

A Theory of Expressiveness in Mechanisms

Michael Benisch, Norman Sadeh, and Tuomas Sandholm

December 2007
CMU-ISRI-07-122

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This report also appears as Computer Science Department Technical Report CMU-CS-07-178.

This work has been supported by the National Science Foundation under IGERT grant 9972762, ITR grants 0205435 and IIS-0427858, and Cyber Trust grant CNS-0627513.

Keywords: mechanism design, expressiveness, message space, efficiency, Vickrey-Clarke-Groves mechanism, combinatorial auction, multi-attribute auction, search keyword auction, sponsored search, catalog offers, menus, web commerce.

Abstract

A key 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, and elicitation of opinions from large and diverse communities. Intuitively, one would think that this increase in expressiveness would lead to more efficient mechanisms (e.g., due to better matching of supply and demand). However, until now we have lacked a general way of characterizing the expressiveness of these mechanisms, analyzing how it impacts the actions taken by rational agents—and ultimately the outcome of the mechanism. In this technical report we introduce a general model of expressiveness for mechanisms. Our model is based on a new measure which we refer to as the *maximum impact dimension*. The measure captures the number of different ways that an agent can impact the outcome of a mechanism. We proceed to uncover a fundamental connection between this measure and the concept of *shattering* from computational learning theory.

We also provide a way to determine an upper bound on the expected efficiency of any mechanism under its most efficient Nash equilibrium which, remarkably, depends only on the mechanism's expressiveness. We show that for any setting and any prior over agent preferences, the bound on efficiency of a mechanism designed optimally under a constraint on expressiveness increases *strictly* as more expressiveness is allowed (until the bound reaches full efficiency). In addition, we prove that a small increase in expressiveness can potentially lead to an arbitrarily large increase in the efficiency bound, depending on the prior.

We conclude with a study of a restricted class of mechanisms which we call *channel based*. The restriction is that these mechanisms take expressions of value through channels from agents to outcomes, and select the outcome with the largest sum. (Channel-based mechanisms subsume most combinatorial and multi-attribute auctions, any Vickrey-Clarke-Groves mechanism, etc.) In this class, a natural measure of expressiveness is the number of channels allowed (this generalizes the k -wise dependence measure of expressiveness traditionally used in the combinatorial auction literature). As a sanity check of our general domain-independent measure of expressiveness, we show that it appropriately relates to the number of channels when applied to channel-based mechanisms. This allows us to transfer all of our results regarding efficiency to this domain.

1 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 imposes 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 opinions. This trend has already manifested itself in combinatorial auctions, multi-attribute auctions, and generalizations thereof, which are used to trade tens of billions of dollars worth of items annually [28, 55, 56, 39, 19]. 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. It is also emerging in the context of security and privacy interfaces developed in different application domains (e.g., [53, 52]). In Web 2.0 parlance, this demand for increasingly diverse offerings is called the Long Tail [4].

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 (e.g., [57, 28, 55, 56, 19]). 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 (e.g., [57, 55, 56, 19]).

Intuitively, one would think that more expressiveness leads to higher efficiency (sum of the agents’ utilities) of the outcome (e.g., better matching of supply and demand). Efficiency improvements have indeed been reported from combinatorial and multi-attribute auctions (e.g., [54, 55, 56, 37, 19]). However, until now, 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. It was not even known whether, in any settings, more expressiveness always leads to more efficiency. (In fact, on the contrary, it has been observed that in certain settings additional expressiveness can give rise to additional equilibria of poor efficiency [40].)

Short of empirical tweaking, participants in the scenarios we described lacked results they can rely on to determine how much—and what forms of—expressiveness they need. These questions have vexed mechanism design theorists, but are not only theoretical. Answers could ensure that ballots are expressed in a form that matches the issues voters care about, that companies are able to identify suppliers that best match their needs, that supply and demand are better matched in B2C and C2C markets, that users of a cell phone-based Friend Finder application can express those privacy preferences that really matter, and so on.

In this paper we introduce a general model of expressiveness for mechanisms (Section 2). This includes a new expressiveness measure, which we refer to as the *maximum impact dimension*. The

measure captures the number of different ways that an agent can impact the outcome of a mechanism. We proceed to uncover a fundamental connection between this measure and the concept of *shattering* from computational learning theory. (We say that a mechanism allows an agent to shatter a set of outcomes if the agent can cause every possible (ordered) combination of those outcomes to be chosen.)

We proceed, in Section 3, to describe perhaps the most important property of our domain-independent measure of expressiveness: how it relates to the efficiency of the mechanism’s outcome. We provide a way to determine an upper bound on the expected efficiency of any mechanism under its most efficient Nash equilibrium which, remarkably, depends only on the mechanism’s expressiveness. This enables us to sidestep two major classic hurdles in studying the relationship between expressiveness and efficiency: 1) it can be analyzed without having to solve for an equilibrium of the mechanism (something that has proved extremely difficult for inexpressive mechanisms [50, 64, 61, 46, 66, 43]), and 2) since it bounds the most efficient equilibrium it can be meaningfully applied to mechanisms with multiple (or infinite) equilibria, e.g., CAs [9]. We show that for any setting and any prior over agent preferences, the bound on efficiency of a mechanism designed optimally under a constraint on expressiveness increases *strictly* as more expressiveness is allowed (until the bound reaches full efficiency). In addition, we prove that a small increase in expressiveness can potentially lead to an arbitrarily large increase in the efficiency bound, depending on the prior.

Finally, in Section 4, we instantiate our model of expressiveness for a restricted class of mechanisms which we call *channel based*. The restriction is that these mechanisms take expressions of value through channels from agents to outcomes, and select the outcome with the largest sum. (Channel-based mechanisms subsume most combinatorial and multi-attribute auctions, any Vickrey-Clarke-Groves [63, 14, 23] mechanism, etc.) In this class, a natural measure of expressiveness is the number of channels allowed (this generalizes the k -wise dependence measure of expressiveness traditionally used in the combinatorial auction literature). As a sanity check of our general domain-independent measure of expressiveness, we show that it appropriately relates to the number of channels when applied to channel-based mechanisms. By studying these mechanisms within our framework we are able to prove that increasing their expressiveness by a small amount (i.e., adding a single channel) cannot decrease our bound on expected efficiency for the mechanism, and under some preference distributions leads to an arbitrarily large increase in this bound.

We conclude with a discussion of related work (Section 5) and a summary of our results (Section 6).

1.1 Preliminaries

The basic 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 settings where $u_i = u_i(t^n, O)$, i.e., an agent's utility depends on the others' private signals. These settings are called *interdependent values* settings. In both types of settings, agents report expressions to the mechanism, denoted θ_i , and based on them the mechanism chooses an outcome.

The mechanism itself consists of an outcome function, $f(\theta^n)$, which aggregates the expressions of all the agents and chooses an outcome from \mathcal{O} . It also consists of a payment function, $\pi(\theta^n)$, which determines how much each agent must pay. 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 and the parameters of the mechanism, $b_i(t_i, f, \pi)$. We do not restrict these equilibrium reports to be deterministic: we allow b_i to be a random variable where the agent specifies a probability distribution over possible reports.

To summarize, we use the following notation.

- $t_i \in T_i$ is the true type of an agent i . θ_i is the expression that agent i reports to the mechanism, the subscript θ_{-i} is used to denote the set of expressions by all the agents other than i , and the superscript θ^n is used to denote a collection of n expressions.
- $O \in \mathcal{O}$ is an outcome from the set of all possible outcomes imposable by the mechanism, \mathcal{O} .
- $u_i : T_i, \mathcal{O} \rightarrow \mathbb{R}$ is agent i 's utility function. It takes as input the agent's true type and an outcome and returns the real-valued utility of the agent if that outcome were to be chosen. (We also discuss results that apply interdependent values settings where $u_i = u_i(t^n, O)$, i.e., an agent's utility also depends on the others' private signals.)
- $f : \Theta^n \rightarrow \mathcal{O}$ is the outcome function of the mechanism. It takes as input the expression of each agent and returns an outcome from the set of all possible outcomes.
- $\pi : \Theta^n \rightarrow \mathbb{R}^n$ is the payment function of the mechanism. It takes as input the expression of each agent and returns the payment to be made by each agent.
- $b_i : T_i, f, \pi \rightarrow P(\theta_i)$ is the expression of agent i in a particular equilibrium. It takes as input agent i 's true type, the outcome function and payment function of the mechanism. It returns a (potentially randomized) expression, in the case of a mixed equilibrium, b_i is a random variable with an underlying probability distribution. (Note that this function cannot depend on the private types of the other agents, even if agent i 's utility does.)

Using this formalism we can describe the expected efficiency, $\mathcal{E}(f, \pi)$, of a mechanism (where expectation is taken over the true types of the agents, and their randomized equilibrium expressions) as

$$(1) \quad E[\mathcal{E}(f, \pi)] = \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(b(t^n, f, \pi) = \theta^n) \sum_i u_i(t^n, f(\theta^n))$$

The following example shows how this formalism can be used to model a combinatorial auction.

Example 1. *In a fully expressive combinatorial auction with m items, each of the agents is a bidder whose type represents his or her private valuation for each of the 2^m different combinations of items. The outcome space includes all of the n^m different ways the goods can be allocated amongst the bidders. Agents are allowed to express their entire type to the mechanism and the outcome function chooses the allocation that maximizes the sum of the bidders' valuations.*

The payment function can charge each agent its bid (aka. the first-price payment rule) or the difference in utility of the other agents had the agent in question not participated (aka. the Vickrey-Clarke-Groves (VCG) payment rule). Under the VCG payment rule, each agent has a (weakly) dominant strategy to tell the truth, so one equilibrium distribution, b , over expressions is a point mass on the agents' true valuations.

2 Characterizing the expressiveness of mechanisms

The primary goal of this technical report is to better understand the tradeoffs associated with making mechanisms more or less expressive. In order to accomplish this, we must first come up with meaningful (and general) definitions of a mechanism's expressiveness. First we will demonstrate that two seemingly natural ways of characterizing the expressiveness of different mechanisms, the dimensionality of their expressions and the granularity of their outcomes, do not capture the fundamental difference between expressive and inexpressive mechanisms.

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 (this is one key difference between fully expressive CAs and auctions that only allow per-item-bids, for example). However, not only does this limit our notion of expressiveness to mechanisms with real-valued expressions, it also does not adequately differentiate between expressive and inexpressive mechanisms. This is because the cardinality of \mathfrak{R} is the same as the cardinality of \mathfrak{R}^n as proved by Cantor in 1890 [13].

Proposition 1. *For any mechanism that allows multi-dimensional real-valued expressions, (i.e., where $\Theta \subseteq \mathbb{R}^d$), there exists an equivalent mechanism that only allows the expression of a single real value (i.e., where $\Theta = \mathbb{R}$).¹*

This illustrates that the fundamental difference between expressive and inexpressive mechanisms cannot be captured simply by the dimensionality of the expressions they allow. 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 wireless spectra) have coarser outcome spaces than those which can allocate them 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 [29, 44].

¹All proofs can be found in the appendix at the end of this technical report.

