

Mechanism Design with a Restricted Action Space

Liad Blumrosen*

Michal Feldman

Microsoft Research
and The Hebrew University

The Hebrew University

Abstract

While traditional mechanism design typically assumes isomorphism between the type space of the players and their action space, behavioral, technical or regulatory factors can severely restrict the set of actions that are actually available to players. We devise a general framework for the study of mechanisms in single-parameter environments with restricted action spaces. Our contribution is threefold. First, we characterize sufficient conditions under which the information-theoretically optimal solution can be implemented in equilibrium. Second, we prove that for a wide family of social-choice rules the optimal mechanisms with k actions incur an expected loss of $O(\frac{1}{k^2})$ compared to the optimal mechanisms with unrestricted action space. Finally, we fully characterize the optimal mechanisms in environments with two players and two alternatives. We apply our general results to signaling games, public-good models and information transmission in networks.

1 Introduction

In standard mechanism-design settings, a social planner wishes to implement some social-choice rule that chooses an alternative based on the private information of the players. Since social planners cannot observe the private information of the players (their *types*), they design mechanisms that make decisions by observing the *actions* of the players. Each player determines his action in the mechanism according to his type in order to maximize his own utility. The challenge of the social planner is to elicit information that will allow him to implement system-wise goals although such goals may conflict with the objectives of the individual players.

Much of the literature on mechanism design restricts attention to *direct revelation* mechanisms, in which the action spaces of the players are identical to their type spaces. This focus is owing to the

*The authors are grateful to Itzhak Gilboa, Zvika Neeman, Noam Nisan, Ady Pauzner, Yuval Salant, Ilya Segal and the participants of several workshops and seminars for valuable discussions.

*revelation principle*¹, which asserts that every mechanism can be transformed into an equivalent incentive-compatible direct-revelation mechanism that implements the same social choice function.

Nonetheless, in most practical settings, direct-revelation mechanisms are not viable since the number of actions available to the players is significantly smaller than their preference space. The most straightforward example are posted-price mechanisms for a single object; in such mechanisms users have only two available actions (to buy the object or not), and the seller clearly lacks the information that would enable her to implement the efficient allocation rule. Another example is the *signaling* model for the labor market by Spence (1973), where employees send signals about their skills to potential employers by the education level they acquire. Although there is a continuum of skill levels, it is unreasonable to expect more than a few education levels in practice (e.g., PhD, M.A., and B.A.). The screening model by Rothschild and Stiglitz (1976) is another example where one might expect a small number of actions; consider an insurance firm that wishes to sell different types of policies to different drivers based on their privately known caution levels. In this model, drivers may have a continuum of possible caution levels, but insurance companies offer only a small number of policies (e.g., a small number of deductible amounts in case of a claim) since it is probably infeasible to market and sell more than a few types of policies. More complicated rules for generating policies may be feasible, but they are rarely used in practice.

Mechanisms with a small, manageable set of choices are widespread in practice, and the main reason for this phenomenon is probably their simplicity. This claim is also supported experimentally, e.g., by Iyengar and Lepper (2000), who showed that a *choice overload* can hamper the willingness of the players to participate in the game, and can degrade their performance in a given transaction. Iyengar et al. compared decision making under a small set of choices and under larger choice sets (not unusually large) and showed that such phenomena are significant even when the number of possible actions is increased from 6 to around 24 or 30. In fact, in many real-life mechanisms the players are required to map their complex preferences into discrete, often dichotomic, decisions. For instance, many mechanisms avoid negotiations and simply post prices for packages or services, and the players are left to decide whether they buy or not under the posted prices. In other settings, players decide whether they participate in or abstain from some transaction, vote for or against some issue in a referendum, and many other similar examples.

Additionally, there are clear evidences for the rare practical use of direct-revelation mechanisms, most prominently VCG mechanisms. One major reason for this fact relates to the *price discovery* process; players usually do not know their exact types and the discovery process may be prohibitively costly (hiring consultants, etc.) or even computationally intractable to compute (see,

¹The work of Myerson (1981), Green and Laffont (1977) and Dasgupta, Hammond, and Maskin (1979) discusses the foundations of the revelation principle.

e.g., Larson and Sandholm (2001)). A well-designed mechanism with limited actions will guide the attention of the players to the information that is most relevant for the decision making. Another critical flaw of direct-revelation mechanisms is that players are typically unwilling to reveal their exact types, even if it is beneficial for them in the short run, worrying that this might harm them in future transactions. A small action space allows the players to preserve some degree of privacy. Papers by Rothkopf, Teisberg, and Kahn (1990) and Ausubel and Milgrom (2006) provide more details on why VCG mechanisms are indeed rare.

Mechanisms with a small action space, in specific models, were studied in several earlier papers. Wilson (1989) measured the effect of discrete “priority classes” of buyers on the efficiency of electricity markets and found that a few priority classes can realize most of the efficiency gains. In a related work, McAfee (2002) showed that in matching and rationing problems at least half of the social value created by optimal complex schemes can be obtained using very coarse action schemes. Dow (1991) considered a simplified decision problem of a single agent searching for a low price with a limited memory; the memory restrictions force the player to divide the set of possible histories into a limited number of categories. It turns out that the optimal partition of the history is obtained, as in our paper, by dividing the range of prices into disjoint intervals. Compared to the above work, our paper incorporates incentives issues in general multi-player domains and also characterizes the exact effect of the expressiveness level allowed in the system. A similar result was obtained in a different setting, studied in Bergemann and Pesendorfer (2001). There, a revenue-maximizing seller faces a set of bidders, who do not know their private types, and he needs to determine the accuracy level by which they learn their types. On the one hand, more information increases efficiency and thus the seller’s revenue, but on the other hand, it increases the information rent of the bidders, thus decreases the seller’s revenue. Once again, partitioning the information range into disjoint intervals is shown to maximize the seller’s revenue. The work of Blumrosen, Nisan, and Segal (2006) is the closest in spirit to our work. They studied single-item auctions with severely-restricted action space, and showed that nearly-optimal social welfare can be achieved even with very strict limitations on the action space. An earlier paper in a similar spirit is by Harstad and Rothkopf (1994) who analyzed discrete bid levels in English auctions.

We next present our framework and results.

1.1 Our Framework

We consider 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 distributed² on real intervals, and a social planner who wishes to implement a *social-choice function* c that maps every profile of types to a chosen alternative. We stress that although we explore the properties of Bayesian-Nash implementation in this paper, all the mechanisms that we construct possess the even stronger dominant-strategy equilibrium. Due to the limited expressiveness that is implied by the restricted action space, the social planner will typically have uncertainty about the desired alternative. That is, for some realizations of the players' types, the decision of the social planner will unavoidably be incompatible with the social-choice function c . In order to quantify how well bounded-action mechanisms can approximate the original social-choice function, we assume that the social-choice function is derived from a *social-value function* g , which assigns a real value to every combination of alternative A and realization $\vec{\theta} = (\theta_1, \dots, \theta_n)$ of the players' types. The social-choice function c will maximize the social value, i.e., $c(\vec{\theta}) \in \operatorname{argmax}_A \{g(\vec{\theta}, A)\}$.³ 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 players gain from it exceeds its construction cost C . There are two alternatives in this model: "build" and "do not build". The social value functions in a 2-player game is given by: $g(\theta_1, \theta_2, \text{"build"}) = \theta_1 + \theta_2 - C$, and $g(\theta_1, \theta_2, \text{"do not build"}) = 0$.
- *Single-item auctions.* Consider a 2-bidder auction where the auctioneer wishes to allocate the item to the bidder with the highest value. The social-choice function is given by $g(\theta_1, \theta_2, \text{"player 1 wins"}) = \theta_1$ and $g(\theta_1, \theta_2, \text{"player 2 wins"}) = \theta_2$.
- *Message delivery over networks.* A message can be delivered over a network composed of two parallel edges. Each edge is owned by a selfish player that has a privately-known probability q_i of delivering the 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, $g(q_1, q_2, \text{"send over network"}) = 1 - (1 - q_1) \cdot (1 - q_2)$ and $g(q_1, q_2, \text{"send over the alternate network"}) = 0.9$. Note that in this example the social-choice function is not welfare maximizing.

