OPTIMAL COLLUSION-RESISTANT MECHANISMS WITH VERIFICATION*

Paolo Penna Carmine Ventre

Abstract

We study the so called *mechanisms with verification* introduced by Nisan and Ronen [1999]. We show that these mechanisms can optimize any "reasonable" cost function and that they can be resistant to any coalition of colluding agents in the strongest model which allows agents to exchange *side payments*. These mechanisms, which are also known as *n-truthful* mechanisms, are stronger than truthful mechanisms. We apply our main result to a number of problems previously studied in the algorithmic mechanism design literature and obtain n-truthful $(1+\varepsilon)$-approximation mechanisms with verification. By contrast, no truthful mechanism *without verification* can achieve such an approximation ratio. Moreover, every mechanism without verification which is collusion-resistant must have an *unbounded* approximation factor.

Key words: Game Theory, Algorithmic Mechanism Design

*Dipartimento di Informatica ed Applicazioni "R.M. Capocelli", Università di Salerno, via Ponte Don Melillo, I-84084 Fisciano (SA), Italy. E-mail: {penna, ventre}@dia.unisa.it Work supported by the European Union under IST FET Integrated Project AEOLUS (IST-015964).

1 Introduction

A large body of the literature studies ways to incorporate economic and game-theoretic considerations in the design of algorithms and protocols. One of the most studied and acknowledged paradigms is *mechanism design* (see e.g. [NR01, Pap01, AT01, FPSS05, CKV07, MS07]). Distributed computations over the Internet often involve self-interested parties (*selfish agents*) which may manipulate the protocol by misreporting a fundamental piece of information they hold (their own *type*). The protocol runs some algorithm which, because of the misreported information, is no longer guaranteed to return a "globally optimal" solution (optimality is naturally expressed as a function of agents' types) [NR01, Pap01]. Since agents can manipulate the algorithm by misreporting their types, one augments algorithms with carefully designed payment functions which make disadvantageous for an agent to do so. A *mechanism* consists of an algorithm (also termed *social choice function*) and payment rule which associates a payment to every agent. Each agent derives a *utility* which depends on the solution computed by the algorithm, on the type of the agent, and on the payment that the agent receives from the mechanism (the solution and the payment depend on the *reported* types). A mechanism is *truthful* if truth-telling is a *dominant strategy* for all agents. That is, the utility of any agent is maximized when this agent reports his/her type truthfully, no matter which strategy the other agents follow. We stress that the construction of a truthful mechanism is a challenging problem since the type of each agent is *unknown* to the mechanism (thus, the utility of the this agent cannot be directly computed by the mechanism). The only known general technique for designing truthful mechanisms are the classical Vickrey-Clarke-Groves (VCG) mechanisms [Vic61, Cla71, Gro73]. Unfortunately, these mechanisms suffer from two main limitations:

- **Not all problems can be optimally solved.** The underlying algorithm must optimize the so called *utilitarian* cost functions, that is, the weighted sum of all agents costs (see Section 2). In particular, truthful VCG mechanisms can achieve only a very poor approximation factor for certain optimization problems [NR01]. Although these problems admit polynomial-time $(1+\varepsilon)$-approximation algorithms, none of these algorithms can be used in a truthful mechanism. By contrast, truthful VCG mechanism are only n-approximate, with n being the number of agents. Nisan and Ronen [NR01] conjectured that such a bound cannot be improved by *any* truthful mechanism for these problems. In partial support of this conjecture, several negative results have been recently proved [CKV07, MS07]. Noticeably, these inapproximability results do not make any computational assumption since they hold also for algorithms (mechanisms) that run in *exponential time* and that use *randomization*.

- **Coalitions of agents can profit by misreporting their types.** It is well-known that agents can improve their own utilities if they collude and exchange side payments. Mechanisms that are resistant to coalitions of up to c agents are called *c-truthful* [GH05]. This is a much stronger requirement, since we want that no coalition of at most c agents can raise the utilities of its members even in the case of *side payments*, i.e., an agent can offer money to other agents in the coalition to induce him/her to lie [Sch00]. Schummer [Sch00] and Goldberg and Hartline [GH05] showed that c-truthful mechanisms are "unfeasible" already for $c = 2$: any such mechanism must essentially output a fixed solution [Sch00, GH05] and therefore its approximation ratio is *unbounded* for any non-trivial problem.

Motivated by the inapproximability results on truthful mechanisms, Nisan and Ronen [NR01] suggested a rather innovative approach called *mechanisms with verification*. The "classical" approach in mechanism design is to award the payment associated to an agent unconditionally, i.e., without performing any kind of verification and solely based on the agents reported types. On the contrary, mechanisms with verification provide the payments after the chosen solution has been implemented. This gives in some cases the possibility to detect that an agent misreported his/her type and thus to penalize him/her by denying the payment. In the following paragraph, we briefly review the main ideas of this approach.

Mechanisms with verification. Nisan and Ronen [NR01] introduced mechanisms with verification in the following scenario. We have a number of tasks and a number of machines. Each machine i is of some *type* t_i meaning that the execution of a task j on a machine i requires time t_i^j. Each machine corresponds to a *selfish agent* and the type of machine i is only known to agent i. We would like to find an assignment minimizing the *makespan*, that is, an assignment x minimizing the expression $\max_i t_i(x)$, where $t_i(x)$ is the sum of the execution times of the tasks that x allocates to machine i. The main idea proposed by Nisan and Ronen [NR01] is that machine i cannot release task j before t_i^j time steps. Therefore, if agent i reports a type b_i and a solution x is implemented, the mechanism is able to detect that b_i is *not* the true type of machine i if $b_i(x) < t_i(x)$. If that is the case, and only in this case, agent i is *caught lying* by the mechanism which can penalize this agent by not providing any payment to him/her. Observe that there are still many ways in which an agent can lie to the mechanism without being caught lying. Indeed, he/she can always report *higher* execution times for the tasks that the mechanism assign to his/her machine. Moreover, the mechanism cannot gain any extra information on the task that are *not* allocated to this agent/machine. So, the agent can lie on any of these tasks without being caught lying by the mechanism. In Section 1.2 we formally define these concepts for a very general setting: For every feasible solution x, an agent of type t_i has a cost $t_i(x)$ associated to this solution. After a solution x is implemented, the mechanism observes a *lower bound* on the cost $t_i(x)$ and provides agent i a suitable payment if and only if $b_i(x) \geq t_i(x)$. For the task scheduling problem above, the mechanism can observe the time machine i releases all tasks assigned to it: Agent i can of course release the tasks at any time *after* they have been executed, but it cannot release them *before*. So, agent i can pretend to have a machine of type b_i whenever $b_i(x) \geq t_i(x)$. The cost for agent i is always the time $t_i(x)$ that this agent actually spent for executing the tasks (the idle time has no cost since the agent can use his/her resources for other purposes). However, reporting a type b_i could be better for agent i since the chosen solution and the payment that this agent will receive depend on b_i (see Section 1.2 for a formal definition). We can regard to mechanisms with verification in a more general and abstract way. Agents have different capabilities/costs associated to each feasible solution. An agent that has a “low” capability incurs a “high” cost since it is more demanding for him/her to “contribute” to a particular solution. The mechanism gains only a *limited* information on the true agents capabilities/costs after a chosen solution is implemented: an agent can always simulate “lower” capabilities/“higher” costs and the observation refers to the implemented solution only (there is no extra information on the costs associated to other solutions that could have been implemented). A mechanism with verification provides a payment to agent i if and only if $b_i(x) \geq t_i(x)$, where x is the chosen solution, b_i is the type reported by i and t_i is his/her true type. In contrast, mechanisms without verification *always* provide the payment to agent i.