Proposition 2. *In any private values setting, the expected efficiency of the VCG mechanism with outcome space \mathcal{O} (when all agents play the weakly dominant truthful equilibrium) is greater than or equal to the expected efficiency of the VCG mechanism with outcome space $\mathcal{O}' \subset \mathcal{O}$.*

In contrast, we are interested in studying the impact of a mechanism’s expressiveness on its efficiency—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. Furthermore, in many practical applications there is no reason to restrict the outcome space,² but there may be a prohibitive burden on agents if they are asked to express a huge amount of information; thus it is limited expressiveness that is the crucial issue.

2.1 A measure of expressiveness: the maximum impact dimension

In order to properly differentiate between expressive and inexpressive mechanisms with the same outcome space, we propose to measure the extent to which an agent can impact the outcome that is chosen. In this technical report we will limit ourselves to studying the mechanism’s outcome function rather than also studying the payment function. In our view, the outcome function is primarily responsible for determining an agent’s expressiveness level. In settings where agents do not care about each others’ payments, this is basically without loss of generality because if an agent could choose between paying more or less for the same outcome, choosing to pay more would be a dominated strategy. Thus that extra expressiveness has no value. (In settings where agents care about each others’ payments, expressiveness related to payments may be an interesting area for future research.)

The fundamental way in which a mechanism allows an agent to express different preferences is by allowing it to cause different *impact vectors* of outcomes to be chosen. An impact vector captures the impact of a particular expression by an agent under each of the joint types that the other agents may have.

Definition 1 (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, we order the joint types in T_{-i} from $t_{-i}^{(1)}$ to $t_{-i}^{(|T_{-i}|)}$; then g_i can be represented as a vector of outcomes $[o_1, o_2, \dots, o_{|T_{-i}|}]$.*

In some cases an agent, i , may wish to impact the mechanism differently under each of i ’s own types. However, i can only actually express an impact vector if there exists some *pure strategy profile* of the other agents such that i can cause the mechanism to choose the correct mapping from the others’ types to outcomes.

Definition 2 (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*

²This is the case as long as the mechanism designer’s goal is efficiency, but this is not always the case for revenue maximization, for example.

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,

$$h_I(t_I) = [h_1(t_1), h_2(t_2), \dots, h_{|I|}(t_{|I|})] = [\theta_1, \theta_2, \dots, \theta_{|I|}]$$

We say that agent i can *express* an impact vector against a pure strategy profile of the other agents if there exists some expression by i that causes each of the outcomes in the impact vector to occur when paired with the expressions made by the other agents under the pure strategy mapping.

Definition 3 (expressability). *Agent i can express an impact vector, g_i , if*

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

Figure 1 illustrates how an agent can express certain impact vectors against a particular pure strategy profile of the other agents. In this example, the agents other than i are playing the pure strategy profile, $[\theta_{-i}^{(x)}, \theta_{-i}^{(y)}]$. Against this pure strategy profile, agent i can express the impact vectors $[A, B]$ and $[C, D]$ by choosing between expressions $\theta_i^{(1)}$ and $\theta_i^{(2)}$.




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

Agent i can only distinguish among a set of impact vectors if it can express each of them against the same pure strategy profile of the other agents. If no pure strategy profile exists that allows the agent to simultaneously express each of the impact vectors in the set, then we say that the agent cannot *distinguish* between those impact vectors. In other words, the agent should be able to express each of the vectors by altering its own expression.

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

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

When this is the case, we write

$$D_i(G_i)$$

Intuitively, more expressive mechanisms allow agents to distinguish among larger sets of impact vectors. We will now define our primary measure of expressiveness based on this intuition. Since mechanisms can allow different levels of expressiveness for different agents, we will define our expressiveness measure in terms of one agent. The measure captures the number of different impact vectors the agent can distinguish among. Since this depends on what the others express, we measure the best case where the others happen to 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,³*

$$d_i = \max_{G_i} \left\{ |G_i| \mid D_i(G_i) \right\}$$

We will show in Section 3 that every agent’s maximum impact dimension ties directly to an upper bound on the expected efficiency of the mechanism’s most efficient Nash equilibrium. In particular, the upper bound increases *strictly* monotonically as the maximum impact dimension for any agent i increases from 1 to d_i^* , where d_i^* is the smallest maximum impact dimension needed by the agent in order for the bound to reach full efficiency.

The maximum impact dimension also has some drawbacks as a measure. First, it does not capture the way in which an agent’s impact vectors are distributed. For example, it is possible that a mechanism that allows a smaller maximum impact dimension can be designed to allow an agent to distinguish among a more important (e.g., for efficiency) set of impact vectors than a mechanism that allows the agent a larger maximum impact dimension. Second, it is not clear that the maximum impact dimension can be measured, numerically or analytically, in settings where even a single agent has an infinite type space.

2.2 Shattering-based expressiveness

We will now proceed to discussing another kind of measure of expressiveness which we will call the *shatterable outcome dimension*. As we will discuss later, it has somewhat different uses than does the maximum impact dimension. The two are closely related, however, as we will discuss.

The shatterable outcome dimension is based on a notion called *shattering* which we adapt from the field of computational learning theory (c.f., pp. 215-216 [41],[62, 10]). The shatterable outcome dimension measure addresses both of the concerns with maximum impact dimension that we raised at the end of the previous section. Unlike the maximum impact dimension, which provides no information as to how the distinguishable impact vectors are distributed, the shatterable outcome dimension measures the *number of different outcomes* that an agent can shatter (i.e., express every possible impact vector over). In addition, as we will show at the end of this section, it has the advantage that we can rule out the shatterability of a set of outcomes over any number of expressions (of the other agents) by merely ruling out the existence of any *pair* of expressions (of the other agents) that allow the agent to shatter the set of outcomes. This enables us to analyze the

³For any agent i , $d_i \leq |\mathcal{O}|^{|T-i|}$ since this is the maximum number of impact vectors over \mathcal{O} of length $|T-i|$.

measure even when agents have infinite type spaces (and can report infinitely many different expressions), and may help one operationalize expressiveness for automated mechanism design [15] in the future.

We will begin by defining what it means for an agent to be able to shatter a set of outcomes. In learning theory, a class of binary classification functions⁴ is said to shatter a set of k instances if there is at least one function in the class that assigns each of the possible 2^k dichotomies of labels to the set of instances. Intuitively, a class of functions that can shatter larger sets of instances is more expressive. To illustrate this idea consider the following example taken from Mitchell pp. 215-216 [41].

Example 2. Consider the class of binary classification functions that assign a 1 to points only if they fall in an interval on the real number line between two constants a and b . Now we can ask whether or not this class of functions has enough expressive power to shatter the set of instances $S = \{3.1, 5.7\}$? Yes, for example the four functions $(1 < x < 2)$, $(1 < x < 4)$, $(4 < x < 7)$ and $(1 < x < 7)$ will assign all possible labels to the instances in S .

Our adaptation of shattering for mechanisms 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 i can distinguish among each of the $|\mathcal{O}'|^{T-i}$ different possible impact vectors that involve only outcomes from \mathcal{O}' . Formally, let

$$G_i^{\mathcal{O}'} = \{g_i | g_i = [o_1, o_2, \dots, o_{T-i}], o_j \in \mathcal{O}'\}$$

Then, i can shatter outcomes \mathcal{O}' if $D_i(G_i^{\mathcal{O}'})$.

Example 3. If agent i can distinguish among the following set of impact vectors, G_i , then it can shatter a set of outcomes, $\{A, B, C, D\}$, over a set of two different joint types of the other agents, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$ (note that rows and columns have no particular significance, they are used for presentation only):

$$G_i = \left\{ \begin{array}{cccc} [A, A], & [B, A], & [C, A], & [D, A], \\ [A, B], & [B, B], & [C, B], & [D, B], \\ [A, C], & [B, C], & [C, C], & [D, C], \\ [A, D], & [B, D], & [C, D], & [D, D] \end{array} \right\}$$

We now introduce a slightly weaker adaptation of shattering for mechanisms. It will be a pairwise notion that involves an agent being able to cause every pair of outcomes to be chosen under every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., which outcome happens against which type)⁵. We call this *semi-shattering* an outcome space.

⁴Binary classification functions are functions that assign each possible input a binary output label of either 0 or 1.

⁵There are many ways to generalize the shattering notion to functions that can return more than two outcomes, c.f. [8]. We have adapted the two most natural ones for our work on expressiveness in mechanism design.

As we will show in Section 3, semi-shattering is more important when designing mechanisms for private values settings (i.e., settings where an agent's utility for each outcome depends only on its own type, and not on the private types of the other agents). In such settings the mechanism designer can build a fully efficient mechanism while only allowing agents to semi-shatter the outcome space; it is not necessary to allow them to fully shatter any of the outcomes (Lemma 3).

Definition 7 (outcome semi-shattering). *A mechanism allows agent i to semi-shatter a set of outcomes, $\mathcal{O}' \subseteq \mathcal{O}$, if i can distinguish among a set of impact vectors that assigns each of the $\binom{|\mathcal{O}'|+1}{2}$ unordered pairs of outcomes (with replacement) to each pair of the other agents' types. Formally, i can semi-shatter \mathcal{O}' if i can distinguish among a set of impact vectors, $G_i^{\mathcal{O}'}$, such that*

$$\forall \left\{ \left\{ t_{-i}^{(1)}, t_{-i}^{(2)} \right\} \mid t_{-i}^{(1)} \neq t_{-i}^{(2)} \right\}, \forall \left\{ \left\{ o_1, o_2 \right\} \mid o_1, o_2 \in \mathcal{O}' \right\}, \exists g_i \in G_i^{\mathcal{O}'},$$

$$\left[g_i \left(t_{-i}^{(1)} \right) = o_1 \text{ and } g_i \left(t_{-i}^{(2)} \right) = o_2 \right] \text{ or } \left[g_i \left(t_{-i}^{(1)} \right) = o_2 \text{ and } g_i \left(t_{-i}^{(2)} \right) = o_1 \right]$$

Example 4. *If agent i can distinguish among the following set of impact vectors, G_i , then it can semi-shatter a set of outcomes, $\{A, B, C, D\}$, over a set of two different joint types of the other agents, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$ (note that the order of the pairs that are included does not matter, for example AB could be replaced with BA):*

$$G_i = \left\{ \begin{array}{l} [A, A], \\ [A, B], [B, B], \\ [A, C], [B, C], [C, C], \\ [A, D], [B, D], [C, D], [D, D] \end{array} \right\}$$

Since semi-shattering is a pairwise notion, it does not always include the entire bottom left half of a matrix of impact vectors as in the previous example. For example, the following set of impact vectors constitutes semi-shattering a set of 3 outcomes.

Example 5. *If agent i can distinguish among the following set of impact vectors, G_i , then it can semi-shatter the set of outcomes $\{A, B, C\}$ over a set of three different joint types of the other agents, $t_{-i}^{(1)}$, $t_{-i}^{(2)}$, and $t_{-i}^{(3)}$:*

$$G_i = \left\{ \begin{array}{l} [A, A, A], \\ [A, A, B], \\ [A, A, C], \\ [A, B, B], [B, B, B], \\ [A, B, C], [B, B, C], \\ [A, C, C], [B, C, C], [C, C, C] \end{array} \right\}$$

Notice that every pair of outcomes appears in every pair of slots at least once. That is exactly the requirement for semi-shattering.