1.2 Our Contribution

The paper centers on the following question: when the players are only allowed to use k actions, which mechanisms achieve the optimal expected social value, and how do they compare to optimal

²Previous work observed that there are joint distribution functions for which the optimal mechanisms with restrictions on the action space become trivial or non-interesting; see, for instance, the work by Blumrosen, Nisan, and Segal (2006).

³Observe 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.

direct-revelation mechanisms? This question is actually composed of two questions.

1. An *information-theoretic* question: what is the optimal method to elicit information on the private information of the players when the players can only reveal information using k actions (recall that their type space may be continuous)?
2. A *game-theoretic* question: what is the best outcome achievable with k actions, given the additional constraint of implementation in a Bayesian-Nash equilibrium?

These two questions raise the question about the “*price of implementation*”: can the optimal information-theoretic result be always implemented in an equilibrium, and to what extent does the equilibrium requirement degrade the optimal result?

Example 1. Consider a public good model with two players whose types θ_1, θ_2 are uniformly distributed between $[0, 1]$. A social planner would like to build the bridge when $\theta_1 + \theta_2 > C$ where C is the construction cost of the bridge. It is well-known that if direct revelation is allowed, the VCG mechanism provides a socially-efficient solution. Assume now that only two actions are available to the players: “No” and “Yes”. Now, due to the inherent information-theoretic constraints, the social planner is no longer able to build the bridge exactly according to the objective function. What is an optimal 2-action mechanism? Consider the following allocation rule and strategies:

Allocation: the social planner always builds the bridge, unless both players report “No”.

Strategies: both players use the following threshold strategy:

“Report “No” if $\theta_i \leq \frac{2}{3} \cdot C$, otherwise report “Yes””

As will be shown later in the paper, the above solution is the best solution for the information-theoretic problem created by using only 2 actions (when $C \leq 1$); there is no other allocation scheme and no other pair of 2-action strategies that together obtain a higher expected social value. The obvious question is whether this result can be obtained in equilibrium, and the answer is affirmative: it is easy to see that the following payment scheme implies that the above strategies are dominant for both players: “If a player is the only player to report “Yes” he should pay $\frac{2}{3} \cdot C$, and otherwise he pays zero”. Consequently, the optimal information-theoretic solution can be supported with dominant strategies (and thus with Bayesian-Nash equilibrium) with no social-value loss!

In the remainder of this section, we informally survey our three research questions and results.

Our first contribution presents a family of social-value functions for which solving the information-theoretic problem actually also solves the game-theoretic problem. The following theorem holds for any number of alternatives, any number of players, and any profile of distribution functions.

Theorem 1: *For all multilinear single-crossing social-value functions, the information-theoretically optimal social-choice rule is dominant-strategy implementable.*

The theorem assumes two properties of the social-value functions – multilinearity and single crossing. *Multilinear* social-value functions are polynomials where each variable has a degree of at most one in each monomial. They capture many important and well-studied models, and include, for instance, any social-*welfare* maximization setting where the valuations are linear in the types (like public-good and auction models), and other models like the above message-delivery example. *Single crossing* is a stronger property than monotonicity, where the latter is required to guarantee the implementability of social-choice functions in the absence of restrictions on the actions. The reason for this stronger requirement is that action-bounded mechanisms will typically not be able to exactly implement the original social-choice function; therefore, the social value of all the alternatives should behave “monotonically,” not only for those alternatives that are chosen by the desired social-choice function (and thus maximize the social value). A formal definition will be given in the next section.

For proving Theorem 1, we prove a useful lemma that presents an alternative characterization of social-choice functions whose “price of implementation” is zero. We show that for every social-choice function, the implementability of the best information-theoretic solution is equivalent to the property that the optimal expected social value is achieved with *threshold* (or *non-decreasing*) strategies.⁴ This lemma actually implies that one can always implement in dominant strategies the optimal social-choice rule that is achievable with threshold strategies.

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 every number of players or alternatives, and for every profile of independent distribution functions, we construct mechanisms that are nearly optimal – up to an additive difference of $O(\frac{1}{k^2})$. This is the same asymptotic rate proved for specific environments by Wilson (1989), Harstad and Rothkopf (1994) and Blumrosen, Nisan, and Segal (2006). Moreover, a better general upper bound cannot be obtained as the work of Blumrosen, Nisan, and Segal (2006) shows that in some auction settings the optimal loss is exactly proportional to $\frac{1}{k^2}$. Note that there are social-choice functions that can be implemented with k actions with no loss at all (for example, the rule “always choose alternative A ”).

⁴The restriction to non-decreasing strategies is very common in the literature. One remarkable result by Athey (2001) 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. Another related result is by Dow (1991), who showed that the optimal way of an agent with limited memory to partition a given set of possible histories into a fixed number of categories is to use thresholds.

Our asymptotic result holds for any Lipschitz-continuous social-value function, i.e., functions for which the effect of local changes in the types on the social value is limited. In particular, all polynomials, including multilinear functions, are Lipschitz continuous.

Theorem 2: *For all single-crossing Lipschitz-continuous social-value functions, the optimal k -action mechanism incurs an expected social loss of $O(\frac{1}{k^2})$ compared with mechanisms with unrestricted action space.*

The proof for this theorem is constructive. We present mechanisms that never exceed this loss. Note that social planners can utilize this characterization to optimize the number of actions when this decision is under their control; that is, they should add an action only if its cost is smaller than the marginal contribution of the action to the expected social value. Due to the above result, we can bound the marginal contribution from an additional action by a value that is proportional to $\frac{1}{k^2} - \frac{1}{(k+1)^2}$, which is in the order of $\frac{1}{k^3}$.

Our final result concerns the problem of finding *the* mechanisms that maximize the expected social value. We fully characterize the optimal mechanisms in environments with two players and two alternatives for every number of actions k and every pair of distribution functions of the players' types. We present them in two parts: we first show that the optimal *allocation* scheme is "diagonal" in the sense that in its matrix representation one alternative will be chosen in, and only in, entries below one of the main diagonals. We then characterize the optimal *strategies* – strategies that are "mutually maximizers". Counter-intuitively (and in contrast to the results obtained in Blumrosen, Nisan, and Segal (2006) in the context of auctions), the optimal "diagonal" mechanism may not utilize all the k available actions for some non-trivial social-value functions.

Theorem 3: *For all multilinear single-crossing social-value functions over two alternatives, the 2-player k -action mechanism that maximizes the social value is diagonal and it possesses dominant strategies that are mutually maximizers.*

Pinpointing the optimal action-bounded mechanism for multi-player or multi-alternative environments seems to be harder and remains an open question. The hardness stems from the fact that the number of diagonal mechanisms is growing exponentially in the number of players.

Finally, we present our results in the context of several natural applications. First, we provide 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 in Blumrosen, Nisan, and Segal (2006)), and show how the optimal allocation scheme depends on the construction cost C . Then, we study the celebrated *signaling* model for the labor market, which is a natural application in our context since education levels are often discrete. Lastly, we present our results in the context

of message delivery in networks. The latter example illustrates how our results apply to settings where the objective function of the social planner is other than welfare maximization.

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. Asymptotic analysis of the social-value loss is given in Section 4. In Section 5 we fully characterize the optimal mechanisms in 2-player environments with two alternatives. Section 6 presents applications of our general results. Some of the proofs are deferred to the appendix.

2 Model and Preliminaries

We first describe a general mechanism-design model for players with one-dimensional types. Later, in Subsection 2.2, we impose limitations on the action space. Consider 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 $v_i : [\underline{\theta}_i, \bar{\theta}_i] \times \mathcal{A} \rightarrow \mathbb{R}$. 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 type 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* $g : \Theta \times \mathcal{A} \rightarrow \mathbb{R}$. That is, $c(\vec{\theta}) \in \mathit{argmax}_{A \in \mathcal{A}} \{g(\vec{\theta}, A)\}$

We assume that for every alternative $A \in \mathcal{A}$, the function $g(\cdot, A)$ is continuous and differentiable with respect to every type. The players reveal information about their types by choosing an *action*, from an action set B .⁵

A strategy of each player is a function $s_i : [\underline{\theta}_i, \bar{\theta}_i] \rightarrow B$, mapping each possible type to an action. 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} . We assume that players have quasi-linear utility functions. Thus, 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 Implementation

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

⁵We assume that the action space is symmetric for all players, and this assumption can easily be relaxed (except for the characterization results in Section 5).