1.1 Our Contribution

In this work we provide the following extension of the VCG theorem. We show that certain VCG mechanisms can be turned into mechanisms with verification which *simultaneously* achieve the following two desiderata:

- It is possible to optimize a wide class of cost functions which includes utilitarian ones as a special case.

- The mechanism is n-*truthful* which is the strongest solution concept known in the literature [GH05].

Our mechanisms are constructive and they apply to the case of agents with *arbitrary bounded domains*. That is, the only requirement is that there is an upper and a lower bound on the cost that an agent can associate to each feasible outcome. All of the afore mentioned impossibility results [NR01, AT01, CKV07, MS07, Sch00, GH05] apply to *finite* domains which is a special case of our domains (actually some lower bounds holds for an even more restricted scenario in which agents are *single-parameter* [AT01, GH05]). Also, it is reasonable to assume that in practice costs cannot be too large or too small.

We apply our main positive result to a number of optimization problems previously considered in the algorithmic game theory community: *inter-domain routing* [FPSS05], its *workload-minimization* variant in [MS07], and *scheduling* tasks on servers with different capabilities [NR01, CKV07, MS07]. For several of these problems, the best known mechanisms without verification is the VCG which guarantees only n-approximations. Also, lower bounds are known for any truthful mechanism (without verification). By contrast, our mechanisms with verification are n-truthful and achieve an approximation factor of $(1+\varepsilon)$, for any $\varepsilon > 0$. Finally, with a similar approach, we show that truthful mechanisms with verification can achieve *exact* solutions for arbitrary monotone non decreasing costs functions. This result extends a recent result in [Ven06] which was proved for finite domains.

Roadmap. Preliminary definitions and results are given in Section 1.2. In Section 2 we consider utilitarian problems and some natural extensions for which we obtain exact truthful mechanisms with verification. Here we present both positive and negative results, showing that our assumptions are "minimal". In Section 3, we present the main result which is the construction of optimal n-truthful mechanisms with verification based on the VCG payments. We apply this result to several optimization problems in Section 4. Due to lack of space some of the proofs are given in Appendix A.

1.2 Preliminaries

We have a finite set A of feasible alternative *solutions* (or *outcomes*), with $a := |A|$. Without loss of generality, we assume that $A = \{1, \ldots, a\}$ and sometimes write $x \leq y$ for two outcomes x and y. There are n selfish agents, each of them having a so called *type*

$$t^i : A \to \mathbb{R}$$

which associates a monetary *cost* to every feasible outcome.[1] If an agent i receives a payment equal to r^i and an outcome x is selected, then his/her *utility* is equal to

$$r^i - t^i(x). \tag{1}$$

Each type t^i belongs to a so called *domain* D^i which consists of all admissible types, that is, a subset of all functions $u : A \to \mathbb{R}$. The type t^i is *private knowledge*, that is, it is known to agent i only. Everything else, including each domain D^i, is *public knowledge*. Hence, each agent i can *misreport* his/her type to *any* other element b^i in the domain D^i. We sometimes call such b^i the *bid* or *reported type* of agent i. We let D being the cross product of all agents domains, that is, D contains all bid vectors $\mathbf{b} = (b^1, \ldots, b^n)$ with b^i in D^i. An *algorithm* f (or *social choice function*) is a function

$$f : D \to A$$

[1] The fact that we consider costs, instead of valuations, is without loss of generality. Indeed, agents with a positive valuation associated to an outcome x can be modeled by means of a "negative cost" $t^i(x) < 0$.

which maps all agents (reported) types $\mathbf{b}$ into a feasible outcome $x = f(\mathbf{b})$. A *mechanism* is a pair (f, p), where f is a social choice function and $p = (p^1, \ldots, p^n)$ is a vector of suitable *payment functions*, one for each agent, where each payment function

$$p^i : D \rightarrow \mathbb{R}$$

associates some amount of money to agent i. We say that D is a *bounded domain* if there exist $\ell_{\min}$ and $\ell_{\max}$ such that $b^i(x)$ belongs to the interval $[\ell_{\min}, \ell_{\max}]$, for all outcomes x, for all b^i in D^i, and for all agents i. If no further restriction is made on D, then we have (social choice functions over) *arbitrary bounded domains*. Throughout the paper we consider only type vectors $\mathbf{v}$ in the domain D and we denote by v^i the type corresponding to agent i.

We say that an agent i is *truthtelling* if he/she reports his/her type, that is, the bid b^i coincides with his/her type t^i. Given a social choice function f and bids $\mathbf{b} = (b^1, \ldots, b^i, \ldots, b^n)$, we say that agent i is *caught lying by the verification* if the following inequality holds:

$$t^i(f(\mathbf{b})) > b^i(f(\mathbf{b})).$$

A mechanism (f, p) is a *mechanism with verification* if, on input bids $\mathbf{b}$, every agent that is caught lying by the verification does not receive any payment, while every other agent i receives his/her associated payment $p^i(\mathbf{b})$. On the contrary, we say that (f, p) is a *mechanism without verification* if every agent receives *always* his/her associated payment, that is, the payment received by agent i, on input bids $\mathbf{b}$, is the associated payment $p^i(\mathbf{b})$. In the sequel, whenever no confusion arises, we denote mechanisms without verification simply to as mechanisms. Given a mechanism (with verification) (f, p), we denote by

$$\mathrm{Utility}^i_{f,p}(\mathbf{b}|t^i)$$

the utility derived by agent i given that t^i is his/her type and that agents report $\mathbf{b}$ to the mechanism. We will omit the subscript 'f, p' whenever this is clear from the context. We say that (f, p) satisfies the *voluntary participation* constraint if truthtelling agents have always a non-negative utility. We remark that the utility function above depends only on the bid vector $\mathbf{b}$ and on his/her type t^i of agent i (in particular, whether agent i receives the payment $p^i(\mathbf{b})$ depends only on the type t^i and on the chosen outcome $f(\mathbf{b})$. So, the utility $\mathrm{Utility}^i_{f,p}(\mathbf{b}|t^i) = -t^i(f(\mathbf{b}))$ if

For any two type vectors $\mathbf{t}$ and $\mathbf{b}$, we say that a coalition C *can misreport* $\mathbf{t}$ *to* $\mathbf{b}$ if the vector $\mathbf{b}$ is obtained by changing the type of some of the agents in C, i.e., $t^i = b^i$ for every agent i *not* in the coalition C. For any two type vectors $\mathbf{t}$ and $\mathbf{b}$, we say that *verification does not catch* $\mathbf{t}$ *misreported to* $\mathbf{b}$ if $t^i(f(\mathbf{b})) \leq b^i(f(\mathbf{b}))$ for every agent i. Conversely, we say that *verification catches* $\mathbf{t}$ *misreported to* $\mathbf{b}$ if $t^i(f(\mathbf{b})) > b^i(f(\mathbf{b}))$ for some agent i.

We consider mechanisms satisfying (at least) the following:

Definition 1 *A mechanism (with verification) is* truthful *if truthtelling is a dominant strategy for all agents. That is, the utility of any agent i is maximized when this agent reports $b^i = t^i$, no matter the types reported by the other agents.*

A much stronger requirement is to have mechanisms which are even resistant to coalitions of $c > 1$ colluding agents that can exchange side payments.