Now we can define a notion of expressiveness based on the size of the largest outcome space that an agent can (semi-)shatter.⁶ It captures the number of outcomes that the mechanism can support full expressiveness over for that agent. We call this the agent’s *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, $\mathcal{O}_i^* \subseteq \mathcal{O}$, has size k_i .*

The shatterable and semi-shatterable outcome dimension measures are closely related to the maximum impact dimension. For example, whenever an agent’s (semi-)shatterable outcome dimension goes up, so does its maximum impact dimension.

Proposition 3. *The most expressive mechanism for agent i (i.e., the mechanism allowing the agent the largest maximum impact dimension) when it has (semi-)shatterable outcome dimension $k_i < |\mathcal{O}|$, has a strictly greater maximum impact dimension than that of any mechanism where agent i has (semi-)shatterable outcome dimension $k_i - 1$.*

While the two measures are related, the shatterable outcome dimension can be thought of as more of a measure of the breadth of an agent’s expressiveness. The maximum impact dimension necessary for an agent to shatter k outcomes increases geometrically in the number of types of the other agents. This illustrates the relationship between expressiveness and uncertainty, since the number of types that the other agents have can be thought of as a support-based measure of agent i ’s uncertainty. The more uncertainty an agent has about the other agents, the more expressiveness the agent needs to shatter a given set of outcomes.

Proposition 4. *Any mechanism that allows agent i to shatter k_i outcomes has maximum impact dimension at least $|T_{-i}|^{k_i}$ for i .*

Because shattering (and semi-shattering) require agents to have a greater amount of control, we have been able to analyze these measures more easily in domains where agents have infinitely many types. In particular, we have derived the following necessary pairwise condition, which can be checked analytically or experimentally to determine whether a mechanism allows an agent to (semi-)shatter a set of outcomes. We actually use this insight throughout our study of channel-based mechanisms in Section 4.

Proposition 5. *Agent i can (semi-)shatter an outcome space \mathcal{O}' only if there exists at least one pair of expressions by the other agents, $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, which allows i to (semi-)shatter \mathcal{O}' . (In other words, there exists a pair of fixed expressions by the other agents such that i can cause any (un-ordered) pair of outcomes from \mathcal{O}' to be chosen.)*

2.3 Uses of the expressiveness measures

The expressiveness measures introduced above enable us to understand mechanisms from a new perspective. The measures being so new, we undoubtedly fail to see all of their possible uses at this time. However, we already see two uses.

⁶Our measure deals with the size of this space, rather than the specific outcomes it contains, because a designer can always re-label the outcomes in the set to transform it into any other set of the same size.

First, we can measure the expressiveness of an existing mechanism, and thereby bound how well the mechanism can do in terms of the designer’s objective. For example, in the next section, we show how our expressiveness measures directly determine an upper bound on the efficiency of any mechanism.

Second, one may be able to use the expressiveness measures in designing new mechanisms. For example, if there are some constraints on what—and how much—information the agents can submit to the mechanism (for example, in a CA, allowing bids on packages of no more than k items), then our measures can be used to design the most expressive mechanism subject to those constraints. This, in turn, hopefully maximizes the mechanism designer’s objective subject to the constraints. For example, this approach can be used to yield the highest upper bound on efficiency.

We can also ask which of the expressiveness measures—maximum impact dimension, shatterable outcome dimension, and semi-shatterable outcome dimension—are most appropriate under which settings and purposes. If the designer knows which impact vectors are (most) important, then the maximum impact dimension is the measure of choice. If, instead, the designer knows which outcomes are (most) important but not which impact vectors are (most) important, then the other two measures can be used to make sure that the agents can fully express themselves over those outcomes. As we will show in Section 3, in private values settings the appropriate measure is semi-shatterable outcome dimension (for one, full semi-shatterability is enough to guarantee that lack of expressiveness will not limit the mechanism’s efficiency at all), and in interdependent values settings the appropriate measure is shatterable outcome dimension. Also, we will show that less than full (semi-)shatterability necessarily leads to inefficiency in any setting under some prior over agent preferences.

Another use of the semi-shatterable outcome dimension is to analyze a broad subclass of mechanisms which we will call channel based. This will be discussed in Section 4.

3 Relationship between expressiveness and efficiency

Perhaps the most important property of our domain-independent measures of expressiveness is how they relate to efficiency of the mechanism’s outcome. We will now present an upper bound on the expected efficiency of the mechanism’s *most efficient* equilibrium, which remarkably depends only on the extent to which agents can impact the mechanism’s outcome. Using this bound allows us to sidestep two of the major roadblocks in analyzing the relationship between expressiveness and efficiency: 1) it can be studied without having to solve for any of the mechanism’s equilibria (attempts at doing this have proved extremely difficult for inexpressive mechanisms [50, 64, 61, 46, 66, 43]), 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 [9].

We achieve these goals by making an optimistic assumption that leads to easier analysis and guarantees that the result is an upper bound. Specifically, we assume 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 when agents need to speculate and counter-speculate about each other. It also implies that we can restrict our analysis to *pure* strategies rather than considering the infinite space of mixed strategies. This is because under our assumption, a pure

strategy always exists that achieves at least as much expected efficiency as any mixture. (This is analogous to the fact that there exists a pure strategy providing at least as much utility as any mixture in general games.) For convenience, let $W(t^n, o)$ denote the total social welfare of outcome o when agents have private types (or private signals) t^n ,

$$W(t^n, o) = \sum_i u_i(t^n, o)$$

Proposition 6. *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 ,*

$$(2) \quad E[\mathcal{E}(f)]^+ = \max_{\hat{B}(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) W\left(t^n, f(\{\hat{B}_1(t_1), \hat{B}_2(t_2), \dots, \hat{B}_n(t_n)\})\right)$$

The maximum is taken over $\hat{B}(\cdot)$, a pure strategy profile that maps every joint type vector to an expression for each agent.⁷

We will now demonstrate that the bound from Equation 2 is closely tied to our notions of expressiveness. First we will prove that the bound *strictly* increases for the best outcome function that can be designed with maximum impact dimension d_i for agent i , as d_i goes from 1 to d_i^* (where d_i^* is the maximum impact dimension needed by the agent for the bound to reach full efficiency). Since we prove this for the *best* outcome function, it also holds true for an upper bound on any mechanism that allows maximum impact dimension d_i for any agent i .

The way we approach this problem is to consider calculating the bound from the fixed perspective of a particular agent i (the value of the bound does not depend on which agent we choose to consider). Based on our assumption, we know that 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 distinguishable impact vectors that lead to the highest expected efficiency.

Observe that there is an impact vector, $g_i^{t_i}$, for each of agent i 's types, t_i , that represents the vector of efficient outcomes when t_i is matched with each of the joint types of the other agents. In order to achieve full efficiency, agent i must be able to distinguish among all of these vectors. We call a set that contains all of these vectors a *fully efficient set*.

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

$$\forall t_i, \exists g_i \in G_i^*, \forall \{t_{-i} \mid P(t_i, t_{-i}) > 0\}, \quad W(\{t_i, t_{-i}\}, g_i(t_{-i})) = \max_{o \in \mathcal{O}} W(\{t_i, t_{-i}\}, o)$$

The first two results regarding our efficiency bound address the conditions under which the mechanism has enough expressiveness for it to reach full expected efficiency. (Our bound never exceeds full efficiency.)

⁷Recall that an agent's strategy can only depend on its own private type, even if its utility depends on the private signals of others.

Proposition 7. *The upper bound, $E[\mathcal{E}(f)]^+$, for any outcome function f reaches full expected efficiency iff f allows at least one agent to distinguish among each of the impact vectors in at least one of its fully efficient sets.*

Our next result demonstrates that allowing any one agent enough expressiveness to make the bound achieve full efficiency, is functionally equivalent to doing so for all agents.

Proposition 8. *If any agent i can distinguish among each of the impact vectors in at least one of its fully efficient sets, then every other agent j can also distinguish among each of the impact vectors in at least one of its fully efficient sets.*

In settings where upon learning its own type an agent knows for sure what the types of the other agents are, the agent only needs an impact dimension of $|\mathcal{O}|$ to bring the bound to full efficiency. (Note that this is slightly more general than assuming the agent has perfect information about the types of the other agents *a priori*, since it need only have this information once its own type is revealed.)

Proposition 9. *If agent i has full information about the types of the other agents based on its own private type, it has a fully efficient set of size $\leq |\mathcal{O}|$. Formally, let G_i^* be agent i 's smallest fully efficient set,*

$$(\forall t_i, \exists t_{-i} \mid P(t_i, t_{-i}) = 1) \Rightarrow |G_i^*| \leq |\mathcal{O}|$$

Corollary 1. *In any setting where an agent i has full information about the types of the other agents based on its own type, there exists an outcome function for which the upper bound reaches full efficiency while limiting agent i to maximum impact dimension $d_i \leq |\mathcal{O}|$.*

3.1 The efficiency bound increases strictly with expressiveness

We will now present the main result relating our notion of expressiveness, maximum impact dimension, to our upper bound on expected efficiency. It demonstrates that a mechanism designer can *strictly* increase the upper bound on expected efficiency by allowing any agent more expressiveness (until the bound reaches full efficiency). The result applies to the mechanism that maximizes the bound subject to the constraint that the agent's expressiveness is less than or equal to that level. The bound attained by such a mechanism also serves as an upper bound on the expected efficiency that is attainable by any outcome function with that expressiveness level.

Theorem 1. *The upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function that limits agent i 's expressiveness to a 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.*

From this result we can also derive the following corollary related to our other two measures of expressiveness.

Corollary 2. *The upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function that limits agent i 's expressiveness to a (semi-)shatterable outcome dimension k_i increases strictly monotonically as k_i goes from 1 to k_i^* , where k_i^* is the (semi-)shatterable outcome dimension necessary for the bound to reach full efficiency.*

3.2 Inadequate expressiveness can lead to arbitrarily low efficiency in any setting

The next three lemmas provide the foundation for our second main theorem regarding the efficiency bound. They demonstrate that in *any* setting there are distributions over agent preferences under which any increase in allowed expressiveness leads to an arbitrary improvement in the upper bound on expected efficiency. We prove that the arbitrary increase is possible by constructing an example under which it is inevitable. We keep these constructions as general as possible: our constructions allow for any number of outcomes, any number of agents, and any number of types.

Lemma 1. *Consider a setting where an agent's utility for any outcome may depend on the private signals of the other agents. For any agent i , in any such setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents) there exist priors over preferences under which the upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function that limits agent i 's expressiveness to a maximum impact dimension of d_i , such that $2 \leq d_i \leq |\mathcal{O}|^{|T-i|}$, is arbitrarily larger than that of any outcome function that limits i 's expressiveness to $d_i - 1$.*

The next lemma deals with the arbitrary improvement that can be achieved by allowing an agent to shatter a single additional outcome. Here we distinguish between an increase in shatterable outcome dimension for interdependent values settings (where an agent's utility for any outcome can depend on its own type and the signals of the other agents), and semi-shatterable outcome dimension for private values settings. As we will see, in private values settings there is no need to allow full shattering in order to achieve efficiency.

