A Theory of Expressiveness in Mechanisms

Michael Benisch and Tuomas Sandholm

Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

We develop a theory that ties the expressiveness of mechanisms to their efficiency. We introduce two notions of expressiveness, *impact dimension* and *outcome shattering*. We show that an upper bound on the expected efficiency of any mechanism increases *strictly* as more expressiveness is allowed, and prove that a small increase in expressiveness can lead to an arbitrarily large increase in the bound. We also show that in any private values setting the bound can be implemented with a budget-balanced mechanism in Bayes-Nash equilibrium. However, without full expressiveness dominant-strategy implementation is not always possible. We then discuss the relationship between expressiveness and communication complexity, and conclude with a study of a mechanism class we call *channel based*, which subsume most combinatorial and multi-attribute allocation mechanisms and any Vickrey-Clarke-Groves mechanism. We show how our expressiveness notions can be used to characterize the often-studied exposure problem faced by participants in channel-based mechanisms that do not allow rich combinatorial expressions. Using our theoretical framework, we are able to prove that this problem can result in arbitrary inefficiency.

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

This work has been supported by a Siebel Scholarship, and the National Science Foundation under grants 9972762, 0205435, 0427858 and 0905390, and Cyber Trust grant CNS-0627513.

1 Introduction

A recent trend in the world—especially in electronic commerce—is a demand for higher levels of expressiveness in mechanisms that mediate interactions, such as the allocation of resources or matching of peers. The most famous expressive mechanism is a *combinatorial auction (CA)*, which allows participants to express valuations over *packages* of items [23, 58]. CAs have the recognized benefit of removing the *exposure problem*, which bidders face when they have preferences over packages but, in traditional auctions, are allowed to submit bids on individual items only. In such traditional inexpressive auctions, when bidding on an item, a given bidder needs to speculate how others will bid on other items, since this will affect which bundles the bidder can construct and at what prices. As we will show in this paper, uncertainty about the others' preferences does, in fact, increase the need for expressiveness, and, conversely, at a fixed level of expressiveness, greater uncertainty leads to reduced efficiency. CAs also have other acknowledged benefits, and preference expression forms significantly more compact and more natural than package bidding have been developed (e.g., [23, 33, 63, 65, 66]).

Expressiveness also plays a key role in *multi-attribute* settings where 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., [23, 63, 65, 66]). CAs, multi-attribute auctions, and generalizations thereof are used to trade tens of billions of dollars worth of items annually [23, 33, 47, 63, 65]. The trend toward more expressiveness is also reflected in the richness of preference expression now offered by businesses as diverse as retail sites, like Amazon and Netflix, and advertising services such as those in sponsored search auctions (e.g., [12, 26]) and display advertising mechanisms (e.g., [17, 73]).

Intuitively, one might think that more expressiveness leads to higher efficiency (sum of the agents' utilities)—for example, due to better matching of supply and demand. Efficiency improvements have indeed been reported from combinatorial multi-attribute auctions (e.g., [23, 62, 63, 65]) and expressive auctions for advertisement allocation [12, 17, 26]. However, until now we have lacked a general way of characterizing the expressiveness of different mechanisms, the impact the expressiveness has on the agents' strategies, and, thereby, the outcome of the mechanism. For example, it was not previously known whether there existed a single setting where more expressiveness could always be used to achieve a more efficient outcome. In fact, on the contrary, it has been observed that in certain settings additional expressiveness can give rise to additional equilibria of poor efficiency [48]. Finally, it is not obvious that more expressiveness should always yield benefits, since certain natural increases in an auction's expressiveness can lead to a reduction in the auctioneer's revenue.[1]

Short of empirical tweaking, mechanism designers lack 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 in nature. Answers could ensure that ballots are expressed in

[1] For example, consider an auction for an apple and an orange with two bidders, one who wants the apple and another who wants the orange. Running a fully expressive combinatorial auction would yield no revenue because the two bidders would not need to compete. On the other hand, running an inexpressive auction where bids are considered only for the bundle of the two items would induce competition, and thus yield some revenue.

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, and so on.

In this paper, we develop a theory that ties the expressiveness of mechanisms to their efficiency in a domain-independent manner. We begin in Section 2 by introducing two notions of expressiveness: i) *impact dimension*, which captures the extent to which an individual agent can impact the mechanism's outcome, and ii) *outcome shattering*, which is based on the concept of *shattering*, a measure of functional complexity from computational learning theory. We refer to increases (or decreases) in these measures as increases (or decreases) in expressiveness.

In Section 3, we derive an upper bound on the expected efficiency of any mechanism under its most efficient Bayes-Nash equilibrium. We show that in any setting the bound of an optimally designed mechanism increases *strictly* as more expressiveness is allowed and, for some distributions over agent valuations, by an arbitrarily large amount via a small increase in expressiveness. We also prove that in any private values setting (i.e., where an agent's utility depends only on its own private information and not the private information of any other agent) the bound is tight in that it is always possible to achieve its efficiency with a budget balanced mechanism in Bayes-Nash equilibrium. Taken together, these results imply that for any private values setting the expected efficiency of the best Bayes-Nash equilibrium increases strictly as more expressiveness is allowed. Interestingly, unlike with full expressiveness, implementing this bound is not always possible in dominant strategies.

In Section 4, we explore the relationship between our expressiveness measures and more traditional notions of communication complexity, such as the amount of information agents must transmit. Specifically, we show that our expressiveness measures can be used to derive both upper and lower bounds on the number of bits needed by the best multi-party communication protocol for computing a given outcome function.

Finally, we study a class of mechanisms which we call *channel based*. They subsume most combinatorial allocation mechanisms (of which CAs and multi-attribute auctions are a subset) and any *Vickrey-Clarke-Groves (VCG)* scheme [19, 28, 72]. We show that our domain-independent measures of expressiveness appropriately relate to the natural measure of expressiveness for channel-based mechanisms: the number of channels allowed (which itself generalizes a classic measure of expressiveness in CAs called k-wise dependence [22]). Using this bridge, our general results yield interesting implications. For example, we prove that for any (channel-based) combinatorial allocation mechanism that does not allow rich combinatorial bids there exist distributions over agent valuations (even distributions satisfying the *free disposal* condition, i.e., where the utility of winning an extra item is always non-negative), for which the mechanism cannot achieve 95% of optimal efficiency. This 5% inefficiency is an order of magnitude greater than a related inefficiency previously proven for combinatorial allocation mechanisms with sub-exponential communication (*Journal of Economic Theory* [54]).

1.1 Preliminaries

The setting we study is that of standard mechanism design. In the model there are n agents. Each agent, i, has some private information (not known by the mechanism or any other agent) denoted by a type, t_i (e.g., the value of the item to the agent in an auction; or, in a CA, a vector of values, potentially one for each bundle). The space of an agent's possible types is denoted T_i. We use the notation t^n to refer to a collection of n types (we occasionally omit the n superscript when it is clear that the entity is a collection of n types). Agent i's types are drawn according to some distribution, $P(t_i)$, that we assume is known to the mechanism designer and to agent i, but not necessarily to all agents.

Each agent has a valuation function, $v_i(o, t_i)$, that indicates its valuation under type t_i, or how much utility the agent gets when it draws type t_i and outcome $o \in \mathcal{O}$ is chosen. We call the distribution of utilities defined by a valuation function and a corresponding probability distribution over types a *preference distribution.* Settings where each agent's valuation function depends only on its own type and the outcome chosen by the mechanism (e.g., the allocation of items to the agent in a CA) are called *private values* settings. We also discuss more general *interdependent values* settings, where $v_i = v_i(o, t^n)$ (i.e., an agent's valuation depends on the others' private signals). In both settings, agents report expressions to the mechanism, denoted θ_i, based only on their own types. We use the notation θ^n to refer to a collection of n expressions. A mapping from types to expressions is called a *pure strategy.*

Definition 1 (pure strategy)**.** *A* pure strategy *for an agent i is a mapping, $h_i : T_i \to \Theta_i$, that selects an expression for each of i's types. A* pure strategy profile *for a subset of agents, I, is a list of pure strategies, one strategy per agent in I, i.e., $h_I \equiv \left[h_1, h_2, \ldots, h_{|I|}\right]$. For shorthand, we often refer to h_I as a mapping from types of the agents in I to an expression for each agent, $h_I(t_I) = \left[\theta_1, \theta_2, \ldots, \theta_{|I|}\right]$.*

We also consider *mixed strategies*, or mappings from types to random variables specifying probability distributions over possible expressions.

Definition 2 (mixed strategy)**.** *A* mixed strategy *for agent i is a mapping, $h_i : T_i \to P(\Theta_i)$, that selects a probability distribution over expressions for each of i's types. A* mixed strategy profile *is a list of mixed strategies, one strategy per agent.*

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

In Section 3, we discuss results pertaining to the implementation of a mechanism under two different solution concepts: Bayes-Nash and dominant strategy equilibria. We do not restrict our

[2]In Section 2, we define our measures of expressiveness based only on the mechanism's outcome function. For our purposes, this is without loss of generality as long as agents do not care about *each others'* payments. We later discuss the payment function in more depth when we examine issues related to incentives in Section 3.

attention to mechanisms with truthful equilibria (i.e., where agents are incentivised to report their true types in equilibrium).[3]

Definition 3 (Bayes-Nash equilibrium)**.** *A mechanism is in* Bayes-Nash equilibrium *when agents play a strategy profile such that no agent can gain expected utility by unilaterally deviating (i.e., assuming the expressions of all the other agents remain fixed).*

Definition 4 (dominant-strategy equilibrium)**.** *A mechanism is in* dominant-strategy equilibrium *when agents play a pure strategy profile such that no agent can gain utility by deviating, regardless of what the other agents do.*

During some of our analysis, we consider the widely studied class of mechanisms in which the set of expressions available to an agent corresponds directly with its types. These are called *direct-revelation mechanisms*.

Definition 5 (direct-revelation mechanism)**.** *A* direct-revelation mechanism *is a mechanism in which each agent's expression space is equivalent to its type space (i.e., $T_i = \Theta_i$, for all i).*

To summarize, we use the following notation.

