dot product valuations, this may be rewritten as:

$$(f(\mathbf{x}_1) - f(\mathbf{x}_2)) \cdot (\mathbf{x}_1 - \mathbf{x}_2) \geq 0. \quad (5)$$

It is known that WMON is a necessary and sufficient condition for truthfulness when the type space is one-dimensional [Mirrlees 1971; Spence 1974; Myerson 1981; Rochet 1987; Archer and Tardos 2001]. Recently Saks and Yu proved that it is necessary and sufficient when the type space is multidimensional but the outcome set is finite [Saks and Yu 2005], a result that was also implicit in work of Jehiel et al. [1999]. However, it is known that WMON is not a sufficient condition for truthfulness when the set of outcomes is infinite, even if the type space is only two-dimensional; see [Saks and Yu 2005] for a counterexample. Thus, characterizing truthfulness in two-dimensional type spaces is inherently more complicated than characterizing truthfulness in one dimension. One of the findings of our work can be summarized as follows: although truthfulness is more complicated in two dimensions than in one, it is not significantly more complicated in infinitely many dimensions than in two.

One of the reasons that the CMON characterization of truthfulness is difficult to apply is that it is not *local*, i.e., it depends on comparing the values of f at points in the type space which may be arbitrarily far from each other. Since it is often much easier to compare an algorithm's behavior on one input to its behavior on a small perturbation of that input, it is desirable to identify a characterization of truthfulness that is *local*, i.e., depends only on the function's behavior in small neighborhoods. One of the findings of our work is that truthfulness can be characterized by local properties: formally, if every point $\mathbf{x}$ in a convex type space has an open neighborhood U such that the restriction $f|_U$ is truthful, then f is truthful.

2. OUR CHARACTERIZATION THEOREM

Our characterization theorem provides two criteria, *local weak monotonicity* and *vortex-freeness*, which together are necessary and sufficient for truthfulness. In this section, we define the two criteria and state the main theorem.

Definition 1. An allocation function f satisfies *local weak monotonicity (local WMON)* if and only if every point $\mathbf{x}_1 \in \mathcal{T}$ has an open neighborhood U such that (4) holds for all $\mathbf{x}_2 \in U$.

For convex type spaces it turns out that local WMON is equivalent to WMON, although this equivalence does not hold in general, even for simply-connected non-convex type spaces.

The definition of vortex-freeness stems from the interpretation of the left side of (3) as a Riemann sum approximating the line integral $\oint f(\mathbf{x}) \cdot d\mathbf{x}$ of the vector field $f(\mathbf{x})$ around a loop in type space that passes through the points $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k$ and finally returns to $\mathbf{x}_1$.

Thus, CMON implies that any loop integral $\oint f(\mathbf{x}) \cdot d\mathbf{x}$ must be non-negative.[3] Since we must still get a non-negative integral when the orientation of the loop is reversed, we see that CMON actually implies that all loop integrals of f are zero. In the finite-dimensional case with dot-product valuations, the vanishing of loop integrals of f is expressed by saying that f is a *conservative vector field*, a criterion whose relation to truthfulness was noted previously in [Jehiel et al. 1999; Müller et al. 2007]. The definition of vortex-freeness captures the simplest non-trivial special case of this necessary condition.

Definition 2. An allocation function $f : \mathcal{T} \to \mathcal{O}$ is *vortex-free* if for every $\mathbf{x} \in \mathcal{T}$ and every 2-dimensional plane Π through $\mathbf{x}$, there exists an open neighborhood U about $\mathbf{x}$ such that the path integral $\oint_{\triangle} f(\mathbf{x}) \cdot d\mathbf{x}$ vanishes for every triangle $\triangle$ in $\Pi \cap U$ with one corner at $\mathbf{x}$.

Having defined local WMON and vortex-freeness, we can now state our main characterization theorem.

Theorem 1. *Let $\mathcal{T}$ be a convex type space and let $f : \mathcal{T} \to \mathcal{O}$ be an allocation function. Then f is truthful if and only if it is vortex-free and satisfies local WMON.*

It is instructive to consider the meanings of local WMON and vortex-freeness when we have finite-dimensional type and outcome spaces with dot product valuations, and the allocation function f is continuously differentiable. In that case, f is vortex-free if and only if it is the gradient of some function[4] $\Phi : \mathcal{T} \to \mathbb{R}$, and f satisfies local WMON if and only if this function Φ is convex. Moreover, both criteria have nice interpretations in terms of the matrix Df of derivatives of f: f is vortex-free if and only if Df is symmetric, and f satisfies local $WMON$ if and only if Df is positive semidefinite. These observations about truthfulness of smooth functions on finite-dimensional convex type spaces are essentially contained in [McAfee and McMillan 1988], though not in the form presented here.

