Strong Implementation of Social Choice Functions in Dominant Strategies

Sven O. Krumke and Clemens Thielen

Abstract

We consider the classical mechanism design problem of strongly implementing social choice functions in a setting where monetary transfers are allowed. In contrast to *weak* implementation, where only *one* equilibrium of a mechanism needs to yield the desired outcomes given by the social choice function, *strong* implementation (also known as *full* implementation) means that a mechanism is sought in which *all* equilibria yield the desired outcomes. For strong implementation, one cannot restrict attention to incentive compatible direct revelation mechanisms via the Revelation Principle, so the question whether a given social choice function is strongly implementable cannot be answered as easily as for weak implementation.

When considering Bayes Nash equilibria, the *Augmented Revelation Principle* states that it suffices to consider mechanisms in which the set of types of each agent is a subset of the set of her possible bids. Moreover, given some additional data, such a mechanism can be constructed by an iterative procedure via *selective elimination* of undesired equilibria in finitely (but possible exponentially) many steps. For dominant strategies as the equilibrium concept, however, no such results have been known so far. We close this gap by showing a variant of the Augmented Revelation Principle for dominant strategies and a selective elimination procedure for constructing the desired mechanisms in *polynomially many* steps. Using these results, we then show that strong implementability in dominant strategies can be decided in nondeterministic polynomial time. This complements the results obtained in the companion paper by Thielen and Westphal [7], where an efficient polynomial time algorithm for the problem is given when one restricts to strong implementation by incentive compatible direct revelation mechanisms.

1 Introduction

Mechanism design is a classical area of noncooperative game theory and microeconomics, which studies how privately known preferences of several people can be aggregated towards a social choice. Applications include the design of procedures for elections and for deciding upon public projects. Recently, the study of the Internet has fostered the interest in algorithmic aspects of mechanism design [5].

In the classical social choice setting considered in this paper, there are n selfish agents, which must make a collective decision from some finite set X of possible social choices. Each agent i has a private value $\theta_i \in \Theta_i$ (called the agent's *type*), which influences the preferences of all agents over the alternatives in X. Formally, this is modeled by a *valuation function* $V_i : X \times \Theta \to \mathbb{Q}$ for each agent i, where $\Theta = \Theta_1 \times \cdots \times \Theta_n$. Every agent i reports some information s_i from a set S_i of possible *bids* of i to the mechanism designer who must then choose an alternative from X based on these bids. The goal of the mechanism designer is to implement a given social choice function $f : \Theta \to X$, that is, to make sure that the alternative $f(\theta)$ is always chosen in equilibrium when the vector of true types is $\theta = (\theta_1, \dots, \theta_n)$. To achieve this, the mechanism designer hands out a payment $P_i(\theta)$ to each agent i, which depends on the bids. Each agent then tries to maximize the sum of her valuation and payment by choosing an appropriate bid depending on her type. A *mechanism* $\Gamma = (S_1, \dots, S_n, g, P)$ is defined by the sets $S_1, \dots, S_n$ of possible bids of the agents, an *outcome function* $g : S_1 \times \cdots \times S_n \to X$, and the payment scheme $P = (P_1, \ldots, P_n)$.

In the most common concept called *weak* implementation, the mechanism Γ is said to *implement* the social choice function f if *some* equilibrium of the noncooperative game defined by the mechanism yields the outcomes specified by f. An important result known as the *Revelation Principle* (cf. [2, p. 884]) states that a social choice function is weakly implementable if and only if it can be *truthfully implemented* by an incentive compatible *direct revelation mechanism*, which means that f can be implemented by a mechanism with $S_i = \Theta_i$ for all i and truthful reporting as an equilibrium that yields the outcome specified by f. As a result, the question whether there exists a mechanism that weakly implements a given social choice function f can be easily answered in time polynomial in $|\Theta|$ by checking for negative cycles in complete directed graphs on the agents' type spaces with changes of valuations as edge weights (cf. [1, 4, 6]).

The more robust concept of implementation called *strong implementation* (also known as *full implementation*) requires that not only one, but *all* equilibria of a mechanism yield the desired outcomes. Hence, a strong implementation does not rely on the implicit assumption that the agents always play the "desired" equilibrium if there is more than one. For strong implementation, the Revelation Principle does not hold, so one cannot, in general, restrict attention to direct revelation mechanisms and truthful implementations when trying to decide whether a social choice function is strongly implementable.

When considering Bayes Nash equilibria as the equilibrium concept, a generalization of the Revelation Principle called the *Augmented Revelation Principle* [3] states that it suffices to consider *augmented revelation mechanisms*, in which the set Θ_i of types of each agent i is a subset of the set S_i of her possible bids. Moreover, it was shown in [3] that one can always obtain an augmented revelation mechanism that strongly implements a strongly implementable social choice function f via the *selective elimination procedure* that starts with an incentive compatible direct revelation mechanism and some additional data on its equilibria and iteratively eliminates all the finitely many equilibria that do not yield the outcomes specified by f. To do so, one of the agents is given a new bid, so her set of possible bids is enlarged by one element. Since the procedure always stops after finitely many iterations, this also implies that the sets S_i can always be chosen to be finite.

For dominant strategies as the equilibrium concept, however, no such results have been known so far and it has not even been clear that one can restrict to finite sets of possible bids or polynomially sized payments. Hence, also the complexity of deciding whether a given social choice function f is strongly implementable in dominant strategies has remained open.

2 Our Contribution

We prove a variant of the Augmented Revelation Principle for dominant strategy equilibria. Our result implies that, as in the case of Bayes Nash equilibria, one can always restrict to augmented revelation mechanisms when trying to decide strong implementability of social choice functions in dominant strategies. Moreover, we present a selective elimination procedure for constructing augmented revelation mechanisms in finitely many steps when dominant strategies are considered. In contrast to the case of Bayes Nash equilibria, where the number of steps needed for selective elimination of all undesired equilibria of an incentive compatible direct revelation mechanism can be exponential, we show that our procedure for dominant strategies always terminates after *polynomially many* steps, which implies that only a polynomial number of possible bids for each agent is needed. Based on this result, we show that the payments in a strong implementation can always be chosen to be of polynomial encoding length and present a method for deciding strong implementability of a given social choice function in nondeterministic polynomial time. Doing so, we prove the first upper bound on the computational complexity of this classical mechanism design problem. We suspect that a matching lower bound can be proved as well, i.e., that deciding strong implementability of a social choice function in dominant strategies is NP-complete.

