

A Theory of Expressiveness in Mechanisms

Michael Benisch

School of Computer Science
Carnegie Mellon University

Norman Sadeh

School of Computer Science
Carnegie Mellon University

Tuomas Sandholm

Computer Science Department
Carnegie Mellon University

Abstract

A key trend in (electronic) commerce is a demand for higher levels of expressiveness in the mechanisms that mediate interactions. We develop a theory that ties the expressiveness of mechanisms to their efficiency in a domain-independent manner. We introduce two new expressiveness measures, 1) *maximum impact dimension*, which captures the number of ways that an agent can impact the outcome, and 2) *shatterable outcome dimension*, which is based on the concept of *shattering* from computational learning theory. We derive an upper bound on the expected efficiency of any mechanism under its most efficient Nash equilibrium. Remarkably, it depends only on the mechanism's expressiveness. We prove that the bound increases strictly as we allow more expressiveness. We also show that in some cases a small increase in expressiveness yields an arbitrarily large increase in the bound. Finally, we study *channel-based* mechanisms, which subsume most combinatorial auctions, multi-attribute mechanisms, and the Vickrey-Clarke-Groves scheme. We show that our domain-independent measures of expressiveness appropriately relate to the natural measure of expressiveness of channel-based mechanisms: the number of channels allowed. Using this bridge, our general results yield interesting implications. For example, any (channel-based) multi-item auction that does not allow rich combinatorial bids can be arbitrarily inefficient—unless agents have no private information.

Introduction

Mechanism design is the science of generating rules of interaction so that desirable outcomes result despite the participating agents (human or computational) acting based on rational self-interest. A *mechanism* takes as input some expressions of preference from the agents, and based on that information computes an *outcome* (such as an allocation of items and potentially also payments). By carefully crafting mechanisms, it is possible to design better auctions, exchanges, catalog offers, voting systems, and so on.

A recent trend in the world—especially in electronic commerce—is a demand for higher levels of expressiveness in the mechanisms that mediate interactions such as the allocation of resources, matching of peers, or elicitation of privacy and security preferences. This trend has already manifested itself in combinatorial auctions, multi-attribute auc-

tions, and generalizations thereof, which are used to trade tens of billions of dollars worth of items annually (Sandholm 2007; Cramton, Shoham, & Steinberg 2006). It is also reflected in the richness of preference expression offered by businesses as diverse as matchmaking sites, sites like Amazon and Netflix, and services like Google's AdSense. In Web 2.0 parlance, this demand for increasingly diverse offerings is called the Long Tail (Anderson 2006).

The most famous expressive mechanism is a *combinatorial auction (CA)*, which allows participants to express valuations over *packages* of items. CAs have the recognized benefit of removing the “exposure” problems that bidders face when they have preferences over packages but in traditional auctions are allowed to submit bids on individual items only. They also have other acknowledged benefits, and preference expression forms significantly more compact and more natural than package bidding have been developed. Expressiveness also plays a key role in *multi-attribute* settings where the participants can express preferences over vectors of attributes of the item—or, more generally, of the outcome. Some market designs are both combinatorial and multi-attribute (Sandholm & Suri 2006; Sandholm 2007; Cramton, Shoham, & Steinberg 2006).

Intuitively, one would think that more expressiveness would lead to higher efficiency (sum of the agents' utilities) of the mechanism's outcome (e.g., due to better matching of supply and demand). Efficiency improvements have indeed been reported from combinatorial and multi-attribute auctions (e.g., (Sandholm 2002; 2007; Cramton, Shoham, & Steinberg 2006)) and expressive auctions for banner advertisement allocation (Boutilier *et al.* 2008). However, we have lacked a general way of characterizing the expressiveness of different mechanisms, the impact that it has on the agents' strategies, and thereby ultimately the outcome. Until now, it was not even known whether, in any settings, more expressiveness can be used to design more efficient mechanisms.¹

In this paper, we develop a theory that ties the expressiveness of mechanisms to their efficiency in a domain-independent manner. We begin by introducing two new

¹In fact, on the contrary, it has been observed that in certain settings additional expressiveness can give rise to additional equilibria of poor efficiency (Milgrom 2007).

expressiveness measures, 1) *maximum impact dimension*, which captures the number of ways that an agent can impact the outcome, and 2) *shatterable outcome dimension*, which is based on the concept of *shattering* from computational learning theory.