- $t_i \in T_i$ is the true type of an agent i. The subscript t_{-i} is used to denote a set of types for all the agents other than i, and the superscript t^n is used to denote a set of n types.
- $\theta_i \in \Theta_i$ is the expression that agent i reports to the mechanism. The subscript θ_{-i} is used to denote a set of expressions for all the agents other than i, and the superscript θ^n is used to denote a set of n expressions.
- $o \in \mathcal{O}$ is an outcome from the set of all possible outcomes imposable by the mechanism.
- $v_i : \mathcal{O}, T_i \rightarrow \mathbb{R}$ is agent i's valuation 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 to interdependent values settings, where $v_i = v_i(o, t^n)$, i.e., an agent's utility also depends on 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.

For convenience, we will let $W(o, t^n)$ denote the total social welfare of outcome o when agents have private types (or private signals) t^n, i.e., $W(o, t^n) = \sum_i v_i(o, t^n)$. Occasionally, we use the

[3] The *revelation principle* of mechanism design states that any outcome function that can be implemented by any mechanism under a non-truthful equilibrium can also be implemented by some mechanism under a truthful equilibrium [45]. However, we do not restrict our analysis to mechanisms with truthful equilibria because in mechanisms without full expressiveness it can be impossible for agents to express their true types.

shorthand W_I, where I refers to some subset of the agents, to denote the total social welfare of only the agents in I. Assuming the agents play a mixed strategy equilibrium denoted by m, the expected efficiency, $\mathcal{E}(f)$, of an outcome function, f, (where expectation is taken over the types of the agents and their randomized equilibrium expressions) is given by,

$$E\left[\mathcal{E}(f)\right] = \int_{t^n} P(t^n) \int_{\theta^n} P(m(t^n) = \theta^n)\ W(f(\theta^n), t^n). \tag{1}$$

2 Characterizing the expressiveness of mechanisms

The primary goal of this paper is to better understand the impact of making mechanisms more or less expressive. In order to achieve this goal, we must first develop meaningful (and general) measures of a mechanism's expressiveness.

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

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

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

Another natural way in which mechanisms can differ is in the granularity of their outcome spaces. For example, auction mechanisms that are restricted to allocating certain items together (e.g., blocks of neighboring frequency bands) have coarser outcome spaces than those that can allocate them individually. 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 less coarse outcome spaces always have more efficient dominant-strategy equilibria [34, 53].

In contrast, we are interested in studying the impact of a mechanism's *expressiveness* on its efficiency. We do this by comparing more and 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 the 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,[5] but there may be

[4] Proofs of all technical claims are located in the Appendix at the end of this paper.

[5] This 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.

a prohibitive burden on agents if they are asked to express a large amount of information; thus it is limited expressiveness that is the crucial issue.

2.1 Impact-based expressiveness

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. We define an *impact vector* to capture the impact of a particular expression by an agent under the different possible types of the other agents. (The subscript $-i$ refers to all the agents other than agent i.)

Definition 6 (impact vector)**.** *An* impact vector *for agent i is a function, $g_i : T_{-i} \to \mathcal{O}$. To represent the function as a vector of outcomes, we order the joint types in T_{-i} from 1 to $|T_{-i}|$; then g_i can be represented as $\left[o_1, o_2, \ldots, o_{|T_{-i}|}\right]$.*

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

Definition 7 (express)**.** *Agent i can* express *an impact vector, g_i, if*

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

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

Definition 8 (distinguish)**.** *Agent i can* distinguish *among a set of impact vectors, G_i, if*

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

when this is the case, we say $D_i(G_i)$ is true.

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

Intuitively, more expressive mechanisms allow agents to distinguish among larger sets of impact vectors. Our first expressiveness measure captures this intuition; it measures the number of different impact vectors that an agent can distinguish among. Since this depends on what the others express, we measure the best case for a given agent, where the others submit expressions that maximize the agent's control. We call this the agent's *maximum impact dimension.*

Definition 9 (maximum impact dimension)**.** *Agent i has* maximum impact dimension *d_i if the largest set of impact vectors, G_i^*, that i can distinguish among has size d_i. Formally,*

$$d_i = \max_{G_i} \left\{ |G_i| \;\middle|\; D_i(G_i) \right\}.$$

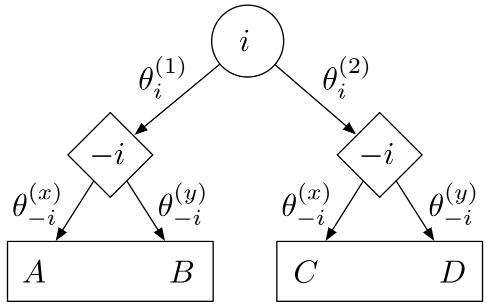

Figure 1: *By choosing between two expressions, $\theta_i^{(1)}$ and $\theta_i^{(2)}$, agent i can distinguish between the impact vectors $[A, B]$ and $[C, D]$ (enclosed in rectangles). The other agents are playing the pure strategy profile* $\left[\theta_{-i}^{(x)}, \theta_{-i}^{(y)}\right]$.

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.

However, 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 let an agent distinguish among a more important (e.g., for efficiency) set of impact vectors. 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 discuss a related notion of expressiveness, which we call *outcome shattering*. As we will show later, it has somewhat different uses than the maximum impact dimension.

Outcome shattering is based on a notion called *shattering*, a measure of functional complexity that we have adapted from the field of computational learning theory [15, 71]. In learning theory, a class of binary classification functions[6] 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 [49].

Example 1. *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*

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

$S = \{3.1, 5.7\}$? *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 involving outcomes from a given set, $\mathcal{O}'$.

Definition 10 (outcome shattering)**.** *A mechanism allows agent* i *to shatter a set of outcomes,* $\mathcal{O}' \subseteq \mathcal{O}$, *over a set of joint types for the other agents,* T_{-i}, *if* $D_i(G_i^{\mathcal{O}'})$, *where,*

$$G_i^{\mathcal{O}'} = \left\{ g_i \middle| g_i = \left[o_1, o_2, \ldots, o_{|T_{-i}|} \right], o_j \in \mathcal{O}' \right\} .$$

Example 2. *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 two joint types of the other agents:*

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

We also use a slightly weaker adaptation of shattering for analyzing the more restricted setting where agents have private values. It captures an agent's ability to cause each of the $\binom{|\mathcal{O}'|+1}{2}$ unordered pairs of outcomes (with replacement) to be chosen for every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., under which of the other agents' types each of the outcomes is chosen). We call this *semi-shattering.* [7]

Definition 11 (outcome semi-shattering)**.** *A mechanism allows agent* i *to semi-shatter a set of outcomes,* $\mathcal{O}'$, *over a set of joint types for the other agents,* T_{-i}, *if* i *can distinguish among a set of impact vectors,* G_i, *such that for every pair of types* $\{\{x, y\} \in T_{-i} | x \neq y\}$, *and every pair of outcomes,* $\{\{o_1, o_2\} \in \mathcal{O}' | o_1 \neq o_2\}$,

$$[(\exists g_i \in G_i, g_i(x) = o_1 \ \wedge g_i(y) = o_2) \ \wedge \ \neg(\exists g_i \in G_i, g_i(x) = o_2 \ \wedge g_i(y) = o_1)] \ \vee$$
$$[(\exists g_i \in G_i, g_i(x) = o_2 \ \wedge g_i(y) = o_1) \ \wedge \ \neg(\exists g_i \in G_i, g_i(x) = o_1 \ \wedge g_i(y) = o_2)] \, .$$

The notion of outcome semi-shattering is best illustrated by the following simple examples.

Example 3. *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 two joint types of the other agents (the order of the pairs that are included does not matter, for example* $[A, B]$ *could be replaced with* $[B, A]$*):*

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

[7] There are many ways to generalize the shattering notion to functions that can return more than two outcomes, c.f. [11]. We have adapted the two most natural ones for our work on expressiveness in mechanism design—in Definitions 10 and 11, respectively. Definition 11 has been slightly altered compared to the version presented at a conference in order to be able to also prove ties to communication complexity.

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

Example 4. *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 three joint types of the other agents:*

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

Notice that every pair of outcomes appears in every pair of slots, in the same order, and at least once, which is exactly the requirement for semi-shattering.

Our second measure of expressiveness is based on the size of the largest outcome space that an agent can shatter or semi-shatter.[8] It captures the number of outcomes that the mechanism can support full expressiveness over for that agent. We call it the *(semi-)shatterable outcome dimension.*

Definition 12 ((semi-)shatterable outcome dimension)**.** *Agent i has (semi-)shatterable outcome dimension k_i if the largest set of outcomes that i can (semi-)shatter has size k_i.*

The (semi-)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 (semi-)shatterable outcome dimension measures the *number of different outcomes* for which an agent has full expressiveness. In addition, it has the advantage that we can rule out the (semi-)shatterability of a set of outcomes by merely ruling out the existence of a *pair* of expressions by the other agents that allows the agent to (semi-)shatter the set.

Observation 1. *Agent i can (semi-)shatter an outcome space $\mathcal{O}'$ only if there exists at least one pair of expressions by the other agents that 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 (unordered) pair of outcomes from $\mathcal{O}'$ to be chosen.)*

This observation will allow us to analyze the measure even when agents have infinite type spaces, and may help one operationalize expressiveness for automated mechanism design [20] in the future. We use this insight throughout the study of channel-based mechanisms in Section 5.

[8] The 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.

The next two results illustrate the close relationship between the shatterable outcome dimension measures and the maximum impact dimension measure. While the two measures are related, the shatterable outcome dimension can be thought of as a measure of expressiveness breadth.

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

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

This result states that the maximum impact dimension necessary for an agent to shatter k outcomes increases geometrically in the number of types of the other agents, which illustrates the relationship between expressiveness and uncertainty. As uncertainty goes up (the number of types that the other agents have can be thought of as a support-based measure of uncertainty), more expressiveness is needed to shatter a given set of outcomes.

2.3 Uses of the expressiveness measures

The expressiveness measures introduced above enable us to understand mechanisms from a new perspective. Because the measures are so new, we undoubtedly fail to see all of their possible uses at this time, however we already see several.

First, we can measure the expressiveness of an existing mechanism, and thereby bound how well the mechanism can do in terms of a designer's objective. For example, in the next section, we show how our expressiveness measures directly relate to 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 (e.g., 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, our results presented in the next section imply that this approach can be used to yield the most efficient possible Bayes-Nash equilibrium in any private values setting.