Lemma 2. *For any agent i , in any setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents) there exist priors over preferences under which the upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function that limits agent i 's expressiveness to*

- shatterable outcome dimension k_i for interdependent values settings, or
- semi-shatterable outcome dimension k_i for private values settings

such that $2 \leq k_i \leq |\mathcal{O}|$, is arbitrarily larger than that of any outcome function that limits i 's expressiveness to $k_i - 1$.

Private values settings place restrictions on the utility functions that agent's can have and therefore on the outcomes that maximize efficiency under different combinations of types. We will now prove that in such settings it is never necessary for an agent to have the ability to fully shatter any set of outcomes in order to achieve full efficiency.

Lemma 3. *In a private values setting, for any agent i , any pair of outcomes, o_1 and o_2 , and any pair of types for the other agents, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, if there is **some** type of agent i , t_i , where it is strictly more efficient for o_1 to happen under type $t_{-i}^{(1)}$ and o_2 to happen under type $t_{-i}^{(2)}$ than the other way around (i.e., o_1 for $t_{-i}^{(2)}$ and o_2 for $t_{-i}^{(1)}$) then it cannot be more efficient for the outcomes to happen in the other order for **any** type of agent i .*

We conclude this section with a result that integrates the three lemmas above. The theorem adds the fact that an arbitrary loss in efficiency can *only* happen if the shatterable (for interdependent values) or semi-shatterable (for private values) outcome dimension is less than the number of outcomes in the mechanism. Thus these dimensions can be used to provide a guarantee that a mechanism has enough expressiveness to avoid arbitrary inefficiency in any setting under any prior over preferences.

Theorem 2. *For any agent i , in any setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents) there exist priors over preferences for which the upper bound, $E[\mathcal{E}(f)]^+$, of the best outcome function is arbitrarily lower than full expected efficiency iff*

- *agent i 's shatterable outcome dimension, k_i , in an interdependent values setting, or*
- *agent i 's semi-shatterable outcome dimension, k_i , in a private values setting*

is less than the number of outcomes, $k_i < |\mathcal{O}|$.

4 An instantiation to illustrate expressiveness: channel-based mechanisms

We will now instantiate our measure 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 sum⁸. Formally, a channel-based mechanism has the following properties:*

- *The expression space of agent i is a vector of real numbers with dimension c_i , (i.e., $\Theta_i \equiv \mathbb{R}^{k_i}$). Each dimension is called a channel.*
- *For each agent i there is a set of channels associated with each outcome o , S_i^o , such that the mechanism's outcome function chooses the outcome with associated channels that have the greatest reported sum:*

$$f(\theta) = \arg \max_{O \in \mathcal{O}} \sum_i \sum_{j \in S_i^O} \theta_{ij}$$

⁸We 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.*

Many different mechanisms for trading goods, information, and services, such as CAs, exchanges, and multi-attribute settings can be cast as channel-based mechanisms. (This class is even more general than 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 [17].

As a sanity check of our domain-independent measures of expressiveness, we will now demonstrate that they appropriately relate to the number of channels allowed in channel-based domains. Our first result deals with the number of channels an agent needs to shatter an outcome space when it has full information about the other agents.

Proposition 10. *If agent i has full information about the types of the other agents (based on its own private type), in a channel-based mechanism it needs only $\lceil \log_2(|\mathcal{O}|) \rceil$ channels to semi-shatter the entire outcome space. Furthermore, fewer channels do not suffice.*

The intuition behind this result is that when an agent knows exactly what the other agents want (and thus what they are going to say) then the agent does not have to express what it would want to happen if they were to say something different. The most important takeaway of this is that perfect information about the other agents' types basically does away with the need for expressiveness. This implies that in prior research that shows that in certain settings even quite inexpressive mechanisms yield full efficiency (e.g., [1]), the assumption that the agents have no private information is essential.

If an agent has fewer than $\lceil \log_2(|\mathcal{O}|) \rceil$ channels, it will be unable to express a preference for at least one outcome no matter how much the agent likes that outcome. Thus there exist distributions

over agent preferences that lead to an arbitrary loss of efficiency (even if the agent has perfect information about the others).

On the other end of the spectrum, the VCG mechanism, which is fully efficient even in private values settings, can be emulated by a channel-based mechanism with $|\mathcal{O}| - 1$ channels per agent.

Proposition 11. *A channel-based mechanism can emulate the VCG mechanism iff it provides each agent with at least $|\mathcal{O}| - 1$ channels.*

The next result shows that in a channel-based mechanism an agent cannot fully shatter any set of two or more outcomes if the agent has even slightly less than perfect information (i.e., when at least one other agent has more than one type). However, these mechanisms are typically used in private values settings where (as demonstrated by Lemma 3) semi-shattering is more important than full shattering for efficiency.

Proposition 12. *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 agents to shatter outcomes, our results from Section 3 imply that in some interdependent values settings any channel-based mechanism, even the VCG mechanism, will be arbitrarily inefficient. That such mechanisms cannot always get full efficiency in interdependent values settings is already known [33].

Corollary 3. *In any interdependent values setting there exists preference distributions for which any channel-based mechanism (even one that emulates the VCG mechanism) results in arbitrarily less than full expected efficiency.*

Our next result deals with a configuration of channels that prevents an agent from being able to *semi-shatter* a set of outcomes. When this configuration is present in a mechanism it can lead to arbitrary inefficiency even in private values settings.

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 the channels that differ between S_i^A and S_i^C are the same as those that differ between S_i^B and S_i^D . Formally, agent i cannot semi-shatter either pair of outcomes if,*

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

The channel configuration discussed in Theorem 3 generalizes one that appears in the channel-based representation of a multi-item auction where bids are allowed on items only. In fact, it is present in any multi-item auction whenever it is assumed that an agent's bid for a bundle is the sum of its bid on two other non-overlapping bundles (e.g., sub-bundles that compose the full bundle). This is true even if the bids on the sub-bundles are complex themselves, (i.e., assumed to be the sum of bids on other bundles). This fact, along with our results from Section 3, imply that such auctions can be arbitrarily inefficient even in private values settings.

Proposition 13. *Any multi-item auction which can be represented as a channel-based mechanism that treats agent i 's bid on a bundle Q to be the sum of its bids on some two other non-overlapping bundles, q_1 and q_2 , (additive valuations are a special case of this), does not allow the agent to semi-shatter the set of outcomes under which it wins Q , q_1 , q_2 or nothing.*

Corollary 4. *Any multi-item auction which can be represented as a channel-based mechanism that treats agent i 's bid on a bundle Q to be the sum of its bids on some two other non-overlapping bundles, q_1 and q_2 , (additive valuations are a special case of this), has some prior over preferences even in private values settings such that the mechanism is arbitrarily inefficient.*

Finally, we will show that as the number of allowed channels for an agent increases, the agent's expressiveness in the most expressive channel-based mechanism *strictly* increases as well (until full expressiveness is reached). Each time a channel is added, the agent can semi-shatter over at least one more outcome. From Lemma 2 we know that this can lead to an arbitrary increase in our upper bound on expected efficiency, even in private values settings.

Proposition 14. *For any agent i , its semi-shatterable outcome dimension, k_i , and maximum impact dimension, d_i , in the most expressive channel-based mechanism strictly increase (until $k_i = |\mathcal{O}|$) as the number of channels assigned to the agent, c_i , increases.*

We conclude with the following corollary to this result (and the results in Section 3) regarding the upper bound on the expected efficiency of the most efficient Nash equilibria in channel-based mechanisms.

Corollary 5. *The upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best channel-based mechanism that allows c_i channels for agent i is greater than or equal to (and can be arbitrarily larger than) that of any mechanism that allows that agent $c'_i < c_i$ channels.*

5 Related work

There has been relatively little work on expressiveness specifically. We discussed some related papers in the body of this technical report. Here we will briefly summarize other work on the most closely related topics. This work started in economics and has more recently been studied in computer science.

5.1 Informational complexity

At a high level, related questions go back at least to the 1940s when Hayek argued that in distributed resource allocation, it is not practical to communicate all the distributed information to a central decision maker [27]. In the 1970s, Mount and Reiter [42] and Hurwicz [30, 31] formalized this in their theory of *informational complexity*, which asked the question: at a minimum, how much information must a mechanism's message space be able to carry in order to accomplish some design goal? That work focused primarily on the number of real-valued dimensions that were needed. Our Proposition 1 shows that, trivially, the number is always one. To get around Cantor's

theorem that begets Proposition 1, the economists made some technical assumptions (such as local threadedness [42] or Lipschitz continuity [32]) that precluded a general mapping between \mathfrak{R}^n and \mathfrak{R}^m . Under these assumptions Proposition 1 does not apply, and the economists proceeded to compare the informational requirements in different economic settings by comparing the number of dimensions in each agent’s expression. In contrast, our work does not rely on such assumptions. In fact, one of our key points is that the dimensionality of the message space is not the essence of expressiveness. Rather, the essence is how the mechanism is wired to use the different inputs.

5.2 Work based on finding or characterizing equilibria

Another thread of related work has tried to characterize the equilibrium behavior in inexpressive mechanisms in specific settings. 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.

For example, Rosenthal and Wang [50] examined an auction setting where a series of globally interested (with nonlinear preferences over different items) and locally interested bidders (with linear preferences for different items) participate in a set of simultaneous first-price sealed-bid auctions where each auction is about a single item. Taken together, the auctions constitute an inexpressive mechanism. The authors were able to construct an equilibrium for each of two regions of the space of parameter values for the bidder type distributions in their model. They found that these equilibria were inefficient for most of their model parameter space. However, they were not able to rule out the possibility that other equilibria exist (although they have not found any) and they were unable to construct equilibria for some parameter values of their model.

Another example is work by Szentes and Rosenthal [61]. They characterized simple efficient equilibria in large inexpressive mechanisms when bidders are identical and each wants to win a specified fraction (more than a half) of the items. The simplicity of this domain illustrates the difficulty in finding equilibria in inexpressive mechanisms. Problems must typically be severely simplified in order to gain traction with analytical or computational techniques.

As further illustration of the difficulty of equilibrium finding, Wilenius and Andersson [64] described a heuristic method for computing approximate equilibrium strategies in first-price sealed bid CAs when bidders either bid on all combinations of items, or on one specific combination and the remaining items individually. They demonstrated the difficulty in finding equilibrium strategies for CAs when they are not dominant-strategy implementable.

All of the work discussed here suggests that there is little hope for a clear general characterization of equilibrium strategies in inexpressive mechanisms.

5.3 Expressiveness issues in dominant-strategy mechanisms

There has been some research related to expressiveness issues in dominant-strategy mechanisms.