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.⁶
- $p : B^n \rightarrow \mathbb{R}^n$ is the payment scheme (i.e., $p_i(b)$ is the payment to the i th player given a vector of actions b).

In this paper we study mechanisms with Bayesian-Nash Equilibria. All the mechanisms that we construct in this paper, however, admit the stronger equilibrium concept of dominant strategies. This is not surprising, given the existing characterization of Bayesian-Nash implementable functions (see below, and also Mookherjee and Reichelstein (1992)). Following are the standard definitions of those equilibrium concepts.

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} . That is, for every type θ_i and action b'_i and b_{-i} , we have that

$$v_i(\theta_i, t(s_i(\theta_i), b_{-i})) - p_i(s_i(\theta_i), b_{-i}) \geq v_i(\theta_i, t(b'_i, b_{-i})) - p_i(b'_i, b_{-i})$$

We say that a strategy s_i is a *best response* to the strategies s_{-i} of the other players, if it gains player i the highest expected utility given that the other players play s_{-i} , i.e., for all types θ_i and bids b'_i :

$$E_{\theta_{-i}} [v_i(\theta_i, t(s_i(\theta_i), s_{-i}(\theta_{-i}))) - p_i(s_i(\theta_i), s_{-i}(\theta_{-i}))] \geq \tag{1}$$

$$E_{\theta_{-i}} [v_i(\theta_i, t(b'_i, s_{-i}(\theta_{-i}))) - p_i(b'_i, s_{-i}(\theta_{-i}))] \tag{2}$$

A profile of strategies s_1, \dots, s_n is a *Bayesian-Nash equilibrium* if for every $1 \leq i \leq n$, s_i is a best response to s_{-i} .

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 Bayesian-Nash 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., for every θ we have $t(s_1(\theta_1), \dots, s_n(\theta_n)) = h(\vec{\theta})$.

Fundamental results in the mechanism-design literature state that under a “single-crossing” condition, *monotonicity* of the social-choice function is a sufficient and necessary condition for implementability (in single-parameter environments). The single-crossing condition (in its different

⁶We will show in the proof of Lemma 1 that, w.l.o.g., we can focus on deterministic allocation schemes.

variants, like the Spence-Mirrlees condition, see Spence (1973) and Mirrlees (1971), or the Milgrom-Shannon condition specified in Milgrom and Shannon (1994)) appears, very often implicitly, in almost every paper on mechanism design in domains with one-dimensional types. Throughout this paper, we assume that the preferences of each one of the players are single-crossing. Our definition of single-crossing valuations may be considered as a hybrid of the differential Spence-Mirrlees single-crossing property and the order-theoretic Single-Crossing property (discussions on these two variants can be found in the work of Milgrom and Shannon (1994) and Edlin and Shannon (1998)). We find this variant convenient as it captures a multitude of models where the type space is a continuous real interval and the space of alternatives, for which players may have individual ordering, is discrete.

A valuation function for player i is single-crossing if for every two non-equivalent alternatives, the effect of an increment in θ_i is greater for one of these alternatives for every θ_i . The single-crossing condition actually defines an order on the alternatives for each one of the players. For example, if the value of player i for alternative A increases more rapidly than his value for alternative B , we can denote it by $A \succ_i B$. This definition rules out preferences where the value for an alternative increases more rapidly (compared to another alternative) on some parts of the support, and slower than the other alternative on different parts of the support. Later on, we will use these orders on the alternatives for defining monotonicity of social-choice functions.

Definition 3. A valuation function $v_i : [\underline{\theta}_i, \bar{\theta}_i] \times \mathcal{A} \rightarrow \mathbb{R}$ is single crossing if there is a partial order \succ_i on the alternatives, such that for every two alternatives $A_j \succ_i A_l$ we have that for every θ_i ,

$$\frac{\partial v_i(\theta_i, A_j)}{\partial \theta_i} > \frac{\partial v_i(\theta_i, A_l)}{\partial \theta_i}$$

and if neither $A_l \succ_i A_j$ nor $A_j \succ_i A_l$ (denoted by $A_j \sim_i A_l$) the functions are identical, i.e., $v_i(\theta_i, A_j) = v_i(\theta_i, A_l)$ for every θ_i . We also denote $A_j \succeq_i A_l$ if either $A_j \sim_i A_l$ or $A_j \succ_i A_l$.

Example 2. Consider a single-item auction among 3 players, with 3 alternatives: $A_1 =$ "1 wins", $A_2 =$ "2 wins", and $A_3 =$ "3 wins". For each player i , $v_i(A_i, \theta_i) = \theta_i$ and for $j \neq i$, $v_i(A_j, \theta_i) = 0$. Indeed, for player 1 the slope of $v_1(A_1, \theta_1)$ is greater than the slope of $v_1(A_2, \theta_1)$ and therefore $A_1 \succ_1 A_2$. The losing alternatives for player 1 gains her the same value therefore $A_2 \sim_1 A_3$.

The definition of monotone social-choice functions requires an order on the actions as well. This order is implicit in most standard settings where, for example, it is defined by the 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 of the actions 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 $<, >.$)

Definition 4. *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 every $b_{-i} \in \{0, \dots, k-1\}^{n-1}$ if $b'_i > b_i$ then $t(b'_i, b_{-i}) \succeq_i t(b_i, b_{-i})$.*

Following is a classic result regarding the implementability of social-choice functions in single-parameter environments. The formal argument is given, for example, in Mookherjee and Reichelstein (1992) and Segal (2003) and also in the survey Hermalin (2005).

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 Bayes-Nash implementable if and only if c is monotone.*

2.2 Restrictions on the Action Space

We study environments where the action space B is restricted. We define a k -action mechanism to be a mechanism in which the number of possible actions for each player is k , i.e., $|B| = k$. In k -action mechanisms, 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*: $E_{\vec{\theta}} \left[g \left(\vec{\theta}, t(s_1(\theta_1), \dots, s_n(\theta_n)) \right) \right]$. We next define social-choice functions that can be achieved by k -action mechanisms. Note that, as opposed to Definition 2, this is an information-theoretic definition that does not involve strategic arguments.

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 every 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.*

Example 3. *Consider an environment with two alternatives $\mathcal{A} = \{A, B\}$, and the following desired social-choice function: $\tilde{c}(\theta_1, \theta_2) = A$ iff $\{\theta_1 > 1/2 \text{ and } \theta_2 > 1/2\}$. \tilde{c} is informationally achievable with two actions: if both players report “1” when their value is greater than 1/2 and “0” otherwise, then the allocation rule “choose alternative A iff both players report 1” results in exactly the desired allocation for every profile of types. Conversely, 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.*

Given a social-value function, we would like to determine mechanisms that maximize the expected social value, given the information-theoretic constraints.

Definition 6. *A social-choice function is k -action informationally optimal with respect to the social-value function g , if it is k -action informationally achievable, and it achieves the maximal expected social value among all the k -action informationally achievable social-choice functions.⁷*

As we will show later, it turns out that the monotonicity of the social-choice function will not suffice for ensuring the monotonicity of the k -action mechanisms. While monotonicity describes the structure of the choices that maximize the social value, mechanisms with discrete action spaces will also take into account the social value obtained by other alternatives. Therefore, the social-value of all the alternatives should similarly be aligned with the preferences of the players. Therefore, we define a single-crossing property on the social-value function g which is stronger than monotonicity.

Definition 7. *Let \succ_1, \dots, \succ_n be the orders on the alternatives implied by the single-crossing condition on the valuations of the players. We say that the social-value function $g(\vec{\theta}, A)$ exhibits the single-crossing property if the following condition is met for every player i :*

for every two alternatives such that $A_j \succ_i A_l$ we have that for every $\vec{\theta} \in \Theta$,

$$\frac{\partial g(\vec{\theta}, A_j)}{\partial \theta_i} > \frac{\partial g(\vec{\theta}, A_l)}{\partial \theta_i}$$

and if $A_j \sim_i A_l$ then for every $\vec{\theta}$ we have $\frac{\partial g(\vec{\theta}, A_j)}{\partial \theta_i} = \frac{\partial g(\vec{\theta}, A_l)}{\partial \theta_i}$

Note that, unlike Definition 3, we do not require that the social value of equivalent alternatives will be identical, but we only require identical slopes.⁸

Example 4. *Consider the auction setting from Example 2. $A_1 \succ_1 A_2$, and indeed the social welfare in A_1 is v_1 and has a greater slope than the welfare in A_2 , v_2 , as a function of v_1 . $A_2 \sim_1 A_3$, and indeed, v_2 and v_3 have identical slopes for all values of θ_1*

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.