3 Problem Definition

We are given n *agents* identified with the set $N = \{1, \ldots, n\}$ and a finite set X of possible *social choices*. For each agent i, there is a finite set Θ_i of possible *types* and we write $\Theta = \Theta_1 \times \cdots \times \Theta_n$. The true type θ_i of agent i is known only to the agent herself. Each agent i has a *valuation function* $V_i : X \times \Theta \to \mathbb{Q}$, where $V_i(x, \theta)$ specifies the value that agent i assigns to alternative $x \in X$ when the types of the agents are $\theta \in \Theta$. A *social choice function* in this setting is a function $f : \Theta \to X$ that assigns an alternative $f(\theta) \in X$ to every vector θ of types.

Definition 1. *A* mechanism $\Gamma = (S_1, \ldots, S_n, g, P)$ *consists of a set* S_i *of possible* bids *for each agent* i*, an* outcome function $g : S \to X$ *and a* payment scheme $P : S \to \mathbb{Q}^n$*, where* $S := S_1 \times \cdots \times S_n$.

A strategy *for agent* i *in the mechanism* Γ *is a function* $\alpha_i : \Theta_i \to S_i$ *that defines a bid* $\alpha_i(\theta_i) \in S_i$ *for every possible type* θ_i *of agent* i*. A* strategy profile *(in the mechanism* Γ*) is an* n*-tuple* $\alpha = (\alpha_1, \ldots \alpha_n)$ *containing a strategy* α_i *for each agent* i.

Definition 2. *Given a mechanism* $\Gamma = (S_1, \ldots, S_n, g, P)$*, a vector* $\theta \in \Theta$ *of types of all agents, and a vector* $s_{-i} \in S_{-i}$ *of bids of all agents except* i*, the* utility *from a bid* $s_i \in S_i$ *for agent* i *is defined as*

$$U_i^\Gamma(s_{-i}, s_i | \theta) := V_i(g(s_{-i}, s_i), \theta) + P_i(s_{-i}, s_i).$$

A bid $\bar{s}_i \in S_i$ *of an agent* i *is called a* dominant bid *for type* $\theta_i \in \Theta_i$ *if it maximizes the utility of an agent* i *of type* θ_i *for every possible vector* $s_{-i} \in S_{-i}$ *of bids of the other agents and every possible vector* $\theta_{-i} \in \Theta_{-i}$ *of types of the other agents, i.e., if*

$$U_i^\Gamma(s_{-i}, \bar{s}_i | \theta) \geq U_i^\Gamma(s_{-i}, s_i | \theta) \quad \forall s_{-i} \in S_{-i}, \theta_{-i} \in \Theta_{-i}, s_i \in S_i.$$

A pair $(\theta, s) \in \Theta \times S$ *of a type vector* $\theta \in \Theta$ *and bid vector* $s \in S$ *is called a* dominant pair *if* s_i *is a dominant bid for* θ_i *for every* $i \in N$*. The strategy profile* α *is a* dominant strategy equilibrium *of* Γ *if* $(\theta, \alpha(\theta))$ *is a dominant pair for every* $\theta \in \Theta$.

Definition 3. *The mechanism* $\Gamma = (S_1, \ldots, S_n, g, P)$ strongly implements *the social choice function* f *if* Γ *has at least one equilibrium and* every *equilibrium* α *of* Γ *satisfies* $g \circ \alpha = f$*. The social choice function* f *is called* strongly implementable *if there exists a mechanism* Γ *that strongly implements* f.

Definition 4. *A mechanism* $\Gamma = (S_1, \ldots, S_n, g, P)$ *is called a* direct revelation mechanism *if* $S_i = \Theta_i$ *for all* $i \in N$*. The direct revelation mechanism* $(\Theta_1, \ldots, \Theta_n, f, P)$ *defined by a social choice function* f *and a payment scheme* P *will be denoted by* $\Gamma_{(f,P)}$*. A direct revelation mechanism* $\Gamma_{(f,P)}$ *is called* incentive compatible *if truthful reporting is a dominant strategy equilibrium of* $\Gamma_{(f,P)}$.

Definition 5. *A mechanism* $\Gamma = (S_1, \ldots, S_n, g, P)$ *is called* augmented revelation mechanism *if* $S_i = \Theta_i \cup T_i$ *for all* $i \in N$ *and arbitrary sets* T_i.

Definition 6 (Strong Implementability Problem)**.**

INSTANCE:	*The number* n *of agents, the set* X *of possible social choices, the sets* Θ_i *of possible types of the agents, the valuation functions* $V_i : X \times \Theta \to \mathbb{Q}$*, and the social choice function* $f : \Theta \to X$.
QUESTION:	*Is* f *strongly implementable in dominant strategies?*

To encode an instance of Strong Implementability, we need to do the following: For every valuation function $V_i : X \times \Theta \to \mathbb{Q}$, we need to store $|X| \cdot |\Theta|$ rational numbers. The social choice function $f : \Theta \to X$ has encoding length $|\Theta| \cdot \log(|X|)$. Thus, the encoding length of an instance of Strong Implementability is in $\Omega(|X| \cdot |\Theta| \cdot n)$.

4 The Augmented Revelation Principle for Dominant Strategies

In this section, we prove the Augmented Revelation Principle for dominant strategies and present our selective elimination procedure that, given an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ and some data on it equilibria, constructs an augmented revelation mechanism that strongly implements f by an iterative procedure that stops after polynomially many steps.

Theorem 1 (Augmented Revelation Principle for dominant strategies)**.** *If a social choice function $f : \Theta \to X$ is strongly implementable in dominant strategies, then f can be strongly implemented in dominant strategies by an augmented revelation mechanism in which truthful reporting is an equilibrium.*

Proof. Given a mechanism $\Gamma = (S_1, \dots, S_n, g, P)$ that strongly implements f in dominant strategies, we construct an augmented revelation mechanism $\bar{\Gamma} = (\bar{S}_1, \dots, \bar{S}_n, \bar{g}, \bar{P})$ that strongly implements f similar to the proof of the Augmented Revelation Principle for Bayes Nash equilibria given in [3]. Additionally, we have to define the new payment scheme $\bar{P}$ in terms of the given payment scheme P since the proof in [3] focused on the case without payments.