Definition 2 (c-truthfulness [GH05]) *A mechanism (with verification) is c-*truthful *if, for any coalition of size at most c and any bid of agents not in the coalition, the sum of the utilities of the agents in the coalition is maximized when all agents in the coalition are truthtelling.*

Throughout the paper we make use of the following standard notation. Given a type vector $\mathbf{v} = (v^1, \ldots, v^n)$, we let $\mathbf{v}^{-i}$ being the vector of length $n-1$ obtained by removing v^i from $\mathbf{v}$, i.e., the vector $(v^1, \ldots, v^{i-1}, v^{i+1}, \ldots, v^n)$. We also let $(w, \mathbf{v}^{-i})$ be the vector $(v^1, \ldots, v^{i-1}, w, v^{i+1}, \ldots, v^n)$, which is obtained by replacing the i-th entry of $\mathbf{v}$ with w. Similarly, we consider subsets C of agents and define vectors $\mathbf{v}^C$, $\mathbf{v}^{-C}$, and $(\mathbf{v}^C, \mathbf{w}^{-C})$ in a similar natural way. We can rephrase the definition of c-truthful mechanism as follows. A mechanism (f, p) is c-truthful if, for every C of size at most c, for every $\mathbf{t}$ (the agents types), for every bid vector $\mathbf{b}^{-C}$ of the agents not in the coalition, and for every bid vector $\mathbf{b}^C$ of the agents in the coalition, the following inequality holds

$$\sum_{i \in C} \mathrm{Utility}^i((\mathbf{t}^C, \mathbf{b}^{-C})|t^i) \geq \sum_{i \in C} \mathrm{Utility}^i((\mathbf{b}^C, \mathbf{b}^{-C})|t^i).$$

Note that if there exist C, $\mathbf{v}$, and $\mathbf{b}$, that violate the above condition, then the vectors $\mathbf{t} = (\mathbf{v}^C, \mathbf{v}^{-C})$ and $\mathbf{b} = (\mathbf{b}^C, \mathbf{v}^{-C})$ satisfy

$$\sum_{i \in C} \mathrm{Utility}^i(\mathbf{t}|t^i) < \sum_{i \in C} \mathrm{Utility}^i(\mathbf{b}|t^i).$$

(Indeed, $t^i = v^i$ for every $i \in C$.)

This simple argument can be used to prove the following fact (see Appendix A.1 for the details).

Fact 3 *A mechanism (with verification) is c-truthful if and only if, for any coalition C of size at most c and for any two type vectors $\mathbf{t}$ and $\mathbf{b}$ such that C can misreport $\mathbf{t}$ to $\mathbf{b}$, the corresponding agents' utilities satisfy*

$$\sum_{i \in C} \mathrm{Utility}^i(\mathbf{t}|t^i) \geq \sum_{i \in C} \mathrm{Utility}^i(\mathbf{b}|t^i). \tag{2}$$

2 From Utilitarian to More General Optimization Problems

Let us consider the well-studied *utilitarian* cost functions, that is, the case

$$\mathrm{Cost}(x, \mathbf{t}) = \sum_{i=1}^{n} \alpha_i \cdot t^i(x),$$

where $\alpha_i > 0$ is a (potentially different) constant for each agent i. A social choice function f minimizing such cost function is called an *utilitarian social choice function*.[2]

We next generalize in two different ways utilitarian social choice functions by considering arbitrary cost functions that are non-decreasing in the agents' costs. Formally:

Definition 4 ((strongly) MIR social choice functions) *Let* $\mathrm{Cost} : A \times D \to \mathbb{R}$ *be a function of the form*

$$\mathrm{Cost}(x, \mathbf{t}) = \mathrm{Cost}(t^1(x), \ldots, t^n(x)),$$

which is monotone non-decreasing *in each $t^i(x)$. We say that f is* minimal in its range (MIR) *if, for any two type vectors $\mathbf{t}$ and $\mathbf{b}$, it holds that* $\mathrm{Cost}(f(\mathbf{t}), \mathbf{t}) \leq \mathrm{Cost}(f(\mathbf{b}), \mathbf{t})$. *Moreover, we say that f* uses a fixed tie breaking rule *if, for any two type vectors $\mathbf{t}$ and $\mathbf{b}$,* $\mathrm{Cost}(f(\mathbf{t}), \mathbf{t}) = \mathrm{Cost}(f(\mathbf{b}), \mathbf{t})$ *implies that* $f(\mathbf{t}) \preceq f(\mathbf{b})$. *Finally, we say that f is a* strongly MIR social choice function *if it is MIR and it uses a fixed tie breaking rule.*

[2]These social choice functions are also called *affine maximizer* in the literature. The term denotes the fact that these social choice functions maximize the "overall content" that is represented by the quantity '-$\mathrm{Cost}(x, \mathbf{t})$'. In the standard game-theoretic terminology, this quantity is called the *(weighted) social welfare*.

The following result will be key ingredient in our proofs. Informally speaking, it says that if no agent is caught lying by the verification, then the global cost of the solution evaluated on the agents' bids is minimal when all agents are truthtelling:

Lemma 5 (essentially due to [Ven06]) *For every MIR social choice function f and every two type vectors $\mathbf{t}$ and $\mathbf{b}$ for which verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$ it holds that* $\mathrm{Cost}(f(\mathbf{t}), \mathbf{t}) \leq \mathrm{Cost}(f(\mathbf{b}), (t^i, \mathbf{b}^{-i}))$.

PROOF. Since verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$, then $t^j(f(\mathbf{b})) \leq b^j(f(\mathbf{b}))$ for every agent j. Thus, from the monotonicity of $\mathrm{Cost}()$, we obtain

$$\begin{aligned}
\mathrm{Cost}(f(\mathbf{b}), \mathbf{t}) &= \mathrm{Cost}(t^1(f(\mathbf{b})), \ldots, t^i(f(\mathbf{b})), \ldots, t^n(f(\mathbf{b}))) \\
&\leq \mathrm{Cost}(b^1(f(\mathbf{b})), \ldots, t^i(f(\mathbf{b})), \ldots, b^n(f(\mathbf{b}))) \\
&= \mathrm{Cost}(f(\mathbf{b}), (t^i, \mathbf{b}^{-i})).
\end{aligned}$$

Since f is MIR we have $\mathrm{Cost}(f(\mathbf{t}), \mathbf{t}) \leq \mathrm{Cost}(f(\mathbf{b}), \mathbf{t})$. This and the inequality above imply the lemma. □

In Section C we show that truthful mechanisms with verification can optimize any strongly MIR social choice function:

Theorem 6 *Every strongly MIR social choice function admits an optimal truthful mechanism with verification over an arbitrary bounded domain (with $\ell_{\min} \geq 0$).*

We next observe that one cannot extend the above result by relaxing the definition of strongly MIR social choice function. Indeed, the "non-decreasing cost function", the "fixed tie breaking rule" and the "minimal in range" assumptions are necessary in order to guarantee the existence of exact truthful mechanisms with verification without introducing other conditions. MIR is necessary because of the optimality of the mechanisms. As for the other two assumptions we prove the next two theorems (whose proofs can be found in Appendix A.2 and A.3).

Theorem 7 *For any cost function that is* not *monotone nondecreasing there exists a bounded domain such that no social choice function that minimizes such a cost function admits a truthful mechanism with verification.*

We next remove the fixed tie breaking rule from our assumption and show that there exists an optimal social choice function not admitting truthful mechanisms with verification.

Theorem 8 *There exists a social choice function which is minimal in its range with respect to some monotone non-decreasing cost function that does not admit any truthful mechanism with verification.*

3 An Extension of the VCG Theorem

We consider a generalization of utilitarian problems for which we can provide optimal n-truthful mechanisms with verification. The basic idea is to consider the cost function as the sum of two nondecreasing cost functions with one of them being utilitarian:

$$\mathrm{Cost}(x, \mathbf{t}) = \left(\sum_i \alpha_i \cdot t^i(x) \right) + \mathrm{NonUtilitarianCost}(x, \mathbf{t}), \tag{3}$$

where each $\alpha_i > 0$ is a known constant and $\mathrm{NonUtilitarianCost}(x, \mathbf{t})$ is monotone non decreasing and is of the form $\mathrm{NonUtilitarianCost}(t^1(x), \ldots, t^n(x))$. A MIR social choice function minimizing a cost function of this type is a *generalized utilitarian social choice function*.