⁷By results shown later in the paper, this maximum is attained and the optimal function is well defined. This holds since the optimal results are achieved by threshold strategies, hence every allocation scheme defines a compact set of social values, and there are finite number of different allocation schemes.

⁸This difference can be demonstrated in multi-item auctions: two allocations in which player i receives the same bundle of items are clearly identical with respect to this player, but their social welfare may differ since the items may be allocated differently among the other players. However, the social welfare changes at the same rate as the value of player i increases.

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 every 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 a 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 lemma that provides a sufficient and necessary condition for the implementability of the optimal solution achievable with k actions: this condition says that the informationally optimal social-choice rule is achieved when all the players use non-decreasing strategies. The intuition behind it 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 alternative increases faster; therefore, monotonicity must hold. The result holds for every number of players and alternatives, and for every profile of distribution functions on the players' types, as long as they are statistically independent ⁹.

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 threshold strategies.

The proof for this lemma can be found in the appendix. Theorem 1 below is based on one direction of the lemma (optimum with threshold strategies implies implementability); the other direction is given for completeness.¹⁰

Next, we show that for a wide family of social-value functions – multilinear functions – the information-theoretically optimal rule is implementable. This family of functions captures many common settings from the literature.

Definition 10. A multilinear function is a polynomial in which the degree of every variable in each monomial is at most 1.¹¹ We say that a social-value function g is multilinear, if $g(\cdot, A)$ is multilinear for every alternative $A \in \mathcal{A}$.

⁹One can easily verify that this result does not hold if the players' types are dependent.

¹⁰We currently do not have a concrete example for social choice functions that achieve an optimum with strategies other than threshold strategies and this remains an open problem.

¹¹For example, $f(x, y, z) = xyz + 5xy + 7$.




Figure 1: Each linear function in the diagram corresponds to the expected social value as a function of the type of Player 1 when she chooses a particular action. Since two linear functions cross at most once, the maximum of k linear functions has at most $k - 1$ breaking points (e.g., x_1, x_2, x_3). Therefore, in order to maximize the social value, Player 1 uses a threshold strategy where those breaking points are the thresholds.

The basic idea behind the proof of the following theorem is as follows: for every player, we show that the expected social welfare for every action he chooses (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 breaking points of the piecewise linear function). Figure 1 graphically illustrates this argument. Since the optimum is achieved with threshold strategies, we can apply Lemma 1 to show the monotonicity of the social-choice rule. Note that in this argument we characterize the players' strategies that maximize the social value rather than 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.*

Proof. We will show that for every k -action mechanism, the optimal expected social value is achieved when all players use threshold strategies. This will be shown by proving that for every player i and for every action b_i of player i , the expected welfare when she chooses the action b_i (fixing the strategies of the other players) 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.

Let t be the allocation function of the mechanism, and let $s_{-i}(\theta_{-i})$ be the strategy profile of all players other than i . For a fixed action b_i of player i , let q_A denote the probability that alternative A is allocated, i.e.,

$$q_A = Pr_{\theta_{-i}} \left[t(s(\vec{\theta})) = A \mid s_i(\theta_i) = b_i \right]$$

Due to the linearity of expectation, the expected social value when player i with type θ_i reports

b_i is:

$$E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i})))] = \sum_{A \in \mathcal{A}} q_A E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, A) | t(b_i, s_{-i}(\theta_{-i})) = A] \quad (3)$$

$$= \sum_{A \in \mathcal{A}} q_A \int_{\theta_{-i}} g(\theta_i, \theta_{-i}, A) f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \quad (4)$$

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

Since g is multilinear, $g(\theta_i, \theta_{-i}, A)$ is a linear function in θ_i for every alternative, where the coefficients depend on the values of θ_{-i} . Denote this function by $g(\theta_i, \theta_{-i}, A) = \lambda_{\theta_{-i}} \theta_i + \beta_{\theta_{-i}}$. Thus, we can write Equation 4 as:

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

In this expression, each integral is a constant independent of θ_i when the strategies of the other players 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.

Finally, we must handle one subtle issue. Showing that the informationally optimal k -action social-choice rule is monotone is actually not enough. We should also show that the same amount of actions also suffices for determining the *prices* that support the (dominant-strategy) implementation of this rule. This clearly holds in our setting. Formally, we can apply Proposition 1 from Segal and Fadel (2006) that claims that in any simultaneous mechanism, the information that allows computing some implementable social-choice function is also sufficient for computing the supporting prices. \square

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 in 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) should be 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.

Also 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.*

4 Asymptotic Analysis of the Social-Value Loss

In this section we prove an upper bound on the social-value loss as a function of the number of actions k . In particular, we show that the social value loss diminishes quadratically with the number of possible actions, k . This result holds for any social value function that is Lipschitz-continuous, and includes, among others, all the bounded-degree polynomials. The main challenge here, compared to earlier results, is in dealing with general social-value functions and any number of players and alternatives. In particular, the social-value function may be asymmetric with respect to the players' types and social-value loss may a-priori occur in every profile of actions.

The basic intuition for the proof is that 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-value functions, even without the Lipschitz-continuous property. Then, Lipschitz-continuity implies that the social-value loss will always be $O(\frac{1}{k})$ in the mechanisms we construct. Taken together, the expected loss becomes $O(\frac{1}{k^2})$. We present an explicit construction for mechanisms that exhibit the desired loss in dominant strategies. The expected social-value loss clearly depends on the length of the support of the type space. Here, we assume that the type space is normalized to $[0, 1]$, that is, for every player i , $\underline{\theta}_i = 0$ and $\bar{\theta}_i = 1$.

Theorem 2. *Assume that the type spaces are normalized to $[0, 1]$. For every number of players and alternatives, and for every set of distribution functions of the players' types, if the social-value*

function g is single crossing and Lipschitz-continuous, then the informationally-optimal k -action social-choice function (with respect to g) incurs an expected social-value loss of $O(\frac{1}{k^2})$.

Proof. 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 pair of consecutive thresholds is $O(\frac{1}{k})$. Using these strategies, we define a mechanism that achieves an $O(\frac{1}{k^2})$ loss. For simplicity, we assume that k is even.

Construction of the threshold strategies:

For each player i let $Y^i = \{y_0^i = 0, y_1^i, \dots, y_{\frac{k}{2}-1}^i, y_{\frac{k}{2}}^i = 1\}$ be a set of thresholds that divide the density function of player i to $\frac{k}{2}$ equi-mass intervals. That is, for every j we have $F_i(y_{j+1}^i) - F_i(y_j^i) = F_i(y_j^i) - F_i(y_{j-1}^i) = \frac{2}{k}$.

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

Now, let $X^i = Y^i \cup Z^i$ be the set of thresholds for player i . That is, player i uses strategy s_i based on the thresholds X^i . Clearly, using a threshold strategy based on X^i (when the thresholds are ordered in an 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 an alternative that maximizes the expected social value when the players use the threshold strategies s based on the vectors X^i defined above. That is,

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

Analysis:

We say that an action profile b is *decisive* if one alternative maximizes the social value for every profile of types (otherwise the profile is *indecisive*). Formally, an action profile 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 profile $\vec{\theta}$, such that $s_i(\theta_i) = b_i$ for every player i . Similarly, the profile 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 players choose the actions b .

We will prove that the above mechanism incurs an expected loss of $O(\frac{1}{k^2})$ using the following two claims. Claim 1 shows that the number of indecisive action profiles is $O(k^{n-1})$. Since the player chooses each action with probability $O(\frac{1}{k})$, each indecisive action profile is chosen with probability $O(\frac{1}{k^n})$, and therefore an indecisive profile 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 profile. Taken together, the expected social-value loss in the above k -action mechanism is $O(\frac{1}{k^2})$.