We now give a brief sketch of the proof of Theorem 1, referring the interested reader to [Archer and Kleinberg 2008] for details. By Rochet's Theorem, truthfulness is equivalent to CMON. We have already seen that CMON implies that f is vortex-free and satisfies local WMON. It remains to prove the converse. Let $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k$ be any k-tuple of types in $\mathcal{T}$, and let P be the polygonal path consisting of line segments from $\mathbf{x}_i$ to $\mathbf{x}_{i+1}$ for $i = 1, 2, \ldots, k$, once again interpreting subscripts mod k. (Here we are using the assumption that $\mathcal{T}$ is convex; otherwise P need not be contained in $\mathcal{T}$.) To establish CMON one proves that

$$\sum_{i=1}^{k} \mathbf{x}_i(f(\mathbf{x}_i)) - \mathbf{x}_{i-1}(f(\mathbf{x}_i)) \geq \oint_P f(\mathbf{x}) \cdot d\mathbf{x} = 0. \tag{6}$$

The fact that the sum is greater than or equal to the integral in (6) is a consequence of local WMON. The proof that the integral equals zero uses only the fact that f is vortex-free. First

[3] Actually, in order for the loop integral to even be well-defined, we must assume that f is *locally path integrable*, meaning that every point of $\mathcal{T}$ has an open neighborhood in which the integral of f over every line segment is well-defined and finite. Note that the integral of f over a line segment Γ from $\mathbf{x}_0$ to $\mathbf{x}_1$ is well-defined, even when $\mathcal{T}$ is infinite-dimensional, using the formula $\int_\Gamma f(\mathbf{x}) \cdot d\mathbf{x} \triangleq \int_0^1 (\mathbf{x}_1(f(\mathbf{x}_t)) - \mathbf{x}_0(f(\mathbf{x}_t)))\, dt$, where $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$. If f satisfies WMON, then this last integrand is an increasing function of t on $[0, 1]$, so the integral exists.

[4] Here we use the assumption that $\mathcal{T}$ is convex, hence simply-connected.

one proves that $\oint_P f(\mathbf{x})d\mathbf{x}$ vanishes whenever P is a triangle, using an argument based on subdividing P into a finite number of sufficiently small triangles and using topological compactness to apply the vortex-free property to each of these smaller triangles. Then one deals with arbitrary polygons P by subdividing them into triangles. This completes the proof of (6), and with it the proof of Theorem 1.

2.1 Corollaries

Our characterization theorem has two easy corollaries that deserve mention.

Corollary 2 (**Local-to-global truthfulness).** *If $\mathcal{T}$ is convex and every $\mathbf{x} \in \mathcal{T}$ has an open neighborhood U such that the restriction to f to U is truthful, then f is truthful.*

Proof: The definitions of local WMON and vortex-free depend only on the restrictions of f to sufficiently small neighborhoods of every point in $\mathcal{T}$. ■

Corollary 3 (**2-dimensional to infinite-dimensional truthfulness).** *If $\mathcal{T}$ is convex and the restriction of f to $\Pi \cap \mathcal{T}$ is truthful for every 2-dimensional affine subspace Π, then f is truthful.*

Proof: The definitions of local WMON and vortex-free depend only on the restrictions of f to sets of the form $\Pi \cap \mathcal{T}$ where Π is an affine subspace of dimension at most 2. (In the case of local WMON, in fact, it suffices to consider 1-dimensional affine subspaces Π.) ■

Note that WMON is equivalent to the assertion that the restriction of f to $\Pi \cap \mathcal{T}$ is truthful for every 1-dimensional affine subspace Π. Thus, the criterion in Corollary 3 can be interpreted as a 2-dimensional generalization of WMON. A special case of Corollary 3 was obtained independently by Vohra [2007].

Our characterization theorem also yields an easy proof of the theorem of Saks and Yu [2005] that WMON suffices for truthfulness on convex domains when the outcome set is finite. Viewing the Saks-Yu theorem as a criterion for truthfulness of piecewise-constant functions, we can generalize it to a criterion for truthfulness of piecewise-truthful functions. Our characterization theorem implies that in order to verify the truthfulness of a function f obtained by "stitching together" truthful functions on subsets of $\mathcal{T}$, it suffices to check that f satisfies local WMON on the boundaries between pieces. The precise statement of this "truthful stitching" theorem is given in Theorem 4 below. A very similar result was established by different means in [Jehiel et al. 1999].