We can also ask which of the expressiveness measures—maximum impact dimension, shatterable outcome dimension, or semi-shatterable outcome dimension—is most appropriate under different settings and for different 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 have full expressiveness over those outcomes. As we will show in Section 3, in private values settings the appropriate measure is semi-shatterable outcome dimension (for one, 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 arbitrary inefficiency under some preference distributions.

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 5.

3 Relationship between expressiveness and efficiency

Perhaps the most important property of our domain-independent measures of expressiveness is how they relate to the efficiency of the mechanism's outcome. In this section, we will discuss a cooperative upper bound on the expected efficiency of a mechanism's *most efficient* equilibrium that is tied directly to the expressiveness of an optimally designed mechanism and can always be implemented by a budget-balanced mechanism in Bayes-Nash equilibrium (in private values settings).[9]

The bound measures the efficiency of the outcome function under the optimistic assumption that the agents play strategies which, taken together, attempt to maximize expected efficiency. Studying this bound allows us to sidestep two of the major roadblocks faced by many prior attempts at analyzing the relationship between expressiveness and efficiency: 1) we do not have to solve for any of the mechanism's equilibria (attempts at doing this have proved difficult for many expressive and inexpressive mechanisms [52, 55, 60, 70, 74, 77]) 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 [13].

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

$$E\left[\mathcal{E}(f)\right]^{+} = \max_{\hat{h}(\cdot)} \int_{t^n} P(t^n)\ W\Big(f(\hat{h}(t^n)), t^n\Big)\,. \tag{2}$$

The bound holds for mixed strategies, but the maximum in the equation need only be taken over the space of pure-strategy profiles, $\hat{h}(\cdot)$.

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

3.1 Conditions under which the bound is fully efficient

There is an impact vector for each of agent i's types that represents the vector of the most efficient outcomes when it is matched with each of the joint types of the other agents. We call a set of such

[9]The upper bound we derive represents a cooperative equilibrium that could be used to bound the value of any objective that depends only on the agents' types and the outcome chosen by the mechanism. By extension, all of our subsequent theory (except for the implementability of the bound discussed in Section 3.4) also applies for any such objective.

impact vectors a *fully efficient set*. Such a set must be distinguishable by each agent for the bound to reach full efficiency.

Definition 13 (fully efficient set)**.** *G_i^* is a* fully efficient set *if*

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

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

In full information settings, whereupon learning its own type an agent knows the types of the other agents for sure, the agent is guaranteed to have a fully efficient set of size $\leq |\mathcal{O}|$. (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 6. *Let G_i^* be agent i's smallest fully efficient set,*

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

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

The implication of this result is that perfect information about the other agents' types essentially eliminates the need for expressiveness. Thus, in prior research showing that in certain settings even quite inexpressive mechanisms yield full efficiency in ex-post Nash equilibrium (e.g., [1]), the assumption that the agents know each other's types is essential.

3.2 The efficiency bound increases strictly with expressiveness

The following results demonstrate that a mechanism designer can *strictly* increase the upper bound on expected efficiency by giving any agent more expressiveness (until the bound reaches full efficiency). The result applies to the outcome function that maximizes the bound subject to the constraint that agent i's expressiveness be less than or equal to a particular level. The bound attained by such an optimal outcome function is also an upper bound for any outcome function at that expressiveness level or lower.

Theorem 1. *For any setting and any distribution over agent preferences, the upper bound on expected efficiency, $E\left[\mathcal{E}(f)\right]^+$, for the best outcome function limiting agent i to maximum impact dimension d_i increases* strictly *monotonically as d_i goes from 1 to d_i^* (where d_i^* is the size of agent i's smallest fully efficient set).*

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

Corollary 2. *The upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function limiting agent i's expressiveness to (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).*

3.3 Inadequate expressiveness can lead to arbitrarily low efficiency in any setting, for some preference distributions

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 an 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 try to keep these constructions as general as possible: we allow for any number of outcomes, any number of agents, and any number of types.)

Lemma 1. *For any agent i in an interdependent values setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents), there exist preference distributions under which $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i's maximum impact dimension to d_i (where $2 \leq d_i \leq |\mathcal{O}|^{|T_{-i}|}$) is* arbitrarily larger *than that of any outcome function limiting i's maximum impact dimension 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, and semi-shatterable outcome dimension, for private values settings. As we will see, in a private values setting there is no need to allow full shattering to achieve full 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 preference distributions under which $E[\mathcal{E}(f)]^+$ for the best outcome function limiting 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*

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

Private values settings place restrictions on agents' utility functions and, therefore, on the efficiency-maximizing outcomes under different types. We will use the following lemma to show that in such settings allowing the agents to semi-shatter the outcomes is sufficient for maximizing the efficiency bound. The lemma proves that the most efficient *order* for two outcomes under any pair of opposing types must be the same for all of agent *i*'s types.

Lemma 3. *In any 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 exists* some *type of agent i, t_i, where it is strictly more efficient for o_1 to be chosen under type $t_{-i}^{(1)}$ and o_2 to be chosen under type $t_{-i}^{(2)}$ than vice-versa (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 be chosen 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 expected efficiency loss under any possible preference distribution.

Theorem 2. *For any setting, there exists a distribution over agent preferences such that $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i to*

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

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

3.4 Bayes-Nash implementation of the upper bound is always possible in private values settings

In addition to the results above, we find that the upper bound on expected efficiency can be implemented in Bayes-Nash equilibrium for any outcome function, in any private values[10] setting, as long as the agents have quasi-linear utility functions. Quasi-linearity means that the agent's utility functions are linear in money or some commonly agreed upon currency. Formally, a quasi-linear utility function for agent i takes the form $u_i = v_i - \pi_i$, where v_i is the agent's valuation for the outcome chosen by the mechanism and π_i is the payment from the agent to the mechanism. This is the first point in the paper where we assume quasi-linearity: all the results so far apply with and without that restriction.

Theorem 3. *For any private values setting with quasi-linear preferences and any outcome function, f, there exists a class of payment functions that achieve the upper bound on efficiency, $E[\mathcal{E}(f)]^+$, in a pure-strategy Bayes-Nash equilibrium.*

This implementability of the upper bound implies that, for private values settings, we can recast all of our earlier results that relate expressiveness to the bound as relating expressiveness to the efficiency of the most efficient implementable Bayes-Nash equilibrium.

[10] Implementing efficient allocations in Bayes-Nash equilibrium for interdependent values settings is impossible even with full expressiveness [40]. The difficulty stems from the need for the mechanism designer to know the beliefs of the agents about *each others'* private information.

3.4.1 Individual rationality and budget balance

In this section, we will discuss individual rationality and budget balance, and how they are related to expressiveness. First, in Bayes-Nash equilibrium, we can always get *strong* budget balance (i.e., the total payments to and from all agents are equal), and we can get *ex ante* individual rationality (i.e., it is always in an agent's best interest to participate in the mechanism prior to learning its own type) as long as agent valuations for outcomes are non-negative.

Proposition 7. *There exists at least one payment function in the class of Theorem 3 that is strongly budget balanced and, if agent valuations for outcomes are non-negative,* ex ante *individually rational.*

These payment functions are derived much like in the expected-form Groves mechanism which is due to d'Aspremont and Gerard-Varet [25] and Arrow [6] (called the dAGVA mechanism). However, as implied by the Myerson-Satterthwaite impossibility theorem [51] for fully expressive mechanisms, there may not exist a payment function in this class that is *ex interim* individually rational (i.e., it may not be in an agent's best interest to participate once the agent knows its own type). Additionally, there exist settings, such as the one described in the following example, where the fully expressive dAGVA mechanism is ex-interim individually rational but a limited-expressiveness variant is not.

Example 5. *Consider an auction for one item run using the dAGVA mechanism. Assume there are two bidders with valuations for the item drawn from the uniform distribution over* $[0, 1]$ *and one auctioneer with zero valuation for the item. Let* $\hat{\theta}$ *represent the bidders' bids, assuming they follow the Bayes-Nash equilibrium and report their valuations truthfully, and let* $f_i(\theta)$ *be an indicator function that returns one if bidder* i *wins the item and zero otherwise. The following reasoning demonstrates that the fully expressive dAGVA mechanism is* ex interim *individually rational for this setting.*

First, we can calculate the dAGVA payment for one of the bidders, i*, for a given set of bids (the payment to the auctioneer is the sum of the payments from the bidders). We begin with the general formula for the dAGVA payment function in a direct-revelation mechanism and then instantiate it for a bidder in this example.*

$$\begin{aligned}
\pi_i(\hat{\theta}) &= -E_{\theta_{-i}}[W_{-i}(f(\hat{\theta}_i, \theta_{-i}), \theta_{-i})] + \frac{1}{n-1}\left(\sum_{j \neq i} E_{\theta_{-j}}\left[W_{-j}(f(\hat{\theta}_j, \theta_{-j}), \theta_{-j})\right]\right) \\
&= -E_{\theta_{-i}}[\theta_{-i} f_{-i}(\hat{\theta}_i, \theta_{-i})] + \frac{1}{2}\left(E_{\theta_i}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] + E_{\theta}[\max(\theta_i, \theta_{-i})]\right) \\
&= \frac{1}{12} + \frac{\hat{\theta}_i^2}{2} - \frac{\hat{\theta}_{-i}^2}{4}
\end{aligned}$$

Next, we can calculate bidder i*'s expected utility when it draws type* θ_i*.*

$$\begin{aligned}
E[u_i|\theta_i] &= E_{\hat{\theta}_{-i}}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] - E_{\hat{\theta}_{-i}}\left[\pi_i(\theta_i, \hat{\theta}_{-i})\right] \\
&= E_{\hat{\theta}_{-i}}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] - E_{\hat{\theta}_{-i}}\left[\frac{1}{12} + \frac{\theta_i^2}{2} - \frac{\hat{\theta}_{-i}^2}{4}\right] \\
&= \frac{\theta_i^2}{2}
\end{aligned}$$

Since θ_i can never be less than zero, $E[u_i|\theta_i]$ can never be negative in this setting.