Claim 1. *For single-crossing social value functions, the number of indecisive action profiles 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 profile with respect to alternatives A and B (assume that $A \succ_1 B$ and $B \succ_2 A$, the other cases are treated analogously). Since the action profile is indecisive, there must be types θ_1, θ_2 for which $s_1(\theta_1) = b_1$ and $s_2(\theta_2) = b_2$, 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 profile 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 $s_1(\theta'_1) = b'_1$ and $s_2(\theta'_2) = b'_2$ 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 action profile 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 profiles for any profile $b_{-\{1,2\}}$ of the other players. Every indecisive action profile is clearly indecisive with respect to some pair of alternatives, thus the number of indecisive action profiles (given $b_{-\{1,2\}}$) is at most $\binom{|A|}{2} \cdot (2k - 1) = O(k)$. Therefore, for every pair of players (out 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 profiles. The total number of indecisive action profiles will therefore be $O(k^{n-2}) \cdot O(k) = O(k^{n-1})$. \square

Claim 2. *For all Lipschitz-continuous social-value functions, the social-value loss incurred when the players play an indecisive action profile is $O(\frac{1}{k})$.*

Proof. Consider an indecisive profile of actions b with respect to a pair of alternatives 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 action profile is also $O(\frac{1}{k})$.

Consider two profiles of types θ and θ' for which the profile of actions b is chosen by the players. We will prove that $g(\theta, A) - g(\theta', B) = O(\frac{1}{k})$.

Since the profile b is indecisive with respect to A, B , and since the social-value function is continuous, we know that there is a profile of types θ^* for which the players choose the actions in b and such that $g(\theta^*, A) = g(\theta^*, B)$. Consider some profile of types θ for which the profile of actions b is chosen. We will show that $|g(\theta, A) - g(\theta^*, A)|$ is at most $O(\frac{1}{k})$, and similarly one can show that $|g(\theta, B) - g(\theta^*, B)|$ is $O(\frac{1}{k})$ and the theorem will follow.

Since the social-value function is Lipschitz-continuous, there exists a non-negative constant α such that for every alternative A , $|g(\theta, A) - g(\theta^*, A)| \leq \alpha \cdot \sum_{i=1}^n |\theta_i - \theta_i^*|$. Since the same action profile is chosen for both θ and θ^* , our construction implies that for every i , $\theta_i - \theta_i^* < \frac{2}{k}$. The claim follows. \square

This concludes the proof of the theorem. \square

Moreover, as proved by Blumrosen, Nisan, and Segal (2006), this bound is asymptotically tight in several environments. That is, there exist a set of distribution functions for the players and social-value functions (e.g., the uniform distribution in auctions and public-good settings) for which *every* mechanism incurs a social-value loss of at least an order of $\frac{1}{k^2}$. Obviously, this claim does not imply that the loss of *every* social-choice function will be proportional to $\frac{1}{k^2}$. For example, in the social-choice function that chooses the same alternative for every type profile, no loss will ever be incurred (even with 0 actions).

5 Optimal Mechanisms for Two Players and Two Alternatives

While in the previous section we presented k -action mechanisms that are asymptotically optimal, we will now consider the problem of finding *the* k -action mechanisms that maximize the expected social value. We will show a full solution for action-bounded environments with two players and two alternatives, when the social-value function is multilinear and single crossing, and for every pair of distribution functions and every number of actions. This solves the problem, for example, for 2-bidder auctions and 2-player public-good games. The characterization of the optimal mechanisms for arbitrary number of alternatives and players remains an open problem, and we will illustrate the intuition behind its hardness.

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. By doing this, we completely

	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 2: 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.

describe which alternative is chosen for every profile of types. It is well known that in monotone mechanisms for one-dimensional environments, the allocation scheme uniquely defines payments that support dominant-strategy implementation. We find this description, which does not explicitly mention the payments, simpler for presentation.

The characterization of the optimal k -action mechanisms is presented in two stages: we first illustrate the allocation scheme in the optimal mechanisms and prove that they must be "diagonal". We then define the optimal strategies in such mechanisms, and prove that they exhibit the "mutually-maximizers" property.

5.1 Diagonal Allocations

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 it is allowed to use, and therefore one might infer that such a mechanism cannot be optimal. Using this property, we then define "diagonal" mechanisms that turns out to exactly characterize the 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., $\forall b_{-i}, t(b_i, b_{-i}) = t(b'_i, b_{-i})$).*

Consider a representation of the allocation scheme using a matrix, where each entry specifies the chosen alternative where the action of one player is choosing a row, and the action of the of the other player is choosing a column. Then, a 2-player mechanism is degenerate with respect to the row player, if there are two rows with identical allocation. We can now define diagonal allocation scheme.

Definition 12. *An allocation scheme for 2-player 2-alternative mechanisms 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. Figure 2 depicts some diagonal 4-action allocation schemes. Simple combinatorial arguments show that diagonal mechanisms may come in very few forms. The direction of the diagonal is determined by whether the players have the same order \succ_i on the alternatives (as in public-good games) or not (like in auctions).

Proposition 2. *Every 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*

$$t(b_1, b_2) = B \text{ iff } b_1 + b_2 \geq k - 1 \quad \text{or} \quad t(b_1, b_2) = B \text{ iff } b_1 + b_2 \geq k$$
2. *If the two players have conflicting preferences (w.l.o.g., $A \succ_1 B$ and $B \succ_2 A$) then either*

$$t(b_1, b_2) = B \text{ iff } b_1 \geq b_2 \quad \text{or} \quad t(b_1, b_2) = B \text{ iff } b_1 > b_2$$
3. *One of the above mechanisms, when one of the fixed-allocation actions is removed for one of the players (i.e., we can subtract the action j of player i such that for any two actions b, b' of the other player we have $t(j, b) = t(j, b')$).*

Proof. Note that in a monotone allocation scheme, there are $k + 1$ possible columns with k alternatives (e.g., for $k = 3$, $[A, A, A]$, $[A, A, B]$, $[A, B, B]$, $[B, B, B]$). Assume that the mechanism is non-degenerate, for example, w.r.t. Player 2 (the column player). If the column $[A, \dots, A]$ appears in the allocation matrix, then clearly the row $[B, \dots, B]$ does not appear there, which leaves k possible distinct rows for the row player. Note that in this case, when we exclude the row $[A, \dots, A]$ of the row player we are still left with k distinct actions for the column player (see Item 3 in the proposition). The actual matrix is defined by the orders on the alternatives, as shown in the proposition. □

We will show that the social-value is maximized in mechanisms with diagonal allocation scheme. Interestingly, one of the possible forms of diagonal mechanisms is degenerate with respect to one of the players (see Item 3 in Proposition 2); that is, it can be described as a mechanism with $k - 1$ actions for this player. For example, the rightmost allocation scheme in Figure 2 will maximize the social value for some 4-action environments, although the row player has only 3 actions. This auction can be viewed as the leftmost mechanism in Figure 2 when the bottom row has been removed.

5.2 Mutually Maximizer Threshold Strategies

In Section 5.1, we provided a characterization of the allocation scheme of the social-value maximizing mechanisms. Here, we complete the characterization of the optimal mechanisms by defining the optimal pricing rules – the pricing rules that support the optimal strategies. We define the notion of *mutually-maximizer* thresholds, and show that threshold strategies based on such thresholds are optimal. The intuition behind it is as follows. Consider some action b (“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, alternative A will be chosen for low actions of Player 2, and alternative B will be chosen for higher actions of Player 2. By determining a threshold for Player 2 that will be used in his threshold strategy, the social planner actually determines the minimal type of Player 2 from which alternative B will be chosen when the row player chooses action b . 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 chooses the action b); for greater values of Player 2, the single-crossing condition ensures that B will be preferred. The diagonal allocation scheme ensures that the value of each threshold follows from those arguments that concern only a single action of the other player.

Definition 13. *Consider a diagonal mechanism, where the players use threshold strategies based on the threshold vectors x, y .¹² 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., $t(i, j) \neq t(i - 1, j)$.

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

Example 5. *Consider the public-good setting in Example 1. The types of the two players are uniformly distributed between $[0, 1]$, each player has 2 actions “0” and “1”, and the mechanism builds the bridge unless both bidders choose the action “0” (this optimal mechanism is illustrated in the left table of Figure 3). Assume that Player 1 chooses the action “0”, and uses a threshold strategy based on the threshold $\frac{2}{3} \cdot C$ (where C is the construction cost of the bridge). What is the minimal type of Player 2 for which the social planner will build the bridge? The expected value of*

¹²For simplicity of presentation, we assume that the mechanism is non-degenerate w.r.t. both players; otherwise, the definition is similar but requires adjusting the indices.