Our main result is the following theorem showing that n-truthful mechanisms with verification can optimize every such cost function. Then, a natural approach suggest itself: Assuming we have a non-utilitarian cost function $\mathrm{NonUtilitarianCost}()$, then we can "approximate" it by adding the "utilitarian" part with α_i's sufficiently small (see Section 4 for a discussion on such an approach).

Theorem 9 *Every generalized utilitarian social choice function maximizer over an arbitrary bounded domain admits an optimal n-truthful mechanism with verification.*

PROOF. We show that, for every generalized utilitarian social choice function f and for any $c \leq n$, there exist payments p such that the mechanism with verification (f, p) is c-truthful. The payment functions are a generalization of the VCG payments:

$$p^i(\mathbf{b}) := \hbar - \frac{1}{\alpha_i}\left(\sum_{j \neq i} \alpha_j \cdot b^j(f(\mathbf{b})) + \mathrm{NonUtilitarianCost}(f(\mathbf{b}), \mathbf{b})\right) \tag{4}$$

with $\hbar$ being a suitable *constant* to be given below.[3]

Let us consider an arbitrary coalition C of size at most c and any two type vectors $\mathbf{t}$ and $\mathbf{b}$ such that C can misreport $\mathbf{t}$ to $\mathbf{b}$. Because of Fact 3, it suffices to prove (2). Either verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$ or verification catches $\mathbf{t}$ misreported to $\mathbf{b}$. We consider the two cases separately.

For the case in which verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$, we make use of the following claim whose proof is given in Appendix A.4.

Claim 10 *For every $\mathbf{b}$ and $\mathbf{t}$ such that agent i receives his/her payment, it holds that* $\mathrm{Utility}^i(\mathbf{b}|t^i) \leq \hbar - \frac{\mathrm{Cost}(f(\mathbf{b}),(t^i,\mathbf{b}^{-i}))}{\alpha_i}$. *Moreover, for $\mathbf{b} = \mathbf{t}$, the inequality holds with '$=$'.*

The claim and Lemma 5 yield

$$\mathrm{Utility}^i(\mathbf{t}|t^i) = \hbar - \frac{\mathrm{Cost}(f(\mathbf{t}), \mathbf{t})}{\alpha_i} \geq \hbar - \frac{\mathrm{Cost}(f(\mathbf{b}), (t^i, \mathbf{b}^{-i}))}{\alpha_i} \geq \mathrm{Utility}^i(\mathbf{b}|t^i),$$

thus implying (2).

If verification catches $\mathbf{t}$ misreported to $\mathbf{b}$, then we have at least one agent $j \in C$ who does not receive any payment for $\mathbf{b}$. We make use of the following claim whose proof is given in Appendix A.5.

Claim 11 *There exist two constants $\rho_{\min}$ and $\rho_{\max} \geq \rho_{\min}$ such such that each agent receives a payment at most $\hbar - \rho_{\min}$, while the payment received by any truthtelling agent is at least $\hbar - \rho_{\max}$.*

From the first part of the claim above we have

$$\sum_{i \in C} \mathrm{Utility}^i(\mathbf{b}|t^i) \quad \leq \quad (|C| - 1)(\hbar - \rho_{\min}) - |C| \cdot \ell_{\min}, \tag{5}$$

[3] Although VCG mechanisms allow $\hbar$ to depend on the bids $\mathbf{b}^{-i}$ of the other agents, here we use a constant in order to achieve collusion-resistant mechanisms.

while the second part of the same claim implies

$$\sum_{i \in C} \text{Utility}^i(\mathbf{t}|t^i) \quad \geq \quad |C|(\hbar - \rho_{\max} - \ell_{\max}). \tag{6}$$

By setting $\hbar := c(\rho_{\max} - \rho_{\min} + \ell_{\max} - \ell_{\min}) + \rho_{\min}$, since $\rho_{\max} \geq \rho_{\min}$, $\ell_{\max} \geq \ell_{\min}$, and $|C| \leq c$, we have $\hbar \geq |C|(\rho_{\max} - \rho_{\min} + \ell_{\max} - \ell_{\min}) + \rho_{\min}$. Hence

$$\begin{aligned}
(|C| - 1)(\hbar - \rho_{\min}) - |C|\ell_{\min} &= |C|(\hbar - \rho_{\min} - \ell_{\min}) - (\hbar - \rho_{\min}) \\
&\leq |C|(\hbar - \rho_{\min} - \ell_{\min}) - |C|(\rho_{\max} - \rho_{\min} + \ell_{\max} - \ell_{\min}) \\
&= |C|(\hbar - \rho_{\max} - \ell_{\max})
\end{aligned}$$

and thus (2) follows from (5) and (6).

Since (2) holds in both cases, Fact 3 implies that the mechanism (f, p) is c-truthful. The theorem thus follows by taking $c = n$. □

Above theorem implies that utilitarian social choice functions admit optimal n-truthful mechanisms with verification: Consider $\text{NonUtilitarianCost}()$ as the constant function of value zero.

Corollary 12 *Every utilitarian social choice function over an arbitrary bounded domain admits an optimal n-truthful mechanism with verification.*

Observe that our n-truthful mechanisms for generalized utilitarian social choice functions use VCG payments. It is well-known that these payment functions are polynomial-time computable if the corresponding social choice function is polynomial-time computable. Hence, the following holds.

Corollary 13 *Every polynomial-time generalized utilitarian social choice function over an arbitrary bounded domain admits a polynomial-time n-truthful mechanism with verification.*

4 Applications

We consider a number of applications for which exact truthful mechanisms *without verification* do not exist. The respective lower bounds apply to the case in which the domain is bounded and holds even if we do not require polynomial running time. By contrast, we obtain n-truthful mechanisms with verification that select $(1 + \varepsilon)$-approximate solutions. Finally, for some restrictions we even obtain *polynomial-time* n-truthful mechanisms and/or *exact* solutions.

4.1 Min-Max Fairness

In this section we study Min-Max social choice functions. While utilitarian social choice functions minimize the "overall cost", Min-Max social choice functions minimize the maximum cost incurred by the agents. More formally, for every n-tuple of types $(t^1, \ldots, t^n)$ the Min-Max function selects the outcome x minimizing the expression $\max_i t^i(x)$. We assume that these functions expect agents bidding a nonnegative cost, i.e. for Min-Max social choice functions $\ell_{\min} \geq 0$. With such an assumption Min-Max functions preserve the essence of their *raison d'être* as many Min-Max problems involving agents supporting costs are NP-hard. Since Min-Max functions are not utilitarian, one would be tempted to plug them in the generalized utilitarian social choice functions definition (3) and use the technique of Theorem 9 to obtain an n-truthful mechanism with verification. While the mechanism is certainly n-truthful we are unaware of which kind of approximation guarantee it has. In the sequel we address the question of bounding the approximation ratio of such a mechanism for the Min-Max function. In detail, next theorem (whose proof is given in Appendix A.6) shows that, by properly setting the values of the α's in (3), it is possible to loose an arbitrarily small factor in the approximation ratio of the non-utilitarian function.