However, if we consider bidder i's expected utility under the best outcome function with less than full expressiveness, we find this is not necessarily true. Say we limit the expressiveness to a maximum impact dimension of one. Now, the best limited-expressiveness outcome function for this setting always allocates the item to the same bidder regardless of the bids. Call that bidder i. Under these assumptions, bidder i's dAGVA payment and expected utility are

$$\begin{aligned} \pi_i(\hat{\theta}) &= -E_{\theta_{-i}}[\theta_{-i} f_{-i}(\hat{\theta}_i, \theta_{-i})] + \frac{1}{2}\Big(E_{\theta_i}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] + E_{\theta}[\theta_i]\Big) \\ &= E_{\theta_i}[\theta_i] \\ E[u_i|\theta_i] &= \theta_i - E_{\theta_i}[\theta_i] \end{aligned}$$

Thus, bidder i's expected utility is negative whenever the valuation it draws is less than its expected valuation.

3.4.2 Impossibility of dominant strategy implementation

While Theorem 3 shows it is always possible to implement the upper bound for private values settings in Bayes-Nash equilibrium, we show below that there exist private values settings for which dominant strategy implementation is impossible without full expressiveness. In other words, it is known that with full expressiveness there is no difference between what is possible in dominant strategies and Bayes-Nash equilibrium (except for issues of individual rationality and budget balance), but we show that with less than fully expressive mechanisms there is a fundamental difference in the power of the two solution concepts.

Theorem 4. *There exist private values settings with quasi-linear preferences where the outcome function that maximizes the upper bound on efficiency, $E[\mathcal{E}(f)]^+$, while limiting agent i to a maximum impact dimension $d_i < d_i^*$ (d_i^* is the size of agent i's smallest fully efficient set), cannot be implemented in dominant strategies.*

The reason for this impossibility is that there exist settings where the best limited-expressiveness outcome function is not guaranteed to satisfy the *weak-monotonicity* property, a condition which has been shown to be necessary for dominant strategy implementation [14]. This property requires that the outcome function react properly to relative changes in an agent's reported preferences for any two outcomes.

4 Relationship between expressiveness and communication complexity

In this section, we consider the relationship between our notions of expressiveness and more traditional notions of communication complexity. Our expressiveness measures quantify how the mechanism uses information, while communication complexity measures how much information has to be communicated (by the agents) to compute it. Although these notions do not measure exactly the same thing, they are closely related. In this section we will begin to formalize this relationship.

One measure of an outcome function's communication complexity for agent i is the size of its expression space, $|\Theta_i|$. As we will show, this determines an upper bound on the amount of information communicated by the agent under any communication procedure that computes the outcome function.

In relating expressiveness to the number of expressions needed for each agent, we consider whether or not a given outcome function can be *emulated* by an outcome function with fewer expressions (essentially losslessly compressed).

Definition 14 (emulate)**.** *An outcome function, f',* emulates *another outcome function, f, if there exists a function, q_i, for each agent, i, that maps from i's expression space under f to i's expression space under f', such that*

$$\forall i,\ \forall \theta_i,\ \forall \theta_{-i},\ f(\theta_i, \theta_{-i}) = f'(q_i(\theta_i'), q_{-i}(\theta_{-i}')).$$

For a given outcome function, f, each agent's maximum impact dimension provides a lower bound on the the number of expressions needed for that agent by any outcome function that emulates f.

Proposition 8. *It is impossible to emulate an outcome function, f, with an outcome function that provides any agent with less expressions than its maximum impact dimension under f.*

Furthermore, for any outcome function that belongs to the widely studied class of direct-revelation mechanisms, an agent's maximum impact dimension is exactly the number of expressions used by the best emulator of the outcome function (i.e., the outcome function that emulates it while minimizing the number of expressions).

Lemma 4. *Under a direct-revelation outcome function, each agent i's impact dimension is maximized when the agents other than i report their types truthfully, (i.e., $h_{-i}(t_{-i}) = t_{-i}$ is the strategy that maximizes i's impact dimension).*

Proposition 9. *Any direct-revelation outcome function, f, can be emulated by another outcome function, f', that provides each agent, i, with exactly d_i expressions, where d_i is agent i's maximum impact dimension under f.*

Given this relationship between our expressiveness measure and the number of expressions needed by any agent, we have the following two Corollaries related to the upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$. Corollary 3 states that increasing a limit on the number of expressions given to an agent strictly increases the bound. Corollary 4 states that some distributions require an agent to have a number of expressions that is exponential in the number of types of the other agents to avoid being arbitrarily less than fully efficient.

Corollary 3. *For any setting and any distribution over agent preferences, $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i to d_i expressions 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.*

Corollary 4. *There exists settings and distributions over agent preferences such that the upper bound on expected efficiency for the best outcome function limiting agent i to less than $|T_{-i}|^{|O|}$ expressions is arbitrarily less than that of the best outcome function.*

While the reasoning above provides an upper bound on an outcome function's communication complexity, it does not account for the possibility of designing clever elicitation protocols, such as protocols that iteratively ask different agents different questions (cf. [64]). To address this, we will also relate our notion of expressiveness to a lower bound on communication complexity. The lower bound is derived by considering the execution of the outcome function as a two-party communication problem, where agent i holds one piece of information (its intended expression) and the agents other than i hold another (their intended joint expression). From this perspective, we can study the outcome function using Yao's model of communication complexity [76], as in Nisan and Segal's seminal work on communication complexity in mechanism design [54].

Yao's model considers the computation of a pre-specified function based on the information held by the agents. It is typical, when using this model, to think of the function being computed as a matrix where the rows represent the possible inputs to the function from one agent, the columns represent the possible inputs from the other agent, and each cell contains the value of the function under the inputs corresponding to its row and column. For a given outcome function, f, and agent, i, we can construct such an input matrix from i's perspective by placing its expressions along the rows and the joint expressions of the other agents along the columns. The cells of the matrix contain the outcome chosen by the outcome function under the corresponding expressions. (Thus, the rows correspond to agent i's possible impact vectors.)

It has been shown that any communication protocol that computes f must involve at least one message for each of the *monochromatic rectangles* (i.e., contiguous rectangles of expressions that result in the same outcome being chosen) in some partitioning of f's input matrix [42]. The following result shows how our notion of semi-shattering is related to the number of monochromatic rectangles needed in any partitioning of f. Specifically, any set of types for the agents other than i over which agent i can semi-shatter a pair of outcomes leads to a corresponding set of expression pairs that cannot be in the same monochromatic rectangle for either outcome.

Lemma 5. *Let T_{-i} be a set of joint types for the agents other than i over which agent i can semi-shatter a pair of outcomes, A and B, under some outcome function, f. There exists a set of $|T_{-i}|-1$ pairs of expressions that cannot be in the same A- or B-monochromatic rectangle of f.*

This leads directly to a lower bound on the number of monochromatic rectangles needed by any partitioning of an outcome function's input matrix and, consequently, a lower bound on the number of messages needed by any communication protocol that computes it.

Theorem 5. *For any outcome function, f, agent, i, and outcome, o, let T^o_{-i} denote the largest set of types over which i can semi-shatter a pair of outcomes containing o. Also, let d_i be i's maximum impact dimension under f. The number of monochromatic rectangles, R, needed by any partitioning* *of the input matrix of f, and the number of messages needed by the best communication protocol that computes f, $M(f)$, satisfy the following inequality,*

$$\min_i d_i \geq M(f) \geq R \geq \max_i \sum_{o \in O} |T^o_{-i}| - 1. \tag{3}$$

This result bounds the number of bits needed by any discrete communication protocol that computes f, since a function that requires $M(f)$ messages to compute must communicate at least $\log_2(M(f))$ bits (i.e., the depth of a binary tree with $M(f)$ leaves). Our bound is also consistent with earlier results showing that combinatorial allocation mechanisms can require the communication of a number of bits that is exponential in the number of items [54], since the number of types an agent has in a combinatorial allocation setting is typically doubly-exponential in the number of items: if there are m items, and an agent has k possible values for each bundle, then the agent has k^{2^m} types. Thus, according to Theorem 5, a combinatorial allocation mechanism that allows an agent to semi-shatter even a single pair of outcomes over the other agents' entire type space would require at least $\log_2(|T_{-i}| - 1)$ bits, which is on the order of 2^m bits.

5 An instantiation to illustrate expressiveness: Channel-based mechanisms

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

Definition 15 (channel-based mechanism)**.** *Each outcome is assigned a set of channels potentially coming from a number of different agents (e.g., outcome A may be assigned channels x_1 and y_1 from Agent 1 and x_2 from Agent 2). Each agent, simultaneously with the other agents, reports real values on each of its channels to the mechanism. The mechanism chooses the outcome whose channels have the largest sum.*[11] *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:*

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

[11] We 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.

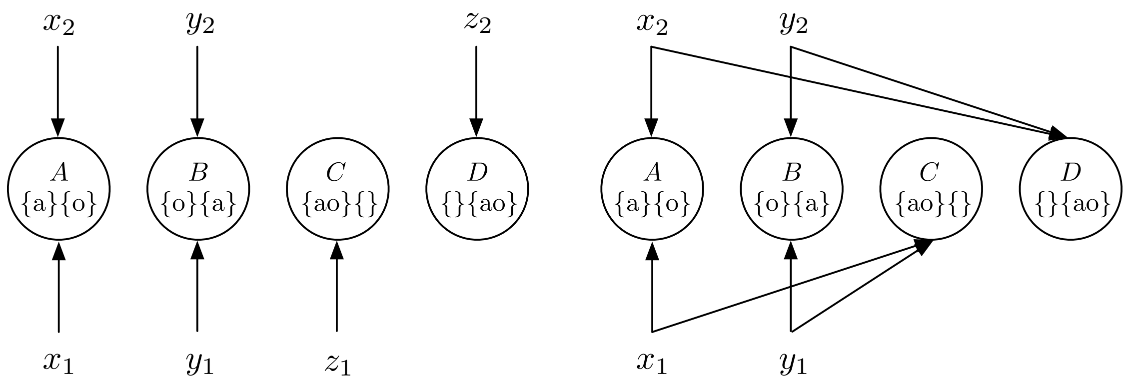

Fully expressive combinatorial auction. Auction that only allows bids on items.