Next, we derive an upper bound on the expected efficiency of any mechanism under its most efficient Nash equilibrium. We show that, remarkably, this bound depends only on the mechanism’s expressiveness. This allows us to sidestep two of the major roadblocks in analyzing the relationship between expressiveness and efficiency: 1) the bound can be studied without having to solve for any of the mechanism’s equilibria (which tends to be extremely difficult for inexpressive mechanisms (e.g., (Rosenthal & Wang 1996))), and 2) since it bounds the *most efficient* equilibrium it can be used to study mechanisms with multiple—or an infinite number of—equilibria, e.g., first price CAs (Bernheim & Whinston 1986). We show that in any setting the bound of an optimally designed mechanism increases *strictly* as more expressiveness is allowed, and in some settings the bound can increase arbitrarily via a small increase in expressiveness.

Finally, we study a class of mechanisms which we call *channel based*. They subsume most combinatorial auctions, multi-attribute mechanisms, and any *VCG* scheme (Vickrey 1961; Clarke 1971; Groves 1973). We show that our domain-independent measures of expressiveness appropriately relate to the natural measure of expressiveness of channel-based mechanisms: the number of channels allowed (which itself generalizes a classic measure of expressiveness in CAs called *k-wise dependence* (Conitzer, Sandholm, & Santi 2005)). Using this bridge, our general results yield interesting implications. For example, any (channel-based) multi-item auction that does not allow rich combinatorial bids can be arbitrarily inefficient—unless agents have no private information.

Preliminaries

The setting we study is that of standard mechanism design. In the model there are n agents. Each agent i has some private information (not known by the mechanism or any other agent) denoted by a type, t_i , (e.g., the value of the item to the agent in an auction; or, in a CA, a vector of values, potentially one value for each package of items) from the space of the agent’s possible types, T_i .

Settings where each agent has a utility function, $u_i(t_i, O)$, that depends only on its own type and the outcome, $O \in \mathcal{O}$, chosen by the mechanism (e.g., the allocation of items to agents in a CA) are called *private values* settings. We also discuss more general *interdependent values* settings, where $u_i = u_i(t^n, O)$, i.e., an agent’s utility depends on the others’ private signals. In both settings, agents report expressions to the mechanism, denoted θ_i , based only on their own types. A mapping from types to expressions is called a *pure strategy*.

Definition 1 (pure strategy). A pure strategy for an agent i is a mapping, $h_i : T_i \rightarrow \Theta_i$, that is, it selects an expression for each of i ’s types. A pure strategy profile is a list of pure strategies, one strategy per agent, i.e., $h_I \equiv$

$[h_1, h_2, \dots, h_{|I|}]$. For shorthand, we often refer to h_I as a mapping from types of the agents in I to an expression for each agent, $h_I(t_I) = [\theta_1, \theta_2, \dots, \theta_{|I|}]$.

Based on these expressions the mechanism computes the value of an outcome function, $f(\theta^n)$, which chooses an outcome from \mathcal{O} . The mechanism may also compute the value of a payment function, $\pi(\theta^n)$, which determines how much each agent must pay or get paid.²

For analysis purposes, we assume that the expression of each agent in a Nash equilibrium can be described by a function that takes as input its type, $m_i(t_i)$. We do not restrict these equilibrium reports to be deterministic pure strategies: we allow m_i to be a *mixed strategy*, i.e., a random variable specifying a probability distribution over possible reports. We also do not restrict our analysis to mechanisms with truthful equilibria (i.e., where agents are incentivised to report their true types in equilibrium).³

For convenience, we will let $W(t^n, o)$ denote the total social welfare of outcome o when agents have private types (or private signals) t^n , $W(t^n, o) = \sum_i u_i(t^n, o)$. Using this formalism we can describe the expected efficiency, $\mathcal{E}(f, \pi)$, of a mechanism (where expectation is taken over the types of the agents, and their randomized equilibrium expressions),

$$E[\mathcal{E}(f, \pi)] = \int_{t^n} P(t^n) \int_{\theta^n} P(m(t^n) = \theta^n) W(t^n, f(\theta^n)).$$

Characterizing mechanism expressiveness

The primary goal of this paper is to better understand the impact of making mechanisms more or less expressive. First we must come up with meaningful (and general) definitions of a mechanism’s expressiveness.

