

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**IMPLEMENTATION WITH A
BOUNDED ACTION SPACE**

by

**LIAD BLUMROSEN
and MICHAL FELDMAN**

Discussion Paper # 412

December 2005

מרכז לחקר הרציונליות
**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Implementation with a Bounded Action Space

Liad Blumrosen and Michal Feldman*

Abstract

While traditional mechanism design typically assumes isomorphism between the agents' type- and action spaces, in many situations the agents face strict restrictions on their action space due to, e.g., technical, behavioral or regulatory reasons. We devise a general framework for the study of mechanism design in single-parameter environments with restricted action spaces. Our contribution is threefold. First, we characterize sufficient conditions under which the information-theoretically optimal social-choice rule can be implemented in dominant strategies, and prove that any multilinear social-choice rule is dominant-strategy implementable with no additional cost. Second, we identify necessary conditions for the optimality of action-bounded mechanisms, and fully characterize the optimal mechanisms and strategies in games with two players and two alternatives. Finally, we prove that for any multilinear social-choice rule, the optimal mechanism with k actions incurs an expected loss of $O(\frac{1}{k^2})$ compared to the optimal mechanisms with unrestricted action spaces. Our results apply to various economic and computational settings, and we demonstrate their applicability to signaling games, public-good models and routing in networks.

1 Introduction

Mechanism design is a sub-field of game theory that studies how to design rules of games resulting in desirable outcomes, when the players are rational. In a standard setting, players hold some private information – their “types” – and choose “actions” from their action spaces to maximize their utilities. The social planner wishes to implement a social-choice function, which maps each possible state of the world (i.e., a profile of the players' types) to a single alternative. For example, a government that wishes to undertake a public-good project (e.g., building a bridge) only if the total benefit for the players exceeds its cost.

Much of the literature on mechanism design restricts attention to *direct revelation* mechanisms, in which a player's action space is identical to his type space. This focus is owing to the *revelation principle* that asserts that if some mechanism achieves a certain result in an equilibrium, the same result can be achieved in a truthful one – an equilibrium where each agent simply reports his private type [15].

Nonetheless, in many environments, direct-revelation mechanisms are not viable since the actions available for the players have a limited expressive power. Consider, for example, the well-studied “screening” model, where an insurance firm wishes to sell different types of policies to different drivers based on their caution levels, which is their private information. In this model, drivers may have a continuum of possible caution levels, but insurance companies offer

*Email: {liad,mfeldman}@cs.huji.ac.il. School of Engineering and Computer Science. The Hebrew University of Jerusalem, Jerusalem, Israel.

only a few different policies since it might be either infeasible or illegal to advertise and sell more than few contracts.

There are various reasons for such strict restrictions on the action spaces. In some situations, players might be reluctant to reveal their accurate type, but willing to disclose partial information about it. For example, agents who are engaged in repeated-games will typically be unwilling to reveal their type, even if it is beneficial for them in the short run, since it might harm them in future transactions. Agents may also not trust the mechanism to keep their valuations private [16], or not even know their exact type while computing it may be expensive [11]. Limitations on the action space can also be caused by technical constraints, such as severe restrictions on the communication lines [5] or from the the need to perform quick transactions (e.g., discrete bidding in English auctions [13]).

Restrictions on the action space, for specific models, were studied in several earlier papers. The work of Blumrosen, Nisan and Segal [4, 6, 5] is the closest in spirit to this paper. They studied single-item auctions where bidders are allowed to send messages with severely bounded size. They characterized the optimal mechanisms under this restriction, and showed that nearly optimal results can be achieved even with very strict limitations on the action space. Other work studied similar models for the analysis of discrete-bid ascending auctions [13, 10, 8, 7], take-it-or-leave-it auctions [17], or for measuring the effect of discrete “priority classes” of buyers on the performance of electricity markets [19, 14]. Our work generalizes the main results of Blumrosen et al. to a general mechanism-design framework that can be applied to a multitude of models. We show that some main properties proved by Blumrosen et al. are preserved in more general frameworks (for example, that dominant-strategy equilibrium can be achieved with no additional cost, and that the loss diminishes with the number of possible actions in a similar rate), where some other properties do not always hold (for example, that asymmetric mechanisms are optimal and that players must always use all their action space).

A standard mechanism design setting is composed of agents with private information (their “types”), and a social planner, who wishes to implement a *social choice function*, c – a function that maps any profile of the agents’ types into a chosen alternative. A classic result in this setting says that under some monotonicity assumption on the agents’ preferences – the “single-crossing” assumption (see definition below) – a social-choice function is implementable in dominant strategies if and only if it is monotone in the players’ types. However, in environments with restricted action spaces, the social planner cannot typically implement every social-choice function due to inherent informational constraints. That is, for some realizations of the players’ types, the decision of the social planner will be incompatible with the social-choice function c . In order to quantitatively measure how well bounded-action mechanisms can approximate the original social-choice functions, we follow a standard assumption that the social choice function is derived from a *social-value* function, g , which assigns a real value for every alternative and realization of the players’ types. The social-choice function c will therefore choose an alternative that maximizes the social value function, given the type vector $\vec{\theta} = (\theta_1, \dots, \theta_n)$, i.e., $c(\vec{\theta}) = \operatorname{argmax}_A \{g(\vec{\theta}, A)\}$. Observe that the social-value function is not necessarily the social *welfare* function – the social welfare function is a special case of g in which g is defined to be the sum of the players’ valuations for the chosen alternative. Following are several simple examples of social-value functions:

- *Public goods.* A government wishes to build a bridge only if the sum of the benefits that agents gain from it exceeds its construction cost C . The social value functions in a 2-player game will therefore be: $g(\theta_1, \theta_2, \text{“build”}) = \theta_1 + \theta_2 - C$, and $g(\theta_1, \theta_2, \text{“do not build”}) = 0$.

- *Routing in networks.* Consider a network that is composed of two links in parallel. Each link has a secret probability p_i of transferring a message successfully. A sender wishes to send his message through the network only if the probability of success is greater than, say, 90 percent - the known probability in an alternate network. That is, $g(p_1, p_2, \text{“send in network”}) = 1 - (1 - p_1) \cdot (1 - p_2)$, and $g(p_1, p_2, \text{“send in alternate network”}) = 0.9$.
- *Single-item auctions.* Consider a 2-player auction, where the auctioneer wishes to allocate the item to the player who values it the most. The social choice function is given by: $g(\theta_1, \theta_2, \text{“player 1 wins”}) = \theta_1$ and for the second alternative is $g(\theta_1, \theta_2, \text{“player 2 wins”}) = \theta_2$.

1.1 Our Contribution

In this paper, we present a general framework for the study of mechanism design in environments with a limited number of actions. We assume a Bayesian model where players have one-dimensional private types, independently distributed on some real interval.

The main question we ask is: when agents are only allowed to use k different actions, which mechanisms achieve the optimal expected social-value? Note that this question is actually composed of two separate questions. The first question is an information-theoretic question: what is the optimal result achievable when the players can only reveal information using these k actions (recall that their type space may be continuous). The other question involves game-theoretic considerations: what is the best result achievable with k actions, where this result should be achieved in a dominant-strategy equilibrium. These questions raise the question about the “*price of truthfulness*”: can the optimal information-theoretic result always be implemented in a dominant-strategy equilibrium? And if not, to what extent does the dominant-strategy requirement degrade the optimal result? What we call “the price of truthfulness” was also explored in other contexts in game theory where computational restrictions apply: for example, is it always true that the optimal polynomial-time approximation ratio (for example, in combinatorial auctions) can be achieved in equilibrium? (The answer for this interesting problem is still unclear, see, e.g., [3, 2, 12].)

Our first contribution is the characterization of sufficient conditions for implementing the optimal information-theoretic social-choice rule in dominant strategies. We show that for the family of *multilinear* social-value functions (that is, polynomials where each variable has a degree of at most one in each monomial) the dominant-strategy implementation incurs no additional cost.

Theorem: *Given any multilinear single-crossing social-value function, and for any number of alternatives and players, the social choice rule that is information-theoretically optimal is implementable in dominant strategies.*

Multilinear social-value functions capture many important and well-studied models, and include, for instance, the routing example given above, and any *social welfare* function in which the players’ valuations are linear in their types (such as public-goods and auctions).

The implementability of the information-theoretically optimal mechanisms enables us to use a standard routine in Mechanism Design and first determine the optimal social-choice rule, and then calculate the appropriate payments that ensure incentive compatibility. To show this result, we prove a useful lemma that gives another characterization for social-choice functions whose “price of truthfulness” is zero. We show that for any social-choice function, incentive compatibility in action-bounded mechanisms is equivalent to the property that the optimal

expected social value is achieved with *non-decreasing* strategies (or *threshold strategies*).¹ In other words, this lemma implies that one can always implement, with dominant strategies, the best social-choice rule that is achievable with non-decreasing strategies.

Our second contribution is in characterizing the optimal action-bounded mechanisms. We identify some necessary conditions for the optimality of mechanisms in general, and using these conditions, we fully characterize the optimal mechanisms in environments with two players and two alternatives. The optimal mechanisms turn out to be “diagonal” – that is, in their matrix representation, one alternative will be chosen in, and only in, entries below one of the main diagonals (this term extends the concept of “Priority Games” used in [5] for bounded-communication auctions). We complete the characterization of the optimal mechanisms with the depiction of the optimal strategies – strategies that are “mutually maximizers”. Since the payments in a dominant-strategy implementation are uniquely defined by a monotone allocation and a profile of strategies, this also defines the payments in the mechanism. We give an intuitive proof for the optimality of such strategies, generalizing the concept of optimal “mutually-centered” strategies from [4]. Surprisingly, as opposed to the optimal auctions in [4], for some non-trivial social-value functions, the optimal “diagonal” mechanism may not utilize all the k available actions.