Theorem 14 *Every Min-Max social choice function over an arbitrary bounded domain admits a $(1+\varepsilon)$-approximation n-truthful mechanism with verification, for every $\varepsilon > 0$.*

4.2 Inter-Domain Routing

We are given an undirected graph $G = (N, L)$ (called the *AS graph*) where each node is an Autonomous System (AS) of which the Internet is comprised. The set of nodes N consists of a destination node d and n source nodes. The set of edges L is the set of links between nodes in N. Each source node is a selfish agent. The intensity of the traffic (number of packets) originating from source i for destination d is denoted by τ_i. Let $neighbours(i)$ be the set of nodes in N linked to i. Each source node has a *private* cost function $t_i : neighbours(i) \to \mathbb{R}^+$ that specifies the per-packet cost incurred by this node for carrying traffic. This cost models the load imposed to the AS internal routing for sending a packet to an adjacent AS. In a single parameter version of the problem the cost function t_i does not depend on $neighbours(i)$, i.e., $t_i^{w_1} = t_i^{w_2}$ for any $w_1, w_2 \in neighbours(i)$. We consider the more general multi-parameter case.

The goal is to assign all source nodes routes to d. We do not allow nodes to transfer traffic to two adjacent nodes. Thus *route allocation* should form a confluent tree to the destination d, or in other words, the route allocation should form a spanning tree of G rooted in d. Therefore, the set of alternatives A is the set of all spanning trees of G (which we also call routing tree) rooted in d.

Let x be a routing tree in A. Let $Tree_x(i)$ be the set of nodes in the subtree of x rooted at i and let $parent_x(i)$ be the parent of i in x. Source node i incurs the following cost:

$$t^i(x) = t_i^j \cdot \sum_{k \in Tree_x(i)} \tau_k,$$

with $j = parent_x(i)$.[4] The objective is to minimize the total cost of routing all packets:

$$\mathrm{Cost}(x, \mathbf{t}) = \sum_{i=1}^{n} t^i(x).$$

We refer to this problem as *Lowest Total Cost* in inter-domain routing. A similar problem has been already considered in [FPSS05]. This is an utilitarian problem. More precisely, it asks for a directed minimum spanning tree. A directed minimum spanning tree can be efficiently computed [Edm67]. Next result follows from the algorithm in [Edm67] and Corollaries 12, 13.

Corollary 15 *The lowest total cost problem in inter-domain routing admits an optimal n-truthful polynomial-time mechanism with verification, if nodes costs are bounded by some constant.*

[4]Observe that in our setting nodes should declare a cumulative cost $t^i()$. However, in this case it suffices that agent i declare the cost function t_i as it defines the cumulative cost we need.

In a different version we seek routing trees in which the workload imposed on the busiest node is minimized. Formally, we wish to minimize the cost function:

$$\mathrm{Cost}(x, \mathbf{t}) = \max_{i=1,\ldots,n} t^i(x).$$

We refer to this problem as *Workload Minimization* in inter-domain routing. This problem has been defined in [MS07] that shows an optimal truthful exponential-time mechanism (without verification) in the single parameter case and a polynomial-time n-approximation truthful mechanism (without verification) in the general case. Since this is a Min-Max problem from Theorem 14 we immediately obtain the following:

Corollary 16 *The workload minimization problem in inter-domain routing admits an n-truthful $(1+\varepsilon)$-approximation mechanism with verification, for every $\varepsilon > 0$, if nodes costs are bounded by some constant.*

4.3 Scheduling Selfish Machines

We are given n tasks that have to be scheduled on m machines. Every machine is controlled by a selfish agent. The type of a machine, for a given schedule x, is the time the machine takes to complete the tasks that x assigned her. The objective is to minimize the maximum completion time (makespan). If machines have the same speed for all tasks then they are referred to as related machines and we talk about *Scheduling Related Machines* problem. For this problem [AAS05] provides a $(1+\varepsilon)$-approximate polynomial-time truthful mechanism (without verification) for any fixed number of machines. Theorem 9, Corollary 13 and the social choice function component of the mechanism by [AAS05] imply the following theorem (whose proof is given in Appendix A.7):

Theorem 17 *For any fixed number of machines, the scheduling related machines problem (so to minimize the makespan) admits an n-truthful $(1+\varepsilon)$-approximation* polynomial-time *mechanism with verification, for every $\varepsilon > 0$, if machines execution times are bounded by some constant.*

A more general variant is that of *Scheduling Unrelated Machine* problem. Peculiarity of the unrelated case is that machines have a speed that is task-dependent, i.e., machine i has a speed s^i_j for task j and s^i_k for task k with s^i_j (potentially) different from s^i_k. This problem is very well studied as a mechanism design problem. It is defined by Nisan and Ronen in their seminal paper [NR01]. They show that any approximation better than 2 cannot be achieved by a truthful deterministic mechanism (*without verification*). Such a lower bound have been recently improved to $1+\sqrt{2}$ in [CKV07] for 3 or more machines.

We would like to obtain an optimal truthful mechanism with verification for the scheduling unrelated machines problem. By Theorem 6 we left to provide a strongly MIR social choice function for a cost function of the form of Definition 4. Since the cost of the scheduling unrelated machines problem is in that form and as there is a trivial strongly MIR social choice function we obtain the following:

Corollary 18 *The scheduling unrelated machines problem (so to minimize the makespan) admits an optimal truthful mechanism with verification, if machines execution times are bounded by some constant.*

Moreover, scheduling unrelated machine is a Min-Max problem. Therefore, from Theorem 14 we immediately obtain the following:

Corollary 19 *The scheduling unrelated machines problem (so to minimize the makespan) admits an n-truthful $(1+\varepsilon)$-approximation mechanism with verification, for every $\varepsilon > 0$, if machines execution times are bounded by some constant.*

Acknowledgements. The authors a grateful to Riccardo Silvestri for several useful comments on an earlier version of this work.

References

[AAS05] Nir Andelman, Yossi Azar, and Motti Sorani. Truthful approximation mechanisms for scheduling selfish related machines. In *Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, Lecture Notes in Computer Science, pages 69–82, 2005.

[ADP$^+$06] V. Auletta, R. De Prisco, P. Penna, G. Persiano, and C. Ventre. New constructions of mechanisms with verification. In *Proc. of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 4051 of *Lecture Notes in Computer Science*, pages 596–607. Springer, 2006.

[ADPP04] V. Auletta, R. De Prisco, P. Penna, and G. Persiano. The power of verification for one-parameter agents. In *Proc. of the 31st International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3142 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2004.

[AT01] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 482–491, 2001.

[CKV07] G. Christodoulou, E. Koutsoupias, and A. Vidali. A lower bound for scheduling mechanisms. In *Proc. of Annual ACM Symposium on Discrete Algorithms (SODA)*, 2007. To appear.

[Cla71] E.H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, pages 17–33, 1971.

[Edm67] J. Edmonds. Optimum branchings. *J. Research of the National Bureau of Standards*, 71B:233–240, 1967.

[FPSS05] Joan Feigenbaum, Christos H. Papadimitriou, Rahul Sami, and Scott Shenker. A bgp-based mechanism for lowest-cost routing. *Distributed Computing*, 18(1):61–72, 2005.

[GH05] A. V. Goldberg and J. D. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proc. of the 16th Annual ACM Symposium on Discrete Algorithms (SODA)*, pages 620–629. SIAM, 2005.

[Gro73] T. Groves. Incentive in Teams. *Econometrica*, 41:617–631, 1973.

[MS07] A. Mu'alem and M. Schapira. Setting lower bounds on truthfulness. In *Proc. of Annual ACM Symposium on Discrete Algorithms (SODA)*, 2007. To appear.

[NR01] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35:166–196, 2001. Extended abstract in the Proc. of the 31st Annual ACM Symposium on Theory of Computing (STOC), pages 129–140, 1999.

[NS06] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *JET*, 2006.

[Pap01] C.H. Papadimitriou. Algorithms, games, and the internet. In *Proc. of Annual ACM Symposium on Theory of Computing (STOC)*, 2001.

[Sch00] J. Schummer. Manipulation through bribes. *Journal of Economic Theory*, 91(3):180–198, 2000.