Player 1, given that he chooses "0", is $\frac{C}{3}$. Therefore, the bridge should be built for any θ_2 such that $\frac{C}{3} + \theta_2 \geq C$, that is $\theta_2 \geq \frac{2}{3} \cdot C$. It follows that the threshold strategies based on $\frac{2}{3} \cdot C$ are mutually maximizers in this game. For further discussion on the public-good example, see Section 6.1.

5.3 The Optimal 2-Action 2-Player Mechanisms

It turns out that in 2-player, 2-alternative environments, where the social-value rule is multilinear and single crossing, the optimal expected social value is achieved in diagonal mechanisms with mutually-maximizer strategies.

The proof centers on proving that the allocation scheme is non-degenerate with respect to one of the players. In non-trivial mechanisms, this, together with monotonicity, will also show that the other player will either have non-degenerate allocation, or slightly degenerate allocation (i.e., $k - 1$ distinct actions). We actually show that in an optimal k -action allocation scheme one of the players will always have k distinct strategies, otherwise we can add a new action for this player and strictly increase the expected social welfare. The proof requires dealing with several sub-cases and is deferred to the appendix.

Theorem 3. *In non-trivial¹³ 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, for some distribution functions, degenerate with respect to one of the players (that is, equivalent to a game where one of the players has only $k - 1$ different actions). Moreover, the proof also identifies the following sufficient condition under which the optimal mechanism will be non-degenerate with respect to both players: if the players have the same order on the alternatives (e.g., $B \succ_1 A$ and $B \succ_2 A$), then the optimal alternative must be identical under the profiles $(\underline{\theta}_1, \bar{\theta}_2)$ and $(\bar{\theta}_1, \underline{\theta}_2)$.¹⁴ Similarly, if the players have the opposite order on the alternatives (e.g., $A \succ_1 B$ and $B \succ_2 A$), then the alternative with the higher social value must be identical for $(\underline{\theta}_1, \underline{\theta}_2)$ and $(\bar{\theta}_1, \bar{\theta}_2)$. This condition clearly holds in the public-good model presented in Section 6.1 and in auctions.

The full characterization of the optimal mechanisms in multi-player and multi-alternative environments is still an open question. The hardness stems from the fact that the necessary conditions

¹³Environments are non-trivial if the original social-choice function chooses each alternative with positive probability, and if for both players the single-crossing condition on the alternative is strict (i.e., \succ_i and not \succeq_i). Otherwise, the solution is easy.

¹⁴More precisely, the condition for non-degeneracy when $B \succ_1 A$ and $B \succ_2 A$ is that $sign(g(\underline{\theta}_i, \bar{\theta}_i, A) - g(\underline{\theta}_i, \bar{\theta}_i, B)) = sign(g(\bar{\theta}_i, \underline{\theta}_i, A) - g(\bar{\theta}_i, \underline{\theta}_i, B))$ (when $sign(0)$ is considered both negative and positive).

$C \leq 1$	0	1
0	<i>No</i> $p_1 = p_2 = 0$	Yes $p_1 = 0; p_2 = \frac{2C}{3}$
1	Yes $p_1 = \frac{2C}{3}; p_2 = 0$	Yes $p_1 = p_2 = 0$

$C \geq 1$	0	1
0	<i>No</i> $p_1 = p_2 = 0$	<i>No</i> $p_1 = p_2 = 0$
1	<i>No</i> $p_1 = p_2 = 0$	Yes $p_1 = p_2 = \frac{2}{3}C - \frac{1}{3}$

Figure 3: 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 $C \geq 1$. The bridge is built in entries labeled as "Yes".

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, the number of monotone and non-degenerate mechanisms rapidly increases as the number of players n grows (it can be shown to grow exponentially in n , a proof is given by Blumrosen and Feldman (2005)). Unlike the 2-player 2-alternative case, it seems that pinpointing the best allocation scheme cannot be done independently of finding the optimal strategies, causing a considerable growth in the complexity of determining the solution.

6 Applications

In this section, we demonstrate the applicability of our results to public-good models, signaling games and message delivery in networks.

6.1 Application 1: Public Goods

This section will discuss in more details the public-good model which was discussed above. The model deals with a social planner who 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 drawn from the interval $[0, 1]$. 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, and this cost is 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 to the informational constraints. Yet, since $g(v, \text{Yes})$ and $g(v, \text{No})$ are multilinear and single crossing, we can directly apply Theorem 1. Hence, the information-theoretically optimal k -action mechanism is implementable.

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 public-good game we can fully characterize the optimal k -action 2-player mechanisms. As mentioned in Section 5.3, when $g(\bar{\theta}_1, \underline{\theta}_2, A) = g(\bar{\theta}_1, \underline{\theta}_2, B)$ 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. An immediate corollary from Theorem 3 is a full characterization of the optimal mechanisms in this setting:

Corollary 3. *The optimal expected welfare in a 2-player k -action public-good game is achieved with one of the following mechanisms (where $x_0 = y_0 = 0$ and $x_k = y_k = 1$):*

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 $1 \leq i \leq k-1$,

$$x_i = C - E[v_2 | v_2 \in [y_{k-i}, y_{k-i+1}]]$$

$$y_i = C - E[v_1 | v_1 \in [x_{k-i}, x_{k-i+1}]]$$

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 $1 \leq i \leq k-1$:

$$x_i = C - E[v_2 | v_2 \in [y_{k-i-1}, y_{k-i}]]$$

$$y_i = C - E[v_1 | v_1 \in [x_{k-i-1}, x_{k-i}]]$$

The construction cost determines which of the two mechanisms above obtains a better result. Recall that we define the optimal mechanisms by their (monotone) allocation scheme and by the optimal strategies for the players. The payments that support dominant-strategy implementation 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 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 apply Corollary 3 for the specific case of the uniform distribution. The example shows how the optimal mechanism is determined by the cost C : a mechanism of type 1 is optimal for construction costs smaller than 1, while a mechanism of type 2 is optimal for higher costs. Note that the optimal mechanisms are symmetric, unlike the solution for auctions in Blumrosen, Nisan, and Segal (2006).

Example 6. *Suppose that the types of both players are uniformly distributed on $[0, 1]$. Figure 3 illustrates the optimal mechanisms for $k = 2$, and how they depend on the construction cost C . For every number of actions k , the welfare-maximizing mechanisms are:*

- 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 $i = \{1, \dots, k-1\}$), $x_i = \frac{2iC}{2k-1}$

6.2 Application 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 wishes 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 productivity levels 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. The standard 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 latter requirement is exactly equivalent to the single-crossing property. 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, a continuum of education levels is somewhat unrealistic. 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 workers to play according to specific threshold strategies. It is easy to verify that the standard condition, the single-crossing condition on the cost function, suffices for ensuring that these threshold strategies will be dominant for the players. We can now apply Theorem 2, and show that if the decision function f of the firm is Lipschitz-continuous (i.e., the decisions are made to maximize a set of Lipschitz-continuous functions), then the firm can design the education system such that the expected loss will be $O(\frac{1}{k^2})$, with a dominant-strategy equilibrium. Note that while in the classic example of the job market it is unreasonable for each firm to select the education level, in other reasonable applications the social planners may be able to determine the thresholds, e.g., by fixing the levels of qualifying exams or other means for the players to demonstrate their skills.

Corollary 4. *Consider a Lipschitz-continuous 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 original decision function f .*

6.3 Application 3: Message Delivery in Networks

Lastly, we show the applicability of our results to settings where messages should be delivered over lossy communication networks. Different parts of the networks, i.e., edges in their graph, are owned by rational players who possess a privately known probability of successfully delivering a message (or completing another task) over this edge. Each player owns at most one edge. A sender knows the topology of the networks, and has to devise a mechanism for deciding which network has the highest success probability. It is natural to assume that the players (i.e., links) may not be able to report (or to figure out) their accurate probabilities of success, but only, e.g., whether these are “low”, “intermediate”, or “high”.

Consider a set of networks, where each network is composed of multiple parallel paths from a given source to a given destination. An example for such a network appears in Figure 4; in this example, the probability that a message will be transmitted successfully in the upper path, for instance, is $q_1 \cdot q_2$. The sender wishes to send the message through the network with the highest success probability.