For example, Blumrosen and Feldman [11] studied the problem of designing a dominant-strategy mechanism with a limited number of discrete actions. They showed a tradeoff between the efficiency of the best possible dominant-strategy mechanism and the number of discrete actions

available to the designer. Similarly, Ronen [49] described methods for achieving near efficiency with limited bidding languages in dominant strategies.

Holzman *et al.* [29] studied CAs where bidders can only bid on restricted sets of bundles. (This is the restricted outcome setting mentioned in Section 2.) Their work shows that truthful bidding is a dominant strategy if and only if the restricted bundle set that agents can bid on forms a quasi-field (and VCG payments are used). They defined a worst-case measure of the economic inefficiency that may result from restricting bids to smaller and smaller quasi-fields. Parkes [48] and Nisan and Segal [45] showed that in order to implement VCG payments, a mechanism must elicit enough information to verify the corresponding universal competitive equilibrium prices.

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 not have dominant strategies (and to settings where they do). Also, our results are not specific to any application, such as a CA.

5.4 Applications of expressiveness in mechanisms

One of the first applications to benefit from expressiveness was strategic sourcing. Sandholm [55, 56] described how building more expressive mechanisms—that generalize both CAs and multi-attribute auctions—for supply chains has saved billions of dollars that would have been lost due to inefficiency. Success with expressive auctions in sourcing has also been reported by others [28, 39, 19]. Schoenherr and Marbert [59] discussed the difficulty faced by business-to-business auction participants in choosing bundles to put up for auction ahead of time. This is a problem that exists because these mechanisms are typically inexpressive: they allow bids on predetermined lots only. If a CA were used instead, the sellers would not have to choose bundles *a priori*: the mechanism would determine the bundles based on the (expressive) bids.

Some work on expressiveness has begun to appear in the context of search keyword auctions (aka sponsored search). Even-Dar, Kearns and Wortman examined an extension of sponsored search auctions, whereby bidders can purchase keywords associated with specific contexts [21]. Under certain probabilistic assumptions they are able to prove that the system becomes more efficient when this extra level of expressiveness is allowed. In a working paper, Milgrom explores the equilibria of sponsored search auctions with limited expressive power (specifically, where bidders submit a single bid to indicate how much they will pay for an ad spot regardless of where it appears on the page) [40]. He finds that by *limiting* expressiveness the auction excludes some bad equilibria. This raises an important counterpoint to our work. We hope that our framework will help us better understand the circumstances under which expressiveness actually helps and when it does not. In another recent paper on sponsored search auctions, Abrams *et. al.* studied the impact of inexpressive bids on efficiency [1]. They found that in a specific auction mechanism, inexpressiveness can lead to an arbitrary amount of inefficiency when all bidders are assumed to play the same pure strategy (regardless of what the strategy is). They proceed to show that the same inexpressive mechanism has an efficient *full information* Nash equilibrium even when bidder valuations are

more complex. They consider this surprising, but it is consistent with our general result that very little expressiveness is needed for efficiency when agents have no uncertainty (Proposition 9).

Another application area that has received recent attention with regard to expressiveness is wireless spectrum trading. For example, Gandhi *et al.* [22] described a prototype wireless spectrum market mechanism. They stressed the importance of allowing spectrum bidders enough expressiveness to communicate their needs, and demonstrated—using synthetic demand distributions and various *ad hoc* bidder behavior models—that their mechanism has good efficiency properties.

The concept of expressiveness has been studied in single-agent applications as well. For example, results from recent studies of user security and privacy policies showed that, in many cases, these policies can be extremely rich and that it is unrealistic to expect users to fully specify them (e.g., [53, 18]). Tradeoffs between expressiveness and simple ease of use are therefore important as well.

5.5 A specific related sub-literature: bundle pricing

There is an extensive literature on bundle pricing. Allowing a seller to price bundles, rather than just individual items, can be seen as increasing the seller’s expressiveness. This is also related to our work on expressiveness. In this subsection we will briefly review some of the bundling literature.

The first mention of being able to increase revenue via bundling is attributed to Stigler in his 1963 discussion of anti-trust Supreme Court rulings over price discrimination via bundling [60]. Bundle pricing in economics has often focused on analyzing two-product settings to provide insight into the way monopolies can improve profits by offering goods in bundles [2, 20, 24, 38, 58]. (One exception is that Armstrong examined n -product settings, but placed severe restrictions on buyers’ utility functions [5].) This work provided sufficient conditions on when bundling is profitable and optimal pricing strategies under various assumptions. However, it did not provide generalized algorithms for determining how to price the bundles. Nor did it typically answer the question of how the increase in expressiveness affects the buyers utility or the efficiency of the market as a whole. There have also been some human subject experiments that explored how people actually perceive savings in bundles [65].

Some work on bundle pricing has been done from an operations research perspective as well. For example, Hason and Martin [26] presented a mixed integer program for optimizing bundle prices for a handful of market segments. They assumed that each of the segments can be described by a single value for each bundle, and that the value of every bundle for every market segment is known in advance. They also did not describe how their bundle pricing strategy compared to using item prices. Rusmevichientong *et al.* investigated the problem of pricing different car configurations based on data collected by GM’s Auto Choice Advisor web site [51].

There has also been work on pricing bundles of information goods, where it is usually assumed that customers care only about how many goods are bundled together (i.e., their valuation for a bundle depends only on its size, not its contents). For example, Kephart *et al.* [35] and Brooks and Durfee [12] described online approaches to pricing in this domain. Additionally, Bakos and Brynjolfsson provided an analytical treatment of this problem with some valuable insights about when bundling is profitable [6].

Finally, computer science work on pricing has focused primarily on pricing items rather than bundles, and for “single-minded” customers that desire only one bundle. For example, Balcan and Blum [7] provided online and approximate algorithms for this setting, and Guruswami et. al. [25] showed that finding the optimal pricing is \mathcal{APX} -Hard. Some work from this community, such as the work by Aggarwal *et al.* [3], considered a more restrictive class of pricing problems called MAX-BUYING, where customers buy the most expensive goods they can afford. Such restricted classes have been shown to be solvable in polynomial time.

Related to bundle pricing, there has recently also been significant work on designing high-revenue CAs (e.g., [47, 16, 36, 37, 34]). Designing for revenue turns out to be much more difficult than designing for efficiency.

6 Conclusions and future research

A recent trend 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 opinions. In this paper 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 perhaps the most important property of our domain-independent expressiveness notions: how they relate to the efficiency of the mechanism’s outcome. We derived an upper bound on the expected efficiency of a mechanism’s most efficient Nash equilibrium which depends only on the extent to which agents can impact the mechanism’s outcome. This bound enables us to study the relationship between expressiveness and efficiency by avoiding two major classic hurdles: 1) our bound can be analyzed without having to solve for an equilibrium of the mechanism, 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 instantiated our model of expressiveness for 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 for trading goods, information, and services—such as combinatorial auctions, exchanges, and multi-attribute auctions—can be cast as channel-based mechanisms. As a sanity check, 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 [17]). Using our general measures of expressiveness and our results on how they relate to efficiency, we were able to prove that in channel-based mechanisms 1) increasing expressiveness by adding a single

channel cannot decrease our upper bound on expected efficiency for the mechanism, and 2) under some preference distributions this leads to an arbitrarily large increase in the bound.

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 in the design of new mechanisms, whether the design is done by hand or by computer.

References

- [1] Zoe Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. In *Proceedings of Workshop on Internet Economics (WINE)*, 2007.
- [2] William James Adams and Janet L. Yellen. Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, 90(3):475–498, Aug 1976.
- [3] Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for multi-product pricing. In *ICALP*, pages 72–83, 2004.
- [4] Chris Anderson. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, July 2006.
- [5] Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, Jan 1996.
- [6] Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, Dec 1999.
- [7] Maria-Florina Balcan and Avrim Blum. Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 29–35, New York, NY, USA, 2006. ACM Press.
- [8] Shai Ben-David, Nicolo Cesa-Bianchi, David Haussler, and Philip M. Long. Characterizations of learnability for classes of $\{0, \dots, n\}$ -valued functions. *Journal of Computer and Systems Sciences*, 50(1):74–86, 1995.
- [9] B. Douglas Bernheim and Michael D. Whinston. Menu auctions, resource allocation, and economic influence. *Quarterly Journal of Economics*, 101(1):1–32, 1986.
- [10] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):865–929, 1989.
- [11] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 62–71, New York, NY, USA, 2006. ACM Press.

- [12] Christopher H. Brooks and Edmund H. Durfee. Toward automated pricing and bundling of information goods. In *AAAI Workshop on Knowledge-based Electronic Markets*, pages 8–13, 2000.
- [13] Georg Cantor. Über ein elementare Frage der Mannigfaltigkeitslehre. *German Mathematical Union*, I:75–88, 1890.
- [14] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [15] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, Edmonton, Canada, 2002.
- [16] Vincent Conitzer and Tuomas Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 132–141, New York, NY, 2004.
- [17] Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial auctions with k -wise dependent valuations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 248–254, Pittsburgh, PA, 2005. Draft in Oct., 2003.
- [18] Jason Cornwell, Ian Fette, Gary Hsieh, Madhu Prabhaker, Jinghai Rao, Karen Tang, Kami Vaniea, Lujio Bauer, Lorrie Cranor, Jason Hong, Bruce McLaren, Mike Reiter, and Norman Sadeh. User-controllable security and privacy for pervasive computing. In *Proceedings of the 8th IEEE Workshop on Mobile Computing Systems and Applications*, 2007.
- [19] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
- [20] Robert E. Dansby and Cecilia Conrad. Commodity bundling. *The American Economic Review*, 74(2):377–381, May 1984.
- [21] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Workshop on Internet Economics (WINE)*, 2007.
- [22] Sorabh Gandhi, Chiranjeev Buragohain, Lili Cao, Haitao Zheng, and Subhash Suri. A general framework for clearing auctions of wireless spectrum. In *IEEE DySPAN*, 2007.
- [23] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [24] Joseph P. Gaultinan. The price bundling of services: A normative framework. *Journal of Marketing*, 51(2):74–85, Apr 1987.
- [25] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

- [26] Ward Hanson and R. Kipp Martin. Optimal bundle pricing. *Manage. Sci.*, 36(2):155–174, 1990.
- [27] Friedrich Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.
- [28] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.
- [29] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, April 2004.
- [30] Leonid Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*. Amsterdam: North Holland, 1972.
- [31] Leonid Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*, chapter 14, pages 297–336. University of Minnesota Press, 1986. Second edition.
- [32] Leonid Hurwicz and Thomas Marschak. Discrete allocation mechanisms: Dimensional requirements for resource-allocation mechanisms when desired outcomes are unbounded. *Journal of Complexity*, 1(2):264–303, 1985.
- [33] Philippe Jehiel and Benny Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, Sept 2001.
- [34] Philippe Jehiel, Moritz Meyer ter Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 127(1):494–512, May 2007.
- [35] Jeffrey O. Kephart, Christopher H. Brooks, Rajarshi Das, Jeffrey K. MacKie-Mason, Robert Gazzale, and Edmund H. Durfee. Pricing information bundles in a dynamic environment. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 180–190, New York, NY, USA, 2001. ACM Press.
- [36] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 232–237, San Jose, CA, 2004.
- [37] Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, 2005.
- [38] R. Preston McAfee, John McMillan, and Michael D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383, May 1989.