[Ven06] C. Ventre. Mechanisms with verification for any finite domain. In *Proc. of the 2nd Workshop on Internet and Network Economics (WINE)*, volume 4286 of *Lecture Notes in Computer Science*, pages 37–49. Springer, 2006.

[Vic61] W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.

A Postponed Proofs

A.1 Proof of Fact 3

In this section, we prove Fact 3. To this end, we extend the notation $(\alpha, \mathbf{b}^{-i})$ to subsets C of agents in the natural way. Hence $(\mathbf{u}^C, \mathbf{b}^{-C})$ denotes the vector whose entry i is u^i for $i \in C$ and b^i otherwise.

PROOF. Let (f, p) be an arbitrary mechanism (with verification). By Definition 24, this mechanism is c-truthful if and only if the following condition holds:

(c-truthfulness) For every C of size at most c, for every $\mathbf{t}$, and for every $\mathbf{b}$, it holds that

$$\sum_{i \in C} \text{Utility}^i_{f,p}((\mathbf{t}^C, \mathbf{b}^{-C})|t^i) \quad \geq \quad \sum_{i \in C} \text{Utility}^i_{f,p}(\mathbf{b}|t^i). \tag{7}$$

The condition in Fact 3 can be rewritten as follows:

(condition in Fact 3) For every C of size at most c, for every $\mathbf{v}$, and for every $\mathbf{b}$ such that C can misreport $\mathbf{v}$ to $\mathbf{b}$, it holds that

$$\sum_{i \in C} \text{Utility}^i_{f,p}(\mathbf{v}|v^i) \quad \geq \quad \sum_{i \in C} \text{Utility}^i_{f,p}(\mathbf{b}|v^i). \tag{8}$$

We have to show that the two conditions above are equivalent.

(c-truthfulness)$\Rightarrow$(condition in Fact 3). Recall that 'C can misreport $\mathbf{v}$ to $\mathbf{b}$' means that $\mathbf{v}$ and $\mathbf{b}$ agree in all entries not in C, that is, $\mathbf{v} = (\mathbf{v}^C, \mathbf{b}^{-C})$. For $\mathbf{t} = \mathbf{v}$, we have that $t^i = v^i$ for all $i \in C$ and thus (8) follows from (7).

(c-truthfulness)$\Leftarrow$(condition in Fact 3). Take $\mathbf{v} := (\mathbf{t}^C, \mathbf{b}^{-C})$ and observe that C can misreport $\mathbf{v}$ to $\mathbf{b}$. By construction, $t^i = v^i$ for all $i \in C$ and thus (7) follows from (8).

This completes the proof. $\square$

A.2 Proof of Theorem 7

PROOF. Since the cost function is not monotone nondecreasing then there exist $\alpha^1, \ldots, \alpha^n, \bar{\alpha}^i$ such that $\alpha^i < \bar{\alpha}^i$ and $\text{Cost}(\alpha^1, \ldots, \alpha^{i-1}, \alpha^i, \alpha^{i+1}, \ldots, \alpha^n) > \text{Cost}(\alpha^1, \ldots, \alpha^{i-1}, \bar{\alpha}^i, \alpha^{i+1}, \ldots, \alpha^n)$. We next show that there exists a bounded domain D for which it is not possible to define a payment function leading to a truthful mechanism with verification. The domain D is the cross product of the following single agents' domains: $D^i = \{u : A \to \{\alpha_i, \bar{\alpha}_i\}\}$ and, for $j \neq i$, $D^j = \{u : A \to \{\alpha_j\}\}$. On such a domain the only agent that can be lying is the agent i. Assume, for sake of contradiction, that a payment function P^i exists. Also assume, without loss of generality, that $A = \{x, \bar{x}\}$ (a greater alternative set would only hurt the readability of the rest of the proof). Consider the following two declarations (belonging to D^i): (i) b^i such that $b^i(x) = \bar{\alpha}^i$ and $b^i(\bar{x}) = \alpha^i$ and (ii) $\bar{b}^i$ such that $\bar{b}^i(x) = \alpha^i$ and $\bar{b}^i(\bar{x}) = \bar{\alpha}^i$, where, by construction, x (respectively $\bar{x}$) is the optimum for b^i (respectively $\bar{b}^i$). (Indeed, since the cost is minimized for $\bar{\alpha}_i$ – and not for α_i – then, when agent i declares b^i, a social choice function minimizing the cost function under consideration selects outcome x.) Since $\bar{\alpha}^i > \alpha^i$ then an agent whose type is b^i (respectively $\bar{b}^i$) can declare to be a type $\bar{b}^i$ (respectively b^i) without being caught by the verification. Thus truthfulness condition implies that $P^i(b^i) - P^i(\bar{b}^i) \geq \bar{\alpha}^i - \alpha^i$ as otherwise if i's true valuation function is b^i then she would have an incentive to declare her valuation to be $\bar{b}^i$. [5] Similarly, it must hold that $P^i(\bar{b}^i) - P^i(b^i) \geq \bar{\alpha}^i - \alpha^i$. Since $\bar{\alpha}^i > \alpha^i$ the sum of the above inequalities leads to a contradiction. □

A.3 Proof of Theorem 8

PROOF SKETCH. We are given three selfish related machines and two tasks and want to minimize the makespan. This problem has a monotone nondecreasing cost function. Consider the social choice function f that assigns the heaviest task to the fastest machine and the lightest task to the slower machine if her completion time does not exceed the completion time of the fastest machine. In such a case f assigns the lightest task as an optimal assignment does. It is easy to see that f is optimal (and thus MIR) and to provide an instance of the problem in which f is not monotone (see Definition 21). Theorem follows since monotonicity is necessary to obtain truthful mechanisms with verification when problems involve one-parameters agents [ADPP04] (as in this case). □

A.4 Proof of Claim 10

PROOF. If agent i receives his/her payment, then his/her utility is as follows:

$$\begin{aligned}
\text{Utility}^i(\mathbf{b}|t^i) &= p^i(\mathbf{b}) - t^i(f(\mathbf{b})) \\
&= \hbar - \frac{1}{\alpha_i}\left(\alpha_i \cdot t^i(f(\mathbf{b})) + \left(\sum_{j\neq i} \alpha_j \cdot b^j(f(\mathbf{b}))\right) + \text{NonUtilitarianCost}(f(\mathbf{b}),\mathbf{b})\right) \\
&= \hbar - \frac{1}{\alpha_i}\Bigg(\text{Cost}(f(\mathbf{b}),(t^i,\mathbf{b}^{-i})) - \\
&\qquad \text{NonUtilitarianCost}(f(\mathbf{b}),(t^i,\mathbf{b}^{-i})) + \text{NonUtilitarianCost}(f(\mathbf{b}),\mathbf{b})\Bigg), \qquad (9)
\end{aligned}$$

where last equality follows from (3). By taking $\mathbf{b} = \mathbf{t}$, this implies $\text{Utility}^i(\mathbf{t}|t^i) = \hbar - \frac{\text{Cost}(f(\mathbf{t}),\mathbf{t})}{\alpha_i}$, and thus the second part of the claim holds.