Theorem: *For any multilinear single-crossing social-value function over two alternatives, the informationally optimal 2-player k -action mechanism is diagonal, and the optimal dominant strategies are mutually-maximizers.*

Achieving a full characterization of the optimal action-bounded mechanism for multi-player or multi-alternative environments seems to be harder. To support this claim, we observe that the number of mechanisms that satisfy the necessary conditions above is growing exponentially in the number of players.

Our next result compares the expected social-value in k -action mechanisms to the optimal expected social value when the action space is unrestricted. For any number of players or alternatives, and for any profile of independent distribution functions, we construct mechanisms that are nearly optimal – up to an additive difference of $O(\frac{1}{k^2})$. This result is achieved in dominant strategies.

Theorem: *For any multilinear social-value function, the optimal k -action mechanism incurs an expected social loss of $O(\frac{1}{k^2})$.*

This is the same asymptotic rate proved for specific environments in [19, 13, 5]. Note that there are social-choice functions that can be implemented with k actions with no loss (for example, the rule “always choose alternative A ”). However, we know that in some settings (e.g., auctions [5]) the optimal loss may be proportional to $\frac{1}{k^2}$, thus a better general upper bound is impossible.

Finally, we present our results in the context of several natural applications. First, we give an explicit solution for a public-good game with k -actions. We show that the optimum is achieved in symmetric mechanisms (in contrast to action-bounded auctions [5]), and that the optimal allocation scheme depends on the value of the construction cost C . Then, we study the celebrated *signaling* model, in which potential employees send signals about their skills to their potential employers by means of the education level they acquire. This is a natural application

¹The restriction to non-decreasing strategies is very common in the literature. One remarkable result by Athey [1] shows that when a non-decreasing strategy is a best response for any other profile of non-decreasing strategies, a pure Bayesian-Nash equilibrium must exist.

in our context since education levels are often discrete (e.g., B.A, M.A and PhD). Lastly, we present our results in the context of routing in networks, where it is reasonable to assume that links report whether they have low or high loss rates, but less reasonable to require them to report their accurate loss rates.

The rest of the paper is organized as follows: our model and notations are described in Section 2. We then describe our general results regarding implementation in multi-player and multi-alternative environments in Section 3, including the asymptotic analysis of the social-value loss. In Section 4, we fully characterize the optimal mechanisms for 2-player environments with two alternative. In Section 5, we conclude with applying our general results to several well-studied models. All proofs can be found in the appendix.

2 Model and Preliminaries

We first describe a standard mechanism-design model for players with one-dimensional types. Then, in Subsection 2.2, we impose limitation of the action space. The general model studies environments with n players and a set $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$ of m alternatives. Each player has a privately known type $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$ (where $\underline{\theta}_i, \bar{\theta}_i \in \mathbb{R}$, $\underline{\theta}_i < \bar{\theta}_i$), and a type-dependent valuation function $v_i(\theta_i, A)$ for each alternative $A \in \mathcal{A}$. In other words, player i with type θ_i is willing to pay an amount of $v_i(\theta_i, A)$ for alternative A to be chosen. Each type θ_i is independently distributed according to a publicly known distribution F_i , with an always positive density function f_i . We denote the set of all possible types' profiles by $\Theta = \times_{i=1}^n [\underline{\theta}_i, \bar{\theta}_i]$.

The social planner has a *social-choice function* $c : \Theta \rightarrow \mathcal{A}$, where the choice of alternatives is made in order to maximize a *social-value function* $g(\vec{\theta}) : \Theta \times \mathcal{A} \rightarrow \mathbb{R}$. That is, $c(\vec{\theta}) \in \operatorname{argmax}_{A \in \mathcal{A}} \{g(\vec{\theta}, A)\}$

We assume that for every alternative $A \in \mathcal{A}$, the function $g(\cdot, A)$ is continuous and differentiable in every type. Since the players' types are private information, in order to choose the optimal alternative, the social planner needs to get the players' types as an input. The players reveal information about their types by choosing an *action*, from an action set B .

Each player uses a *strategy* for determining the action he plays for any possible type. A strategy for player i is therefore a function $s_i : [\underline{\theta}_i, \bar{\theta}_i] \rightarrow B$. We denote a profile of strategies by $s = s_1, \dots, s_n$ and the set of the strategies of all players except i by s_{-i} . The utility of player i of type θ_i from alternative A under the payment p_i is $u_i = v_i(\theta_i, A) - p_i$.

2.1 Dominant-Strategy Implementation

Following is a standard definition of a mechanism. The action space B is traditionally implicit, but we mention it explicitly since we later examine limitations on B .

Definition 1. *A mechanism with an action set B is a pair (t, p) where:*

- $t : B^n \rightarrow \mathcal{A}$ is the allocation rule.²
- $p : B^n \rightarrow \mathbb{R}^n$ is the payment scheme (i.e., $p_i(b)$ is the payment to the i th player given a vector of actions b).

²We will show that, w.l.o.g., we can focus on deterministic mechanisms.

The main goal of this paper is to optimize the expected social value (in action-bounded mechanisms) while preserving a dominant-strategy equilibrium.

We say that a strategy s_i is *dominant* for player i in mechanism (t, p) if player i cannot increase his utility by reporting a different action than $s_i(\theta_i)$, regardless of the actions of the other players b_{-i} .³

Definition 2. We say that a social-choice function h is implementable with a set of actions B if there exists a mechanism (t, p) with a dominant-strategy equilibrium s_1, \dots, s_n (where for each i , $s_i : [\underline{\theta}_i, \bar{\theta}_i] \rightarrow B$) that always chooses an alternative according to h , i.e., $t(s_1(\theta_1), \dots, s_n(\theta_n)) = h(\vec{\theta})$.

A fundamental result in the mechanism-design literature states that under reasonable conditions, monotonicity of the social-choice function is a sufficient and necessary condition for dominant-strategy implementability (in single-parameter environments). For defining monotonicity, the preferences of the players must exhibit some order on the alternatives. Namely, each player has a complete, weak transitive order \succeq_i over the alternatives in \mathcal{A} . If $A \succeq_i B$ but not $B \succeq_i A$, we use the notation $A \succ_i B$. If both $A \succeq_i B$ and $B \succeq_i A$ we use the notation $A \sim_i B$.⁴

Monotonicity also requires defining an order on the actions. In standard settings, the order is defined by an order on the real numbers (e.g., in direct revelation mechanisms where each type is drawn from a real interval). When the action space is discrete, the order can be determined by the names of the actions, for example, “0”, “1”, ..., “ $k-1$ ” for k -action mechanisms. (We therefore describe this order with the standard relation on natural numbers $<, >$.)

Given these orders, we can now define the notion of monotonicity:

Definition 3. A deterministic mechanism is monotone if when player i raises his reported action, and fixing the actions of the other players, the mechanism never chooses an inferior alternative for i . That is, for any $b_{-i} \in \{0, \dots, k-1\}^{n-1}$ if $b'_i > b_i$ then $t(b'_i, b_{-i}) \succeq_i t(b_i, b_{-i})$.

The last ingredient in the characterization of incentive compatibility in the classic model requires that the valuations of the players will exhibit the *single-crossing* property (also known as Spence-Mirrlees condition). In our model, the single-crossing property implies that for every player, the effect of an increment in the player’s type on the player’s valuation is greater as the alternative is higher in this player’s order \succeq_i . Throughout the paper we assume that the players have single-crossing valuations.

Definition 4. A function $h : \Theta \times \mathcal{A} \rightarrow \mathbb{R}$ is single crossing with respect to θ_i if for any two alternatives $A_j \succ_i A_i$ we have,

$$\frac{\partial h(\vec{\theta}, A_j)}{\partial \theta_i} > \frac{\partial h(\vec{\theta}, A_i)}{\partial \theta_i}$$

and if $A_j \sim A_i$ then $h(\cdot, A_j) \equiv h(\cdot, A_i)$ (i.e., the functions are identical).

Following is a classic result regarding the implementability of social-choice functions in single-parameter environments. This result can be found in different forms, very often implicit, in

³That is, for every type θ_i and every action b'_i , $v_i(\theta_i, t(s_i(\theta_i), b_{-i})) - p_i(s_i(\theta_i), b_{-i}) > v_i(\theta_i, t(b'_i, b_{-i})) - p_i(b'_i, b_{-i})$

⁴For example, in an auction model the alternatives may be A =”player 1 wins”, B =”player 2 wins” and C =”player 3 wins”. In this case, $A \succ_1 B$ and $B \succ_2 A$ but $B \sim_1 C$.

almost every paper on mechanism design in one-dimensional domains. This characterization assumes single-crossing preferences; without this assumption, general sufficient conditions for implementability are not known (for a survey on this topic see [9]). Note, however, that this characterization does not hold when the action space is bounded.

Proposition 1. *Assume that the valuation functions $v_i(\theta_i, A)$ are single crossing and that the action space is unrestricted. A social-choice function c is dominant-strategy implementable if and only if c is monotone.*

2.2 Action-Bounded Mechanisms

The set of actions B is usually implicit in the literature, and it is assumed to be isomorphic to the type space. In this paper, we study environments where this assumption does not hold. We define a k -action game to be a game in which the number of possible actions for each player is k , i.e., $|B| = k$. In k -action games, the social planner typically cannot always choose an alternative according to the social choice function c due to the informational constraints. Instead, we are interested in implementing a social-choice function that, with k actions, maximizes the *expected social value*: $E_{\vec{\theta}} g(\vec{\theta}, t(s_1(\theta_1), \dots, s_n(\theta_n)))$.