- [39] Theresa Metty, Rob Harlan, Quentin Samelson, Tom Moore, Thomas Morris, Ron Sorensen, Avner Schneur, Olga Raskina, Rina Schneur, Joshua Kanner, Kevin Potts, and Jeffrey Robbins. Reinventing the supplier negotiation process at Motorola. *Interfaces*, 35(1):7–23, 2005.
- [40] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions. Working paper, 2007.
- [41] Tom Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [42] Kenneth Mount and Stanley Reiter. The informational size of message spaces. *Journal of Economic Theory*, 8(2):161–192, 1974.
- [43] Victor Naroditskiy and Amy Greenwald. Using iterated best-response to find bayes-nash equilibria in auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, Canada, 2007. Poster.
- [44] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.
- [45] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.
- [46] Shmuel Oren and Michael Rothkopf. Optimal bidding in sequential auctions. *Operations Research*, 23:1080–1090, 1975.
- [47] Thomas R. Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2):463–483, Mar 1983.
- [48] David C. Parkes. Price-based information certificates for minimal-revelation combinatorial auctions. In *AAMAS '02: Revised Papers from the Workshop on Agent Mediated Electronic Commerce on Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems*, pages 103–122, London, UK, 2002. Springer-Verlag.
- [49] Amir Ronen. Mechanism design with incomplete languages. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 105–114, 2001.
- [50] Robert W. Rosenthal and Ruqu Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behavior*, 17:32–55, 1996.
- [51] Paat Rusmevichientong, Benjamin Van Roy, and Peter W. Glynn. A nonparametric approach to multiproduct pricing. *Oper. Res.*, 54(1):82–98, 2006.
- [52] Norman Sadeh, Fabien Gandon, and Oh Buyng Kwon. Ambient intelligence: The MyCampus experience. In T. Vasilakos and W. Pedrycz, editors, *Ambient Intelligence and Pervasive Computing*. ArTech House, 2006. Also available as Carnegie Mellon University Technical Report CMU-ISRI-05-123.

- [53] Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. Understanding and capturing people’s privacy policies in a people finder application. In *Proceedings of the Ubicomp Workshop on Privacy*, 2007.
- [54] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002. Earlier versions: ICE-98 keynote, Washington U. tech report WUCS-99-01 Jan. 1999, IJCAI-99.
- [55] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
- [56] Tuomas Sandholm, David Levine, Michael Concordia, Paul Martyn, Rick Hughes, Jim Jacobs, and Dennis Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.
- [57] Tuomas Sandholm and Subhash Suri. Side constraints and non-price attributes in markets. *Games and Economic Behavior*, 55:321–330, 2006. Extended version in IJCAI-2001 Workshop on Distributed Constraint Reasoning.
- [58] Richard Schmalensee. Gaussian demand and commodity bundling. *The Journal of Business*, 57(1):S211–S230, Jan 1984.
- [59] Tobias Schoenherr and Vincent A. Marbert. The use of bundling in B2B online reverse auctions. *Journal of Operations Management*, 2008. Forthcoming.
- [60] George J. Stigler. United States v. Loew’s Inc.: A note on block-booking. *The Supreme Court Review*, 1963:152–157, 1963.
- [61] Balazs Szentes and Robert W. Rosenthal. Beyond chopsticks: Symmetric equilibria in majority auction games. *Games and Economic Behavior*, 45:278–295, 2003.
- [62] Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [63] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [64] Jim Wilenius and Arne Andersson. Discovering equilibrium strategies for a combinatorial first price auction. In *Proceedings of IEEE Joint Conference on E-Commerce Technology (CEC’07) and Enterprise Computing, E-Commerce and E-Services (EEE’07)*, 2007.
- [65] Manjit S. Yadav and Kent B. Monroe. How buyers perceive savings in a bundle price: An examination of a bundle’s transaction value. *Journal of Marketing Research*, 30(3):350–358, Aug 1993.

[66] Weili Zhu and Peter Wurman. Structural leverage and fictitious play in sequential auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Edmonton, Canada, 2002.

7 Appendix

Proof of Proposition 1. Given a mechanism with reportable type space in \mathcal{R}^d we can construct an equivalent mechanism with reportable type space \mathcal{R} by constructing an injective mapping from \mathcal{R}^d to \mathcal{R} . Then, when an agent makes a report in \mathcal{R} , we use the reverse mapping and act as if the agent had expressed the corresponding point in \mathcal{R}^d in the original mechanism.

One way to construct the injective mapping is as follows. Let σ_i^j be the i th bit (or digit) of the real number that the agent expresses for dimension $j \in \{1, 2, \dots, n\}$. Let p_k be the k th prime number. Our desired number in \mathcal{R} is

$$\prod_i \prod_j (p_{(i-1)n+j})^{\sigma_i^j}$$

□

Proof of Proposition 2. This follows trivially from the fact that both mechanisms have welfare maximizing truthful dominant-strategy equilibria and for any particular group of participants the welfare maximizing alternative in the larger set has efficiency equal to or greater than the welfare maximizing alternative in the smaller set.

Let $M = \langle f, \pi \rangle$ be the mechanism with the larger outcome space and $M' = \langle f', \pi' \rangle$ be the one with the smaller outcome space. Equation 1 describes the expected efficiency of a mechanism, thus we wish to show that for population with types drawn from any distribution $P(T^n)$ with any utility functions the following inequality holds,

$$\begin{aligned} \mathcal{E}(f, \pi) &= \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(b(t^n, f, \pi) = \theta^n) \sum_i u_i(t_i, f(\theta^n)) \geq \\ &\quad \mathcal{E}(f', \pi') = \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(b(t^n, f', \pi') = \theta^n) \sum_i u_i(t_i, f'(\theta^n)) \end{aligned}$$

We know that the equilibria of the two mechanisms are truthful since they are both VCGs therefore the above inequality simplifies to,

$$\int_{t^n \in T^n} P(T^n = t^n) \sum_i u_i(t_i, f(t^n)) \geq \int_{t^n \in T^n} P(T^n = t^n) \sum_i u_i(t_i, f'(t^n))$$

Since both mechanisms are choosing welfare maximizing outcomes and $\mathcal{O}' \subseteq \mathcal{O}$ for any particular type vector, t^n , for the agents we have,

$$\forall t^n \quad \sum_i u_i(t_i, f(t^n)) \geq \sum_i u_i(t_i, f'(t^n))$$

□

Proof of Proposition 3. Every time we allow agent i to (semi-)shatter a new outcome the most expressive mechanism allows the agent to distinguish among all of the impact vectors it had previously distinguished between plus one additional impact vector (the impact vector that was preventing it from (semi-)shattering that outcome). □

Proof of Proposition 4. This is fairly straight forward. The number of impact vectors over $|T_{-i}|$ involving k different outcomes is $|T_{-i}|^k$. Shattering requires that an agent be able to distinguish among each of these vectors, thus its maximum impact dimension must be greater than or equal to this amount. □

Proof of Proposition 5. If there exists a pair of types, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, that agent i cannot (semi-)shatter over, then there is at least one (un-ordered in the case of semi-shattering) pair of outcomes, A and B , that agent i cannot force the mechanism to choose when the other agents have types $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$. This means that agent i cannot express any impact vector where A and B (in either order for semi-shattering) happen against $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$ (i.e., where $g_i(t_{-i}^{(1)}) = A$ and $g_i(t_{-i}^{(2)}) = B$). Without being able to express these preference vectors agent i cannot fully (semi-)shatter outcome space \mathcal{O} . □

Proof of Proposition 6. The following reasoning demonstrates that Equation 2 is a valid upper bound on the maximum attainable expected efficiency by any mechanism using the outcome function f in equilibrium:

$$\begin{aligned} E_{t^n} [\mathcal{E}(f, \pi)]^+ &= \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(b(t^n, f, \pi) = \theta^n) W(t^n, f(\theta^n)) \\ &\leq \max_{B(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(B(t^n) = \theta^n) W(t^n, f(\theta^n)) \\ &= \max_{\hat{B}(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) W(t^n, f(\hat{B}(t^n))) \\ &= \max_{\hat{B}(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) W(t^n, f(\{\hat{B}_1(t_1), \dots, \hat{B}_n(t_n)\})) \end{aligned}$$

The step between the second and third equations follows from the fact that one of the maxima of the function in the second equation must have each entry of $B(\cdot)$ (a function that maps every type vector to a mixed strategy profile) as a point mass. This is because there is at least one single pure strategy combination for each type vector that leads to the outcome with highest welfare, so there is no reason to consider mixed strategies in this bound. The last step is valid because the strategy of each agent can depend only on its own private type. □

Proof of Proposition 7. First we will prove the forward implication, namely that the upper bound reaches full efficiency if any agent i can distinguish among each of the impact vectors in at least one of its fully efficient sets.

The fact that some agent i can distinguish among each of the impact vectors in some fully efficient set, G_i^* , implies that there is a pure strategy for agent i , h_i , which is a mapping from its types to expressions, and a pure strategy profile for the agents other than i , h_{-i} , mapping from each of their types to expressions that causes the most efficient outcome to be chosen by the mechanism for every possible combination of types. If we set $\hat{B}(t^n) = \{h_i(t_i), h_{-i}(t_{-i})\}$ then $E[\mathcal{E}(f)]^+$ will reach full efficiency.

Now we will prove the backwards implication, namely that if any agent i cannot distinguish among each of the impact vectors in at least one of its fully efficient sets then the upper bound cannot be fully efficient.

Let agent i be an agent that cannot distinguish among each of its impact vectors in any of its fully efficient sets. Consider any set of impact vectors that agent i can distinguish among, G_j . Based on the predicate of the proposition, at least one of the impact vectors, g_i^* corresponding to fully efficient outcomes when agent i has type t_i^* , in any fully efficient set cannot be expressed by agent i .

This means that no matter what strategies the agents other than i choose, agent i will not be able to express some g_i^* and at least one of the outcomes chosen by the mechanism when agent i has type t_i^* will be less than fully efficient. \square

Proof of Proposition 8. This proof is relatively straightforward. We know that the predicate implies there is some pure strategy for agent i , h_i that achieves full efficiency when played against some pure strategy profile, h_{-i} for the other agents. Let h_j be agent j 's pure strategy in the profile h_{-i} . Construct a new pure strategy profile, h_{-j} , by starting with h_{-i} and removing agent j 's pure strategy. Now add agent i 's pure strategy h_i to complete the profile. Since we have not changed the strategies played in any circumstances h_j will achieve full efficiency against h_{-j} , thus completing our proof. \square

Proof of Proposition 9. In these settings, as soon as agent i knows its own type it knows for certain the single most efficient outcome. It never needs to distinguish between more than one-dimensional preference vectors and there are only $|\mathcal{O}|$ such vectors. \square

Proof of Corollary 1. This follows directly from Proposition 9 and Proposition 7 \square

Proof of Theorem 1. The set of mechanisms allowing agent i maximum preference dimension d_i is a super-set of the mechanisms allowing agent i maximum preference dimension $d_i' < d_i$. Thus the fact that the bound for the best mechanism increases weakly monotonically is trivially true for any increase in d_i . The challenge is proving the strictness of the monotonicity.