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 combinatorial allocation mechanisms, CAs, exchanges, and multi-attribute mechanisms can be cast as channel-based mechanisms.

A natural measure of expressiveness in channel-based mechanisms is the number of channels allowed. For 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 [22]. First, we will demonstrate that our domain-independent expressiveness measures relate appropriately to the number of channels allowed in a channel-based mechanism.

Proposition 10. *For any agent i, its semi-shatterable outcome dimension, k_i, in the most expressive channel-based mechanism strictly increases (until $k_i = |\mathcal{O}|$) as the number of channels assigned to the agent increases.*

Based on Theorems 1 and 2, we know that an increase in expressiveness will always yield an increase in our efficiency bound and can lead to an arbitrary large increase, even in private values settings.

Corollary 5. *For any setting, the upper bound on expected efficiency of the best channel-based mechanism that allows c_i channels for agent i is strictly greater than, and can be arbitrarily larger than, that of the best mechanism that allows agent i to have c'_i channels, where $c_i < c'_i \leq c^*_i$ and c^*_i is the number of channels needed for full efficiency.*

However, if an agent has full information it only needs a logarithmic number of channels to bring the bound to full efficiency. (This also happens to be the number of channels in any multi-item auction that allows item bids only.)

Proposition 11. *If agent i has full information about the other agents, in a channel-based mechanism it needs only* $\lceil log_2(|\mathcal{O}|)\rceil$ *channels to shatter the entire outcome space.*

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

Proposition 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 full shattering, our results from the previous section imply that in some interdependent values settings any channel-based mechanism, even one that emulates the VCG mechanism, will be arbitrarily inefficient. (That such mechanisms cannot always get full efficiency in interdependent values settings is already known [40].)

Corollary 6. *In any interdependent values setting, there exist preference distributions for which any channel-based mechanism (even one that emulates the VCG mechanism) results in arbitrarily less than full expected efficiency.*

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

Proposition 13. *A channel-based mechanism can emulate the VCG mechanism if and only if it provides each agent with at least* $|\mathcal{O}| - 1$ *channels.*

Our next two results deal with a configuration of channels that prevents an agent from being able to semi-shatter a set containing two particular pairs of outcomes.

Lemma 6. *For any sets, A, B, C, and D, the following bi-directional implication holds,*

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

Lemma 7. *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 any set of outcomes containing both pairs* $\{A, B\}$ *and* $\{C, D\}$ *(i.e., there is no fixed pair of expressions by the other agents allowing i to cause the mechanism to select A and B with one expression, and C and D with another) if,*

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

The channel configuration discussed in Lemma 7 generalizes one that appears in the channel-based representation of a CA where bids are allowed on items only. In fact, it is present in any combinatorial allocation mechanism 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).

Based on this insight, we can prove that for any combinatorial allocation mechanism where an agent's bid on any bundle is the sum of its bid on two other non-overlapping bundles, there exists a preference distribution satisfying free disposal (i.e., an agent's valuation for a bundle of items is greater than or equal to its valuation for any sub-bundle) where the mechanism cannot achieve expected efficiency within 5% of the maximum. While 5% may seem like a relatively small gap, it can be arbitrarily large in absolute terms. Furthermore, it is ten times larger than the expected efficiency gap found by Nisan and Segal [54] in their prior work on communication complexity in combinatorial allocation mechanisms. Their result pertains to mechanisms that communicate less than an exponential number of bits and involves a single prior over preferences. Our result pertains to limited expressiveness and potentially uses a different prior for each mechanism.

Theorem 6. *Consider a combinatorial allocation mechanism, M, which can be represented as a channel-based mechanism that treats agent i's bid on any bundle Q to be the sum of its bids on some two other non-overlapping sub-bundles, q_1 and q_2. There exists a distribution over agent valuations, that satisfies the private values and free disposal assumptions, such that M cannot achieve expected efficiency within 5% of the maximum possible for the setting.*

The setting that is used to prove Theorem 6 provides some insight into the circumstances under which limited expressiveness can be particularly problematic. It involves two agents who care only about the items with limited expressiveness (i.e., the items in q_1 and q_2). Each of the agents has two equally probable types: a *complementarity* type and a *substitutability* type. Under the complementarity type, the agent only derives utility from winning the super-bundle (i.e., the bundle containing both q_1 and q_2). Under the substitutability type, the agent derives no additional utility from winning more than one sub-bundle (either q_1 or q_2). In other words, we find that expressiveness is important for combinatorial allocation mechanisms when agents may have either complementarity or substitutability for the same items.

6 Connections to related research

There has been relatively little work on expressiveness specifically. We discussed some related papers in the body of this paper. 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.

6.1 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 [32]. In the 1970s, Mount and Reiter [50] and Hurwicz [35, 36] 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 (cf. [38])? That work focused primarily on the number of real-valued dimensions that were needed. It was well known that in general, as our Proposition 1 shows, 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 [50] or Lipschitz continuity [37]) that precluded a general mapping between $\Re^n$ and $\Re^m$. 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.

6.2 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 [60] 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 for 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 [70], who 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 [74] 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.

6.3 Expressiveness issues in dominant-strategy mechanisms

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

For example, Blumrosen and Feldman [16] studied the problem of designing a dominant-strategy mechanism with a limited number of discrete actions. They showed a trade off between the efficiency of the best possible dominant-strategy mechanism and the number of discrete actions available to the designer. Similarly, Ronen [59] described methods for achieving near efficiency with limited bidding languages in dominant strategies.

Holzman *et al.* [34] 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 [57] and Nisan and Segal [54] 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 the types of questions about expressiveness that can be addressed. 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.

6.4 Applications of expressiveness in mechanisms

One of the first applications to benefit from expressiveness was strategic sourcing. Sandholm [63, 65] 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 [23, 33, 47]. Schoenherr and Marbert [68] 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 [26]. Under certain probabilistic assumptions they are able to prove that the system becomes more efficient when this extra level of expressiveness is allowed. It has also been shown that increasing the expressiveness of today's sponsored search auctions very slightly by allowing for a premium bid for premium slots removes most of the inefficiency of today's design [12]. Also, highly expressive mechanisms have been designed for trading entire advertising campaigns [17, 73]. 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) [48]. 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 6).

Another application area that has received recent attention with regard to expressiveness is wireless spectrum trading. For example, Gandhi *et al.* [27] 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.

6.5 Bundle pricing and revenue-maximizing CAs

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 [69]. 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, 24, 29, 46, 67]. (One exception is that Armstrong examined n-product settings, but placed severe restrictions on buyers' utility functions [4].) 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 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 [75].

Some work on bundle pricing has been done from an operations research perspective as well. For example, Hason and Martin [31] 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 only 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 [61].

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.* [41] and Brooks and Durfee [18] 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 [8].

Finally, computer science work on pricing has focused primarily on pricing items rather than bundles and on "single-minded" customers that desire only one bundle. For example, Balcan *et al.* [9, 10] provided online and approximate algorithms for this setting, and Guruswami et. al. [30] showed that finding the optimal pricing is computationally complex ($\mathcal{APX}$-hard). Some work from that 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 quickly even in the worst case (in polynomial time).

Related to bundle pricing, there has recently also been significant work on designing high-revenue CAs (e.g., [5, 7, 39, 56]). The design problem is NP-hard [21], meaning that it is very unlikely that a short characterization of the answer exists or that any algorithm can compute the answer quickly in the worst case (in polynomial time). There has been recent work on algorithms for automatically designing high-revenue CAs (e.g., [43, 44]). This work has followed two paths: designing algorithms that provide logarithmic approximations to revenue while running fast even in the worst case, and designing algorithms that yield very high revenue in practice, but do not provide worst-case guarantees on solution quality or run time.

7 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 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 equilibrium that 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) the bound can be analyzed without having to solve for an equilibrium of the mechanism, and 2) the 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 lead to arbitrarily large increases in the efficiency bound, depending on the prior over agents' preferences. We ended the discussion with proof that the bound is tight in private values settings: it is always possible to build a strongly budget balanced payment function that achieves the efficiency of the bound in Bayes-Nash equilibrium (with *ex ante* but not necessarily *ex interim* individual rationality). This implies that for any private values setting, the expected efficiency of the best Bayes-Nash equilibrium increases strictly as more expressiveness is allowed. However, we showed that unlike with full expressiveness, implementing the bound is not always possible in dominant strategies with less than full expressiveness.

Next, we explored the relationship between our expressiveness measures and communication complexity. We showed that the expressiveness measures can be used to derive both upper and lower bounds on the number of bits used by the best communication protocol for running any mechanism.

Finally, we instantiated our model of expressiveness for a class of mechanisms, called 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. 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 [22]). Using our general measures of expressiveness and the results on how they relate to efficiency, we proved that in channel-based mechanisms 1) increasing expressiveness by allowing an additional channel leads to an increase in the upper bound on expected efficiency for the mechanism, and 2) under some preference distributions this leads to an arbitrarily large increase in the bound. We also used our theoretical framework to prove that for any (channel-based) multi-item allocation mechanism that does not allow rich combinatorial bids, there exist distributions over agent preferences that satisfy the free disposal condition for which the mechanism cannot achieve 95% of optimal efficiency. This inefficiency is ten times larger than a related expected efficiency gap found by Nisan and Segal [54] in their prior work on communication complexity in combinatorial allocation mechanisms.

The framework we developed enables one to understand mechanisms from a new perspective and opens the door for a possible new avenue of research within mechanism design. In addition to the theoretical contributions of this work, we already see two practical uses for our expressiveness measures. First, they can be used to bound the efficiency—and therefore provide a lower bound on inefficiency—of existing mechanisms. Second, they can be used to aid 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 the Workshop on Internet Economics (WINE)*, 2007.

[2] William James Adams and Janet L. Yellen. Commodity bundling and the burden of monopoly. *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. *Lecture Notes in Computer Science*, pages 253–266, 2004.

[4] Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, Jan 1996.

[5] Mark Armstrong. Optimal multi-object auctions. *Review of Economic Studies*, 67:455–481, 2000.

[6] Kenneth Arrow. The property rights doctrine and demand revelation under incomplete information. In M Boskin, editor, *Economics and human welfare*. New York Academic Press, 1979.