Given an arbitrary equilibrium $\alpha = (\alpha_1, \dots, \alpha_n)$ of Γ, we define $\bar{S}_i := \Theta_i \cup T_i$, where

$$T_i := \{s_i \in S_i \mid s_i \notin \operatorname{image}(\alpha_i)\},$$

and $\operatorname{image}(\alpha_i) = \{\alpha_i(\theta_i) \mid \theta_i \in \Theta_i\}$ denotes the image of the function $\alpha_i : \Theta_i \to S_i$. We consider the functions $\phi_i : \bar{S}_i \to S_i$ given by

$$\phi_i(\bar{s}_i) := \begin{cases} \alpha_i(\theta_i) & \text{if } \bar{s}_i = \theta_i \text{ for } \theta_i \in \Theta_i \\ \bar{s}_i & \text{if } \bar{s}_i \in T_i \end{cases}$$

and define the outcome function $\bar{g} : \bar{S} \to X$ as $\bar{g} := g \circ \phi$, where $\phi = (\phi_1, \dots, \phi_n)$. The payment scheme $\bar{P} : \bar{S} \to \mathbb{Q}$ is defined analogously as $\bar{P} := P \circ \phi$.

To show that $\bar{\Gamma}$ strongly implements f in dominant strategies, suppose that $\bar{\alpha} = (\bar{\alpha}_1, \dots, \bar{\alpha}_n)$ is an equilibrium of $\bar{\Gamma}$ and again consider the strategy profile $\alpha^* = (\alpha_1^*, \dots, \alpha_n^*)$ in Γ given by $\alpha_i^* := \phi_i \circ \bar{\alpha}_i$. As before, we then have $g \circ \alpha^* = g \circ \phi \circ \bar{\alpha} = \bar{g} \circ \bar{\alpha}$ and $P \circ \alpha^* = P \circ \phi \circ \bar{\alpha} = \bar{P} \circ \bar{\alpha}$ and claim that α^* is an equilibrium of Γ.

Since every $\phi_j : \bar{S}_j \to S_j$ is surjective, we can choose $\bar{s}_j \in \bar{S}_j$ with $\phi_j(\bar{s}_j) = s_j$ for each $j \in N$ and each $s_j \in S_j$. Then, for all $i \in N, \theta \in \Theta, s_{-i} \in S_{-i}$, and $s_i \in S_i$,

$$\begin{aligned} U_i^{\Gamma}(s_{-i}, \alpha_i^*(\theta_i)|\theta) &= V_i(g(s_{-i}, \alpha_i^*(\theta_i)), \theta) + P_i(s_{-i}, \alpha_i^*(\theta_i)) \\ &= V_i(\bar{g}(\bar{s}_{-i}, \bar{\alpha}_i(\theta_i)), \theta) + \bar{P}_i(\bar{s}_{-i}, \bar{\alpha}_i(\theta_i)) \\ &\geq V_i(\bar{g}(\bar{s}_{-i}, \bar{s}_i), \theta) + \bar{P}_i(\bar{s}_{-i}, \bar{s}_i) \\ &= V_i(g(s_{-i}, s_i), \theta) + P_i(s_{-i}, s_i) \\ &= U_i^{\Gamma}(s_{-i}, s_i|\theta), \end{aligned}$$

where the inequality follows since $\bar{\alpha}$ is an equilibrium of $\bar{\Gamma}$. Thus, α^* is an equilibrium of Γ as claimed. So since Γ strongly implements f, it follows that $f = g \circ \alpha^* = \bar{g} \circ \bar{\alpha}$, i.e., the equilibrium $\bar{\alpha}$ yields the outcomes specified by f. Hence, it just remains to show that truthful bidding is an equilibrium of $\bar{\Gamma}$. But this follows easily since, for every $\theta \in \Theta$, we have $\bar{g}(\theta) = (g \circ \phi)(\theta) = g(\alpha(\theta))$ and $\bar{P}(\theta) = (P \circ \phi)(\theta) = P(\alpha(\theta))$ and α is an equilibrium of Γ. □

We now present our selective elimination procedure for dominant strategies. To this end, we need the following definition:

Definition 7. *A dominant bid $\bar{\theta}_i \in \Theta_i$ for type $\tilde{\theta}_i \in \Theta_i$ of agent $i \in N$ in a direct revelation mechanism $\Gamma_{(f,P)}$ can be* selectively eliminated *if there exists a nonempty subset $\bar{N} \subseteq N \setminus \{i\}$ of the other agents such that the following holds: For $\bar{S}_j := \Theta_j \cup \{\bar{s}_j\}$ for $j \in \bar{N}$, $\bar{S}_j := \Theta_j$ for $j \in N \setminus \bar{N}$, and $\bar{S} := \bar{S}_1 \times \cdots \times \bar{S}_n$, there exist functions $h : \bar{S} \to X$ and $\bar{P}_j : \bar{S} \to \mathbb{Q}, j \in N$, with $h_{|\Theta} = f$ and $(\bar{P}_j)_{|\Theta} = P_j$ such that:*

1. *For some $\tilde{\theta}_{-i} \in \Theta_{-i}$ and some bid vector $\bar{\theta}_{-(\bar{N}\cup\{i\})} \in \Theta_{-(\bar{N}\cup\{i\})}$ of the agents not in $\bar{N} \cup \{i\}$*

$$
\begin{aligned}
& V_i(h(\bar{s}_{\bar{N}}, \bar{\theta}_{-(\bar{N}\cup\{i\})}, \tilde{\theta}_i), \tilde{\theta}) + \bar{P}_i(\bar{s}_{\bar{N}}, \bar{\theta}_{-(\bar{N}\cup\{i\})}, \tilde{\theta}_i) \\
> \quad & V_i(h(\bar{s}_{\bar{N}}, \bar{\theta}_{-(\bar{N}\cup\{i\})}, \bar{\theta}_i), \tilde{\theta}) + \bar{P}_i(\bar{s}_{\bar{N}}, \bar{\theta}_{-(\bar{N}\cup\{i\})}, \bar{\theta}_i).
\end{aligned}
$$

2. *For all $j \in N, \theta \in \Theta, s \in \bar{S} \setminus \Theta$*

$$
V_j(h(s_{-j}, \theta_j), \theta) + \bar{P}_j(s_{-j}, \theta_j) \geq V_j(h(s_{-j}, s_j), \theta) + \bar{P}_j(s_{-j}, s_j).
$$

A dominant pair $(\theta, \theta') \in \Theta^2$ can be selectively eliminated if the dominant bid $\theta'_i \in \Theta_i$ for type $\theta_i \in \Theta_i$ can be selectively eliminated for some $i \in N$.