Definition 5. *We say that a social-choice function $h : \Theta \rightarrow \mathcal{A}$ is informationally achievable with a set of actions B if there exists a profile of strategies s_1, \dots, s_n (where for each i , $s_i : [\underline{\theta}_i, \bar{\theta}_i] \rightarrow B$), and an allocation rule $t : B^n \rightarrow \mathcal{A}$, such that t chooses the same alternative as h for any type profile, i.e., $t(s_1(\theta_1), \dots, s_n(\theta_n)) = h(\vec{\theta})$. If $|B| = k$, we say that h is k -action informationally achievable.*

Note that this definition does not take into account strategic considerations. For example, consider an environment with two alternatives $\mathcal{A} = \{A, B\}$, and the following social-choice function: $\tilde{c}(\theta_1, \theta_2) = A$ iff $\theta_1 > 1/2$ and $\theta_2 > 1/2$. \tilde{c} is informationally achievable with two actions: if both players bid “0” when their value is greater than 1/2 and “1” otherwise, then the allocation rule “choose alternative A iff both players report 1” derives exactly the same allocation for every profile of types. In contrast, it is easy to see that the function $\hat{c}(\theta_1, \theta_2) = A$ iff $\theta_1 + \theta_2 > 1/2$ is not informationally achievable with two actions.

We now define a social-choice rule that maximizes the social value under the information-theoretic constraints that are implied by the limitations on the number of actions.

Definition 6. *A social-choice function is k -action informationally optimal with respect to the social-value function g , if it achieves the maximal expected social value among all the k -action informationally achievable social-choice functions.⁵*

Earlier in this section, we defined the single-crossing property for the players valuations. We now define a single-crossing property on the social-value function g . This property clearly ensures the monotonicity of the corresponding social choice rule, and we will later show that it is also useful for action-bounded environments.

Definition 7. *We say that the social-choice rule $g(\vec{\theta}, A)$ exhibits the single-crossing property if for every player i , g exhibits the single-crossing property with respect to θ_i .*

⁵The optimal function is well defined since there is a finite number of such functions.

Note that the definition above requires that g will be single crossing with respect to every player i , given her individual order \succeq_i on the alternatives.

Finally, we call attention to a natural set of strategies – “non-decreasing” strategies, where each player reports a higher action as her type increases. Equivalently, such strategies are *threshold strategies* – strategies where each player divides his type support into intervals, and simply reports the interval in which her type lies.

Definition 8. A real vector $x = (x_0, x_1, \dots, x_k)$ is a vector of threshold values if $x_0 \leq x_1 \leq \dots \leq x_k$.

Definition 9. A strategy s_i is a threshold strategy based on a vector of threshold values $x = (x_0, x_1, \dots, x_k)$, if for any action j it holds that $s_i(\theta_i) = j$ iff $\theta_i \in [x_j, x_{j+1}]$. A strategy s_i is called a threshold strategy, if there exists a vector x of threshold values such that s_i is threshold strategy based on x .

3 Implementation with a Limited Number of Actions

In this section, we study the general model of action-bounded mechanism design. Our first result is a sufficient and necessary condition for the implementability of the optimal solution achievable with k actions: this condition says that the optimal social-choice rule is achieved when all the players use non-decreasing strategies. The basic idea is that with non-decreasing strategies (i.e., threshold strategies), we can apply the single-crossing property to show that when a player raises his reported action, the expected value for his high-priority alternatives increases faster; therefore, monotonicity must hold. The result holds for any number of players and alternatives, and for any profile of distribution functions on the players’ types, as long as they are statistically independent. (It is easy to illustrate that this result does not hold if the players’ types are dependent.)

Lemma 1. Consider a single-crossing social-value function g . The informationally optimal k -action social-choice function c^* (with respect to g) is implementable if and only if c^* achieves its optimum when the players use non-decreasing strategies.

Next, we show that for a wide family of social-value functions – multilinear functions – the “price of truthfulness” is zero. That is, the information-theoretically optimal rule is dominant-strategy implementable. This family of functions captures many common settings from the literature. In particular, it generalizes the auction setting studied by Blumrosen et al. [4, 6].

Definition 10. A multilinear function is a polynomial in which the degree of every variable in each monomial is at most 1.⁶ We say that a social-choice rule g is multilinear, if $g(\cdot, A)$ is multilinear for every alternative $A \in \mathcal{A}$.

The basic idea behind the proof of the following theorem is as follows: for every player, we show that the expected social welfare when he chooses any action (fixing the strategies of the other players) is a linear function of his type. This is a result of the multilinearity of the social-value function and of the linearity of expectation. The maximum over a set of linear functions is a piecewise-linear function, hence the optimal social value is achieved when the player uses threshold strategies (the thresholds are the switching points). Since the optimum is achieved

⁶For example, $f = xyz + 5xy + 7$.

with threshold strategies, we can apply Lemma 1 to show the monotonicity of this social-choice rule. Note that in this argument we characterize the players’ strategies that maximize the social value, and not the players’ utilities.

Theorem 1. *If the social-value function is multilinear and single crossing, the informationally optimal k -action social-choice function is implementable.*

Observe that the proof of Theorem 1 actually works for a more general setting. For proving that the information-theoretically optimal result is achieved with threshold strategies, it is sufficient to show that the social-choice function exhibits a *single-crossing condition on expectation*: given any allocation scheme, and fixing the behavior of the other players, the expected social value in any two actions (as a function of θ_i) is single crossing. Theorem 1 shows that this requirement holds for multilinear functions, but we were not able to give an exact characterization of this general class of functions.

The implementability of the information-theoretically optimal solution makes the characterization of the optimal incentive-compatible mechanisms significantly easier: we can apply the standard mechanism-design technique and first calculate the optimal allocation scheme and then find the “right” payments.

Observe that if the valuation functions of the players are linear and single crossing, then the social-welfare function (i.e., the sum of the players’ valuations) is multilinear and single-crossing. This holds since the single-crossing conditions on the valuations are defined with a similar order on the alternatives as in the social-value function. Therefore, an immediate conclusion from Theorem 1 is that the optimal social welfare, which is achievable with k actions, is implementable when the valuations are linear.

Corollary 1. *If the valuation functions $v_i(\cdot, A)$ are single crossing and linear in θ_i for every player i and for every alternative, then the informationally optimal k -action social welfare function is implementable.*

3.1 Asymptotic Analysis

In this section we show that the social value loss of multilinear social-value rules diminishes quadratically with the number of possible actions, k . This is the same asymptotic ratio presented in the study of specific models in the same spirit [19, 5, 18, 13]. The main challenge here, compared to earlier results, is in dealing with the general mechanism-design framework, that allows a large family of social-value functions for any number of players and alternatives. As opposed to the specific models, the social-value function may be asymmetric with respect to the players’ types; for instance, the social-value loss may a-priori occur in any “entry” (i.e., profile of actions).

The basic intuition for the proof is that even for this general framework, we can construct mechanisms where the probability of having an allocation that is incompatible with the original social-choice function is $O(\frac{1}{k})$. (This fact holds for all single-crossing social-choice functions, not only for multilinear functions.) Then, we can use the multilinearity to show that the social-value loss will always be $O(\frac{1}{k})$ in the mechanisms we construct. Taken together, the expected loss becomes $O(\frac{1}{k^2})$. Our proof is constructive – we present an explicit construction for a mechanism that exhibits the desired loss in dominant strategies.

Theorem 2. *For any number of players and alternatives, and for any set of distribution functions of the players’ types, if the social-value function g is single crossing and multilinear, then*

the informationally optimal k -action social-choice function (with respect to g) incurs an expected social-value loss of $O(\frac{1}{k^2})$.

Moreover, as discussed in [4], this bound is asymptotically tight. That is, there exists a set of distribution functions for the players (the uniform distribution in particular) and there are social-value functions (e.g., auctions) for which *any* mechanism incurs a social-value loss of at least $\Omega(\frac{1}{k^2})$. In that sense, auctions are the hardest problems with respect to the incurred loss. Yet, note that this claim does not imply that the loss of *any* social-choice function will be proportional to $\frac{1}{k^2}$. For example, in the social choice function that chooses the same alternative for any type profile, no loss will be incurred (even with 0 actions).

4 Optimal Mechanisms for Two Players and Two Alternatives

In this section, we present a full characterization of the optimal mechanisms in action-bounded environments with two players and two alternatives, where the social-choice functions are multilinear and single crossing.

Note that in this section, as in most parts of this paper, we characterize monotone mechanisms by their *allocation* scheme and by a profile of *strategies* for the players. Doing this, we completely describe which alternative is chosen for every profile of types of the players. It is well known that in monotone mechanisms for one dimensional environments, the allocation scheme uniquely defines the payments in the dominant-strategy implementation. We find this description, which does not explicitly mention the payments, easier for the presentation.

A key notion in our characterization of the optimal action-bounded mechanism, is the notion of *non-degenerate* mechanisms. In a degenerate mechanism, there are two actions for one of the players that are identical in their allocation. Intuitively, a degenerate mechanism does not utilize all the action space he is allowed to use, and therefore it cannot be optimal. Using this property, we then define “diagonal” mechanisms that turns out to exactly characterize the set of optimal mechanisms.

Definition 11. *A mechanism is degenerate with respect to player i if there exist two actions b_i, b'_i for player i such that for all profiles b_{-i} of actions of the other players, the allocation scheme is identical whether player i reports b_i or b'_i (i.e., $\forall b_{-i}, t(b_i, b_{-i}) = t(b'_i, b_{-i})$).*

For example, a 2-player mechanism is degenerate with respect to the “rows” player, if there are two rows with identical allocation in the matrix representation of the game.