[7] Christopher Avery and Terrence Hendershott. Bundling and optimal auctions of multiple products. *Review of Economic Studies*, 67:483–497, 2000.

[8] Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, Dec 1999.

[9] 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.

[10] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 50–59, Chicago, IL, 2008.

[11] Shai Ben-David, Nicolo Cesa-Bianchi, David Haussler, and Philip M. Long. Characterizations of learnability for classes of {0,...,n}-valued functions. *Journal of Computer and Systems Sciences*, 50(1):74–86, 1995.

[12] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. Methodology for designing reasonably expressive mechanisms with application to ad auctions. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 2009. Earlier version in proceedings of ACM EC-08 Workshop on Advertisement Auctions.

[13] B. Douglas Bernheim and Michael D. Whinston. Menu auctions, resource allocation, and economic influence. *Quarterly Journal of Economics*, 101(1):1–32, 1986.

[14] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.

[15] 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.

[16] 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.

[17] Craig Boutilier, David Parkes, Tuomas Sandholm, and William Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2008.

[18] Chirstopher 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.

[19] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[20] 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.

[21] 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.

[22] 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.

[23] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[24] Robert E. Dansby and Cecilia Conrad. Commodity bundling. *The American Economic Review*, 74(2):377–381, May 1984.

[25] Claude d'Aspremont and Louis-Andre Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[26] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Proceedings of the Workshop on Internet Economics (WINE)*, 2007.

[27] Sorabh Gandhi, Chiranjeeb Buragohain, Lili Cao, Haitao Zheng, and Subhash Suri. A general framework for clearing auctions of wireless spectrum. In *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2007.

[28] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[29] Joseph P. Guiltinan. The price bundling of services: A normative framework. *Journal of Marketing*, 51(2):74–85, Apr 1987.

[30] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.

[31] Ward Hanson and R. Kipp Martin. Optimal bundle pricing. *Management Science*, 36(2):155–174, 1990.

[32] Friedrich Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.

[33] 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.

[34] 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.

[35] L. Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*. Amsterdam: North Holland, 1972.

[36] L. 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.

[37] 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.

[38] Leonid Hurwicz and Stanley Reiter. *Designing Economic Mechanisms*. Number 9780521724104 in Cambridge Books. Cambridge University Press, 2006.

[39] Philippe Jehiel, Moritz Meyer-Ter-Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 127(1):494–512, 2007.

[40] Philippe Jehiel and Benny Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, Sept 2001.

[41] 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.

[42] E Kushilevitz and N Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[43] 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.

[44] Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, 2005.

[45] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[46] R. Preston McAfee, John McMillan, and Michael D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, 104(2):371–383, May 1989.

[47] 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.

[48] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions. *Games and Economic Behavior*, 70:62–70, 2010.

[49] Tom Mitchell. *Machine Learning*. McGraw-Hill, 1997.

[50] Kenneth Mount and Stanley Reiter. The informational size of message spaces. *Journal of Economic Theory*, 8(2):161–192, 1974.

[51] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[52] Victor Naroditsky 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. Short paper.

[53] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.

[54] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.

[55] Shmuel Oren and Michael Rothkopf. Optimal bidding in sequential auctions. *Operations Research*, 23:1080–1090, 1975.

[56] Thomas Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51:463–484, 1983.

[57] David Parkes. Price-based information certificates for minimal-revelation combinatorial auctions. In Julian Padget, Onn Shehory, David Parkes, Norman Sadeh, and William Walsh, editors, *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Computer Science*, pages 253–276. Springer Berlin / Heidelberg, 2002.

[58] Stephen Rassenti, Vernon Smith, and Robert Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2):402–417, 1982.

[59] Amir Ronen. Mechanism design with incomplete languages. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 105–114, 2001.

[60] Robert W. Rosenthal and Ruqu Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behavior*, 17:32–55, 1996.

[61] Paat Rusmevichientong, Benjamin Van Roy, and Peter W. Glynn. A nonparametric approach to multiproduct pricing. *Operations Research*, 54(1):82–98, 2006.

[62] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002.

[63] 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.

[64] Tuomas Sandholm and Craig Boutilier. Preference elicitation in combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, pages 233–263. MIT Press, 2006. Chapter 10.

[65] 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.

[66] 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.

[67] Richard Schmalensee. Gaussian demand and commodity bundling. *The Journal of Business*, 57(1):S211–S230, Jan 1984.

[68] Tobias Schoenherr and Vincent A. Marbert. The use of bundling in B2B online reverse auctions. *Journal of Operations Management*, 2008.

[69] George J. Stigler. United States v. Loew's Inc.: A note on block-booking. *The Supreme Court Review*, 1963:152–157, 1963.

[70] Balazs Szentes and Robert W. Rosenthal. Beyond chopsticks: Symmetric equilibria in majority auction games. *Games and Economic Behavior*, 45(2):278–295, 2003.

[71] Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

[72] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[73] W. Walsh, C. Boutilier, T. Sandholm, R. Shields, G. Nemhauser, and D. Parkes. Automated channel abstraction for advertising auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2010.

[74] Jim Wilenius and Arne Andersson. Discovering equilibrium strategies for a combinatorial first price auction. In *Proceedings of E-Commerce Technology (CEC)*, 2007.

[75] 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.

[76] A. C. Yao. Some complexity questions related to distributed computing. In *Proceedings of the 11th ACM Symposium on Theory of Computing (STOC)*, pages 209–213, 1979.

[77] 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.

Appendix: Proofs of all technical claims

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

One way to construct the injective mapping is as follows. Let σ_i^j be the ith bit (or digit) of the real number that the agent expresses for dimension $j \in \{1, 2, \ldots, n\}$. Let p_k be the kth prime number. Our desired number in $\Re$ is given by,

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

□

Proof of Proposition 2. Every time we allow agent i to (semi-)shatter a new outcome, the most expressive possible 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 3. The number of possible impact vectors for agent i with k different outcomes when the other agents have types T_{-i} 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 4. 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.

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

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. There is at least one single pure strategy combination for each type vector that leads to the outcome with highest welfare, thus 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 5. First we will prove the forward implication, namely that the upper bound reaches full efficiency if any agent 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 backward implication, namely that if any agent cannot distinguish among each of the impact vectors in at least one of its fully efficient sets, then the upper bound cannot reach full efficiency.

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_i. Based on the predicate of the proposition, at least one of the impact vectors, g_i^*, corresponding to the fully efficient outcomes when agent i has type t_i^*, cannot be expressed by agent i. This means that no matter which strategies the agents other than i choose, at least one of the outcomes chosen by the mechanism when agent i has type t_i^* will be less than fully efficient. □

Proof of Proposition 6. 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 among more than one-dimensional impact vectors and there are only $|\mathcal{O}|$ such vectors. □

Proof of Corollary 1. This follows directly from Propositions 5 and 6. □

Proof of Theorem 1. The set of mechanisms allowing agent i maximum impact dimension d_i is a super-set of the mechanisms allowing agent i maximum impact 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 (for efficiency) set of impact vectors that agent i can distinguish among when restricted to $d_i^{(1)}$ vectors (i.e., the set of $d_i^{(1)}$ impact vectors that maximize the upper bound on expected efficiency). We know that there are at least $d_i^* - d_i^{(1)} \geq 1$ impact vectors corresponding to fully efficient sets of outcomes that cannot be expressed by agent i, and thus at least that many fully-efficient impact vectors 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, the new mechanism with maximum impact dimension $d_i^{(2)}$ has a strictly higher expected efficiency bound. □

Proof of Corollary 2. This follows directly from Theorem 1 and Proposition 2. □

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, $t_i^{g^{(j)}}$. Set the total welfare of all agents, as shown below, to an arbitrarily large number for every combination of joint types corresponding to the impact vectors in G_i (in an interdependent values setting their are no restrictions on the agent's utility functions, so the welfare function for each set of joint types can be constructed arbitrarily):

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

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

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.

To prove the implication for private values, 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 the 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, without loss of generality, that the outcomes in $\mathcal{O}'$ are the only outcomes that any of the agents derive utility from. We will also 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 describe how to set the utility of agent j for every outcome under every one of its types. Our construction sets agent j's utility for any 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. For a fixed one of agent j's types, all of the differences in utility for successive outcomes will be the same size. However, the gap amount (i.e., the difference in utility between o_m and o_{m-1}) will increase by an arbitrary amount for each successive successive type. This results 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 for each successive type. Formally, we will set agent j's utility function in the following way (let M be an arbitrarily large number),

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

Now for each of the $\binom{|\mathcal{O}'|}{2}$ unordered 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 (this value does not have to be negative infinity, just 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 the types in $T_i^{(a,b)}$ 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, (i.e., $(b-a) \times 2 \times M$). In other words, o_a will provide successively more utility to agent i as its type 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,

$$\left(\forall m \mid t_i^{(m)} \in T_i^{(a,b)}\right) \quad v_i(o_b, t_i^{(m)}) = M$$

$$\left(\forall m \mid t_i^{(m)} \in T_i^{(a,b)}\right) \quad v_i(o_a, t_i^{(m)}) = (m-1) \times (b-a) \times 2 \times M$$

$$\left(\forall o_j \in \mathcal{O} \setminus \mathcal{O}', \forall m \mid t_i^{(m)} \in T_i^{(a,b)}\right) v_j(o_j, t_i^{(m)}) = -\infty.$$

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.

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

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 that requires o_b to happen against $t_j^{(m)}$ and o_a against $t_j^{(m')}$ to avoid an arbitrary loss in efficiency (because any other outcome would lead to at least M less welfare).

Now, we can repeat this process for each pair of outcomes in $\mathcal{O}'$ by constructing types for agent i that select the pair. This guarantees that agent i must be able to make every pair of outcomes happen against every pair of agent j's types in the same order, or else face 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. □

Proof of Lemma 3. Let agent i's utility for outcomes o_1 and o_2 under type $t_i^{(1)}$ and $t_i^{(2)}$ be denoted as X and Y, respectively. 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 implied by this assumption held we would have the following,

$$\begin{aligned} b - a < \quad & X - Y \quad < b' - a' \\ b' - a' < \quad & X' - Y' \quad < b - a, \end{aligned}$$

Contradiction. □

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 5 (since there will always be a fully efficient set that contains every possible impact vector). In the private values setting, the backward implication is implied by Lemma 3, since it proves that it is never necessary for full efficiency in such settings to shatter any set of outcomes (only to semi-shatter them). □

Proof of Theorem 3. Let h_I^* be a pure strategy profile that achieves the expected efficiency of the bound $E[\mathcal{E}(f)]^+$. For shorthand, let $h^*(t_i)$ be agent i's expression under type t_i and the pure strategy profile h^*, and let $h^*(t_{-i})$ denote the expressions of the agents other than i under that profile.