Here, each agent $j \in \bar{N}$ is given a new bid $\bar{s}_j$. The function h extends f to the enlarged set $\bar{S}$ of possible bids by specifying the outcomes chosen when at least one agent chooses a non-type message. Similarly, the functions $\bar{P}_j$ extend the payment functions P_j to $\bar{S}$. The first condition says that, for some type vector $\tilde{\theta}_{-i} \in \Theta_{-i}$ of the other agents and some bid vector $\bar{\theta}_{-(\bar{N}\cup\{i\})}$ of the agents not in $\bar{N} \cup \{i\}$, agent i can increase her utility by bidding her true type $\tilde{\theta}_i$ instead of $\bar{\theta}_i$ in the case that the agents in $\bar{N}$ choose their new non-type messages. Thus, $\bar{\theta}_i$ is not a dominant bid for type $\tilde{\theta}_i$ anymore. On the other hand, the second condition ensures that all pairs $(\theta, \theta) \in \Theta^2$ stay dominant pairs, so truthful reporting is preserved as an equilibrium.

Definition 8. *A dominant pair $(\theta, \theta') \in \Theta^2$ in the direct revelation mechanism $\Gamma_{(f,P)}$ is called* bad *if $f(\theta) \neq f(\theta')$. $\Gamma_{(f,P)}$ satisfies the* selective elimination condition *if every bad dominant pair can be selectively eliminated.*

The idea behind Definition 8 is the following observation, which follows immediately from the definitions:

Observation 1. *A direct revelation mechanism $\Gamma_{(f,P)}$ with at least one dominant strategy equilibrium has a* bad *dominant strategy equilibrium if and only if there exists a bad dominant pair $(\theta, \theta') \in \Theta^2$ in $\Gamma_{(f,P)}$. In particular, an* incentive compatible *direct revelation mechanism $\Gamma_{(f,P)}$ has a bad dominant strategy equilibrium if and only if there exists a bad dominant pair $(\theta, \theta') \in \Theta^2$ in $\Gamma_{(f,P)}$.*

Hence, selectively eliminating all bad dominant pairs will lead to elimination of all bad equilibria. The difference to the case of Bayes Nash equilibria discussed in [3] is that one does not need to consider complete equilibria α and check whether they can be selectively eliminated. Here, one has to consider only bad dominant pairs $(\theta, \theta') \in \Theta^2$. While there are potentially exponentially many bad equilibria, the number of bad dominant pairs $(\theta, \theta') \in \Theta^2$ is bounded by $|\Theta|^2$, which is *polynomial* in the encoding length of the input. This observation will play a crucial role when we show that Strong Implementability is in NP when considering dominant strategies.

Theorem 2. *Suppose that the social choice function $f : \Theta \to X$ is strongly implementable in dominant strategies. Then there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ that satisfies the selective elimination condition.*

Proof. Theorem 1 states that there exists an augmented revelation mechanism $\Gamma = (S_1, \ldots, S_n, g, P)$ that strongly implements f in dominant strategies and in which truthful reporting is an equilibrium. In particular, this implies that $g_{|\Theta} = f$, and we claim that the direct revelation mechanism $\Gamma_{(f,P_{|\Theta})}$ is as required.

Incentive compatibility of $\Gamma_{(f,P_{|\Theta})}$ follows directly from the fact that truthful reporting is an equilibrium in Γ. To show that $\Gamma_{(f,P_{|\Theta})}$ satisfies the selective elimination condition, consider a bad dominant pair $(\theta, \theta') \in \Theta^2$ in $\Gamma_{(f,P_{|\Theta})}$ (if none exists, we are done). Since Γ strongly implements f, it can have no bad equilibria and, thus, (θ, θ') cannot be a dominant pair in Γ. Hence, there must be an agent i, a vector $\tilde{\theta}_{-i} \in \Theta_{-i}$ of types of the other agents, and a vector $\bar{s} \in S$ of bids such that

$$U_i^\Gamma(\bar{s}_{-i}, \bar{s}_i | \theta_i, \tilde{\theta}_{-i}) > U_i^\Gamma(\bar{s}_{-i}, \theta_i' | \theta_i, \tilde{\theta}_{-i}).$$

Moreover, since truthful reporting is an equilibrium in Γ, we know that

$$U_i^\Gamma(\bar{s}_{-i}, \theta_i | \theta_i, \tilde{\theta}_{-i}) \geq U_i^\Gamma(\bar{s}_{-i}, \bar{s}_i | \theta_i, \tilde{\theta}_{-i}),$$

so we obtain

$$U_i^\Gamma(\bar{s}_{-i}, \theta_i | \theta_i, \tilde{\theta}_{-i}) > U_i^\Gamma(\bar{s}_{-i}, \theta_i' | \theta_i, \tilde{\theta}_{-i}). \tag{1}$$

Moreover, since (θ, θ') is a dominant pair in $\Gamma_{(f,P_{|\Theta})}$, we know that $\bar{s}_j \notin \Theta_j$ for at least one $j \neq i$. Hence, the set $\bar{N} := \{j \neq i : \bar{s}_j \notin \Theta_j\}$ is nonempty. We now set $\bar{S}_j := \Theta_j \cup \{\bar{s}_j\}$ for $j \in \bar{N}$, $\bar{S}_j := \Theta_j$ for $j \in N \setminus \bar{N}$, and $\bar{S} := \bar{S}_1 \times \cdots \times \bar{S}_n$ as in the definition of selective elimination. The function $h : \bar{S} \to X$ is defined as the restriction of g to $\bar{S} \subseteq S$ and it satisfies $h_{|\Theta} = g_{|\Theta} = f$. Analogously, the functions $\bar{P}_j : \bar{S} \to \mathbb{Q}$ are defined as the restrictions of the P_j to $\bar{S} \subseteq S$ and we have $(\bar{P}_j)_{|\Theta} = (P_j)_{|\Theta}$. Defining $\bar{\theta}_j := \bar{s}_j \in \Theta_j$ for all $j \notin (\bar{N} \cup \{i\})$, it is now immediate that the dominant bid θ_i' for type θ_i of agent i can be selectively eliminated, i.e., that Conditions 1 and 2 in the definition of selective elimination are satisfied: Condition 1 follows directly from (1) and the definitions, and Condition 2 follows since truthful reporting is an equilibrium in Γ. Thus, the bad dominant pair $(\theta, \theta') \in \Theta^2$ can be selectively eliminated. □

We are now ready to present our selective elimination procedure for constructing augmented revelation mechanisms. This procedure is used to prove the following theorem, which is states that the selective elimination condition is also sufficient for strong implementability:

Theorem 3. *Suppose that there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ that satisfies the selective elimination condition. Then the social choice function $f : \Theta \to X$ is strongly implementable in dominant strategies.*

Proof. We start with the direct revelation mechanism $\Gamma_{(f,P)}$ and proceed inductively to selectively eliminate all bad dominant pairs $(\theta, \theta') \in \Theta^2$ in $\Gamma_{(f,P)}$ one by one without introducing any new dominant pairs by augmenting the mechanism appropriately. Since there can only be finitely many bad dominant pairs, the procedure stops after a finite number of steps with an augmented revelation mechanism without bad dominant pairs and, thus, without bad equilibria. In fact, the procedure stops after a *polynomial* number of steps since there can only be $|\Theta|^2$ many bad dominant pairs in $\Gamma_{(f,P)}$.