If we consider mechanisms that allow expressions from the set of multi-dimensional real numbers, such as CAs and combinatorial exchanges, one seemingly natural way of characterizing their expressiveness is the dimensionality of the expressions they allow (e.g., this is one difference between CAs and auctions that only allow per-item-bids). However, this notion does not adequately differentiate between expressive and inexpressive mechanisms.

Proposition 1. For any mechanism that allows multi-dimensional real-valued expressions, (i.e., $\Theta_i \subseteq \mathbb{R}^d$), there exists an equivalent mechanism that only allows the expression of one real value (i.e., $\Theta_i = \mathbb{R}$).⁴ (This follows immediately from Cantor (1890): being able to losslessly map between the spaces \mathbb{R}^d and \mathbb{R} .)

²In this paper we only study the mechanism’s outcome function. For our purposes this is basically without loss of generality as long as agents do not care about *each others’* payments.

³The *revelation principle* of mechanism design states that any outcome function that can be implemented by any mechanism under a non-truthful equilibrium can also be implemented by some mechanism under a truthful equilibrium. However, we do not restrict our analysis to mechanisms with truthful equilibria because in inexpressive mechanisms it can be impossible for agents to express their true types.

⁴A more detailed treatment of the work in this paper, including proof of all technical claims, is available as a technical report (Benisch, Sadeh, & Sandholm 2007).

Thus, it is not the number of real-valued questions that a mechanism can ask that truly characterizes expressiveness, it is how the answers are used!

Another natural way in which mechanisms can differ is in the granularity of their outcome spaces. For example, auction mechanisms that are restricted to allocating certain items together (e.g., blocks of neighboring frequency bands) have coarser outcome spaces than those which can allocate them individually to different agents. Some prior work addresses the impact of a mechanism’s *outcome space* on its efficiency. For example, it has been shown that in private values settings VCG mechanisms with finer-grained outcome spaces have more efficient dominant-strategy equilibria (Holzman *et al.* 2004; Nisan & Ronen 2007).

In contrast, we are interested in studying the impact of a mechanism’s *expressiveness* on its efficiency. We do this by comparing more versus less expressive mechanisms with the *same* outcome space (e.g., fully expressive CAs and multi-item auctions that allow bids on individual items only). In our approach the outcome space can be unrestricted or restricted; thus our results can be used in conjunction with those stating that larger outcome spaces beget greater efficiency.

Impact-based expressiveness

In order to properly differentiate between expressive and in-expressive mechanisms with the same outcome space, we propose to measure the extent to which an agent can impact the outcome that is chosen. We define an *impact vector* to capture the impact of a particular expression by an agent under the different possible types of the other agents. (Given a mechanism let the subscript $-i$ refer to to all the agents other than agent i .)

Definition 2 (impact vector). *An impact vector for agent i is a function, $g_i : T_{-i} \rightarrow \mathcal{O}$. To represent the function as a vector of outcomes, we order the joint types in T_{-i} from 1 to $|T_{-i}|$; then g_i can be represented as $[o_1, o_2, \dots, o_{|T_{-i}|}]$.*

We say that agent i can *express* an impact vector if there is some pure strategy profile of the other agents such that one of i ’s expressions causes each of the outcomes in the impact vector to occur.

Definition 3 (express). *Agent i can express an impact vector, g_i , if $\exists h_{-i}, \exists \theta_i, \forall t_{-i}, f(\theta_i, h_{-i}(t_{-i})) = g_i(t_{-i})$.*

We say that agent i can *distinguish* among a set of impact vectors if it can express each of them against the same pure strategy profile of the other agents by changing only its own expression.

Definition 4 (distinguish). *Agent i can distinguish between a set of impact vectors, G_i , if*

$$\exists h_{-i}, \forall g_i \in G_i, \exists \theta_i, \forall t_{-i}, f(\theta_i, h_{-i}(t_{-i})) = g_i(t_{-i}),$$

when this is the case, we write $D_i(G_i) = \top$.

Figure 1 illustrates how an agent can distinguish between two different impact vectors against a pure strategy profile of the other agents.

Intuitively, more expressive mechanisms allow agents to distinguish among larger sets of impact vectors. Our first expressiveness measure captures this intuition; it measures the




Figure 1: By choosing between two expressions, $\theta_i^{(1)}$ and $\theta_i^{(2)}$, agent i can distinguish between the impact vectors $[A, B]$ and $[C, D]$ (enclosed in rectangles). The other agents are playing the pure strategy profile $[\theta_{-i}^{(x)}, \theta_{-i}^{(y)}]$.

number of different impact vectors that an agent can distinguish among. Since this depends on what the others express, we measure the best case, where the others submit expressions that maximize the agent’s control. We call this the agent’s *maximum impact dimension*.