Consider the class of payment functions, π^+, that charges agent i some constant function of the other agent's expressions minus the expected welfare of the other agents, given that agent i expresses θ_i and the other agents play the pure strategies denoted by h^*,

$$\pi_i^+(\theta_i, \theta_{-i}) = C_i(\theta_{-i}) - E_{t_{-i}}[W_{-i}(f(\theta_i, h^*(t_{-i})), t_{-i})].$$

Now, we will prove that under any payment function in the class π^+ the pure strategy profile h^* is a Bayes-Nash equilibrium. The following inequality implies that it is always (weakly) preferable in expectation over the types of the agents other than i for agent i, under any type t_i, to report $h^*(t_i)$ rather than a different expression, assuming that the other agents play according to h^* as well. (In the first equation, agent i's payment is outside of the expectation since it does not depend on the types of the other agents, and we omit the C_i terms since they do not depend on agent i's expression.)

$$E\left[v_i\left(f(h^*(t_i), h^*(t_{-i})), t_i\right)\right] - \pi_i^+\left(h^*(t_i), h^*(t_{-i})\right) \quad \geq \quad E\left[v_i\left(f(\theta_i', h^*(t_{-i})), t_i\right)\right] - \pi_i^+\left(\theta_i', h^*(t_{-i})\right)$$

$$E[v_i(f(h^*(t_i), h^*(t_{-i})), t_i)] + E[W_{-i}(f(h^*(t_i), h^*(t_{-i})), t_{-i})] \quad \geq$$
$$E[v_i(f(\theta_i', h^*(t_{-i})), t_i)] + E[W_{-i}(f(\theta_i', h^*(t_{-i})), t_{-i})].$$

The left-hand side of the final inequality is the expected welfare when the agents play the pure strategy profile h^* and agent i has type t_i. The right-hand side is the expected welfare when agent i deviates from h^* under type t_i. This inequality holds because it is impossible for any deviation from h^* to increase expected welfare. Based on our assumptions, it would have to already be reflected in h^*. □

Proof of Proposition 7. We can set the C_i in term π_i^+ for each agent to be the average expected total payment to all other agents. This amount does not depend on the agent's own expression, and it gives strong budget balance because each agent pays an equal share of the total payments made to the other agents.

The following reasoning proves that the resulting mechanism is *ex ante* individually rational (i.e., that any agent's utility for participating is always positive in expectation, prior to learning its own type). $E[u_i]$, the expected utility of agent i, is given by the following.

$$\begin{aligned}
E_t[u_i] &= E_t\left[v_i\left(f(h^*(t_i), h^*(t_{-i})), t_i\right) - \pi_i^+\left(h^*(t_i), h^*(t_{-i})\right)\right] \\
&= E_t\left[W\left(f(h^*(t)), t\right) - \frac{1}{(n-1)}\left(\sum_{j\neq i}\sum_{k\neq j} v_k\left(f(h^*(t_k), h^*(t_{-k})), t_k\right)\right)\right] \\
0 &\leq E_t\left[W(f(h^*(t)), t) - \left(v_i(f(h^*(t)), t_i) + \frac{(n-2)}{(n-1)}W_{-i}(f(h^*(t)), t_{-i})\right)\right]
\end{aligned}$$

□

Proof of Theorem 4. We will prove this by showing that the outcome function implementing our bound under a limit on expressiveness does not necessarily satisfy the weak-monotonicity (W-Mon) property, which has been shown to be a necessary condition for dominant-strategy implementation [14].

Consider the following example where agent one has three types and agent two has two types. The agents' valuations for each of three different outcomes are given below.

Agent 1:

type	A	B	C
$t_1^{(1)}$	14	0	0
$t_1^{(2)}$	0	0	0
$t_1^{(3)}$	1	0	12

Agent 2:

type	A	B	C
$t_2^{(1)}$	11	13	0
$t_2^{(2)}$	10	0	0

Under the valuations given above, the total social welfare of each outcome is given by the following table (the welfare of the most efficient outcome associated with each joint type is shown in bold).

Outcome	$t_1^{(1)}, t_2^{(1)}$	$t_1^{(1)}, t_2^{(2)}$	$t_1^{(2)}, t_2^{(1)}$	$t_1^{(2)}, t_2^{(2)}$	$t_1^{(3)}, t_2^{(1)}$	$t_1^{(3)}, t_2^{(2)}$
A	**25**	**24**	11	**10**	12	11
B	13	0	**13**	0	**13**	0
C	0	0	0	0	12	**12**

Consider a direct-revelation mechanism with a socially optimal outcome function, f. The impact vectors with the highest social welfare for agent one correspond to $[A, A], [B, A]$, and $[B, C]$ (these are the outcomes with the greatest welfare under each combination of types). If we are forced to design an outcome function that limits agent one to maximum impact dimension $d_1 \leq 2$ and the type $t_1^{(2)}$ is highly unlikely (e.g., $P(t_1^{(2)}) = \epsilon$), then the outcome function with the highest expected welfare will provide agent one with the impact vectors, $[A, A], [A, A]$ and $[B, C]$.

The W-Mon property states that the following inequality must hold for all t_1, t_1', and t_2,

$$v_1(f(t_1, t_2), t_1) - v_1(f(t_1', t_2), t_1) \quad \geq \quad v_1(f(t_1, t_2), t_1') - v_1(f(t_1', t_2), t_1').$$

If we use $t_1^{(2)}$ and $t_1^{(3)}$ for t_1 and t_1', respectively, and $t_2^{(1)}$ for t_2, then we can rewrite the inequality for our limited-expressiveness mechanism as follows,

$$\begin{aligned} v_1(f(t_1^{(2)}, t_2^{(1)}), t_1^{(2)}) - v_1(f(t_1^{(3)}, t_2^{(1)}), t_1^{(2)}) \quad &\geq \quad v_1(f(t_1^{(2)}, t_2^{(1)}), t_1^{(3)}) - v_1(f(t_1^{(3)}, t_2^{(1)}), t_1^{(3)}) \\ v_1(A, t_1^{(2)}) - v_1(B, t_1^{(2)}) \quad &\geq \quad v_1(A, t_1^{(3)}) - v_1(B, t_1^{(3)}). \end{aligned}$$

This inequality is violated by the valuation functions in our example, so the inexpressive mechanism cannot be implemented in dominant strategies. □

Proof of Proposition 8. Suppose the opposite: that agent i has maximum impact dimension d_i under f, and there exists an outcome function, f', that emulates f while providing agent i with less than d_i expressions. Let G_i be one of the largest sets of impact vectors that the agent can distinguish among under f (i.e., $|G_i| = d_i$ and $D_i(G_i)$ is true), let h_{-i} be a strategy by the other agents that allows i to distinguish among the vectors in G_i under f, and let q be the mapping that allows f' to emulate f.

Since d_i expressions are needed to distinguish among each of the different impact vectors in G_i, our assumption implies that there will be at least two distinct impact vectors in G_i that agent i cannot distinguish among under f'. Let these be denoted as $g_i^{(1)}$ and $g_i^{(2)}$. Since $g_i^{(1)}$ and $g_i^{(2)}$ are distinct, there must be at least one joint type for the agents other than i, t_{-i}, such that they map to different outcomes. Furthermore, the impact vectors are expressible under f, so it must be possible for agent i to cause both the outcome mapped by $g_i^{(1)}$ and the outcome mapped by $g_i^{(2)}$ under t_{-i} to be chosen by f (i.e., there exists a $\theta_i^{(1)}$ and $\theta_i^{(2)}$ such that $f(\theta_i^{(1)}, h_{-i}(t_{-i})) \neq f(\theta_i^{(2)}, h_{-i}(t_{-i}))$).

However, since agent i cannot distinguish between $g_i^{(1)}$ and $g_i^{(2)}$ under f', $\theta_i^{(1)}$ and $\theta_i^{(2)}$ must map to the same expression under q. Thus, we get the following starting from the equation above,

$$f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))) \neq f'(q_i(\theta_i^{(2)}), q_{-i}(h_{-i}(t_{-i})))$$
$$f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))) \neq f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))).$$

Contradiction. □

Proof of Lemma 4. If one of the agents other than i lies, the number of impact vectors that agent i can distinguish among can only decrease. Any impact vector that was distinct because of an outcome chosen under the type that is now being reported untruthfully will no longer be distinct, and no new impact vectors can become distinct because of the lie. □

Proof of Proposition 9. We have already shown in Proposition 8 that no outcome function can emulate f using fewer expressions for each agent, i, than its maximum impact dimension under f, d_i. We will now show that if f is a direct-revelation outcome function, it is always possible to emulate it using exactly d_i expressions for each agent. To prove this, we will use Lemma 4, which implies that we can construct a mapping from expressions under f to expressions under f' such that any two types resulting in the same impact vector under f are also mapped to the same expression under f'. We can set the outcomes in f' as follows,

$$f'(q_i(t_i), q_{-i}(t_{-i})) = f(t_i, t_{-i}).$$

This will produce a valid mapping because the types that map to to the same expression will result in the same outcome for all joint types of the other agents. □

Proof of Corollary 3. From Theorem 1, we know that this is true for maximum impact dimension. From Proposition 8, we know that the set of outcome functions that limit agent i to d_i expressions does not include any outcome functions where agent i has maximum impact dimension greater than d_i. A direct-revelation mechanism can always be designed to maximize the bound subject to a limit on expressiveness, and that outcome function can be emulated by one that provides d_i expressions to each agent i. □

Proof of Corollary 4. This follows directly from Theorem 2, which states that an agent may need to shatter its entire outcome space to avoid arbitrary inefficiency, and Proposition 3, which states that $|T_{-i}|^{|O|}$ expressions are needed needed to shatter an outcome space. □

Proof of Lemma 5. Let Θ_{-i} be expressions for the agents other than i that allow agent i to shatter T_{-i} (if f is a direct-revelation mechanism these two sets will be identical). Construct a total ordering of Θ_{-i} so that for any pair, θ^j_{-i} and θ^k_{-i} (where $j < k$), and any expression by agent i, $\theta_i \in \Theta_i$, that causes the mechanism to choose A and B when the other agents express θ^j_{-i} and θ^k_{-i}, A is chosen for θ^j_{-i} and B for θ^k_{-i}. If this condition is not met, we can simply switch j and k. Re-labeling all of the expressions to satisfy this condition is possible because of the semi-shattering requirement that, under any strict ordering of the expressions of the agents other than i, all expressions by agent i in Θ_i that cause outcomes A and B to be chosen by f do so in the same order.