Consider increasing d_i from $d_i^{(1)} < d_i^*$ to $d_i^{(2)} > d_i^{(1)}$. Let $G_i^{(1)}$ be the best set of preference vectors that agent i distinguishes between when restricted to $d_i^{(1)}$ vectors (i.e., the set of preference vectors that maximize the upper bound on expected efficiency). We know that there are at least $d_i^* - d_i^{(1)} \geq 1$ preference vectors corresponding to fully efficient sets of outcomes with non-zero probability that cannot be expressed by agent i , and thus at least that many preference vectors

corresponding to fully efficient sets of outcomes that are absent from $G_i^{(1)}$. When we increase our expressiveness limit from $d_i^{(1)}$ to $d_i^{(2)}$, we can add one of those missing vectors to $G_i^{(1)}$ to get $G_i^{(2)}$. Since $G_i^{(2)}$ allows agent i to distinguish among all the same vectors as $G_i^{(1)}$ and an additional vector which corresponds to a fully efficient set of outcomes (recall that these outcomes must be strictly more efficient), the new mechanism with maximum preference dimension $d_i^{(2)}$ has a strictly higher expected efficiency bound. \square

Proof of Corollary 2. This follows directly from Theorem 1 and Proposition 3. \square

Proof of Lemma 1. Start with any number of outcomes and any number of types for the agents other than i with equal likelihood (and let the probability of any particular set of types for the agents other than i be independent of i 's type). Choose a set, G_i , of unique impact vectors for agent i with size d_i . Construct one non-zero probability type for agent i for each impact vector in G_i , $t_i^{g^{(i)}}$. Set the total welfare of all agents to an arbitrarily large number for every combination of joint types according to the impact vectors corresponding to i 's type, (in an interdependent values setting there are no restrictions on the agent's utility functions parametrized on the full joint type space, so the welfare function for each set of joint types can be constructed arbitrarily):

$$\forall g_i \in G_i, \forall t_{-i}, \quad W(\{t_i^{g_i}, t_{-i}\}, g_i(t_{-i})) = M$$

If agent i cannot distinguish among all of the d_i impact vectors then the efficiency bound will be arbitrarily smaller than if it can. Thus for the best outcome function the move from $d_i - 1$ to d_i necessarily results in an arbitrary increase. \square

Proof of Lemma 2. The part that applies to the interdependent values setting follows directly from Lemma 1 since decreasing k_i by one also decreases d_i by at least 1.

Next we will prove the implication in the private value setting. To prove this we will construct a setting (i.e., utilities, types and outcomes), such that agent i must be able to semi-shatter an outcome space of size k_i in order to avoid our upper bound being arbitrarily lower than full efficiency. Our constructed setting can have any number of outcomes, any number of other agents and any number of joint types for the other agents. However, in order to assign the total utility of the other agents for each of their joint types in an arbitrary way, we will limit every other agent except for one, agent j , to a single type (agent j will have $|T_{-i}|$ types). We will set the utility of every agent other than i and j to 0 in all circumstances and build our construction using only these two agents.

We will start with a set of outcomes \mathcal{O}' that has size k_i (if $k_i = 1$ the rest of this proof is trivial, if every single outcome provides an arbitrary amount of welfare then not being able to make any one of them happen will lead to arbitrary inefficiency). We will assume the outcomes in \mathcal{O}' are the only outcomes that any of the agents derive any utility from. We will assume that there is some strict ordering on the outcomes from o_1 to o_{k_i} and on agent j 's types from $t_j^{(1)}$ to $t_j^{(|T_j|)}$. We will now set the utility of agent j for every outcome under every one of its types. (Recall that in a private value setting the utility of the agents other than i cannot depend on i 's type, and vice versa).

Our construction sets agent j 's utility for outcome o_m under each of its types to be arbitrarily larger than for the outcome preceding it in the strict ordering, o_{m-1} (with the first outcome always leading to utility 0). Under a single type, all of the gaps between successive outcomes will be the

same size, however this gap amount will increase by an arbitrary amount for each successive type. This will result in agent j 's utility under each of its types being a step function over the strictly ordered outcomes in \mathcal{O}' , with the step sizes increasing under each successive type. Formally we will set agent j 's utility function in the following way (let M be an arbitrarily large number),

$$(\forall m, \forall l) \quad u_j(t_j^{(m)}, o_l) = (l - 1 \times ((m - 1) \times 2 \times M))$$

Now for each of the $\binom{|\mathcal{O}'|}{2}$ un-ordered pairs of outcomes, o_a and o_b (where a is always before b in our strict ordering), we will construct a set of $|T_j|$ types for agent i , which we will call $T_i^{(a,b)}$. Agent i 's utility under all of the types in $T_i^{(a,b)}$ will be hugely negative for all outcomes other than o_a and o_b (note that this value does not have to be negative infinity, it just has to be arbitrarily lower than the total welfare of any outcome under any circumstance), thus causing an arbitrary loss of efficiency if either of these outcomes is not chosen. Again, we will assume a strict ordering on the types in $T_i^{(a,b)}$, from 1 to $|T_j|$. Agent i 's utility for o_b under each of these types will be set to the arbitrarily large number M , and for o_a (the typically less preferred outcome by agent j , since it comes earlier in the ordering) will be set to successively increasing multiples of the distance between the outcomes in the strict ordering times twice the arbitrarily large number used above, $(b - a) \times 2 \times M$. In other words, o_a will provide successively more utility to agent i as its type from the pair selecting set increases from 1 to $|T_j|$. Formally we will set agent i 's utility under the types in $T_i^{(a,b)}$ to be the following,

$$\begin{aligned} (\forall m \mid t_i^{(m)} \in T_i^{(a,b)}) \quad & u_i(t_i^{(m)}, o_b) = M \\ (\forall m \mid t_i^{(m)} \in T_i^{(a,b)}) \quad & u_i(t_i^{(m)}, o_a) = (m - 1) \times (b - a) \times 2 \times M \\ (\forall o_j \in \mathcal{O} \setminus \mathcal{O}', \forall m \mid t_i^{(m)} \in T_i^{(a,b)}) \quad & u_j(t_i^{(m)}, o_j) = -\infty \end{aligned}$$

When $t_i^{(m)}$ is matched with $t_j^{(m)}$ the total welfare of outcome o_b will be at least M larger than the total welfare of o_a . However, for all of j 's types smaller than m the opposite will be true.

$$\begin{aligned} W(\{t_i^{(m)}, t_j^{(m)}\}, o_b) &= M + [(b - 1) \times (m - 1) \times 2 \times M] \\ W(\{t_i^{(m)}, t_j^{(m)}\}, o_a) &= [(b - a) \times (m - 1) \times 2 \times M] + [(a - 1) \times (m - 1) \times 2 \times M] \\ &= [(b - 1) \times (m - 1) \times 2 \times M] \end{aligned}$$

By constructing the utility functions in this way we have guaranteed that for any pair of agent j 's types, $t_j^{(m)}$ and $t_j^{(m')}$ (where $m < m'$ in our strict ordering), there is a type for agent i requiring o_b to happen against $t_j^{(m)}$ and o_a against $t_j^{(m')}$ to avoid an arbitrary loss in efficiency (because the second best outcome always leads to at least M less welfare).

Now we can simply repeat this process for every pair of outcomes in \mathcal{O}' by constructing types for agent i that select that pair. We can also construct one type for agent i for each outcome in \mathcal{O}' ,

where it prefers that outcome hugely more than any other outcome. This will guarantee that agent i must be able to make every pair of outcomes happen against every pair of agent j 's types, and must be able to make every single outcome happen against every pair of agent j 's types, in order to avoid an arbitrary loss of efficiency in some non-zero probability combination of types. This is equivalent to saying that agent i must be able to semi-shatter the outcome space \mathcal{O}' in order to avoid an arbitrary decrease in the expected efficiency bound. \square

Proof of Lemma 3. Let agent i 's utility for outcomes o_1 and o_2 under type $t_i^{(1)}$ and be denoted as X and Y . For the agents other than i , let the sum of their utilities for the outcomes o_1 and o_2 under types $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, be denoted as, a and b , and, a' and b' , respectively. We wish to show that the ordering on efficient outcomes imposed by this collection of types cannot be reversed. Formally,

$$(X + a > Y + b) \text{ and } (Y + b' > X + a') \Rightarrow \\ \neg(\exists X', Y') (X' + a < Y' + b) \text{ and } (Y' + b' < X' + a')$$

We will proceed by assuming this is true, namely that there exists an X' and Y' that satisfy the second set of inequalities, and show that it leads to a contradiction. If all of the inequalities held we would have the following,

$$b - a < X - Y < b' - a' \\ b' - a' < X' - Y' < b - a$$

which leads to a contradiction. \square

Proof of Theorem 2. The forward implication in both settings follows directly from Lemma 2. The backward implication in the interdependent values setting follows from Lemma 1 and Proposition 7 (since there will always be a fully efficient set that contains every possible impact vector). In the private value setting the backward implication is implied by Lemma 3, since it proves that it is never necessary for full efficiency in this setting to shatter any pair of outcomes (only semi-shatter them). \square

Proof of Proposition 10. This proof is based on a pigeon hole argument. With fewer than $\lceil \log_2(|\mathcal{O}|) \rceil$ channels there will be at least 2 outcomes connected to the exact same set of channels. If agent i has C_i channels then it has 2^{C_i} sets of channels. When C_i is small the number of sets of channels will be less than the number of outcomes.

$$C_i < \lceil \log_2(|\mathcal{O}|) \rceil \Rightarrow 2^{C_i} < |\mathcal{O}|$$

This will prevent the agent from forcing the mechanism to choose both of those outcomes against with different plays since the agent's own contribution to the two outcomes will always be identical. \square

Proof of Proposition 11. With that many channels we can construct a VCG outcome function in the following manner. For each agent i , connect each of i 's channels to a different outcome, leaving one outcome with no channel from that agent. The agent then reports its utility under each outcome relative to the outcome with no channels. The mechanism chooses outcome whose channels have the largest sum, which is equivalent to choosing the welfare-maximizing outcome. The payment rule will not be affected by the fact that each agent is reporting its utility relative to a particular outcome. To see this consider the VCG (i.e., Clarke tax) payment of any agent i . This payment is equal to the total difference in utility of the other agents, had agent i not participated. Let the outcome with agent i in the mechanism be A and the outcome without agent i be B . Let the outcome with no channels attached be o_j for every agent j . Then we have the payment for agent i as,

$$\begin{aligned}
\pi_i &= \sum_j (u_j(t_j, A) - u_j(t_j, o_j)) - \sum_j (u_j(t_j, B) - u_j(t_j, o_j)) \\
&= \sum_j (u_j(t_j, A) - u_j(t_j, B)) - \sum_j (u_j(t_j, o_j) - u_j(t_j, o_j)) \\
&= \sum_j (u_j(t_j, A) - u_j(t_j, B))
\end{aligned}$$

Since the $u_j(t_j, o_j)$ terms drop out of this equation, having every agent report their utility for every outcome minus their utility for one particular outcome does not effect the payment calculation. This shows that the payment rule can be properly calculated even when each agent is left with a single outcome with no channels.