We describe a representative stage of this iterative procedure. From the previous iteration, we are given an augmentation $\Gamma = (S_1, \ldots, S_n, g, P')$ of $\Gamma_{(f,P)}$ with $S_i = \Theta_i \cup T_i$ for all $i \in N$, and $\theta' \in \Theta$ for every dominant pair (θ, θ') of Γ. Let (θ, θ') be a bad dominant pair of Γ. Let $i \in N$ be such that the dominant bid θ_i' for type θ_i of agent i can be selectively eliminated, and suppose that $\emptyset \neq \bar{N} \subseteq N \setminus \{i\}$, $\bar{S} = \bar{S}_1 \times \cdots \times \bar{S}_n$ with $\bar{S}_j = \Theta_j \cup \{\bar{s}_j\}$ for $j \in \bar{N}$ and $\bar{S}_j = \Theta_j$ for $j \in N \setminus \bar{N}$, $h : \bar{S} \to X$, and $\bar{P}_j : \bar{S} \to \mathbb{Q}$ are as in the definition of selective elimination. Consider the mechanism $\tilde{\Gamma} = (\tilde{S}_1, \ldots, \tilde{S}_n, \tilde{g}, \tilde{P})$ with

$$\begin{aligned} \tilde{S}_j &:= S_j \cup \{\bar{s}_j\} \text{ for } j \in \bar{N}, \\ \tilde{S}_j &:= S_j \text{ for } j \in N \setminus (\bar{N} \cup \{i\}), \\ \tilde{S}_i &:= S_i \cup \{\text{CFL}\}. \end{aligned}$$

Hence, each agent $j \in \bar{N}$ is given a new bid $\bar{s}_j$ (a *flag*), and agent i is given a new *counterflag* CFL. We set $\tilde{g}_{|S} := g$ and $\tilde{P}_{|S} := P'$, i.e., outcomes and payments associated with bids from the previous stages are left unchanged. Outcomes and payments associated with the new bids are defined as follows:

1. If the bid vector is in $\bar{S}$, the outcome and the payments are given by h and the $\bar{P}_j$, respectively, i.e., $\tilde{g}(s) := h(s)$ for $s \in \bar{S}$ and $\tilde{P}_j(s) := \bar{P}_j(s)$ for $s \in \bar{S}, j \in N$. Note that this definition agrees with the outcomes and payments of the previous stages when the bid vector is in Θ since $\tilde{g}_{|\Theta} = f$ and $(\tilde{P}_j)_{|\Theta} = P$.

2. If some agents $\emptyset \neq \tilde{N} \subseteq \bar{N}$ choose their new bids, agent i does *not* choose her new counterflag CFL, but some agents $j \in \tilde{\tilde{N}} \subseteq N \setminus \tilde{N}$ choose bids in $T_j = S_j \setminus \Theta_j$, then outcome and payments are as if each agent $k \in \tilde{N}$ had chosen a fixed type $\theta_k^0 \in \Theta_k$, but each agent $k \in \tilde{N}$ is charged $\epsilon > 0$ for choosing her new bid $\bar{s}_k$ if agent i chooses a bid in $T_i = S_i \setminus \Theta_i$.

3. If *at least two* agents $\tilde{N} \subseteq \bar{N}, |\tilde{N}| \geq 2$, choose their new bids and agent i chooses CFL, then outcome and payments are as if agent i had reported a fixed type $\theta_i^0 \in \Theta_i$.

4. If *no* agent in $\bar{N}$ chooses her new bid, but agent i chooses CFL, then outcome and payments are as if agent i had reported $\theta_i^0 \in \Theta_i$, but agent i is charged $\epsilon > 0$ for choosing CFL.

5. If *exactly one* agent $k \in \bar{N}$ chooses her new bid and agent i chooses CFL, then outcome and payments are as if agent i had reported $\theta_i^0 \in \Theta_i$ and agent k had reported $\theta_k^0 \in \Theta_k$, but agent k is charged $\epsilon > 0$.

We now have to show that truthful reporting is still an equilibrium in $\tilde{\Gamma}$, (θ, θ') is not a dominant pair anymore, and there are no new dominant bid pairs in $\tilde{\Gamma}$ (or, equivalently, no new dominant bids). Note that, since the outcomes and payments associated with bids from the previous stages are left unchanged, any new dominant bid of an agent would have to be one of the agent's new bids.

Claim 1. *Truthful reporting is an equilibrium in* $\tilde{\Gamma}$.

Proof. We consider a fixed agent $j \in N$ and show that truthful reporting is a dominant strategy for j in $\tilde{\Gamma}$.

As long as the other agents bid a vector in $S_{-j} \cup \bar{S}_{-j}$, truthful reporting is always optimal for agent j among all bids in $S_j \cup \bar{S}_j$ by Condition 2 in the definition of selective elimination and since truthful reporting is an equilibrium in Γ. When agent j bids a new bid (if she has one), this can only lead to some agents being charged ϵ as long as the other agents still bid a vector in $S_{-j} \cup \bar{S}_{-j}$. Hence, truthful reporting is always optimal among *all* bids of agent j in this case.

If the vector of bids of the other agents is *not* in $S_{-j} \cup \bar{S}_{-j}$, some agents choose new bids and some agents choose previously added bids. Hence, Case 2 in the definition of $\tilde{\Gamma}$ applies when agent i does not choose CFL, and Case 3,4, or 5 applies when agent i chooses CFL. But outcomes and payments in each of these cases are equivalent to the outcome and payments resulting from some bid vector in $S \cup \bar{S}$, except that some agents are possibly charged ϵ. Hence, truthful reporting is optimal for agent j by the case considered above. □

Claim 2. *(θ, θ') is not a dominant pair in $\tilde{\Gamma}$.*

Proof. Follows immediately from Case 1 in the definition of $\tilde{\Gamma}$ and the definition of selective elimination. □

Claim 3. *There is no dominant bid $\bar{s}_k$ for any type $\theta_k \in \Theta_k$ of any agent $k \in \bar{N}$ in $\tilde{\Gamma}$.*