Definition 5 (maximum impact dimension). *Agent i has maximum impact dimension d_i if the largest set of impact vectors, G_i^* , that i can distinguish among has size d_i . Formally, $d_i = \max_{G_i} \{|G_i| \mid D_i(G_i) = \top\}$.*

Shattering-based expressiveness

We will now discuss a related measure of expressiveness, which we call *shatterable outcome dimension*. As we will show later, it has somewhat different uses than does the maximum impact dimension.

The shatterable outcome dimension is based on a notion called *shattering*, which we have adapted from the field of computational learning theory (Vapnik & Chervonenkis 1971; Blumer *et al.* 1989). Our adaptation captures an agent’s ability to distinguish among each of the $|\mathcal{O}'|^{|T_{-i}|}$ impact vectors that include only outcomes from a given set \mathcal{O}' .

Definition 6 (outcome shattering). *A mechanism allows agent i to shatter a set of outcomes, $\mathcal{O}' \subseteq \mathcal{O}$, if $D_i(G_i^{\mathcal{O}'})$, where $G_i^{\mathcal{O}'} = \{g_i \mid g_i = [o_1, o_2, \dots, o_{|T_{-i}|}], o_j \in \mathcal{O}'\}$.*

We also use a slightly weaker adaptation of shattering for analyzing the more restricted setting where agents have private values. It captures an agent’s ability to cause each of the $\binom{|\mathcal{O}'|+1}{2}$ unordered pairs of outcomes (with replacement) to be chosen for every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., which outcome happens for which type). We call this *semi-shattering*.

Definition 7 (outcome semi-shattering). *A mechanism allows agent i to semi-shatter a set of outcomes, \mathcal{O}' , if it can distinguish among a set of impact vectors, $G_i^{\mathcal{O}'}$, such that*

$$\forall \{x, y\} \mid x, y \in T_{-i} \wedge x \neq y, \forall o_1, o_2 \in \mathcal{O}', \exists g_i \in G_i^{\mathcal{O}'}, [g_i(x) = o_1 \wedge g_i(y) = o_2] \vee [g_i(x) = o_2 \wedge g_i(y) = o_1].$$

Our second measure of expressiveness is based on the size of the largest outcome space that an agent can shatter or semi-shatter. We call it the *(semi-)shatterable outcome dimension*.

Definition 8 ((semi-)shatterable outcome dimension). *Agent i has (semi-)shatterable outcome dimension k_i if the largest set of outcomes that i can (semi-)shatter has size k_i .*

The next two results illustrate the close relationship between the shatterable outcome dimension measures and the maximum impact dimension measure.

Proposition 2. *Increasing an agent's shatterable or semi-shatterable outcome dimension also increases its maximum impact dimension.*

Proposition 3. *In order to shatter k_i outcomes agent i must be able to distinguish among at least $|T_{-i}|^{k_i}$ impact vectors.*

The number of types that the other agents have can be thought of as a support-based measure of agent i 's uncertainty. Thus, the more uncertainty an agent has, the more expressiveness it needs to shatter a given set of outcomes.

Expressiveness and efficiency

We will now present an upper bound on the expected efficiency of a mechanism's *most efficient* equilibrium. We will also show that the upper bound for an optimally designed mechanism is tied directly to its expressiveness.

We derive the bound by making the optimistic assumption that the agents play strategies which, taken together, attempt to maximize social welfare. This allows us to avoid the difficulty involved in calculating equilibrium strategies. It also implies that we can restrict our analysis to *pure* strategies because a pure strategy always exists that achieves at least as much expected efficiency as any mixed strategy.

Proposition 4. *The following quantity, $E[\mathcal{E}(f)]^+$, is an upper bound on the expected efficiency of the most efficient equilibrium in any mechanism with outcome function f ,*

$$E[\mathcal{E}(f)]^+ = \max_{\hat{h}(\cdot)} \int_{t^n} P(t^n) W(t^n, f(\hat{h}(t^n))). \quad (1)$$

The maximum is taken over $\hat{h}(\cdot)$, a pure strategy profile that maps every joint type to an expression for each agent.⁵

To see how this bound is tied to our notions of expressiveness, consider calculating the bound from the fixed perspective of a particular agent i . Based on our assumption, the other agents will choose whatever pure strategies are best for maximizing the mechanism's expected efficiency. Thus, from agent i 's perspective, the maximization problem comes down to finding the set of expressible impact vectors that lead to the highest expected efficiency.