Using a pigeon hole argument we can see that an agent with fewer than $|\mathcal{O}| - 1$ channels will either have at least 2 outcomes sharing a channel, making it impossible for that agent to express arbitrary non-linear utility for every outcome (something that is required in order to implement a VCG), or it will have 2 outcomes without a channel, making it impossible for that agent to express any preference for one of the outcomes (if the agent had only one outcome with no channel, then it could express its preferences relative to that outcome, as described above). \square

Proof of Proposition 12. We will show that no agent can shatter any set of 2 outcomes against any 2 types, even when it has a channel dedicated solely to each of the two outcomes (so that it can place an arbitrary amount of value on either outcome). This implies that it is impossible for any larger set of outcomes or types in any channel-based mechanism.

We will assume for contradiction that there is some agent i that can shatter a pair of outcomes A and B in a channel-based mechanism. Let agent i 's channel value connected to outcome A be X and let its channel value connected to B be Y . Consider two types for the agents other than i , $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, and the reports mapped to them in *any* pure strategy, $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$. Let the sum of the reports by the other agents on the channels connected to A be denoted a_1 and a_2 under the first and second expressions, respectively. Likewise let b_1 and b_2 be the sum of the reports on B . We have assumed (for contradiction) that there exists an X, Y, X' and Y' that satisfy the following

inequalities,

$$A \text{ against } 1, B \text{ against } 2 \begin{cases} X + a_1 > Y + b_1 \\ Y + a_2 > X + b_2 \end{cases}$$

$$B \text{ against } 1, A \text{ against } 2 \begin{cases} Y' + b_1 > X' + a_1 \\ X' + a_2 > Y' + b_2 \end{cases}$$

This leads directly to the contradiction,

$$\begin{aligned} b_1 - a_1 &< X - Y < b_2 - a_2 \\ b_2 - a_2 &< X' - Y' < b_1 - a_1 \end{aligned}$$

□

Proof of Corollary 3. This follows directly from Proposition 12 and Lemmas 1 and 2. □

Proof of Theorem 3. We will first present a lemma regarding an implication of the predicate based on set algebra.

Lemma 4. *For any sets, $A, B, C,$ and $D,$ the following bi-directional implication holds,*

$$(A \setminus C = B \setminus D) \text{ and } (C \setminus A = D \setminus B) \Leftrightarrow (A \setminus D = C \setminus B) \text{ and } (D \setminus A = B \setminus C)$$

Proof. We will prove the forward implication, once that is proved the backward implication is trivial since we can just switch the labels of C and D . From the predicate we know that the only part of A that is not in D must be contained completely in C (since $(A \setminus C) \subseteq D$), in particular we know that,

$$\begin{aligned} A \setminus D &= C \setminus (C \setminus A) \\ &= C \setminus (B \setminus D) \\ &= C \setminus B \end{aligned}$$

The last step is valid because we know that no elements from D can be in the set on the right hand side (since we are removing them from A). Thus it cannot make a difference if we leave them in B before subtracting it from C . This same logic can be repeated for the other side,

$$\begin{aligned} D \setminus A &= B \setminus (B \setminus D) \\ &= B \setminus (C \setminus A) \\ &= B \setminus C \end{aligned}$$

□

From Lemma 4 in addition to our predicate we know that the following must also be true (we drop the i subscript on the channel sets for shorthand, since all sets of channels discussed in this proof belong to agent i),

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

Now we will assume for contradiction that agent i can semi-shatter both pairs of outcomes, $\{A, B\}$ and $\{C, D\}$. From Proposition 5, we know that in order for i to be able to semi-shatter a set of outcomes, it must be able to semi-shatter it for any *pair* of types of the other agents. Thus, there must be at least one pair of reports by the agents other than i , $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, such that agent i can cause all four outcomes to happen (although we are dealing with semi-shattering so the order in which they happen does not matter). 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.

Lets assume (without loss of generality) that $b_1 - a_1 < b_2 - a_2$ and that A will happen against $\theta_{-i}^{(1)}$ and B will happen against $\theta_{-i}^{(2)}$ (if the inequality does not hold, we can reverse the labels on the θ_{-i} 's). In order to cause A to happen against the first opponent profile and B against the second the following inequalities must hold (from here on we use the shorthand S^A to denote the sum of agent i 's report on the channels in S^A , and we assume that ties are broken consistently so that an agent cannot use them to semi-shatter),

$$\begin{aligned} A \text{ happens against 1} & \begin{cases} S^A + a_1 > S^B + b_1 \\ S^A + a_1 > S^C + c_1 \\ S^A + a_1 > S^D + d_1 \end{cases} \\ B \text{ happens against 2} & \begin{cases} S^B + b_2 > S^A + a_2 \\ S^B + b_2 > S^C + c_2 \\ S^B + b_2 > S^D + d_2 \end{cases} \end{aligned}$$

Now let the difference between the sum of channels in $S^A - S^C = S_1$, and notice from the predicate that $S^D - S^B = S_1$. This is because the channels that are in S^A and not S^C are the same as those that are in S^D and not S^B , and also the channels in S^C that are not in S^A are the same as those that are in S^B and not S^D . In addition, let the difference in the sum of the channels in $S^A - S^D = S^C - S^B = S_2$ (this equality is also implied by the predicate). Now the equations above simplify to,

$$\begin{aligned} b_1 - a_1 & < S^A - S^B < b_2 - a_2 \\ c_1 - a_1 & < S_1 < b_2 - d_2 \\ a_1 - d_1 & < S_2 < b_2 - c_2 \end{aligned}$$

In order to semi-shatter C and D with C happening against the first report by the other agents and D against the second we have the following inequalities generated in the same fashion,

$$\begin{aligned}
c_1 - d_1 &< S^C - S^D < c_2 - d_2 \\
b_2 - d_2 &< S_1 < c_1 - a_1 \\
b_1 - c_1 &< S_2 < a_2 - d_2
\end{aligned}$$

In order to semi-shatter over C and D in the opposite direction (with D first and C second) the constraints would change to the following,

$$\begin{aligned}
c_2 - d_2 &< S^C - S^D < c_1 - d_1 \\
b_1 - d_1 &< S_1 < c_2 - a_2 \\
b_2 - c_2 &< S_2 < a_1 - d_1
\end{aligned}$$

Now we can see that our assumption that we can semi-shatter both sets of outcomes under even a single pair of types leads to a contradiction since the following sets of constraints would have to be satisfied,

$$\begin{aligned}
c_1 - a_1 &< b_2 - d_2 \\
b_2 - d_2 &< c_1 - a_1
\end{aligned}$$

or,

$$\begin{aligned}
c_2 - b_2 &< a_1 - d_1 \\
a_1 - d_1 &< c_2 - b_2
\end{aligned}$$

□

Proof of Proposition 13. Let A be an outcome under which agent i is allocated bundle Q , let B be an outcome under which it is allocated q_1 , C for q_2 and D for nothing (also let S^A , S^B , S^C , and S^D be the sets of channels connected to those outcomes for agent i). Since agent i 's bid on Q equals the sum of its bid on q_1 and q_2 , we have that $S^A = S^B \cup S^C$ and its bid for the outcome where it wins nothing is always 0, so we have $S^D = \emptyset$. Notice that these sets of channels meet the conditions of Theorem 3,

$$\begin{aligned}
(S^A \setminus S^C = S^B \setminus S^D) &\quad \text{and} \quad (S^C \setminus S^A = S^D \setminus S^B) \\
((S^B \cup S^C) \setminus S^C = S^B \setminus \emptyset) &\quad \text{and} \quad (S^C \setminus (S^B \cup S^C) = \emptyset \setminus S^B)
\end{aligned}$$

□

Proof of Corollary 4. This follows trivially from Proposition 13 and Lemma 2. □

Proof of Proposition 14. We will prove this statement for the semi-shatterable outcome dimension, k_i , which will imply it is true for maximum impact dimension d_i as well (based on Proposition 3).

Consider any channel-based mechanism that assigns c_i channels to agent i , and allows it a semi-shatterable outcome dimension $k_i < |\mathcal{O}|$. We will assume from here on that $k_i \geq 2$, since if $k_i = 1$ the theorem is trivially true (we can build a fully expressive VCG mechanism over 2 outcomes with a single channel and thus adding a channel will definitely increase k_i to at least 2).

Let the largest set of outcomes that agent i can shatter over in this mechanism be \mathcal{O}' (if there are ties just choose one arbitrarily). Note that there is a non-empty set of outcomes missing from \mathcal{O}' , we will call that $\mathcal{O}^* = \mathcal{O} \setminus \mathcal{O}'$. Now consider adding one channel for agent i to the mechanism and connecting it to one of the outcomes $o^* \in \mathcal{O}^*$. Clearly the agent can still semi-shatter over \mathcal{O}' , since it can just ignore the new channel. However, it can now also semi-shatter a larger set, $\mathcal{O}' \cup \{o^*\}$.

To verify this notice that with the additional channel connected to o^* the agent can control the amount of utility it reports on this outcome arbitrarily (without affecting its reports on any other outcomes). Consider any pair of outcomes in the original set, $o'_1, o'_2 \in \mathcal{O}'$. Agent i can now make o^* happen against any type where either of those outcomes happened in the old mechanism by setting its report on the new channel to be ϵ greater than the sum of its reports on the channels connected to the outcome it chooses. Formally, if C_i is the channel mapping from the original mechanism, then we can translate any report in the old mechanism, θ_i , to a report in the new mechanism, θ_i^* , which causes o^* to happen whenever any o' did previously,

$$\begin{aligned} (\forall j \mid 1 \leq j \leq c_i) \theta_{i,j}^* &= \theta_{i,j} \\ \theta_{i,c_i+1}^* &= \sum_{j \in C_i(o')} \theta_{i,j} + \epsilon \end{aligned}$$

Since agent i can do that with both outcomes from the original semi-shatterable set we have confirmed that it has reports in the new mechanism that make o^* happen with every pair of outcomes in \mathcal{O}' (this is an inductive argument, since each of those outcomes had this property before)⁹ Thus agent i can semi-shatter the new larger outcome set. \square

Proof of Corollary 5. The fact that the bound is weakly monotonic is true because the extra channel can always be ignored. The fact that the increase can be arbitrarily large follows directly from Proposition 14 above and Lemma 2 (since increasing the number of channels by 1 can increase the agent's semi-shatterable outcome dimension). \square

⁹Note that we have assumed the agent was not using the tie-breaking properties of the original mechanism to shatter the outcomes. If this assumption does not hold, the proof is still valid as long as the mechanism always breaks ties consistently (i.e., when the channels connected to outcomes o_1 and o_2 have the same sum it always chooses either o_1 or o_2).