Consider a subset of expressions by agent i, Θ'_i, that allow it to choose between A and B for each of the immediately subsequent, or neighboring, pairs of expressions by the agents other than i. In other words, there exists at least one $\theta'_i \in \Theta'_i$, such that f chooses A and B under the expression pairs $(\theta'_i, \theta^j_{-i})$ and $(\theta'_i, \theta^{j+1}_{-i})$, for all j. Additionally, order Θ'_i so that an expression that chooses between A and B when the other agents express θ^1_{-i} and θ^2_{-i} is the first in the ordering. This expression is then followed in the ordering by one that chooses between A and B when the other agents express θ^2_{-i} and θ^3_{-i}, and so on until the last of the neighboring pairs of θ_{-i}'s is reached.

Under this ordering, none of the $|T_{-i}| - 1$ pairs of expressions along the diagonal of an input matrix corresponding to agent i's expressions in Θ'_i (i.e., pairs of the form $(\theta'^j_i, \theta^j_{-i})$) can be in the same A-monochromatic rectangle. To see this, consider that whenever θ'^j_i is matched against an expression larger than θ^j_{-i}, the outcome function does not choose A (since we ensured that A's always come before B's). Thus, all of the A's in the reduced input matrix corresponding to Θ_{-i} and Θ'_i are to the left of the diagonal (in a triangle pattern) and the diagonal is all A's. This implies that each of the pairs to the left of the diagonal must be in a different A-monochromatic rectangle, since they all result in the same value but are different when crossed with any other member of the set. We can reverse the total ordering of Θ_{-i} and make the same argument for B-monochromatic rectangles. □

Proof of Theorem 5. The upper bound follows directly from Propositions 8 and 9, since they imply that an agent's maximum impact dimension is an upper bound on the number of messages needed by the best communication protocol for running f. The lower bound follows from Lemma 5, since it implies that if agent i can semi-shatter a set of types containing some outcome o, T^o_{-i}, there must be at least $|T^o_{-i}| - 1$ monochromatic rectangles for outcome o in any partitioning of f's input matrix from i's perspective. It has been previously shown that each rectangle requires at least one message in any communication protocol [42]. □

Proof of Proposition 10. We will prove this statement for the semi-shatterable outcome dimension, k_i, which will imply it is true for maximum impact dimension as well (based on Proposition 2).

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 semi-shatter over in this mechanism be $\mathcal{O}'$. 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 $o^* \in \mathcal{O}^*$. 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, $\mathcal{O}' \cup \{o^*\}$.

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 used to select. 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 other outcome, o', did previously.

$$\begin{aligned}(\forall j \mid 1 \leq j \leq c_i)\, \theta^*_{i,j} &= \theta_{i,j} \\ \theta^*_{i,c_1+1} &= \sum_{j \in C_i(o')} \theta_{ij} + \epsilon.\end{aligned}$$

Since agent i can do this 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).[12] Thus, agent i can semi-shatter the new larger outcome set using the additional channel. □

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 10 and Lemma 2 (since increasing the number of channels by one can be used to increase the agent's semi-shatterable outcome dimension). □

Proof of Proposition 11. This proof is based on a pigeon hole argument. With fewer than $\lceil log_2(|\mathcal{O}|)\rceil$ channels there will be at least two 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,

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

This will prevent the agent from forcing the mechanism to choose both of those outcomes with different expressions, since the agent's own contribution to the two outcomes will always be identical. □

Proof of Proposition 12. We will show that no agent can shatter any set of two outcomes against any two 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 to shatter 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 total channel value connected to outcome A be X and let its total channel value connected to B be Y. Consider two types for the agents other

[12] 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).

than i, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, and the reports mapped to them under *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 following.

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

Contradiction. □

Proof of Corollary 6. This follows directly from Proposition 12 and Lemmas 1 and 2. □

Proof of Proposition 13. With $|\mathcal{O}| - 1$ 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,

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

Since the $v_j(o_j, t_j)$ 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 calculation. This shows that the payment rule can be properly calculated even when each agent has 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 two sharing a channel, making it impossible for that agent to express arbitrary non-linear utility for every outcome (a requirement for implementing a VCG mechanism), or it will have two outcomes without a channel, making it impossible for that agent to express any preference for one of the outcomes. □

Proof of Lemma 6. We will prove only the forward implication. Once that is proved, the backward implication will be trivial since we can just switch the labels of C and D. From the predicate, we have $(A \setminus C = D \setminus B)$, which implies that every element, x, in A and not in D is in either B or C, and every element, y, in A and B must also be in C (if y were not in either C or D, or if y were in only D, it would contradict the predicate). Thus, the difference between A and D must be contained completely in C (i.e., $(A \setminus D) \subseteq C$). The following reasoning proves the rest of the claim,

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

The last step is valid because we know that no elements from D can be in the set on the right-hand side, once all the operators are applied (since the left-hand side involves removing all elements in D from A). Thus, it cannot make a difference if we leave them in B before subtracting it from C, since the set minus operator in the parentheses on the right-hand side only serves to maintain the elements from D in the resulting set. This same logic can be repeated for other half of the conjunction in the predicate. □

Proof of Lemma 7. From Lemma 6, 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),

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

We will assume, for contradiction, that agent i can semi-shatter both pairs of outcomes, $\{A, B\}$ and $\{C, D\}$. From Observation 1, 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 considering 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.

We will 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).

$$A \text{ happens against } 1 \begin{cases} S^A + a_1 > S^B + b_1 \\ S^A + a_1 > S^C + c_1 \\ S^A + a_1 > S^D + d_1 \end{cases}$$

$$B \text{ happens against } 2 \begin{cases} S^B + b_2 > S^A + a_2 \\ S^B + b_2 > S^C + c_2 \\ S^B + b_2 > S^D + d_2 \end{cases}$$

Let the difference between the sum of channels in S^A and S^C be denoted S_1 (i.e., $S^A - S^C = S_1$). From the predicate, we have 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. Additionally, 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. Let the difference in the sum of the channels in S^A and S^D be denoted by S_2. This leads the following equality, which is implied by Lemma 6: $S^A - S^D = S^C - S^B = S_2$. Now the equations above simplify to the following.

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

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.

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

Our assumption that agent i could semi-shatter both sets of outcomes under 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,

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

Contradiction. □

Proof of Theorem 6. We will prove this by providing a distribution over valuations such that a channel-based mechanism that treats agent one's bid on any bundle Q to be the sum of its bids on some two other non-overlapping bundles, q_1 and q_2, cannot achieve within 5% of the maximum expected efficiency.

We will first show that, in such a mechanism, agent one cannot choose between the pairs of outcomes where it wins q_1 or q_2, and Q or nothing, since the channels connected to these outcomes overlap in the fashion described in Lemma 7. Let A be an outcome under which agent one is allocated bundle q_1, let B be an outcome under which it is allocated q_2, C for Q 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 the agent). Since agent one's bid on Q equals the sum of its bid on q_1 and q_2, we have that $S^C = S^A \cup S^B$, and its bid for the outcome where it wins nothing is always 0, so we have $S^D = \emptyset$. These sets of channels meet the conditions of Lemma 7.

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

Now, consider the following example with two agents, where each agent has two equally likely types: a "substitutable" type and a "complementary" type. The agents' valuations for the bundles q_1, q_2, and Q are given below. Valuations for all other bundles, including the empty bundle, are assumed to be 0 or their minimum possible value. Since the items other than those in q_1, q_2, and Q provide no utility to either agent, and the bids for these items cannot affect how the items in q_1, q_2, and Q are allocated, we ignore the additional items in the rest of our proof.

Agent 1

type	q_1	q_2	Q
t_1^s	0.5	0.5	0.5
t_1^c	0	0	0.75

Agent 2

type	q_1	q_2	Q
t_2^s	0	0.5	0.5
t_2^c	0.75	0	1

Under the valuations given above, the total social welfare of each outcome is given by the following table (the welfare of the most efficient outcome associated with each joint type is shown in bold).

Outcome	t_1^s, t_2^s	t_1^s, t_2^c	t_1^c, t_2^s	t_1^c, t_2^c
$A : \{q_1, q_2\}$	**1**	0.5	0.5	0
$B : \{q_2, q_1\}$	0.5	**1.25**	0	0.75
$C : \{Q, \emptyset\}$	0.5	0.5	**0.75**	0.75
$D : \{\emptyset, Q\}$	0.5	1	0.5	**1**

The maximum expected efficiency, $E[\mathcal{E}^*]$, is then given by the following. We drop the t_1 and t_2 notation in favor of shorthand where types are simply referred to as s or c. P_{sc} denotes the probability of agent one having type s and agent two having type c.

$$\begin{aligned} E[\mathcal{E}^*] &= P_{ss}W(A, \{s, s\}) + P_{sc}W(B, \{s, c\}) + P_{cs}W(C, \{c, s\}) + P_{cc}W(D, \{c, c\}) \\ &= \frac{1}{4} \times 1 + \frac{1}{4} \times 1.25 + \frac{1}{4} \times 0.75 + \frac{1}{4} \times 1 = 1. \end{aligned}$$

Since agent one cannot choose between the pairs of outcomes $\{A, B\}$ and $\{C, D\}$, the mechanism cannot achieve the expected efficiency of the optimal allocation for some combination of types. In the best case, it will assign the second best outcome for one of the $\{s, c\}$, $\{c, s, \}$, or $\{c, c\}$ types, which will cost at least 6.25% in expected efficiency. □