Proof. If agent k has type θ_k and chooses the bid $\bar{s}_k$, consider the situation in which no other agent $j \in \bar{N} \setminus \{k\}$ chooses her new bid and agent i chooses CFL. Then, by Cases 4 and 5 in the definition of $\tilde{\Gamma}$, agent k could increase her utility by $\epsilon > 0$ by bidding θ_k^0 instead of $\bar{s}_k$ for every possible vector θ_{-k} of types of the other agents. □

Claim 4. *CFL is not a dominant bid for any type $\theta_i \in \Theta_i$ of agent i in $\tilde{\Gamma}$.*

Proof. If agent i has type θ_i and chooses the bid CFL, consider the situation in which no agent $j \in \bar{N}$ chooses her new bid. Then, by Case 4 in the definition of $\tilde{\Gamma}$, agent i could increase her utility by $\epsilon > 0$ by bidding θ_i^0 instead of CFL for every possible vector θ_{-i} of types of the other agents. □

By inductive application of the claims, the final mechanism obtained after eliminating all bad dominant pairs (θ, θ') in $\Gamma_{(f,P)}$ has no bad dominant pairs, but truthful reporting is still an equilibrium. Hence, this mechanism strongly implements f in dominant strategies, which proves the theorem. □

Theorem 4. *The social choice function $f : \Theta \to X$ is strongly implementable in dominant strategies if and only if there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ that satisfies the selective elimination condition.* □

Theorem 4 is the main ingredient needed for the proof of our complexity result on Strong Implementability in the next section. The following lemma resolves one last formal problem resulting from the definition of selective elimination: Giving all agents in $\bar{N}$ a new bid could yield an exponentially large space $\bar{S}$ of possible bids in the definition of selective elimination. This can, however, only happen if some of the agents have only one possible type, and the behavior of such agents cannot impose restrictions on the implementability of a social choice function since the types of these agents are common knowledge.

Lemma 1. *Let $Z \subseteq N$ denote the set of agents whose type space consists of only one element, i.e., $|\Theta_j| = 1$ for every $j \in Z$. Consider the instance of Strong Implementability for the agents in $N \setminus Z$ given by the valuations $V_i^{-Z} : X \times \Theta_{-Z} \to \mathbb{Q}$ defined by $V_i^{-Z}(x, \theta_{-Z}) := V_i(x, \theta_{-Z}, \theta_Z)$ and the social choice function $f^{-Z} : \Theta_{-Z} \to X$ defined by $f^{-Z}(\theta_{-Z}) := f(\theta_{-Z}, \theta_Z)$, where θ_Z is the unique type vector of the agents in Z. Then f is* *strongly implementable in dominant strategies if and only if f^{-Z} is strongly implementable in dominant strategies.*

Proof. Suppose that f is strongly implementable in dominant strategies. Then, by Theorem 1, there exists an augmented revelation mechanism $\Gamma = (S_1, \dots, S_n, g, P)$ that strongly implements f in dominant strategies and in which truthful reporting is an equilibrium. Without loss of generality, we assume that $N \setminus Z = \{1, \dots, z\}$ with $z := |N \setminus Z|$ and consider the mechanism $\Gamma^{-Z} = (S_1^{-Z}, \dots, S_z^{-Z}, g^{-Z}, P^{-Z})$ defined by

$$\begin{aligned} S_i^{-Z} &:= S_i \text{ for } i = 1, \dots, z \\ g^{-Z}(s_1, \dots, s_z) &:= g(s_1, \dots, s_z, \theta_Z) \\ P_i^{-Z}(s_1, \dots, s_z) &:= P_i(s_1, \dots, s_z, \theta_Z). \end{aligned}$$

Then $\alpha^{-Z} = (\alpha_1^{-Z}, \dots \alpha_z^{-Z}) \longmapsto (\alpha_1^{-Z}, \dots, \alpha_z^{-Z}, id_{\Theta_Z}) = \alpha$ defines an injective map from the set of strategy profiles in Γ^{-Z} to the set of strategy profiles in Γ, and α^{-Z} is an equilibrium in Γ^{-Z} if and only if α is an equilibrium in Γ (here, we use that truthful bidding is a dominant strategy in Γ for each agent in Z). Moreover, we have $g^{-Z} \circ \alpha^{-Z} = g^{-Z}$ if and only if $g \circ \alpha = g$. Hence, since Γ strongly implements f, it follows that Γ^{-Z} strongly implements f^{-Z}.

Conversely, assume that f^{-Z} is strongly implementable and denote an augmented revelation mechanism that strongly implements it in dominant strategies and in which truthful reporting is an equilibrium by $\Gamma^{-Z} = (S_1^{-Z}, \dots, S_z^{-Z}, g^{-Z}, P^{-Z})$. We define a mechanism $\Gamma = (S_1, \dots, S_n, g, P)$ for all agents as follows:

$$\begin{aligned} S_i &:= S_i^{-Z} \text{ for } i = 1, \dots, z \\ S_i &:= \Theta_i = \{\theta_i\} \text{ for } i \in Z \\ g(s_1, \dots, s_z, \theta_Z) &:= g^{-Z}(s_1, \dots, s_z) \\ P_i(s_1, \dots, s_z, \theta_Z) &:= P_i^{-Z}(s_1, \dots, s_z) \text{ for } i = 1, \dots, z \\ P_i(s_1, \dots, s_z, \theta_Z) &:= 0 \text{ for } i \in Z \end{aligned}$$

Then $\alpha = (\alpha_1, \dots \alpha_z, id_{\Theta_Z}) \longmapsto (\alpha_1, \dots, \alpha_z) = \alpha^{-Z}$ defines a bijection between the set of strategy profiles in Γ and the set of strategy profiles in Γ^{-Z}, and α is an equilibrium in Γ if and only if α^{-Z} is an equilibrium in Γ^{-Z}. Again, we have $g^{-Z} \circ \alpha^{-Z} = g^{-Z}$ if and only if $g \circ \alpha = g$. Hence, since Γ^{-Z} strongly implements f^{-Z}, Γ strongly implements f. $\square$

Lemma 1 shows that, when trying to decide strong implementability of a social choice function f in dominant strategies, one can disregard all agents that have only one possible type by considering the equivalent problem of strong implementability of the social choice function f^{-Z}. Hence, we may from now on assume that $|\Theta_j| \geq 2$ for every agent $j \in N$. With this assumption, the cardinality of the set $\bar{S}$ in the definition of selective elimination is only quadratic in $|\Theta|$:

$$|\bar{S}| = \prod_{j \in \bar{N}} \underbrace{(|\Theta_j| + 1)}_{\leq 2|\Theta_j|} \cdot \prod_{j \in N \setminus \bar{N}} |\Theta_j| \leq \underbrace{2^{|N|-1}}_{\leq \prod_{j=1}^{n} |\Theta_j|} \prod_{j=1}^{n} |\Theta_j| \leq (\prod_{j=1}^{n} |\Theta_j|)^2 = |\Theta|^2.$$

5 Solving Strong Implementability in Nondeterministic Polynomial Time

In this section, we use our results on augmented revelation mechanisms and selective elimination to show that Strong Implementability can be decided in nondeterministic polynomial time when dominant strategies are considered.