Conditions under which the bound is fully efficient

Observe that there is an impact vector for each of agent i 's types that represents the vector of efficient outcomes when it

⁵Recall that an agent's strategy can only depend on its own private type, even if its utility depends on the private signals of others.

is matched with each of the non-zero probability joint types of the other agents. We call a set that contains such vectors for each of i 's types a *fully efficient set*. Such a set must be distinguishable for the bound to reach full efficiency.

Definition 9 (fully efficient set). G_i^* is a fully efficient set if

$$\forall t_i, \exists g_i \in G_i^*, \forall \{t_{-i} \mid P(t_i, t_{-i}) > 0\},$$

$$W([t_i, t_{-i}], g_i(t_{-i})) = \max_{o \in \mathcal{O}} W([t_i, t_{-i}], o).$$

Proposition 5. *The upper bound, $E[\mathcal{E}(f)]^+$, reaches full expected efficiency iff each agent can distinguish among the impact vectors in at least one of its fully efficient sets.*

In full information settings, where upon learning its own type an agent knows the types of the other agents for sure, the agent is guaranteed to have a fully efficient set of size $\leq |\mathcal{O}|$.

Proposition 6. *Let G_i^* be agent i 's smallest fully efficient set, $(\forall t_i, \exists t_{-i} \mid P(t_i, t_{-i}) = 1) \Rightarrow |G_i^*| \leq |\mathcal{O}|$.*

This implies that in such settings an agent requires less expressiveness to bring the bound to full efficiency.

Corollary 1. *If agent i has full information then there exists an outcome function for which the upper bound reaches full efficiency while limiting i to maximum impact dimension $d_i \leq |\mathcal{O}|$.*

One important takeaway of this is that perfect information about the other agents' types basically does away with the need for expressiveness. Thus, in prior research that shows that in certain settings even quite inexpressive mechanisms yield full efficiency (e.g., (Abrams, Ghosh, & Vee 2007)), the assumption that the agents have no private information is essential.

The efficiency bound increases strictly with expressiveness

Our main result demonstrates that a mechanism designer can *strictly* increase the upper bound on expected efficiency by giving any agent more expressiveness (until the bound reaches full efficiency). The result applies to the outcome function that maximizes the bound subject to the constraint that agent i 's expressiveness be less than or equal to a particular level. The bound attained by such an outcome function also serves as an upper bound for any outcome function which allows that expressiveness level to i .

Theorem 1. *For any setting and any distribution over agent preferences, the upper bound on expected efficiency for the best outcome function limiting agent i to maximum impact dimension d_i increases strictly monotonically as d_i goes from 1 to d_i^* , where d_i^* is the size of agent i 's smallest fully efficient set.*

Proof intuition. The proof is by induction on d_i . Briefly, if agent i 's maximum impact dimension is less than d_i^* , then there is at least one impact vector in any of its fully efficient sets that it cannot express. Increasing i 's maximum impact dimension by one will allow it to express at least one additional fully efficient impact vector and thus strictly increase the efficiency bound. \square

From Proposition 2 we know that any increase in shatterable or semi-shatterable outcome dimension implies an increase in maximum impact dimension; thus Theorem 1 implies that strict monotonicity holds for these measures as well.

Inadequate expressiveness can lead to arbitrarily low efficiency in any setting

In addition to strict monotonicity, we find that in *any* setting there exist distributions over agent preferences under which a small increase in allowed expressiveness leads to an arbitrary improvement in the upper bound.

Theorem 2. *For any setting there exists a distribution over agent preferences such that the upper bound on expected efficiency for the best outcome function limiting agent i to*

- shatterable outcome dimension, $k_i < |\mathcal{O}|$, in an interdependent values setting, or
- semi-shatterable outcome dimension, $k_i < |\mathcal{O}|$, in a private values setting

is arbitrarily less than that of the best outcome function limiting agent i to (semi-)shatterable outcome dimension $k_i + 1$.