In this example we assume that every player has a single-crossing valuation function over the alternatives. That is, each player wishes that the message will be sent through his network, and his benefit is positively correlated with his private data (e.g., the valuation of player i for delivering the task may be exactly q_i). We would like to emphasize that the social planner in this example (the sender) does not aim to maximize the social welfare. That is, the social value is neither the




Figure 4: An example for a parallel-path network, where each link has a probability q_i for transmission success. We show that the overall probability of success in such networks is multilinear in q_i , and thus the optimal k -action social-choice function is dominant-strategy implementable.

sum of the players' types nor any weighted sum of the types ("affine maximizer").

The success probability of sending a message through a parallel-path network is multilinear, since it can be expressed by the following multilinear formula (where \mathcal{P} denotes the set of all paths between the source and the sink):

$$f(\vec{q}) = 1 - \prod_{P \in \mathcal{P}} (1 - \prod_{j \in P} q_j) \quad (5)$$

For example, in the network presented in figure 4, the probability of success is given by

$$f(\vec{q}) = 1 - (1 - q_1 q_2)(1 - q_3)$$

Thus, if all the candidate networks are parallel-path networks, the social-value function is multilinear.¹⁵ We also note that for every edge i , the partial derivative in q_i of the success probability written in Equation 5 is positive where in all the other networks, that do not contain link i , the partial derivative is clearly zero. Therefore, the social-value function is single crossing and our general results apply.

Corollary 5. *For all social-choice functions that maximize the success probability over parallel-path networks, the informationally-optimal k -action social-choice function is implementable in dominant strategies (for every k).*

References

- Athey, S., 2001, "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica*, 69, 861–89.
- Ausubel, L. M., and P. Milgrom, 2006, *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 1. The Lovely but Lonely Vickrey Auction*, MIT Press.

¹⁵The results obtained here hold for all directed networks with no cycles (also known as DAG – directed acyclic graphs).

- Bergemann, D., and M. Pesendorfer, 2001, "Information Structures in Optimal Auctions," Discussion Paper 2991, C.E.P.R. Discussion Papers.
- Blumrosen, L., and M. Feldman, 2005, "Implementation with a Bounded Action Space," Discussion paper no. 412, the center for the study of rationality., The Hebrew University.
- Blumrosen, L., N. Nisan, and I. Segal, 2006, "Auctions with Severely Bounded Communications," *Journal of Artificial Intelligence Research*, 28, 233–266.
- Dasgupta, P., P. Hammond, and E. Maskin, 1979, "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46, 185 – 216.
- Dow, J., 1991, "Search Decisions with Limited Memory," *Review of Economic Studies*, 58, 1–14.
- Edlin, A. S., and C. Shannon, 1998, "Strict Single Crossing and the strict Spence-Mirrlees condition: a comment on monotone comparative statics," *Econometrica*, 68, 1417–1425.
- Green, J., and J. Laffont, 1977, "Characterization of satisfactory mechanism for the revelation of preferences for public goods," *Econometrica*, pp. 427–438.
- Harstad, R. M., and M. H. Rothkopf, 1994, "On the role of discrete bid levels in oral auctions," *European Journal of Operations Research*, pp. 572–581.
- Hermalin, B. E., 2005, "Lecture notes in Economics," University of California, Berkeley.
- Iyengar, S. S., and M. R. Lepper, 2000, "When choice is demotivating : Can one desire too much of a good thing?," *Journal of personality and social psychology*, 79, 995–1006.
- Larson, K., and T. Sandholm, 2001, "Costly valuation computation in auctions," in *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pp. 169–182.
- McAfee, P., 2002, "Coarse Matching," *Econometrica*, 70, 2025–2034.
- Milgrom, P., and C. Shannon, 1994, "Monotone Comparative Statics," *Econometrica*, 62, 157–180.
- Mirrlees, J., 1971, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175–208.
- Mookherjee, D., and S. Reichelstein, 1992, "Dominant strategy implementation of Bayesian incentive compatible allocation rules," *Journal of Economic Theory*, 56, 378–399.
- Myerson, R. B., 1981, "Optimal auction design," *Mathematics of Operations Research*, 6, 58–73.

- Rothkopf, M. H., T. J. Teisberg, and E. P. Kahn, 1990, “Why Are Vickrey Auctions Rare?,” *Journal of Political Economy*, 98, 94–109.
- Rothschild, M., and J. E. Stiglitz, 1976, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *The Quarterly Journal of Economics*, 90, 630–49.
- Segal, I., 2003, “Optimal Pricing Mechanisms with Unknown Demand,” *American Economic Review*, 93, 509–529.
- Segal, I., and R. Fadel, 2006, “The Communication Cost of Selfishness,” .
- Spence, M., 1973, “Job Market Signaling,” *The Quarterly Journal of Economics*, 87, 355 – 374.
- Wilson, R., 1989, “Efficient and Competitive Rationing,” *Econometrica*, 57, 1–40.

A Missing Proofs from Section 3

Proof of Lemma 1:

Proof. We first observe that for every social-value function there exists an informationally optimal k -action mechanism with a *deterministic* allocation scheme. This observation is general and does not require the use of threshold strategies or single-crossing conditions. Consider an optimal k -action mechanism that achieves the optimal result with some set of strategies $s = s_1, \dots, s_n$. At least the same expected social value will clearly be achieved by the following deterministic allocation scheme: for each profile of actions b , the mechanism chooses an alternative that maximizes the expected social value, i.e., $t(b) \in \operatorname{argmax}_{A'} E_{\vec{\theta}} \left[g(\vec{\theta}, A') \mid \forall i s_i(\theta_i) = b_i \right]$. Of course, this procedure may ruin incentive-compatibility properties of the mechanisms, but we will handle the incentive considerations separately.

With this observation in hand, we now turn to prove the two directions of the lemma. By Proposition 1, it is sufficient to show that the optimum is achieved with threshold strategies if and only if the optimal k -action mechanism is monotone. \Leftarrow :

Denote the thresholds used by player i by $x_0^i, x_1^i, \dots, x_k^i$. Namely, when player i reports an action b_i and uses a threshold strategy, her type lies between $[x_{b_i}^i, x_{b_i+1}^i]$. Consider a deterministic choice rule as described above, and consider an action profile $b = (b_1, \dots, b_n)$. Let A and B be two alternatives such that $A \succeq_i B$ (as determined by the single-crossing property). Now consider another action vector $b' = (b'_i, b_{-i})$, where $b'_i > b_i$. An optimal mechanism chooses for each profile of bids the alternative that maximizes the expected social value. Let A be the alternative chosen

by the mechanism under action profile b . For For proving monotonicity, it suffices to show that if A gains a higher pected social value than B for the action profile b , this will also hold for the action profile b' . That is, if

$$E_{\vec{\theta}} \left[g(\vec{\theta}, A) \mid s(\vec{\theta}) = b \right] \geq E_{\vec{\theta}} \left[g(\vec{\theta}, B) \mid s(\vec{\theta}) = b \right]$$

then

$$E_{\vec{\theta}} \left[g(\vec{\theta}, A) \mid s(\vec{\theta}) = b' \right] \geq E_{\vec{\theta}} \left[g(\vec{\theta}, B) \mid s(\vec{\theta}) = b' \right]$$

This will be an immediate conclusion from the following intuitive statement: fixing θ_{-i} , the expected difference in social value between alternatives A and B is greater for b' than for b .¹⁶ Formally,

$$E_{\theta_i} \left[g(\vec{\theta}, A) - g(\vec{\theta}, B) \mid s_i(\theta_i) = b'_i \right] \tag{6}$$

$$= \frac{1}{F_i(x_{b'_i+1}^i) - F_i(x_{b'_i}^i)} \int_{x_{b'_i}^i}^{x_{b'_i+1}^i} \left(g(\vec{\theta}, A) - g(\vec{\theta}, B) \right) f_i(\theta_i) d\theta_i \tag{7}$$

$$\geq \frac{1}{F_i(x_{b'_i+1}^i) - F_i(x_{b'_i}^i)} \int_{x_{b'_i}^i}^{x_{b'_i+1}^i} \left(g(x_{b'_i+1}^i, \theta_{-i}, A) - g(x_{b'_i+1}^i, \theta_{-i}, B) \right) f_i(\theta_i) d\theta_i \tag{8}$$

$$= \frac{1}{F_i(x_{b_i+1}^i) - F_i(x_{b_i}^i)} \int_{x_{b_i}^i}^{x_{b_i+1}^i} \left(g(x_{b_i+1}^i, \theta_{-i}, A) - g(x_{b_i+1}^i, \theta_{-i}, B) \right) f_i(\theta_i) d\theta_i \tag{9}$$

$$\geq \frac{1}{F_i(x_{b_i+1}^i) - F_i(x_{b_i}^i)} \int_{x_{b_i}^i}^{x_{b_i+1}^i} \left(g(\vec{\theta}, A) - g(\vec{\theta}, B) \right) f_i(\theta_i) d\theta_i \tag{10}$$

$$= E_{\theta_i} \left[g(\vec{\theta}, A) - g(\vec{\theta}, B) \mid s_i(\theta_i) = b_i \right] \tag{11}$$

Where inequalities 8 and 10 are due to the single-crossing property of the social-value function, and Equations 8 and 9 are equal since they are the expected value of the same constant value.