Definition 12. *A 2-player 2-alternative mechanism with k -possible actions is called diagonal if it is monotone, and non-degenerate with respect to at least one of the players.*

The term “diagonal” originates from the matrix representation of these mechanisms, in which one of the diagonals determines the boundary between the choice of the two alternatives (see Figure 1). Simple combinatorial considerations show that diagonal mechanisms may come in very few forms. Interestingly, one of these forms is degenerate with respect to one of the players; that is, it can be described as a mechanism with $k - 1$ actions for this player.

Proposition 2. *Any diagonal 2-player mechanism has one of the following forms:*

1. *If both players favor the same alternative (w.l.o.g., $B \succ_i A$ for $i = 1, 2$) then either:*

- (a) $t(b_1, b_2) = B$ iff $b_1 + b_2 \geq k - 1$

	0	1	2	3
0	A	A	A	B
1	A	A	B	B
2	A	B	B	B
3	B	B	B	B

	0	1	2	3
0	A	A	A	A
1	A	A	A	B
2	A	A	B	B
3	A	B	B	B

	0	1	2	3
0	B	B	B	B
1	A	B	B	B
2	A	A	B	B
3	A	A	A	B

	0	1	2	3
0	A	A	A	B
1	A	A	B	B
2	A	B	B	B

Figure 1: The three left tables show all possible diagonal allocation scheme with 4 possible actions for each player. The rightmost table show an example for a diagonal allocation scheme where one of the player has only $k - 1$ possible actions.

(b) $t(b_1, b_2) = B$ iff $b_1 + b_2 \geq k$.

2. If the two players have conflicting preferences (e.g., $A \succ_1 B$ and $B \succ_2 A$) then either:

(a) $t(b_1, b_2) = B$ iff $b_1 \geq b_2$

(b) $t(b_1, b_2) = B$ iff $b_1 > b_2$.

In both cases, the optimal mechanism can also take the form of one of the possibilities described, except one of the players is not allowed to choose the “fixed allocation” action.

To complete the description of the optimal allocation scheme, we now move to determine the optimal strategies in diagonal mechanisms. We define the notion of *mutually-maximizer* thresholds, and show that threshold strategies based on such thresholds are optimal. The reason why mutually-maximizer strategies maximize the expected social value in monotone mechanisms is intuitive: Consider some action i (“row” in the matrix representation) for player 1. In a monotone mechanism, the allocation in such a row will be of the form $[A, A, \dots, B, B]$ (assuming that $B \succ_2 A$). That is, the alternative A will be chosen for low actions of player 2, and the alternative B will be chosen for higher actions of player 2. By determining a threshold for player 2, the social planner actually determines the minimal type of player 2 from which the alternative B will be chosen. For optimizing the expected social value, this type for player 2 should clearly be the type for which the expected social value from A equals the expected social value from B (given that player 1 plays i); for greater values of player 2, the single-crossing condition ensures that B will be preferred.

Definition 13. Consider a monotone 2-player mechanism g that is non-degenerate with respect to the two players, where the players use threshold strategies based on the threshold vectors x, y . We say that the threshold x_i of one player (w.l.o.g. player 1) is a maximizer if

$$E_{\theta_2} (g(x_i, \theta_2, A) \mid \theta_2 \in [y_j, y_{j+1}]) = E_{\theta_2} (g(x_i, \theta_2, B) \mid \theta_2 \in [y_j, y_{j+1}])$$

where j is the action of player 2 for which the mechanism swaps the chosen alternative exactly when player 1 plays i , i.e., $t(i, j) \neq t(i - 1, j)$ (we denote, w.l.o.g., $t(i, j) = A$, $t(i - 1, j) = B$).

The threshold vectors x, y are called mutually maximizers if all their thresholds are maximizers (except the first and the last).

It turns out that in 2-player, 2-alternative environments, where the social-choice rule is multilinear and single crossing, the optimal expected social value is achieved in diagonal mechanisms with mutually-maximizer strategies. In the proof, we start with a $k \times k$ allocation matrix, and show that the mechanism cannot be degenerate with respect to one of the players (we show how

to choose this player). If the player, w.l.o.g., the columns player, is degenerate, then there are two columns with an identical allocation. These two columns can be unified to a single action, and the mechanism can therefore be described as a $k \times k - 1$ matrix. We then show that we can insert a new missing column, and an appropriately chosen threshold, and strictly increase the expected social value in the mechanism. Therefore, the original mechanism was not the optimal k -action mechanism.

Theorem 3. *In environments with two alternatives and two players, if the social-value function is multilinear and single crossing, then the optimal k -action mechanism is diagonal, and the optimum is achieved with threshold strategies that are mutually maximizers.*

A corollary from the proof of Theorem 1 is that the optimal 2-player k -action mechanism may be degenerate for one of the players (that is, equivalent to a game where one of the players has only $k - 1$ different actions). However, the proof identifies the following sufficient condition under which the optimal mechanism will be non-degenerate with respect to both players: if the players' preferences are correlated (e.g., $A \succ_1 B$ and $A \succ_2 B$), then the optimal alternative must be the same under the profiles $(\underline{\theta}_1, \bar{\theta}_2)$ and $(\bar{\theta}_1, \underline{\theta}_2)$. Similarly, if the players' preferences are conflicting (e.g., $A \succ_1 B$ and $B \succ_2 A$), then the optimal alternative must be the same under the profiles $(\underline{\theta}_1, \underline{\theta}_2)$ and $(\bar{\theta}_1, \bar{\theta}_2)$. Examples in which this condition holds are the public good model presented in section 5 and auctions [5].

We do not know how to give an exact characterization of the optimal mechanisms in multi-player and multi-alternative environments. The hardness stems from the fact that the necessary conditions we specified before for the optimality of the mechanisms (i.e., non-degenerate and monotone allocations) are not restrictive enough for the general model. In other words, for $n > 2$ players, the number of monotone and non-degenerate mechanisms becomes exponential in n .

Proposition 3. *The number of monotone non-degenerate k -action mechanisms in an n -player game is exponential in n , even if $|\mathcal{A}| = 2$.*

5 Examples

Our results apply to a variety of economic, computational and networked settings. In this section, we demonstrate the applicability of our results to public-good models, signaling games and routing applications.

5.1 Example 1: Public Goods

The public-good model deals with a social planner (e.g., government) that needs to decide whether to supply a public good, such as building a bridge. Let *Yes* and *No* denote the respective alternatives of building and not building the bridge. $v = v_1, \dots, v_n$ is the vector of the players' types – the values they gain from using the bridge. The decision that maximizes the social welfare is to build the bridge if and only if $\sum_i v_i$ is greater than its cost, denoted by C . If the bridge is built, the social welfare is $\sum_i v_i - C$, and zero otherwise; thus, $g(v, \text{Yes}) = \sum_i v_i - C$, and $g(v, \text{No}) = 0$. The utility of player i under payment p_i is $u_i = v_i - p_i$ if the bridge is built, and 0 otherwise. It is well-known that under no restriction on the action space, it is possible to induce truthful revelation by VCG mechanisms, therefore full efficiency can be achieved. Obviously, when the action set is limited to k actions, we cannot achieve full efficiency due

$c \leq 1$	0	1	$c \geq 1$	0	1
0	No $p_1 = p_2 = 0$	No $p_1 = p_2 = 0$	0	No $p_1 = p_2 = 0$	Yes $p_1 = 0; p_2 = \frac{2c}{3}$
1	No $p_1 = p_2 = 0$	Yes $p_1 = p_2 = \frac{2}{3}c - \frac{1}{3}$	1	Yes $p_1 = \frac{2c}{3}; p_2 = 0$	Yes $p_1 = p_2 = 0$

Figure 2: Optimal mechanisms in a 2-player, 2-alternative, 2-action public-goods game, when the types are uniformly distributed in $[0, 1]$. The mechanism on the left is optimal when $c \leq 1$ and the other is optimal when $c \geq 1$.

to the informational constraints. Yet, since $g(v, Yes)$ and $g(v, No)$ are multilinear and single crossing, we can directly apply Theorem 1. Hence, the information-theoretically optimal k -action mechanism is implementable in dominant strategies.

Corollary 2. *The k -action informationally optimal social welfare in the n -player public-good game is implementable in dominant strategies.*

Moreover, as Theorem 3 suggests, in the k -action 2-player public-good game, we can fully characterize the optimal mechanisms. In the proof of Theorem 3, we saw that when $g(\underline{\theta}_i, \bar{\theta}_i, A) = g(\bar{\theta}_i, \underline{\theta}_i, B)$, the mechanism is non-degenerate with respect to both players. This condition clearly holds here ($1 + 0 - C = 0 + 1 - C$), therefore the optimal mechanisms will use all k actions.

Corollary 3. *The optimal expected welfare in a 2-player k -action public-good game is achieved with one of the following mechanisms:⁷*

1. Allocation: Build the bridge iff $b_1 + b_2 \geq k$.

Strategies: Threshold strategies based on the vectors \vec{x}, \vec{y} where for every $1 \leq i \leq k-1$,

$$x_i = c - E[v_2 | v_2 \in [y_{k-i}, y_{k-i+1}]] \quad ; \quad y_i = c - E[v_1 | v_1 \in [x_{k-i}, x_{k-i+1}]]$$

2. Allocation: Build the bridge iff $b_1 + b_2 \geq k - 1$.

Strategies: Threshold strategies based on the vectors \vec{x}, \vec{y} where for every $1 \leq i \leq k-1$:

$$x_i = c - E[v_2 | v_2 \in [y_{k-i-1}, y_{k-i}]] \quad ; \quad y_i = c - E[v_1 | v_1 \in [x_{k-i-1}, x_{k-i}]]$$

Recall that we define the optimal mechanisms by their allocation scheme and by the optimal strategies for the players. It is well known, that the allocation scheme in monotone mechanisms uniquely defines the payments that ensure incentive-compatibility. In public-good games, these payments satisfy the rule that a player pays his lowest value for which the bridge is built, when the action of the other player is fixed. Therefore, the payments for the players 1 and 2 reporting the actions b_1 and b_2 are as follows: in mechanism 1 from Proposition 3, $p_1 = x_{b_2}$ and $p_2 = y_{b_1}$; in mechanism 2 from Proposition 3, $p_1 = x_{b_2-1}$ and $p_2 = y_{b_1-1}$.