Proof intuition. Construct preference distributions for any setting that require (semi-)shattering $k_i + 1$ outcomes. In interdependent values settings this is trivial, since preferences can depend arbitrarily on the private types of others. In private values settings our construction ensures that at least one of agent i 's types makes every pair of outcomes arbitrarily more efficient than others, under each pair of the other agents' joint types. \square

This implies that a mechanism must allow every agent to shatter (in interdependent values settings) or semi-shatter (in private values settings) its entire outcome space in order to guarantee that it will not be arbitrarily inefficient under some preference distribution.

An application of our expressiveness theory

We will now instantiate our theory of expressiveness for an important class of mechanisms, which we call *channel based*. Channel-based mechanisms are defined by the following (a small example is also presented in Figure 2),

Definition 10 (channel-based mechanism). *Each outcome is assigned a set of channels potentially coming from a number of different agents (e.g., outcome A may be assigned channels x_1 and y_1 from Agent 1 and x_2 from Agent 2). Each agent, simultaneously with the other agents, reports real values on each of its channels to the mechanism. The mechanism chooses the outcome whose channels have the largest sum.*⁶

Many different mechanisms for trading goods, information, and services, such as CAs, multi-attribute mechanisms, and any VCG-based mechanism, can be cast as channel-based mechanisms. (This class is even more general than

⁶We assume that ties are broken consistently according to some strict ordering on the outcomes. This prevents an agent from using the mechanism's tie breaking behavior as artificial expressiveness.




Figure 2: *Channel-based representations of two auctions. The items auctioned are an apple (a) and an orange (o). The channels for each agent i are denoted x_i , y_i , and z_i . The possible allocations are A, B, C, and D. In each one, the items that agent 1 gets are in the first braces, and the items agent 2 gets are in the second braces.*

CAs because it can model settings where agents care about how the items that they do not win get allocated across the other agents.)

A natural measure of expressiveness in channel-based mechanisms is the number of channels allowed. In CAs, it is able to capture the difference between fully expressive CAs, multi-item auctions that allow bids on individual items only (Fig. 2), and an entire spectrum in between. In fact, it generalizes a classic measure of expressiveness in CAs called k -wise dependence (Conitzer, Sandholm, & Santi 2005).

First, we will demonstrate that our domain-independent expressiveness measures relate appropriately to the number of channels allowed in a channel-based mechanism.

Proposition 7. *For any agent i , its semi-shatterable outcome dimension, k_i , in the most expressive channel-based mechanism strictly increases (until $k_i = |\mathcal{O}|$) as the number of channels assigned to the agent increases.*

From Theorem 2 we know that this increase in expressiveness can lead to an arbitrary increase in our efficiency bound, even in private values settings. However, if an agent has full information it only needs a logarithmic number of channels to bring the bound to full efficiency. (This also happens to be the number of channels in any multi-item auction that allows item bids only.)

Proposition 8. *If agent i has full information about the other agents, in a channel-based mechanism it needs only $\lceil \log_2(|\mathcal{O}|) \rceil$ channels to shatter the entire outcome space.*

On the other hand, an agent with less than full information cannot fully shatter any set of two or more outcomes in a channel-based mechanism.

Proposition 9. *No channel-based mechanism allows any agent to shatter any set of two or more outcomes when the other agents have two or more types.*

Since channel-based mechanisms do not allow full shattering, our results from the previous section imply that in some interdependent values settings any channel-based mechanism, even one that emulates the VCG mechanism, will be arbitrarily inefficient. (That such mechanisms cannot always get full efficiency in interdependent values settings is already known (Jehiel & Moldovanu 2001).)

However, full efficiency can be achieved in any private values setting—despite agent uncertainty—by a channel-based mechanism with $|\mathcal{O}| - 1$ channels per agent that emulates the VCG mechanism.

Proposition 10. *A channel-based mechanism can emulate the VCG mechanism iff it provides each agent with at least $|\mathcal{O}| - 1$ channels.*

Our next result deals with a configuration of channels that prevents an agent from being able to even *semi*-shatter a set that contains two particular pairs of outcomes.