Suppose we are given a yes-instance of Strong Implementability, i.e., an instance with a strongly implementable social choice function f. Theorem 4 then tells us that there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ satisfying the selective elimination condition. We denote the set of all bad dominant pairs (θ, θ') by $D \subseteq \Theta^2$. Similarly, for each $i \in N$, we denote the set of all pairs $(\tilde{\theta}_i, \bar{\theta}_i) \in \Theta_i^2$ such that $\bar{\theta}_i$ is a dominant bid for type $\tilde{\theta}_i$ of agent i by $D_i \subseteq \Theta_i^2$.

Since $\Gamma_{(f,P)}$ satisfies the selective elimination condition, we know that each bad dominant pair $(\theta, \theta') \in D$ can be selectively eliminated. Suppose that, for every $(\theta, \theta') \in D$, $\left(i^{(\theta,\theta')}, \bar{N}^{(\theta,\theta')}, h^{(\theta,\theta')}, \tilde{\theta}^{(\theta,\theta')}_{-i^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}\cup\{i^{(\theta,\theta')}\})}\right)$ is the data which, together with appropriate payment functions $\bar{P}_j^{(\theta,\theta')}$ for $j \in N$, can be used to selectively eliminate the bad dominant pair (θ, θ').

Similarly, suppose that, for every $i \in N$ and every pair $(\tilde{\theta}_i, \bar{\theta}_i) \in \Theta_i^2 \setminus D_i$ (i.e., for every pair $(\tilde{\theta}_i, \bar{\theta}_i)$ of types of agent i such that $\bar{\theta}_i$ is *not* a dominant bid for type $\tilde{\theta}_i$), $\left(\tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}\right)$ is a pair of a type vector and a bid vector of the other agents such that

$$U_i^{\Gamma_{(f,P)}}\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i | \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right) > U_i^{\Gamma_{(f,P)}}\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_i | \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right).$$

The possible payment functions $P_j : \Theta \to \mathbb{Q}$ of the mechanism $\Gamma_{(f,P)}$ and the functions $\bar{P}_j^{(\theta,\theta')} : \bar{S}^{(\theta,\theta')} \to \mathbb{Q}$ are then given by the solutions of the system of linear inequalities in the variables $P_j(\theta)$ for $j \in N, \theta \in \Theta$ and $\bar{P}_j^{(\theta,\theta')}(s)$ for $j \in N, (\theta, \theta') \in D$, $s \in \bar{S}^{(\theta,\theta')} \setminus \Theta$ displayed on Page 11. Note that the values $\bar{P}_j^{(\theta,\theta')}(s)$ for $s \in \Theta$ do not need to appear in the system since we require that $(\bar{P}_j^{(\theta,\theta')})_{|\Theta} = P_j$ for all $j \in N, (\theta, \theta') \in D$.

Inequalities (2) and (3) encode exactly which bids $\bar{\theta}_i \in \Theta_i$ are dominant bids for any type $\tilde{\theta}_i$ of an agent i in $\Gamma_{(f,P)}$ (in particular, (3) encodes incentive compatibility of $\Gamma_{(f,P)}$) and (4) corresponds to Condition 1 in the definition of selective elimination. Inequalities (5) and (6) correspond to Condition 2, where (6) is stated separately since it involves the variable $P_j(s_{-j}, \theta_j)$ instead of $\bar{P}_j^{(\theta,\theta')}(s_{-j}, \theta_j)$ as in (5).

Note that there are only *polynomially many* variables and inequalities in this system and all coefficients have polynomial encoding length. Hence, we can find a relative interior point of the polyhedron defined by the system, which corresponds to a solution of the original system with strict inequalities in (2) and (4), in polynomial time (e.g., by using the ellipsoid method). In particular, this shows that all the values $P_j(\theta)$ and $\bar{P}_j^{(\theta,\theta')}(s)$ can be chosen to have polynomial encoding length, which proves the following Theorem:

Theorem 5. *The social choice function $f : \Theta \to X$ is strongly implementable in dominant strategies if and only if there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$* of polynomial encoding length *that satisfies the selective elimination condition. In this case, for every (fixed) bad dominant pair (θ, θ') of $\Gamma_{(f,P)}$, the data $(i, \bar{N}, h, \tilde{\theta}_{-i}, \bar{\theta}_{-(\bar{N}\cup\{i\})})$ needed to selectively eliminate (θ, θ') can be chosen to have polynomial encoding length.* □

Using Theorem 5, we can now state our nondeterministic polynomial time algorithm for Strong Implementability and, thus, prove the main result of this section:

Theorem 6. *Strong Implementability* $\in$ NP.

Proof. Assume that the given social choice function f is strongly implementable in dominant strategies. Then, by Theorem 5, there exists an incentive compatible direct revelation mechanism $\Gamma_{(f,P)}$ of polynomial encoding length that satisfies the selective elimination condition. Moreover, for every bad dominant pair (θ, θ') of $\Gamma_{(f,P)}$, the

For all $i \in N, (\tilde{\theta}_i, \bar{\theta}_i) \in \Theta_i^2 \setminus D_i$:

$$\begin{aligned} & V_i\left(f(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i), \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right) + P_i\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right) \\ & -V_i\left(f(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_i), \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right) - P_i\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_i\right) \quad > \quad 0 \end{aligned} \tag{2}$$

For all $i \in N, (\tilde{\theta}_i, \bar{\theta}_i) \in D_i$ and all $\tilde{\theta}_{-i}, \bar{\theta}_{-i} \in \Theta_{-i}, \theta_i \in \Theta_i$:

$$\begin{aligned} & V_i\left(f(\bar{\theta}_{-i}, \bar{\theta}_i), \tilde{\theta}\right) + P_i\left(\bar{\theta}_{-i}, \bar{\theta}_i\right) \\ & -V_i\left(f(\bar{\theta}_{-i}, \theta_i), \tilde{\theta}\right) - P_i\left(\bar{\theta}_{-i}, \theta_i\right) \quad \geq \quad 0 \end{aligned} \tag{3}$$