We now show a more specific example that assumes uniform distributions. The example shows how the optimal mechanism is determined by the cost C : for low costs, mechanism of type 1 is optimal, and for high costs the optimal mechanism is of type 2. An additional interesting feature of the optimal mechanisms in the example is that they are symmetric with respect to the players. This come as opposed to the optimal mechanisms in the auction model [5] that are asymmetric (even when the players' values are drawn from identical distributions).

⁷We denote $x_0 = y_0 = 0$ and $x_k = y_k = 1$.

Example 1. Suppose that the types of both players are uniformly distributed on $[0, 1]$. Then, the welfare-maximizing mechanisms are (Figure 2 illustrates the optimal mechanisms for $k = 2$):

- If the cost of building is at least 1:

Allocation: Build iff $b_1 + b_2 \geq k$

Strategies: The thresholds of both players are (for $i = \{1, \dots, k-1\}$), $x_i = \frac{2(k-i) \cdot c}{2k-1} - \frac{2k-4i+1}{2k-1}$

- If the cost of building is smaller than 1:

Allocation: Build iff $b_1 + b_2 \geq k - 1$

Strategies: The thresholds of both players are (for $i = \{1, \dots, k-1\}$), $x_i = \frac{2ic}{2k-1}$

5.2 Example 2: Signaling

We now study a signaling model in labor markets. In this model, the type of each worker, $\theta_i \in [\underline{\theta}, \bar{\theta}]$, describes the worker's productivity level. The firm wants to make her hiring decisions according to a decision function $f(\vec{\theta})$. For example, the firm may want to hire the most productive worker (like the auction model), or hire a group of workers only if their sum of productivities is greater than some threshold (similar to the public-good model). However, the worker's productivity is invisible to the firm; the firm only observes the worker's education level e that should convey signals about her productivity level. Note that the assumption here is that acquiring education, at any level, does not affect the productivity of the worker, but only signals about the worker's skills.

A main component in this model, is the fact that as the worker is more productive, it is easier for him to acquire high-level education. In addition, the cost of acquiring education increases with the education level. More formally, a continuous function $C(e, \theta)$ describes the cost to a worker from acquiring each education level as a function of his productivity. The standard assumptions about the cost function are: $\frac{\partial c}{\partial e} > 0$, $\frac{\partial c}{\partial \theta} < 0$, $\frac{\partial c}{\partial e \partial \theta} < 0$, where the last requirement is the single-crossing property (when both variables are differentiable). The utility of a worker is determined according to the education level he chooses and the wage $w(e)$ attached to this education level, that is, $u_i(e, \theta_i) = -C(\theta_i, e) + w(e)$.

An action for a worker in this game is the education level he chooses to acquire. In standard models, this action space is continuous, and then a "fully separating equilibrium" exists (under the single-crossing conditions on the cost function). That is, there exists an equilibrium in which every type is mapped into a different education level; thus, the firm can induce the exact productivity levels of the workers by this signaling mechanism. However, it is hard to imagine a world with a continuum of education levels. It is usually the case that there are only several discrete education levels (e.g., BSc, MSc, PhD).

With k education levels, the firm may not be able to exactly follow the decision function f . For achieving the best result in k actions, the firm may want the worker to play according to a certain threshold strategy based on the thresholds x_0, x_1, \dots, x_k . Our first claim is that the standard condition, the single-crossing condition on the cost function, suffices for a dominant-strategy k -action implementation.

Proposition 4. Consider a worker with a single-crossing cost function, and a k -action threshold strategy s . There are education levels and wages such that s is a dominant strategy.




Figure 3: An example of a series-parallel network, where each link has a probability p_i for transmission success. We show that the overall probability of success in such networks is multilinear and monotone in p_i , and thus the optimal k -action social-choice function is dominant-strategy implementable.

We can now apply Theorem 2, and show that if the decision function f of the firm is multilinear (i.e., the decisions are made to maximize a set of multilinear functions), then the firm can design the education system such that the expected incurred loss will be $O(\frac{1}{k^2})$.

Corollary 4. *Consider a multilinear decision function f , and a single-crossing cost function for the players. With k education levels, the firm can implement in dominant strategies a decision function that incurs a loss of $O(\frac{1}{k^2})$ compared with the decision function f .*

5.3 Example 3: Routing

In our last example, we show the applicability of our results to routing in lossy networked systems. In such systems, a sender needs to decide which network to transmit his message through. In this setting, it is natural to assume that the agents (i.e., links) cannot report their accurate probabilities of success, but may be able to report, e.g., whether it is “low”, “intermediate”, or “high”. In this example, we focus on *parallel-path* networks.

Let N_1 denote an n -edge network that is composed of multiple parallel paths of variable lengths from a given source to a given sink (as in Figure 3), where the edges are controlled by n different selfish agents. Suppose that the sender, who wishes to send a message from the source to the sink, knows the topology of the network, but the probability of success on each link, p_i , is the link’s private information. The problem of the sender is to decide whether to send a message through the network N_1 or through an alternate network, N_2 , with a known success probability of p' . Obviously, the sender wishes to send the message through N_1 only if the total probability of success in N_1 , is greater than p' . Let $f^N(\vec{p})$ denote the probability of success in network N with a success-probability vector \vec{p} .

In this example, the sender’s set of alternatives is $\mathcal{A} = \{N_1, N_2\}$, and we assume that every agent on N_1 wishes the message to be sent, and has a single-crossing valuation function over the alternatives. The social choice function is: $c(\vec{p}) \in \operatorname{argmax}_{A \in \{N_1, N_2\}} \{g(\vec{p}, A)\}$, where: $g(\vec{p}, N_1) = f^{N_1}(\vec{p})$, and $g(\vec{p}, N_2) = p'$.

Proposition 5. *Given a parallel-path network, the social-choice function $c(\vec{p})$ is multilinear and single crossing.*

Based on the above proposition, we can apply Theorem 1 and get the following corollary.

Corollary 5. *Given a parallel-path network, the informationally optimal k -action social-choice function $c(\vec{p})$ is implementable.*

For example, in the network presented in figure 3, the probability of success is given by $f(\vec{p}) = 1 - (1 - p_1 p_2) \cdot (1 - p_3) \cdot (1 - p_4 p_5)$. This function is multilinear and has positive derivatives with respect to all p_i . Therefore, the optimal social-choice function with k actions is implementable.

Acknowledgment. We thank Noam Nisan for helpful discussions.

References

- [1] Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–89, 2001.
- [2] Moshe Babaioff and Liad Blumrosen. Computationally-feasible auctions for convex bundles. In *APPROX 04*, pages 27–38, 2004.
- [3] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Mechanism design for single-value domains. In *AAAI’05*, pages 241–247, 2005.
- [4] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communications. 43th Annual Symposium on Foundations of Computer Science (FOCS 2002), 2002.
- [5] Liad Blumrosen, Noam Nisan, and Ilya Segal. Auctions with severely bounded communications. Working paper, The Hebrew University. Preliminary versions appeared in FOCS 2002 and ESA 03., 2003.
- [6] Liad Blumrosen, Noam Nisan, and Ilya Segal. Multi-player and multi-round auctions with severely bounded communication. *ESA 2003*, 2003.
- [7] M.S.Y. Chwe. The discrete bid first price auction. In *Economics Letters*, volume 31, pages 303–306, 1989.
- [8] E. David, A. Rogers, J. Schiff, S. Kraus, and N. Jennings. Optimal design of english auctions with discrete bid levels. In *EC 05*.
- [9] Benjamin E. Hermalin. *Lecture notes in economics*, 2005.
- [10] A. Kress and C. Boutilier. A study of limited precision, incremental elicitation in auctions. In *AAMAS 04*.
- [11] K. Larson and T. Sandholm. Costly valuation computation in auctions. In *8th conference of theoretical aspects of knowledge and rationality*, 2001.
- [12] R. Lavi, A. Mu’alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. 2003.
- [13] Harstad R. M. and Rothkopf M. H. On the role of discrete bid levels in oral auctions. *European Journal of Operations Research*, pages 572–581, 1994.
- [14] P. McAfee. Coarse matching. *Econometrica*, 70(5):2025–2034, 2002.

- [15] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [16] M. Naor, B Pinkas, and R. Summer. Privacy preserving auctions and mechanism design. In *ACM Conference on Electronic Commerce*, 1999.
- [17] T. Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full-valuation revelation. In *AMEC-V*, 2003.
- [18] Mark A. Satterthwaite and Steven R Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.
- [19] R. Wilson. Efficient and competitive rationing. *Econometrica*, 57:1–40, 1989.

A Missing Proofs from Section 3

Proof of Lemma 1:

Proof. We first show that we can assume, w.l.o.g., that the optimal k -action social-choice function is deterministic.⁸ Consider an optimal k -action mechanism that achieves the optimal result with some set of strategies $s = s_1, \dots, s_n$. Assume that there is an action vector b_1, \dots, b_n for which the mechanism randomizes over alternatives. Consider a similar mechanism that deterministically chooses an alternative that maximizes the expected social value for the action vector \vec{b} , i.e., $t(\vec{b}) \in \operatorname{argmax}_{A'} E_{\vec{\theta}}(g(\vec{\theta}, A') | \forall i s_i(\theta_i) = b_i)$. The expected social value for the designer clearly has not decreased. We can similarly change the allocation for all the actions combinations and get a deterministic mechanism with at least the same expected social value.