Theorem 4. *Suppose that a finite-dimensional convex type space $\mathcal{T}$ is covered by closed sets $\{\mathcal{T}_i : i \in \mathcal{I}\}$ such that:*
(1) *the covering is locally finite;*
(2) *each set $\mathcal{T}_i$ is the closure of its interior;*
(3) *the pairwise intersections $\mathcal{T}_i \cap \mathcal{T}_j$ are piecewise differentiable and have positive codimension in $\mathcal{T}$.*

Suppose that f is a function on $\mathcal{T}$, and that for each $i \in \mathcal{I}$, we have a locally truthful function f_i on $\mathcal{T}_i$, continuous at each point of the boundary $\partial\mathcal{T}_i$, such that $f = f_i$ on the interior of $\mathcal{T}_i$. If f satisfies local WMON, then f is truthful.

Truthful stitching represents a promising tool for designing algorithms to compute truthful allocation functions, since functions computed by a program containing branch points (e.g., if-then-else statements) tend to have the property that they partition the type space into a finite number of pieces (according to the branch outcomes), and on each piece the function has a simple description.

2.2 Non-convex type spaces

Up to this point in the article, we have dealt exclusively with convex type spaces. Indeed, most of the characterization results stated above fail when the type space is non-convex. For example, Corollary 2 fails to hold in the non-convex setting: a locally truthful function is not necessarily truthful, even if the type space is simply connected and two-dimensional. Nevertheless, convex type spaces are in some sense the canonical ones, by the following theorem. This result requires a technical condition that we call *outcome compactness*; see [Archer and Kleinberg 2008] for details.

Theorem 5. *Let $\mathcal{T}$ be any type space, let $\mathcal{T}^\sharp$ denote the convex hull of $\mathcal{T}$, and let $\mathcal{O}$ be an outcome space that satisfies outcome compactness. An allocation function $f : \mathcal{T} \to \mathcal{O}$ is truthful if and only if there exists a truthful allocation function $f^\sharp : \mathcal{T}^\sharp \to \mathcal{O}$ such that f is the restriction of $f^\sharp$ to $\mathcal{T}$.*

The idea behind the proof is that f implicitly defines a payment for each outcome, so we can extend f to $\mathcal{T}^\sharp$ by defining $f^\sharp(\mathbf{x})$ to be the outcome that maximizes the total utility for type $\mathbf{x}$ given this payment. The outcome compactness condition is there to guarantee that there is some type in $\mathcal{T}$ that achieves this maximum.

3. FUTURE DIRECTIONS

One of the focal optimization problems studied in algorithmic mechanism design has been the problem of scheduling jobs on unrelated machines to minimize makespan, often denoted $R||C_{\max}$ in the scheduling literature. In this problem, there are n jobs that we need to schedule on m machines. Each machine has a different processing time that it would require to finish each job, and the mechanism must decide how to split up the jobs amongst the machines, with the goal of minimizing the time at which the last job is completed. The players are the machines, and a machine's type is simply its vector of processing times. We assume that a machine's valuation for any collection of jobs equals the negative of its total processing time for those jobs.

This problem was first proposed in the mechanism design context by Nisan and Ronen [2001]. They note that assigning each job to the machine that can process it the fastest is a truthful allocation, but the makespan it achieves can exceed that of the optimal solution by as much as a factor of m. A big question has been whether this approximation factor can be improved by a better truthful mechanism, or whether a matching lower bound can be obtained. The lower bound of 2 given by Nisan and Ronen [2001] has been improved to 2.414 by Christodoulou et al. [2007], then to 2.618 by Koutsoupias and Vidali [2007]. For the version of the problem where each job can be split fractionally amongst machines, Christodoulou et al. [2007] give a lower bound of $2 - \frac{1}{m}$ and an upper bound of $\frac{m+1}{2}$. Thus, for both versions of the problem, there is still an enormous gap between the upper and lower bounds. For deterministic mechanisms in the unsplittable case, there are only a finite number of outcomes, so the characterization theorem from [Saks and Yu 2005] is the relevant one. However, for randomized mechanisms and for the fractional case our Theorem 1 might be of use either in designing better truthful mechanisms or in proving stronger lower bounds.