For all $(\theta, \theta') \in D$:

$$\begin{aligned} & V_{i^{(\theta,\theta')}}\left(h^{(\theta,\theta')}(\bar{s}_{\bar{N}^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}^{(\theta,\theta')}\cup\{i^{(\theta,\theta')}\})}, \theta_{i^{(\theta,\theta')}}), \tilde{\theta}^{(\theta,\theta')}_{-i^{(\theta,\theta')}}, \theta_{i^{(\theta,\theta')}}\right) + \bar{P}^{(\theta,\theta')}_{i^{(\theta,\theta')}}\left(\bar{s}_{\bar{N}^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}\cup\{i^{(\theta,\theta')}\})}, \theta_{i^{(\theta,\theta')}}\right) \\ & -V_{i^{(\theta,\theta')}}\left(h^{(\theta,\theta')}(\bar{s}_{\bar{N}^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}^{(\theta,\theta')}\cup\{i^{(\theta,\theta')}\})}, \theta'_{i^{(\theta,\theta')}}), \tilde{\theta}^{(\theta,\theta')}_{-i^{(\theta,\theta')}}, \theta_{i^{(\theta,\theta')}}\right) - \bar{P}^{(\theta,\theta')}_{i^{(\theta,\theta')}}\left(\bar{s}_{\bar{N}^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}\cup\{i^{(\theta,\theta')}\})}, \theta'_{i^{(\theta,\theta')}}\right) > \quad 0 \end{aligned} \tag{4}$$

For all $j \in N, (\theta, \theta') \in D$ and all $\theta \in \Theta, s_{-j} \in \bar{S}_{-j}^{(\theta,\theta')} \setminus \Theta_{-j}, s_j \in \bar{S}_j^{(\theta,\theta')}$:

$$\begin{aligned} & V_j\left(h^{(\theta,\theta')}(s_{-j}, \theta_j), \theta\right) + \bar{P}_j^{(\theta,\theta')}\left(s_{-j}, \theta_j\right) \\ & -V_j\left(h^{(\theta,\theta')}(s_{-j}, s_j), \theta\right) - \bar{P}_j^{(\theta,\theta')}\left(s_{-j}, s_j\right) \quad \geq \quad 0 \end{aligned} \tag{5}$$

For all $j \in N, (\theta, \theta') \in D$ and all $\theta \in \Theta, s_{-j} \in \Theta_{-j}, s_j \in \bar{S}_j^{(\theta,\theta')} \setminus \Theta_j$:

$$\begin{aligned} & V_j\left(h^{(\theta,\theta')}(s_{-j}, \theta_j), \theta\right) + P_j\left(s_{-j}, \theta_j\right) \\ & -V_j\left(h^{(\theta,\theta')}(s_{-j}, s_j), \theta\right) - \bar{P}_j^{(\theta,\theta')}\left(s_{-j}, s_j\right) \quad \geq \quad 0 \end{aligned} \tag{6}$$

data $\left(i^{(\theta,\theta')}, \bar{N}^{(\theta,\theta')}, h^{(\theta,\theta')}, \tilde{\theta}^{(\theta,\theta')}_{-i^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}\cup\{i^{(\theta,\theta')}\})}\right)$ needed to selectively eliminate (θ,θ') can be chosen to have polynomial encoding length. Now consider the following nondeterministic algorithm for verifying that f is strongly implementable:

Algorithm 1.

1. *Guess the (polynomially many) values* $P_j(\theta)$.
2. *For every* $i \in N$*, guess the set* D_i *of all pairs* $(\tilde{\theta}_i, \bar{\theta}_i)$ *of types* $\tilde{\theta}_i \in \Theta_i$ *and dominant bids* $\bar{\theta}_i \in \Theta_i$ *for type* $\tilde{\theta}_i$ *in* $\Gamma_{(f,P)}$.
3. *Guess the set* $D \subseteq \Theta^2$ *of all bad dominant pairs in* $\Gamma_{(f,P)}$.
4. *For every* $i \in N$ *and every pair* $(\tilde{\theta}_i, \bar{\theta}_i) \in \Theta_i^2 \setminus D_i$*, guess the pair* $(\tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)})$ *of a type vector and a bid vector of the other agents such that*
$$U_i^{\Gamma_{(f,P)}}\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i | \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right) > U_i^{\Gamma_{(f,P)}}\left(\bar{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \bar{\theta}_i | \tilde{\theta}_{-i}^{(\tilde{\theta}_i,\bar{\theta}_i)}, \tilde{\theta}_i\right).$$
5. *For every* $(\theta,\theta') \in D$*, guess the data* $\left(i^{(\theta,\theta')}, \bar{N}^{(\theta,\theta')}, h^{(\theta,\theta')}, \tilde{\theta}^{(\theta,\theta')}_{-i^{(\theta,\theta')}}, \bar{\theta}^{(\theta,\theta')}_{-(\bar{N}\cup\{i^{(\theta,\theta')}\})}\right)$ *needed to selectively eliminate the bad dominant pair* (θ,θ').
6. *Check all the (polynomially many) inequalities in the system displayed on Page 11.*

Since all the values $P_j(\theta)$ and the data needed for selective elimination of each of the polynomially many bad dominant pairs $(\theta,\theta') \in D$ have polynomial encoding length, Algorithm 1 runs in polynomial time, which proves the claim. □

References

[1] H. Gui, R. Müller, and R. V. Vohra: *Dominant strategy mechanisms with multidimensional types.* In: *Computing and Markets* (2005).

[2] A. Mas-Colell, M. D. Whinston, and J. R. Green: *Microeconomic Theory.* Oxford University Press (1995).

[3] D. Mookherjee and S. Reichelstein: *Implementation via augmented revelation mechanisms.* Review of Economic Studies 57(3): pp. 453–475 (1990).

[4] R. Müller, A. Perea, and S. Wolf: *Weak monotonicity and Bayes-Nash incentive compatibility.* Games and Economic Behavior 61: pp. 344–358 (2007).

[5] N. Nisan and A. Ronen: *Algorithmic mechanism design.* Games and Economic Behavior 35: pp. 166–196 (2001).

[6] J.-C. Rochet: *A necessary and sufficient condition for rationalizability in a quasilinear context.* Journal of Mathematical Economics 16: pp. 191–200 (1987).

[7] C. Thielen and S. Westphal: *A combinatorial algorithm for strong implementation of social choice functions.* COMSOC 2010.

Sven O. Krumke and Clemens Thielen
Department of Mathematics
University of Kaiserslautern
Paul-Ehrlich-Str. 14
67663 Kaiserslautern, Germany
Email: `{krumke,thielen}@mathematik.uni-kl.de`