We now show that when the optimum is achieved with threshold strategies, the optimal mechanism is monotone (and hence incentive compatible, Prop 1). This will follow from the single-crossing condition on g . Denote the thresholds used by player i by $x_0^i, x_1^i, \dots, x_n^i$. Specifically, when player i reports an action b_i and uses a threshold strategy, her type will lie between $[x_{b_i}^i, x_{b_i+1}^i]$. Consider a deterministic choice rule as described above, and consider an action vector b_1, \dots, b_n . Let A and B be two alternatives where player i prefers alternative A to B (i.e., $A \succeq_i B$). Now consider another action vector $b' = (b'_i, b_{-i})$, where $b'_i > b_i$. For proving monotonicity, it suffices to show that if choosing A gains a higher social value than choosing B for the actions vector b , this will also hold for the actions vector b' . That is, if

$$E_{\vec{\theta}} \left(g(\vec{\theta}, A) | \forall j s_j(\theta_j) = b_j \right) \geq E_{\vec{\theta}} \left(g(\vec{\theta}, B) | \forall j s_j(\theta_j) = b_j \right) \quad (1)$$

then

$$E_{\vec{\theta}} \left(g(\vec{\theta}, A) | \forall j s_j(\theta_j) = b'_j \right) \geq E_{\vec{\theta}} \left(g(\vec{\theta}, B) | \forall j s_j(\theta_j) = b'_j \right)$$

To see this, we show that given any profile of types θ_{-i} of the other players, the change in the expected value of $g(\cdot, A)$ will be greater than the change in $g(\cdot, B)$ when player i bids a higher bid.

⁸This result is general and its proof does not require that the players use threshold strategies.

And indeed,

$$\begin{aligned}
& E_{\vec{\theta}} \left(g(\vec{\theta}, A) \mid \forall j \ s_j(\theta_j) = b'_j \right) - E_{\vec{\theta}} \left(g(\vec{\theta}, B) \mid \forall j \ s_j(\theta_j) = b'_j \right) \\
&= E_{\theta_{-i}} \left(E_{\theta_i} \left(g(\vec{\theta}, A) - g(\vec{\theta}, B) \mid \forall s_i(\theta_i) = b'_i \right) \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&= E_{\theta_{-i}} \left(\frac{1}{F_i(x_{b_{i+2}}^i) - F_i(x_{b_{i+1}}^i)} \int_{x_{b_{i+1}}^i}^{x_{b_{i+2}}^i} \left(g(\vec{\theta}, A) - g(\vec{\theta}, B) \right) f_i(\theta_i) d\theta_i \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&> E_{\theta_{-i}} \left(\frac{1}{F_i(x_{b_{i+2}}^i) - F_i(x_{b_{i+1}}^i)} \int_{x_{b_{i+1}}^i}^{x_{b_{i+2}}^i} \left(g(x_{b_{i+1}}^i, \theta_{-i}, A) - g(x_{b_{i+1}}^i, \theta_{-i}, B) \right) f_i(\theta_i) d\theta_i \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&= E_{\theta_{-i}} \left(g(x_{b_{i+1}}^i, \theta_{-i}, A) - g(x_{b_{i+1}}^i, \theta_{-i}, B) \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&= E_{\theta_{-i}} \left(\frac{1}{F_i(x_{b_{i+1}}^i) - F_i(x_{b_i}^i)} \int_{x_{b_i}^i}^{x_{b_{i+1}}^i} \left(g(x_{b_{i+1}}^i, \theta_{-i}, A) - g(x_{b_{i+1}}^i, \theta_{-i}, B) \right) f_i(\theta_i) d\theta_i \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&> E_{\theta_{-i}} \left(\frac{1}{F_i(x_{b_{i+1}}^i) - F_i(x_{b_i}^i)} \int_{x_{b_i}^i}^{x_{b_{i+1}}^i} \left(g(\vec{\theta}, A) - g(\vec{\theta}, B) \right) \mid \forall j \neq i \ s_j(\theta_j) = b_j \right) \\
&= E_{\vec{\theta}} \left(g(\vec{\theta}, A) \mid \forall j s_j(\theta_j) = b_j \right) - E_{\vec{\theta}} \left(g(\vec{\theta}, B) \mid \forall j s_j(\theta_j) = b_j \right) \\
&\geq 0
\end{aligned}$$

The strict inequalities follow from the single-crossing condition on g , and since $A \succeq_i B$. The other equalities hold since θ_i is drawn independently from the other types and due to the linearity of expectation. The last inequality holds since for the action vector b , the alternative A achieves a higher social value than B (Equation 1).

Therefore, when player i reports a higher message, an optimal mechanism will necessarily choose an alternative with higher priority for player i . The monotonicity of the optimal mechanism then follows.

We now prove the other direction of the lemma: if a mechanism is monotone, then the optimum is achieved with threshold strategies. The basic idea: for each player, we consider the expected social value as a function of her type θ_i when he chooses a particular action. We show that for every two actions $j_1 < j_2$ this expected social value is single crossing; it suffices here to show that the single-crossing property holds in the weaker sense – if for some θ_i the expected social value is equal for the two actions j_1, j_2 of player i , then for any higher type the expected value in j_2 will be strictly higher.

Let θ_i^* be the type for player i for which the expected social value is equal either when he chooses j_1 or j_2 , that is (we denote the actions of the players except i when their types are θ_{-i} by $s_{-i}(\theta_{-i})$):

$$E_{\theta_{-i}} \left(g(\theta_i^*, \theta_{-i}, t(j_1, s_{-i}(\theta_{-i}))) \right) = E_{\theta_{-i}} \left(g(\theta_i^*, \theta_{-i}, t(j_2, s_{-i}(\theta_{-i}))) \right) \quad (2)$$

We will show that for every $\epsilon > 0$, the expected social value when player i chooses j_2 is strictly greater than the expected social value in j_1 when player i 's type is $\theta_i^* + \epsilon$.

Given a profile of actions b_{-i} played by the other players, let A be the chosen alternative when player i bids j_1 and let B be the chosen alternative for j_2 (that is, $t(j_1, b_{-i}) = A, t(j_2, b_{-i}) = B$). Since we assumed that the allocation scheme is monotone, then if $A \neq B$ we must have that

$A \succ_i B$. The social value function is single crossing, hence the change in the expected social value when alternative B is chosen should be greater, that is:

$$\begin{aligned} E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, B) \mid s_{-i}(\theta_{-i}) = b_{-i}) &- E_{\theta_{-i}}(g(\theta_i^*, \theta_{-i}, B) \mid s_{-i}(\theta_{-i}) = b_{-i}) > \\ E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, A) \mid s_{-i}(\theta_{-i}) = b_{-i}) &- E_{\theta_{-i}}(g(\theta_i^*, \theta_{-i}, A) \mid s_{-i}(\theta_{-i}) = b_{-i}) \end{aligned}$$

Now, summing over all the possible b_{-i} , we get:⁹

$$\begin{aligned} E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, t(j_2, s_{-i}(\theta_{-i})))) &- E_{\theta_{-i}}(g(\theta_i^*, \theta_{-i}, t(j_2, s_{-i}(\theta_{-i})))) > \\ E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, t(j_1, s_{-i}(\theta_{-i})))) &- E_{\theta_{-i}}(g(\theta_i^*, \theta_{-i}, t(j_1, s_{-i}(\theta_{-i})))) \end{aligned}$$

Since for θ_i^* the expected social value in j_1 and j_2 is equal (Equation 2), our claim follows:

$$E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, t(j_2, s_{-i}(\theta_{-i})))) > E_{\theta_{-i}}(g(\theta_i^* + \epsilon, \theta_{-i}, t(j_1, s_{-i}(\theta_{-i}))))$$

Finally, it is easy to see now that the optimal social value can be achieved with threshold strategies for k -action games: the strategy for player i that maximizes the social value is a maximum over k pairwise single-crossing functions, and such a function must have at most $k - 1$ switching points. \square

Proof for Theorem 1

Proof. We will show that for any k -action mechanism, the optimal expected social value is achieved when all players use threshold strategies. This will be shown by proving that for any player i and for any action b_i of player i , the expected welfare when she chooses the action b_i is a linear function in player i 's type θ_i . Then, it will follow from Lemma 1 that the social choice function is implementable.

For every action b_i of player i , let q_A denote the probability that alternative A is allocated, i.e., $q_A = Pr_{\vec{\theta}} [t(s(\vec{\theta})) = A \mid s_i(\theta_i) = b_i]$. Due to the linearity of expectation, the expected social value when player i with type θ_i reports b_i is:

$$\begin{aligned} &\sum_{A \in \mathcal{A}} q_A E_{\theta_{-i}}(g(\theta_i, \theta_{-i}, A) \mid t(b_i, s_{-i}(\theta_{-i})) = A) \\ &= \sum_{A \in \mathcal{A}} q_A \int_{\theta_{-i}} g(\theta_i, \theta_{-i}, A) f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \end{aligned}$$

where $f_{-i}^A(\theta_{-i})$ equals $\frac{\prod_{j \neq i} f_j(\theta_j)}{q_A}$ for types profiles θ_{-i} such that $t(b_i, s_{-i}(\theta_{-i})) = A$, and 0 otherwise.