Since i receives his/her payment, it must be the case $t^i(f(\mathbf{b})) \leq b^i(f(\mathbf{b}))$. Since the function $\text{NonUtilitarianCost}()$ is monotone nondecreasing, we have $\text{NonUtilitarianCost}(f(\mathbf{b}),(t^i,\mathbf{b}^{-i})) \leq \text{NonUtilitarianCost}(f(\mathbf{b}),\mathbf{b})$ which, combined with (9), implies $\text{Utility}^i(\mathbf{b}|t^i) \leq \hbar - \frac{\text{Cost}(f(\mathbf{b}),(t^i,\mathbf{b}^{-i}))}{\alpha_i}$. This proves the first part of the claim and concludes its proof. □

A.5 Proof of Claim 11

PROOF. Let $\alpha_{\min} := \min_i\{\alpha_i\}$ and $\alpha_{\max} := \max_i\{\alpha_i\}$. We define

$$\rho_{\min} := \frac{\ell_{\min} \cdot \left(\sum_{j\neq i} \alpha_j\right) + \text{NonUtilitarianCost}(\ell_{\min},\dots,\ell_{\min})}{\alpha_{\max}}, \qquad (10)$$

$$\rho_{\max} := \frac{\ell_{\max} \cdot \left(\sum_{j\neq i} \alpha_j\right) + \text{NonUtilitarianCost}(\ell_{\max},\dots,\ell_{\max})}{\alpha_{\min}}. \qquad (11)$$

[5] We removed from P^i notation the dependence from the declaration of all agent but i since they are fixed.

Since $\ell_{\min} \leq \ell_{\max}$, $\alpha_{\min} \leq \alpha_{\max}$ and from the monotonicity of $\mathrm{NonUtilitarianCost}()$ we have $\rho_{\min} \leq \rho_{\max}$. From (4) we have $\hbar - \rho_{\max} \leq p^i(\mathbf{b}) \leq \hbar - \rho_{\min}$. □

A.6 Proof of Theorem 14

PROOF. Since the Min-Max social choice function is not utilitarian then we use it as the $\mathrm{NonUtilitarianCost}()$ component of a generalized utilitarian social choice function (3). Let (f, p) be the mechanism of Theorem 9 with, for a given $\varepsilon > 0$, $\alpha_i = \varepsilon/n$ for any $i = 1, \ldots, n$. Let x be the outcome output of (f, p). Theorem 9 ensures that (f, p) is n-truthful and that x minimizes the following quantity:

$$\varepsilon \cdot \frac{\sum_{i=1}^n t^i(x)}{n} + \max_{i=1,\ldots,n} t^i(x),$$

where $(t^1, \ldots, t^n)$ is the true type vector of the agents. Let x^* be the optimal outcome of the Min-Max social choice function. Let us assume for sake of contradiction that x is not a $(1+\varepsilon)$-approximation of x^*. Formally:

$$\max_{i=1,\ldots,n} t^i(x) > (1+\varepsilon) \cdot \max_{i=1,\ldots,n} t^i(x^*). \tag{12}$$

We next derive simple properties. "Chaining" the inequalities given by these properties and the (12) we will obtain a contradiction. Since $n \cdot \max_{i=1,\ldots,n} t^i(x^*) \geq \sum_{i=1}^n t^i(x^*)$ then we have:

$$(1+\varepsilon) \cdot \max_{i=1,\ldots,n} t^i(x^*) \geq \varepsilon \cdot \frac{\sum_{i=1}^n t^i(x^*)}{n} + \max_{i=1,\ldots,n} t^i(x^*). \tag{13}$$

Because x minimizes the generalized utilitarian social choice function then it holds that:

$$\varepsilon \cdot \frac{\sum_{i=1}^n t^i(x)}{n} + \max_{i=1,\ldots,n} t^i(x) \leq \varepsilon \cdot \frac{\sum_{i=1}^n t^i(x^*)}{n} + \max_{i=1,\ldots,n} t^i(x^*). \tag{14}$$

Since $\ell_{\min} \geq 0$ then $\sum_{i=1}^n t^i(x) \geq 0$ and thus

$$\varepsilon \cdot \frac{\sum_{i=1}^n t^i(x)}{n} + \max_{i=1,\ldots,n} t^i(x) \geq \max_{i=1,\ldots,n} t^i(x). \tag{15}$$

Inequalities (13), (14), (15) and (12) imply that

$$(1+\varepsilon) \cdot \max_{i=1,\ldots,n} t^i(x^*) > (1+\varepsilon) \cdot \max_{i=1,\ldots,n} t^i(x^*).$$

We have reached the contradiction, and thus, x is a $(1+\varepsilon)$-approximation of x^*. □

A.7 Proof of Theorem 17

PROOF SKETCH. The social choice function part of the mechanism in [AAS05] considers a subset A' of alternatives and find the optimum (using a fixed tie breaking rule) in A'. A' is constructed in polynomial-time in such a way that it contains a $(1+\varepsilon)$-approximate solution of the makespan. Such a social choice function is thus MIR: no solution in $A \setminus A'$ is explored and the function returns the optimum in A'. We use this polynomial-time algorithm to construct A' and find the optimum in A' and consider the output of this algorithm as the NonUtilitarianCost() component of (3). Now by mimicking the proof of Theorem 14 we can show that we can set α_i's in (3) in such a way that the final outcome is a $(1+\varepsilon)$-approximation of the makespan (the optimum in A' is a $(1+\varepsilon)$-approximation also for the utilitarian component of (3)). To conclude the proof we stress that social choice function in [AAS05] is polynomial time if the number of machine is fixed and that, in such a case, the payments can be computed in polynomial time as stated in Corollary 13. □

B Collusion-Resistant Mechanisms for Single-Parameter Agents

In this section we consider the case of *single-parameter* agents. Here, each outcome partitions the agents into two sets: those that are *selected* and those that are *not selected*. The value $t^i(x)$ depends *uniquely* on the fact that i selected in x or not and it is completely specified by a *parameter* t_i, which is a real number such that

$$t^i(x) = \begin{cases} t_i & \text{for } i \text{ selected in } x, \\ 0 & \text{for } i \text{ not selected in } x. \end{cases} \tag{16}$$

Whether i selected in x is publicly known, for every outcome x, and thus each agent can only specify (and misreport) the parameter t_i.[6] Also for single-parameter agents we assume agents to have a *bounded domain* with $\ell_{\min} = 0$ and $\ell_{\max} = \ell$. So, each parameter t_i belongs to the interval $[0, \ell]$.

Notice that, mechanisms with verification are not more powerful than those without verification, if we simply require truthfulness (see the discussion in [ADP+06]). As soon as we consider c-truthful mechanisms, instead, verification turns out to be an extremely powerful tool:

Remark 20 *For single-parameter agents, every 2-truthful mechanism must use posted prices (see [Sch00, GH05]). Hence, no generalized utilitarian social choice function admits a 2-truthful mechanism. By contrast, every generalized utilitarian social choice function admits an optimal n-truthful mechanism with verification (see Theorem 9).*

In the sequel we will provide sufficient conditions for the existence of c-truthful mechanisms, for any given $c \leq n$. These conditions do not assume social welfare maximization.

B.1 Sufficient conditions for c-truthfulness

We begin with a *necessary* condition. Observe that in order to have truthful mechanisms for single-parameter agents (even when using verification [ADPP04]) the social choice function must select agents "monotonically":

Definition 21 (monotone) *We say that f is* monotone *if the following holds. Having fixed the bids of all agents but i, agent i is selected if bidding a cost less than a threshold value $b_i^{\oplus}$, and is not selected if bidding a cost more than a threshold value $b_i^{\oplus}$. In particular, for every $\mathbf{b} \in D$ and for every i, there exists a value $b_i^{\oplus}$ which depends only on $\mathbf{b}^{-i}$ and such that (i) i selected in $f(b^i, \mathbf{b}^{-i})$ for $b_i < b_i^{\oplus}$ and (ii) i not selected in $f(b^i, \mathbf{b}^{-i})$ for $b_i > b_i^{\oplus}$.*

From Definition 21 we can easily obtain the following:

Fact 22 *If f is monotone and i selected in $f(\mathbf{b})$, then $b_i \leq b_i^{\oplus}$. Moreover, if i not selected in $f(\mathbf{b})$ then $b_i \geq b_i^{\oplus}$. Hence, for bounded domains the threshold values of Definition 21 are in the interval $[0, \ell]$.*

[6] Single-parameter agents differ from one-parameter agents [AT01] since their types are expressed in a slightly different way. Indeed, in the latter case types have to consider also some kind of "load" that output assigns.