It is well-known that any allocation function that maximizes an affine combination of the players' valuations over any fixed subset of $\mathcal{O}$ is truthful. For single-player mechanisms, the converse is obvious, because the truthful allocation f is an affine maximizer where the additive offset term is just the payment function (using the fact that the payment function is constant on $f^{-1}(\mathbf{a})$, for each $\mathbf{a} \in \mathcal{O}$). For multi-player mechanisms, the converse does not generally hold. However, a remarkable theorem by Roberts [1979] shows that for mechanism design problems where $3 \leq |\mathcal{O}| < \infty$ and the type space is unrestricted (i.e., $\mathcal{T} = \mathbb{R}^{\mathcal{O}}$), the converse does hold; that is, every truthful allocation function is an affine maximizer. For combinational auctions, under some additional restrictions, Lavi et al. [2003] proved a similar theorem. It would be good to understand just how generally this converse holds. For instance, perhaps it is the case that whenever the type space is a linear subspace of $\mathbb{R}^{\mathcal{O}}$ of sufficiently high dimension, the only truthful allocation functions are the affine maximizers.

In the context of mechanisms for maximizing the utilitarian objective function in 1-dimensional type spaces, Mu'alem and Nisan [2002] showed that under certain circumstances, one can apply max operators and if-then-else constructs to combine multiple truthful mechanisms into a single, improved truthful mechanism. This would be an excellent tool to be able to employ more generally. We hope that our truthful stitching theorem will be useful in creating such a tool for multi-dimentional type spaces.

REFERENCES

ARCHER, A. AND KLEINBERG, R. 2008. Truthful germs are contagious: A local-to-global characterization of truthfulness. In *Proceedings of the* 10^{th} *ACM Conference on Electronic Commerce*.

ARCHER, A. AND TARDOS, É. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings of the* 42^{nd} *IEEE Symposium on Foundations of Computer Science*. 482–491.

CHRISTODOULOU, G., KOUTSOUPIAS, E., AND KOVÁCS, A. 2007. Mechanism design for fractional scheduling on unrelated machines. In *Proceedings of the* 34^{th} *International Colloquium on Automata, Languages and Programming*. 40–52.

CHRISTODOULOU, G., KOUTSOUPIAS, E., AND VIDALI, A. 2007. A lower bound for scheduling mechanisms. In *Proceedings of the* 18^{th} *Annual ACM-SIAM Symposium on Discrete Algorithms*. 1163–1170.

JEHIEL, P., MOLDOVANU, B., AND STACCHETTI, E. 1999. Multidimensional mechanism design for auctions with externalities. *J. Economic Theory 85*, 258–293.

KOUTSOUPIAS, E. AND VIDALI, A. 2007. A lower bound of $1 + \phi$ for truthful scheduling mechanisms. In *32nd International Symposium on Mathematical Foundations of Computer Science*. 454–464.

LAVI, R., MU'ALEM, A., AND NISAN, N. 2003. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the* 44^{th} *Annual IEEE Symposium on Foundations of Computer Science*. 574–583.

LAVI, R. AND SWAMY, C. 2007. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In *Proceedings of the* 9^{th} *ACM Conference on Electronic Commerce*. 252–261.

MCAFEE, R. P. AND MCMILLAN, J. 1988. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory 46*, 335–354.

MIRRLEES, J. 1971. An exploration in the theory of optimum income taxation. *Review of Economic Studies 38*, 175–208.

MU'ALEM, A. AND NISAN, N. 2002. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings of the* 18^{th} *National Conference on Artificial intelligence (AAAI)*. 379–384.

MÜLLER, R., PEREA, A., AND WOLF, S. 2007. A network approach to bayes-nash incentive compatible mechanisms. *Games and Economic Behavior 61,* 2 (November), 344–358.

MYERSON, R. B. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1 (February), 58–73.

NISAN, N. AND RONEN, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior 35*, 166–196.

ROBERTS, K. 1979. The characterization of implementable choice rules. In *Aggregation and Revelation of Preferences*, J.-J. Laffont, Ed. North-Holland, Amsterdam, 321–348.

ROCHET, J.-C. 1987. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics 16,* 2 (April), 191–200.

SAKS, M. E. AND YU, L. 2005. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the* 7^{th} *ACM Conference on Electronic Commerce*. 286–293.

SPENCE, M. 1974. Competitive and optimal responses to signals. *Journal of Economic Theory 7*, 196–232.

VOHRA, R. 2007. Paths, cycles and mechanism design. Working paper, updated 2008.