Since g is multilinear, every function $g(\theta_i, \theta_{-i}, A)$ is a linear function in θ_i , where the coefficients depend on the values of θ_{-i} . Denote this function by $g(\theta_i, \theta_{-i}, A) = \alpha_{\theta_{-i}} \theta_i + \beta_{\theta_{-i}}$. Thus, we can write Equation 3 as:

$$\begin{aligned} &\sum_{A \in \mathcal{A}} q_A \int_{\theta_{-i}} (\alpha_{\theta_{-i}} \theta_i + \beta_{\theta_{-i}}) f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \\ &= \sum_{A \in \mathcal{A}} q_A \left(\theta_i \int_{\theta_{-i}} \alpha_{\theta_{-i}} f_{-i}^A(\theta_{-i}) d(\theta_{-i}) + \int_{\theta_{-i}} \beta_{\theta_{-i}} f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \right) \end{aligned}$$

⁹Note that there must be some b_{-i} for which $t(j_2, b_{-i}) \succ_i t(j_1, b_{-i})$ otherwise the allocation scheme is identical in j_1 and j_2 , thus we can ignore one of them,

In this expression, each integral is a constant independent of θ_i when the strategies of the other player are fixed. Therefore, each summand, thus the whole function, is a linear function in θ_i . For achieving the optimal expected social value, the player must choose the action that maximizes the expected social value. A maximum of k linear functions is a piecewise-linear function with at most $k - 1$ breaking points. These breaking points are the thresholds to be used by the player. For all types between subsequent thresholds, the optimum is clearly achieved by a single action; Since linear functions are single-crossing, every action will be maximal in at most one interval.

The same argument applies to all the players, and therefore the optimal social value is obtained with threshold strategies. \square

Proof of Theorem 2:

Proof. For simplicity, we will assume that all the types are drawn from the support $[0, 1]$ (otherwise, the lengths of the supports only affect the constants in the asymptotic analysis), and that k is even.

Given a set of n players, we will define a k -action threshold strategy for each player where each action j is chosen with probability $O(\frac{1}{k})$, and the distance between each consecutive thresholds is $O(\frac{1}{k})$. Using these strategies, we define a mechanism that achieves an $O(\frac{1}{k^2})$ loss.

Construction of the threshold strategies:

For each player i let $Y^i = \{y_0^i = \underline{\theta}, y_1^i, \dots, y_{\frac{k}{2n}-1}^i, y_{\frac{k}{2}}^i = \bar{\theta}\}$ be a set of threshold thresholds that divide the density function of player i to $\frac{k}{2}$ equi-mass intervals. That is, for every j, l we have $F_i(y_{j+1}^i) - F_i(y_j^i) = F_i(y_{l+1}^i) - F_i(y_l^i) = \frac{2}{k}$.

In addition, let $Z^i = \{z_0^i = \underline{\theta}, z_1^i, \dots, z_{\frac{k}{2n}-1}^i, z_{\frac{k}{2}}^i = \bar{\theta}\}$ be a set of thresholds that divides the interval $[0, 1]$ to $\frac{k}{2}$ equi-sized intervals. That is, for every j, l we have $y_{j+1}^i - y_j^i = y_{l+1}^i - y_l^i = \frac{2}{k}$.

Now, let $X^i = Y^i \cup Z^i$ be the set of thresholds for player i . Clearly, using a threshold strategy based on X^i (when the thresholds are ordered in increasing order), player i chooses each action j with probability $O(\frac{1}{k})$, and the distance between each consecutive thresholds is $O(\frac{1}{k})$.

The allocation rule:

For each vector of actions b , the mechanism will choose the alternative that maximize the expected social-value when the players use the threshold strategies s based on the vectors X^i defined above. That is,

$$t(b) = \operatorname{argmax}_A E \left[g(\vec{\theta}, A) \mid s(\vec{\theta}) = b \right]$$

All the definitions and claims below refer to the mechanism above, where each player plays according to the threshold strategy s_i based on the thresholds X^i .

We say that an actions vector b is *decisive* if one alternative maximizes the social value for every profile of types (otherwise the vector is *indecisive*). In other words, if the social planner chooses a particular alternative for this actions' vector then no loss in social-value is incurred. More formally, an actions vector b is *decisive* if there exists an alternative A for which $A \in \operatorname{argmax}_B g(\theta_1, \dots, \theta_n, B)$ for every where profile $\vec{\theta}$ of types such that $s^*(\theta_i) = b_i$ for every player i . Similarly, the vector b is *decisive with respect to a pair of alternatives A, B* , if one of these alternatives is always superior to the other when the player choose the actions b .

We will prove that the mechanism incurs an expected loss of $O(\frac{1}{k^2})$ using the two claims below. Claim 1 shows that the number of indecisive actions vectors is $O(k^{n-1})$. Since the player

choose each action with probability $O(\frac{1}{k})$, each indecisive action vector is chosen with probability $O(\frac{1}{k^n})$, and therefore an indecisive vector will be chosen with probability of $O(k^{n-1} \cdot \frac{1}{k^n}) = O(\frac{1}{k})$. Claim 2 proves that the maximal possible social-value loss, compared to the optimal allocation with unrestricted actions, is $O(\frac{1}{k})$ for each indecisive action vector. Therefore, it follows from the following claims that the expected social-value loss in the k -action mechanism we constructed above is $O(\frac{1}{k^2})$.

Claim 1. *The number of indecisive actions profile is at most $O(k^{n-1})$.*

Proof. Consider a pair of players 1, 2 and a pair of alternatives A, B and fix the actions $b_{-\{1,2\}}$ of the other players. Let $(b_1, b_2, b_{-\{1,2\}})$ be an indecisive vector with respect to alternatives A and B (assume that $A \succ_1 B$ and $B \succ_2 A$, the other cases are treated similarly). Since the action vector is indecisive, there must be types θ_1, θ_2 for which $s(\theta_1) = b_1$ and $s(\theta_2) = b_2$, and also

$$E_{\theta_{-\{1,2\}}} [g(\theta_1, \theta_2, \theta_{-\{1,2\}}, A)] > E_{\theta_{-\{1,2\}}} [g(\theta_1, \theta_2, \theta_{-\{1,2\}}, B)]$$

Now consider an action vector b'_1, b'_2 such that $b'_1 > b_1$ and $b'_2 < b_2$. We will show that for any pair of types θ'_1, θ'_2 for which $S(\theta'_1) = b'_1$ and $s(\theta'_2) = b'_2$ we have:

$$E_{\theta_{-\{1,2\}}} [g(\theta'_1, \theta'_2, \theta_{-\{1,2\}}, A)] > E_{\theta_{-\{1,2\}}} [g(\theta'_1, \theta'_2, \theta_{-\{1,2\}}, B)]$$

The formal argument is proved similarly to the proof in Lemma 1, and it follows from the single-crossing condition: changing the types from θ_1, θ_2 to θ'_1, θ'_2 clearly increases the type of player 1 and decreases the type of player 2 – both changes increase the gap between the social value achieved with the alternative A and the alternative B . We conclude that if $b_1, b_2, b_{-\{1,2\}}$ is indecisive with respect to A, B , then any other indecisive actions vector cannot include a smaller action for one of the players 1, 2 and a higher action for the other. Thus, there are at most $2k - 1$ indecisive vectors for any profile $b_{-\{1,2\}}$ of the other players. Every indecisive actions vector is clearly indecisive with respect to some pair of alternatives, thus the number of indecisive actions vectors (given $b_{-\{1,2\}}$) is at most $\binom{A}{2} \cdot (2k - 1) = O(k)$. Therefore, for any pair players (of $\binom{n}{2}$ pairs), there are k^{n-2} different actions for the other players, each one allows at most a linear number of indecisive action vectors. The total number of indecisive actions vectors will therefore be $O(k^{n-2}) \cdot O(k) = O(k^{n-1})$. \square

Claim 2. *The social-value loss incurred when the players play an indecisive actions vector is $O(\frac{1}{k})$.*

Proof. Consider an indecisive vector of actions b with respect to a pair of alternative A, B . Given that the players choose the actions b , we show that the difference between the social value gained by choosing A and B is always at most $O(\frac{1}{k})$. It will follow immediately that the expected loss incurred given each actions vector is $O(\frac{1}{k})$.

Suppose w.l.o.g that the mechanism chooses the alternative A for the action vector b . Let $\theta_1^A, \theta_2^A \in \operatorname{argmax}_{\theta_1, \theta_2} g(\theta_1, \theta_2, A)$ and let $\theta_1^B, \theta_2^B \in \operatorname{argmin}_{\theta_1, \theta_2} g(\theta_1, \theta_2, B)$. Since the vector b is indecisive with respect to A, B , and since the social value function is continuous, we know that there are types θ_1^*, θ_2^* for which $g(\theta_1^*, \theta_2^*, A) = g(\theta_1^*, \theta_2^*, B)$. We will show that $g(\theta_1^A, \theta_2^A, A) - g(\theta_1^*, \theta_2^*, A)$ is at most $O(\frac{1}{k})$, and similarly one can show that $g(\theta_1^B, \theta_2^B, B) - g(\theta_1^*, \theta_2^*, B)$ is $O(\frac{1}{k})$ and the theorem will follow.

Since the social-value function g is multilinear, we can write $g(\theta_1^A, \theta_2^A, A) = a\theta_1\theta_2 + b\theta_1 + c\theta_2 + d$, where $a, b, c, d \in \mathbb{R}$. The social value will increase, when moving from θ_1^A, θ_2^A to θ_1^*, θ_2^* , by at most

$$\begin{aligned} |a|(\theta_1^A - \theta_1^*)(\theta_2^A - \theta_2^*) + |b|(\theta_1^A - \theta_1^*) + |c|(\theta_2^A - \theta_2^*) &\leq |a|\frac{2n}{k}\frac{2n}{k} + |b|\frac{2n}{k} + |c|\frac{2n}{k} \\ &= O\left(\frac{1}{k}\right) \end{aligned}$$

The inequality holds since in the construction of the threshold strategies, the size of each interval is $O(\frac{1}{k})$.