Theorem 3. *Consider a set of outcomes, $\{A, B, C, D\}$, connected to different sets of channels for agent i , $\{S_i^A, S_i^B, S_i^C, S_i^D\}$, respectively. Agent i cannot semi-shatter both pairs of outcomes $\{A, B\}$ and $\{C, D\}$ if,*

$$(S_i^A \setminus S_i^C = S_i^D \setminus S_i^B) \quad \text{and} \quad (S_i^C \setminus S_i^A = S_i^B \setminus S_i^D).$$

Proof intuition. Let the sum of the reported channels under the first (second) profile for the other agents connected to outcome A be a_1 (a_2), to outcome B be b_1 (b_2), and so on. The assumption that agent i can semi-shatter both pairs of outcomes leads to the following contradiction,

$$(b_2 - d_2 < c_1 - a_1 < b_2 - d_2) \vee (a_1 - d_1 < c_2 - b_2 < a_1 - d_1).$$

□

This channel configuration generalizes one that appears in any channel-based multi-item auction that considers some agent’s bid for a bundle to be the sum of its bids on two other (non-overlapping) bundles. This fact, along with our previous results, implies that if any agent has less than full information, such auctions can be arbitrarily inefficient.

Corollary 2. *Any multi-item auction that*

- *can be represented as a channel-based mechanism, and*
 - *treats some agent i ’s bid for some bundle b to be the sum of i ’s bids on some other two (non-overlapping) bundles,*
- can be (i.e., for some prior) arbitrarily inefficient in any setting (even a private values setting).*

Other related work, briefly

Mount and Reiter (1974) and Hurwicz (1972) asked the question: how many real-valued dimensions must a mechanism’s message space have in order to accomplish some design goal? To get around Cantor’s theorem that begets our Proposition 1, they relied on certain technical assumptions that precluded a general mapping between \mathbb{R}^n and \mathbb{R}^m .

Another thread of related work tries to characterize the equilibrium in inexpressive mechanisms in specific settings (e.g., (Rosenthal & Wang 1996)). The challenge here is that determining equilibrium behavior is usually prohibitively difficult even for the simplest non-trivial mechanisms. Furthermore, when a particular equilibrium is found to have certain properties, one often cannot rule out the possibility of additional equilibria that do not share those properties.

There has been some research related to expressiveness issues in dominant-strategy mechanisms. Blumrosen and Feldman (2006) showed a tradeoff between the efficiency of

the best possible mechanism and the number of discrete actions available to the designer. Similarly, Ronen (2001) described methods for achieving near efficiency with limited bidding languages. The restriction to studying dominant-strategy mechanisms imposes severe limitations on which questions about expressiveness arise. In particular, uncertainty about others’ private information becomes an issue only when considering mechanisms that do not have dominant strategies. As we showed, the larger the possible type space of others, the more expressiveness an agent may need for efficiency. Our results apply to settings where agents do, or do not, have dominant strategies. Also, our results are not specific to any application, such as a CA.

One of the first applications to benefit from expressiveness was strategic sourcing. Sandholm (2007) described how building more expressive mechanisms—that generalize both CAs and multi-attribute auctions—for supply chains has saved billions of dollars through increased efficiency. Some work on expressiveness has begun to appear in the context of search keyword auctions (aka sponsored search) as well. Even-Dar *et al.* (2007) examined auctions where bidders can purchase keywords associated with specific contexts. Under certain probabilistic assumptions they are able to prove that the system becomes more efficient when this extra level of expressiveness is allowed.

Conclusions and future work

A recent trend in (electronic) commerce is a demand for higher levels of expressiveness in mechanisms. We provided the first general model of expressiveness for mechanisms. Our model included a new expressiveness measure, maximum impact dimension, that captures the number of different ways that an agent can impact the outcome of a mechanism. We also introduced two related measures of expressiveness based on the concept of shattering from computational learning theory.

We then described how these measures relate to efficiency. We derived an upper bound on the expected efficiency of a mechanism’s most efficient Nash equilibrium. The bound depends only on the extent to which agents can impact the mechanism’s outcome. This bound enables one to study the relationship between expressiveness and efficiency by avoiding two classic hurdles: 1) our bound can be analyzed without having to solve for equilibrium, and 2) our bound applies to the most efficient equilibrium so it can be used to analyze mechanisms with multiple (or an infinite number of) equilibria. We proved that this bound increases *strictly* monotonically for the best mechanism that can be designed as the limit on any agent’s expressiveness increases (until the bound reaches full efficiency). In addition, we proved that a small increase in expressiveness can potentially lead to arbitrarily large increases in the efficiency bound, depending on the prior over agents’ preferences.