\implies : We now assume that a mechanism possesses a monotone allocation scheme, and prove that the optimum is achieved with threshold strategies.

The basic idea: we consider the expected social value of some player as a function of her type θ_i

¹⁶Note that due to the linearity of expectation,

$$\begin{aligned} & E_{\vec{\theta}} \left[g(\vec{\theta}, A) \mid s(\vec{\theta}) = b \right] - E_{\vec{\theta}} \left[g(\vec{\theta}, B) \mid s(\vec{\theta}) = b \right] \\ &= E_{\theta_{-i}} \left[E_{\theta_i} \left[g(\vec{\theta}, A) - g(\vec{\theta}, B) \mid s_i(\theta_i) = b_i \right] \mid s_{-i}(\theta_{-i}) = b_{-i} \right] \end{aligned}$$

when she chooses a particular action. We show that such functions, for every two actions $b_i < b'_i$, cross at most once; that is, if for some θ_i^* the expected social value is equal when player i chooses either b_i or b'_i , then for any $\theta_i > \theta_i^*$ the expected social value when choosing b_i is at most the expected social value when choosing b'_i . The optimality of threshold strategies for this mechanism will be derived directly from this weak single-crossing property.

Consider two actions $b'_i > b_i$ for player i . Let θ_i^* be a type for which the expected social value is equal either when player chooses b'_i or b_i , that is (we denote the actions of the players except i when their types are θ_{-i} by $s_{-i}(\theta_{-i})$):

$$E_{\theta_{-i}} [g(\theta_i^*, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i})))] = E_{\theta_{-i}} [g(\theta_i^*, \theta_{-i}, t(b'_i, s_{-i}(\theta_{-i})))] \quad (12)$$

We will show that for every $\theta_i > \theta_i^*$, the expected social value when player i chooses b_i is at most the expected social value in b'_i .

Monotonicity implies that $t(b'_i, b_{-i}) \succeq_i t(b_i, b_{-i})$ for every b_{-i} . Consider some profile of actions of the other players b_{-i} , and denote $t(b_i, b_{-i}) = A$ and $t(b'_i, b_{-i}) = B$. Since the social value function is single crossing, the change in the expected social value when alternative B is chosen must be greater, that is:

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

Now, considering all possible b_{-i} and using the linearity of expectation, we get:

$$E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, t(b'_i, s_{-i}(\theta_{-i})))] - E_{\theta_{-i}} [g(\theta_i^*, \theta_{-i}, t(b'_i, s_{-i}(\theta_{-i})))] \quad (13)$$

$$\geq E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i})))] - E_{\theta_{-i}} [g(\theta_i^*, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i})))] \quad (14)$$

Due to Equation 12 and Inequality 13-14, indeed for any $\theta_i > \theta_i^*$ the expected social value in b_i is at most the social welfare in b'_i :

$$E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, t(b'_i, s_{-i}(\theta_{-i})))] \geq E_{\theta_{-i}} [g(\theta_i, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i})))]$$

Finally, we conclude that the optimal social value can be achieved with threshold strategies for k -action games; each player should choose, for every type θ_i , the action that maximizes the expected social value. The maximum over k pairwise single-crossing functions have at most $k - 1$ switching points between the functions, therefore the social value is maximized using a threshold strategy that always chooses the action with the highest social value. (A similar argument is given for

maximum of linear functions, that are also single-crossing in this since, in Theorem 1.) □

B Missing Proofs from Section 5

Proof of Theorem 3:

Proof. Since the social-value function is multilinear and single crossing, the optimal expected social value is achieved by threshold strategies and therefore in a monotone mechanism (Lemma 1 and Theorem 1). To show that the mechanism is diagonal, we should also show that the allocation scheme is non-degenerate with respect to one of the players.

We prove the theorem for the case where the preferences \succ_i of the player are conflicting, and the proof for correlated preferences is similar. We assume, w.l.o.g., that $A \succ_1 B$ and $B \succ_2 A$ and that $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$. For such preferences, we show that the optimal mechanism will be non-degenerate with respect to Player 2. In other words, in the matrix representation of the optimal mechanism there will be no identical columns. Showing this will suffice, as it is easy to see that in a monotone allocation scheme where the column player has k distinct columns, the row player clearly has either $k - 1$ or k players.¹⁷

If Player 2 has two identical columns, then monotonicity derives that these columns will be adjacent, so in an equivalent allocation scheme this player will actually have $k - 1$ possible actions. 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. We therefore assume that the optimal k -action social value is 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} .

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:

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

¹⁷ Assuming that $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$, we can show that the optimal allocation scheme is non-degenerate with respect to Player 2. If the converse is true, we can show in the same way that the optimal allocation scheme is non-degenerate w.r.t. Player 1. Similar arguments also prove that when $g(\bar{\theta}_1, \bar{\theta}_2, A) \geq g(\bar{\theta}_1, \bar{\theta}_2, B)$ the optimal allocation will be non-degenerate w.r.t. Player 2 (otherwise, w.r.t. Player 1). Therefore, a sufficient condition for having an optimal allocation scheme that is non-degenerate w.r.t. both players is having both $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$ and $g(\bar{\theta}_1, \bar{\theta}_2, A) \leq g(\bar{\theta}_1, \bar{\theta}_2, B)$, or when both inequalities are in the opposite direction.

We know that $diff(y_0) > 0$ (since we assumed that $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$ 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 $y^* \in (y_0, y_1)$ for which $diff(y^*) = 0$ ($diff(\cdot)$ is clearly continuous since each both $g(\theta_1, \theta_2, A)$ and $g(\theta_1, \theta_2, B)$ are continuous w.r.t. θ_2). Setting y' to be, for example, $\frac{y_0 + y^*}{2}$ 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 $[B, B, \dots, A, A]$ and only $k - 1$ columns in the allocation matrix, it must be the case that some “internal” column is missing, hence, 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, between actions j and j' in the order on the actions, that is identical to the allocation in column j except $t(i, j') = B$, will strictly increase the expected social value. For the exact construction, we have to consider two different sub-cases: 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} :

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

Due to the (strict) single-crossing condition, clearly

$$E [g(\theta_1, y_{j+1}, A) \mid \theta_1 \in [x_{i+1}, x_{i+2}]] > E [g(\theta_1, y_{j+1}, B) \mid \theta_1 \in [x_{i+1}, x_{i+2}]]$$

Therefore, due to similar intermediate-value considerations, there must be some threshold $y^* > y_{j+1}$ for which

$$E [g(\theta_1, y_{j+1}, A) \mid \theta_1 \in [x_{i+1}, x_{i+2}]] = E [g(\theta_1, y_{j+1}, B) \mid \theta_1 \in [x_{i+1}, x_{i+2}]]$$

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 increases when $\theta_1 \in [x_i, x_{i+1}], \theta_2 \in [y_{j+1}, y']$, while the allocation in all other cases remains unchanged.

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

Let y^* be again the value for which

$$E(g(\theta_1, y^*, A) \mid \theta_1 \in [x_i, x_{i+1}]) = E(g(\theta_1, y^*, B) \mid \theta_1 \in [x_i, x_{i+1}])$$

Clearly, now $y^* < y_{j+1}$. Similar arguments show that adding a new threshold $y' = \frac{y^* + y_{j+1}}{2}$ yields a higher expected social surplus.

Given that the mechanism is diagonal, it is clear that each threshold of a player affects the decision that is made only for one action of the other player. Therefore, it is easy to see that each threshold must be a maximizer, based on the arguments given in Section 5.2. \square