This argument easily extends to any (constant) number of players. Since the proof holds for every two alternatives, the maximal loss is always $O(\frac{1}{k})$. \square

\square

B Missing Proofs from Section 4

Proof of Theorem 3:

Proof. We will show that the optimal mechanism will be non-degenerate with respect to (w.l.o.g.) player 2. In other words, in the matrix representation of the optimal mechanism there will be no identical columns. Denote the two alternatives as A and B and the two players as 1 and 2. We will prove the theorem for the case where the preferences of the players are conflicting, that is $A \succeq_1 B$ and $B \succeq_2 A$. The case where the preferences are correlated ($A \succeq_1 B$ and $A \succeq_2 B$) can be proved similarly. Assume w.l.o.g. that $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$ (recall that $\underline{\theta}_i$ denotes the lower bound of the support of player i). If player 2 has two identical columns, then monotonicity derives that these columns will be adjacent, so this player will actually have $k - 1$ possible actions (note that here we only consider the allocation scheme). We will prove that a mechanism where player 2 has $k - 1$ possible actions cannot be optimal, since we can add a new column and strictly increase the expected social value. Let the optimal k -action social value be achieved when player 1 uses the threshold vector x_0, \dots, x_k and player 2 has $k - 1$ possible actions and uses the threshold vector y_0, \dots, y_{k-1} . (Theorem 1 shows that for multilinear social-choice rules the optimal result is achieved in a monotone mechanism with threshold strategies).

Case 1: the column $[A, A, \dots, A]$ does not appear in the allocation matrix.

We will add this column to the game as the first column (action “0”), and add an additional threshold y' such that the expected social value strictly improves in the new mechanism when player 2 uses the threshold vector $y_0, y', y_1, \dots, y_{k-1}$. Consider the expected difference between the social value of the two alternatives when both players report 0, as a function of the second threshold of player 2:

$$diff(y) = E_{\theta_1, \theta_2} (g(\theta_1, \theta_2, A) - g(\theta_1, \theta_2, B) \mid \theta_1 \in [x_0, x_1], \theta_2 \in [y_0, y])$$

We know that $diff(y_0) > 0$ (since we assumed that $g(a_1, a_2, A) \geq g(a_1, a_2, B)$ and due to the single-crossing property). We also know that $diff(y_1) < 0$, otherwise alternative A would be preferred in this entry and the column $[A, \dots, A]$ would have existed (monotonicity). Due to the Intermediate-Value theorem, there must be some y^* for which $diff(y^*) = 0$ ($diff(\cdot)$ is clearly continuous since each $g(\cdot, \cdot, \text{alt})$ is continuous). Setting y' to be, for example, $\frac{y_{j+1} + y^*}{2}$

ensures that when θ_2 is between $[y_0, y']$ and when player 1 reports “0”, the expected social-value strictly increases. the allocation in all other cases remains unchanged.

Case 2: when the column $[A, A, \dots, A]$ exists.

Since there are $k + 1$ possible columns of the form $[B, B, \dots, A, A]$, it must be the case that some “internal” column is missing, that is, there are actions $i, i + 1$ for player 1 and $j, j + 1$ for player 2 such that $t(i, j) = t(i + 1, j) = A$ and $t(i, j + 1) = t(i + 1, j + 1) = B$. We will show that adding an action (column) j' for player 2 that is identical to the allocation in column j except $t(i + 1, j') = B$, will strictly increase the expected social value. For the exact construction, we have to consider two different subcases: If the expected social value when player 1 reports 0 and player 2’s type is y_{j+1} is greater for alternative A than for B , then we will define a new threshold which is greater than y_{j+1} ; Otherwise, the threshold will be smaller than y_{j+1} :

Case 2.1.: When $E(g(\theta_1, y_{j+1}, A)|\theta_1 \in [x_i, x_{i+1}])) \geq E(g(\theta_1, y_{j+1}, B)|\theta_1 \in [x_i, x_{i+1}]))$:

Due to the (strict) single-crossing condition, clearly $E(g(\theta_1, y_{j+1}, A)|\theta_1 \in [x_{i+1}, x_{i+2}])) > E(g(\theta_1, y_{j+1}, B)|\theta_1 \in [x_{i+1}, x_{i+2}]))$. Therefore, due to similar Intermediate-Value considerations, there must be some threshold $y^* > y_{j+1}$ for which $E(g(\theta_1, y_{j+1}, A)|\theta_1 \in [x_{i+1}, x_{i+2}])) = E(g(\theta_1, y_{j+1}, B)|\theta_1 \in [x_{i+1}, x_{i+2}]))$. Now, let player 2 use the threshold strategy based on the vector $y_0, \dots, y_{j+1}, y', \dots, y_{k-1}$, for example, $y' = \frac{y_{j+1} + y^*}{2}$. The expected social value strictly increased when player 2 reports the new bid (that is when $\theta_2 \in [y_{j+1}, y']$), while the allocation in the other cases remains unchanged.

Case 2.2.: When $E(g(\theta_1, y_{j+1}, A)|\theta_1 \in [x_i, x_{i+1}])) < E(g(\theta_1, y_{j+1}, B)|\theta_1 \in [x_i, x_{i+1}]))$:

Let y^* be again the value for which $E(g(\theta_1, y^*, A)|\theta_1 \in [x_i, x_{i+1}])) = E(g(\theta_1, y^*, B)|\theta_1 \in [x_i, x_{i+1}]))$. Clearly, now $y^* < y_{j+1}$. Similar arguments show that adding a new threshold $y' = \frac{y_{j+1} + y^*}{2}$ yields a higher expected social surplus. \square

Proof of Proposition 3:

Proof. We first prove that when $k = 2$, the number of monotone non-degenerate (MND) mechanisms is exponential in n by induction on the number of players. Suppose the number of MND n -player mechanisms is at least 2^n . We will show that the number of MND $(n+1)$ -player mechanisms is at least 2^{n+1} . Let \mathcal{M} denote the set of MND n -player mechanisms, and suppose that $B \succeq_{n+1} A$. Also, let $t^M(b)$ denote the allocation under mechanism M , given the vector of actions, b .

For each mechanism $M \in \mathcal{M}$, construct two $(n+1)$ -player mechanisms, M_1 and M_2 , as follows. (M_1) : $t^{M_1}(b_1, \dots, b_n, 0) = A$, and $t^{M_1}(b_1, \dots, b_n, 1) = t^M(b_1, \dots, b_n)$. (M_2) : $t^{M_2}(b_1, \dots, b_n, 0) = t^M(b_1, \dots, b_n)$, and $t^{M_2}(b_1, \dots, b_n, 1) = B$. We will show that mechanism M_1 is MND. The case of M_2 can be proved similarly.

Monotonicity: It is easy to see that the monotonicity of M_1 for the initial n players follows from the monotonicity of M . In addition, since $t^{M_1}(b_1, \dots, b_n, 0) = A$, and $B \succeq_{n+1} A$, M_1 must be monotone with respect to player $n + 1$ too.

non-degenerate: From the allocation function of M_1 , it follows that if M is non-degenerate with respect to the initial n players, then the same applies to M_1 . In addition, since M is non-degenerate, it cannot be the case that $\forall b_i, t^M(b_1, \dots, b_n) = A$. But since $\forall b_i, t^{M_1}(b_1, \dots, b_n, 0) = A$, M_1 is non-degenerate w.r.t. player $(n+1)$.

Similar arguments show that mechanism M_2 is MND. To complete the proof, we need to show that no two mechanisms are identical. Since M is non-degenerate, it cannot be the case that for all $b_i, t(b_1, \dots, b_n) = A$ or $t(b_1, \dots, b_n) = B$. But since $t^{M_1}(b_1, \dots, b_n, 0) = A$, and $t^{M_2}(b_1, \dots, b_n, 1) = B$, there does not exist a vector b for which $t^{M_1}(b) = t^{M_2}(b)$. In addition,

since $\forall M \in \mathcal{M}$, M is non-degenerate, two mechanisms that are both constructed according to M_1 or both constructed according to M_2 cannot be identical. Thus, if the number of MND n -player mechanisms is 2^n , then the number of MND $(n+1)$ -player mechanisms is 2^{n+1} . \square

C Missing Proofs from Section 5

Proof of Proposition 5:

Let \mathcal{P} denote the set of all paths from the source to the destination, and denote $P_i = P \in \mathcal{P} : i \in P$. For all i , the success probability function can be expressed as:

$$f(\vec{p}) = 1 - \prod_{P \in \mathcal{P}} (1 - \prod_{j \in P} p_j)$$

It is easy to see that $f(\vec{p})$ is linear in p_i for all i , and therefore multilinear. Therefore g is multilinear. In addition, the derivative of $f(\vec{p})$ with respect to p_i is positive:

$$\frac{\partial f(\vec{p})}{\partial p_i} = \prod_{j \in P_i: j \neq i} p_j > 0$$

Thus, we get:

$$\frac{\partial g(\vec{p}, N_1)}{\partial p_i} = \frac{\partial f(\vec{p})}{\partial p_i} = \prod_{j \in P_i: j \neq i} p_j > 0$$

while:

$$\frac{\partial g(\vec{p}, N_2)}{\partial p_i} = 0$$

Therefore, $\frac{\partial g(\vec{p}, N_1)}{\partial p_i} > \frac{\partial g(\vec{p}, N_2)}{\partial p_i}$ for all i , and g is single crossing.