Finally, we applied our model of expressiveness to a class of mechanisms which we call channel based. This class involves mechanisms that take expressions of value through channels from agents to outcomes, and select the outcome with the largest sum. Many mechanisms—such as combinatorial auctions, multi-attribute mechanisms,

and any Vickrey-Clarke-Groves scheme—can be cast as channel-based mechanisms. We showed that our domain-independent measures of expressiveness appropriately relate to a natural notion of expressiveness in channel-based mechanisms, the number of channels allowed (which already generalizes a traditional measure of expressiveness in CAs called k -wise dependence (Conitzer, Sandholm, & Santi 2005)). Our general measures of expressiveness and our results on how they relate to efficiency then yield interesting results for channel-based mechanisms: 1) allowing one additional channel can yield an arbitrarily large increase in the bound, and 2) any (channel-based) multi-item auction that assumes additivity anywhere in any agent’s bids (e.g., auctions where bids can be submitted on individual items only) can be arbitrarily inefficient—unless agents have no private information.

The framework we developed enables one to understand mechanisms from a new perspective. This opens the door for a possible new avenue of research within mechanism design. On the practical side, we already see two uses of our expressiveness measures. They can be used to bound the efficiency—and therefore provide a lower bound on inefficiency—of existing mechanisms. They can also potentially be used to design new mechanisms, either by hand or by computer (Conitzer & Sandholm 2002).

Acknowledgement

This work has been supported by the National Science Foundation under grants IGERT-9972762, ITR-0205435, IIS-0427858, and CNS-0627513. We would also like to thank George B. Davis for his helpful feedback.

References

- Abrams, Z.; Ghosh, A.; and Vee, E. 2007. Cost of conciseness in sponsored search auctions. In *Proceedings of Workshop on Internet Economics (WINE)*.
- Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion.
- Benisch, M.; Sadeh, N.; and Sandholm, T. 2007. Theory of expressiveness in mechanisms. Technical Report CMU-CS-07-178, School of Computer Science, Carnegie Mellon University.
- Bernheim, B. D., and Whinston, M. D. 1986. Menu auctions, resource allocation, and economic influence. *Quarterly Journal of Economics* 101(1):1–32.
- Blumer, A.; Ehrenfeucht, A.; Haussler, D.; and Warmuth, M. K. 1989. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM* 36(4):865–929.
- Blumrosen, L., and Feldman, M. 2006. Implementation with a bounded action space. In *Proceedings of ACM Conference on Electronic Commerce (ACM-EC)*.
- Boutilier, C.; Parkes, D.; Sandholm, T.; and Walsh, W. 2008. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11:17–33.
- Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Conitzer, V.; Sandholm, T.; and Santi, P. 2005. Combinatorial auctions with k -wise dependent valuations. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*.
- Cramton, P.; Shoham, Y.; and Steinberg, R. 2006. *Combinatorial Auctions*. MIT Press.
- Even-Dar, E.; Kearns, M.; and Wortman, J. 2007. Sponsored search with contexts. In *Proceedings of Workshop on Internet Economics (WINE)*.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Holzman, R.; Kfir-Dahav, N.; Monderer, D.; and Tennenholtz, M. 2004. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior* 47:104–123.
- Hurwicz, L. 1972. On informationally decentralized systems. In McGuire, C., and Radner, R., eds., *Decision and Organization*. Amsterdam: North Holland.
- Jehiel, P., and Moldovanu, B. 2001. Efficient design with interdependent valuations. *Econometrica* 69(5):1237–1259.
- Milgrom, P. 2007. Simplified mechanisms with applications to sponsored search and package auctions. Working paper.
- Mount, K., and Reiter, S. 1974. The informational size of message spaces. *Journal of Economic Theory* 8(2):161–192.
- Nisan, N., and Ronen, A. 2007. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research* 29:19–47.
- Ronen, A. 2001. Mechanism design with incomplete languages. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 105–114.
- Rosenthal, R. W., and Wang, R. 1996. Simultaneous auctions with synergies and common values. *Games and Economic Behavior* 17:32–55.
- Sandholm, T., and Suri, S. 2006. Side constraints and non-price attributes in markets. *Games and Economic Behavior* 55:321–330.
- Sandholm, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135:1–54.
- Sandholm, T. 2007. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Mag.* 28(3):45–58.
- Vapnik, V. N., and Chervonenkis, A. Y. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications* 16(2):264–280.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.