From (16) we immediately get the following:

Fact 23 *For single-parameter agents, it holds that verification does not catch* $\mathbf{t}$ *misreported to* $\mathbf{b}$ *if and only if* $t_i \leq b_i$ *for all* i *selected in* $f(\mathbf{b})$.

We next give a *sufficient* condition for c-truthfulness on single-parameter domains. Obviously this condition takes into account the bound c on the size of the coalition:

Definition 24 (c-resistant) *We say that* $\mathbf{b}$ *is* c-different *from* $\mathbf{t}$ *if these two type vectors differ for at most* c *agents' types. A monotone* f *is* c-resistant *if, for every* $\mathbf{b}$ *which is* c*-different from* $\mathbf{t}$ *and such that verification does not catch* $\mathbf{t}$ *misreported to* $\mathbf{b}$*, it holds that* $t_i^{\oplus} \leq b_i^{\oplus}$ *for all* i *not selected in* $f(\mathbf{b})$.

Theorem 25 *Every* c*-resistant* f *admits an optimal* c*-truthful mechanism with verification for single-parameter agents.*

PROOF. We define the payment functions as follows:

$$p^i(\mathbf{b}) := \begin{cases} \hbar - b_i^{\oplus} & \text{if } i \text{ not selected in } f(\mathbf{b}) \\ \hbar & \text{otherwise} \end{cases}$$

where $\hbar := c \cdot \ell$.

Let us consider an arbitrary coalition C of size at most n and any two type vectors $\mathbf{t}$ and $\mathbf{b}$ such that C can misreport $\mathbf{t}$ to $\mathbf{b}$. Because of Fact 3, it suffices to prove (2). Either verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$ or verification catches $\mathbf{t}$ misreported to $\mathbf{b}$. We consider the two cases separately.

If verification catches $\mathbf{t}$ misreported to $\mathbf{b}$, then we have at least one agent $j \in C$ which does not receive any payment for $\mathbf{b}$. Moreover, the payment received by every other agent i in the coalition is at most $\hbar$. Hence, we have

$$\sum_{i \in C} \text{Utility}^i(\mathbf{b}|t^i) \leq (c-1)\hbar = c\hbar - \hbar.$$

We next show that the utility of every truthtelling agent is at least $\hbar - \ell$. Indeed, the definition of $p^i()$ implies that $\text{Utility}^i(\mathbf{t}|t^i)$ is either $\hbar - t_i^{\oplus}$ for i not selected in $f(\mathbf{t})$, or $\hbar - t_i$ for i selected in $f(\mathbf{t})$. Fact 22 says that $t_i^{\oplus} \leq \ell$ and, if i selected in $f(\mathbf{t})$, then $t_i \leq t_i^{\oplus}$. Hence, $\text{Utility}^i(\mathbf{t}|t^i) \geq \hbar - \ell$. From this and from our choice of $\hbar$, we obtain

$$\sum_{i \in C} \text{Utility}^i(\mathbf{t}|t^i) \geq c(\hbar - \ell) = c\hbar - c\ell = c\hbar - \hbar.$$

The two inequalities above clearly imply (2).

It left to show the case in which verification does not catch $\mathbf{t}$ misreported to $\mathbf{b}$. We will show that for any $i \in C$ it holds

$$\text{Utility}^i(\mathbf{t}|t^i) \geq \text{Utility}^i(\mathbf{b}|t^i),$$

which clearly implies (2). There are four possible cases:

Case 1 (i selected in $f(\mathbf{t})$ and i selected in $f(\mathbf{b})$)**.** In this case nothing changes for i. Indeed, by the definition of $p^i()$, we have

$$\text{Utility}^i(\mathbf{t}|t^i) = \hbar - t_i = \text{Utility}^i(\mathbf{b}|t^i).$$

Case 2 (i not selected in $f(\mathbf{t})$ and i selected in $f(\mathbf{b})$). Fact 22 implies that $t_i^{\oplus} \leq t_i$. This and the definition of $p^i()$ imply

$$\text{Utility}^i(\mathbf{t}|t^i) = \hbar - t_i^{\oplus} \geq \hbar - t_i = \text{Utility}^i(\mathbf{b}|t^i).$$

Case 3 (i not selected in $f(\mathbf{t})$ and i not selected in $f(\mathbf{b})$). Since f is c-resistant, we have that $t_i^{\oplus} \leq b_i^{\oplus}$. This and the definition of $p^i()$ imply

$$\text{Utility}^i(\mathbf{t}|t^i) = \hbar - t_i^{\oplus} \geq \hbar - b_i^{\oplus} = \text{Utility}^i(\mathbf{b}|t^i).$$

Case 4 (i selected in $f(\mathbf{t})$ and i not selected in $f(\mathbf{b})$). Since i selected in $f(\mathbf{t})$, Fact 22 implies $t_i \leq t_i^{\oplus}$. Since i not selected in $f(\mathbf{b})$ and as f is c-resistant, we have that $t_i^{\oplus} \leq b_i^{\oplus}$, thus implying $t_i \leq b_i^{\oplus}$. This and the definition of $p^i()$ imply

$$\text{Utility}^i(\mathbf{t}|t^i) = \hbar - t_i \geq \hbar - b_i^{\oplus} = \text{Utility}^i(\mathbf{b}|t^i).$$

This concludes the proof. □

B.2 Sufficient conditions for n-truthfulness

We next give a *sufficient* condition for obtaining n-truthful mechanisms with verification for single-parameter agents.

Definition 26 (threshold-resistant) *A monotone f is* threshold-resistant *if $t_i^{\oplus} \leq b_i^{\oplus}$ for every $\mathbf{t}$ and every $\mathbf{b}$ obtained by increasing one agent entry of $\mathbf{t}$, where $t_i^{\oplus}$ and $b_i^{\oplus}$ are the threshold values of Definition 21.*

The rest of the section will be devoted to the proof of the next theorem asserting that above property is indeed sufficient for n-truthfulness.

Theorem 27 *Every threshold-resistant strongly MIR social choice function admits an optimal n-truthful mechanism with verification for single-parameter agents.*

Because of Theorem 25 a proof for next theorem actually shows Theorem 27 above.

Theorem 28 *Every threshold-resistant strongly MIR social choice function is n-resistant.*

First, we show the following technical result:

Lemma 29 *Let f be a strongly MIR social choice function for single-parameter agents. Consider any two vectors $\mathbf{b}$ and $\tilde{\mathbf{b}}$ which differ only in the agent i's entry. If it holds i not selected in $f(\mathbf{b})$ and i not selected in $f(\tilde{\mathbf{b}})$, then $f(\mathbf{b}) = f(\tilde{\mathbf{b}})$.*

PROOF. From i not selected in $f(\mathbf{b})$ and i not selected in $f(\tilde{\mathbf{b}})$ and since the two vectors agree in every coordinate other than i, we have the following two identities:

$$\text{Cost}(f(\mathbf{b}), \mathbf{b}) = \text{Cost}(f(\mathbf{b}), \tilde{\mathbf{b}}) \text{ and } \text{Cost}(f(\tilde{\mathbf{b}}), \mathbf{b}) = \text{Cost}(f(\tilde{\mathbf{b}}), \tilde{\mathbf{b}}).$$

From this and since f is MIR we obtain

$$\begin{aligned}\mathrm{Cost}(f(\mathbf{b}),\mathbf{b}) &\leq \mathrm{Cost}(f(\tilde{\mathbf{b}}),\mathbf{b})\\ &= \mathrm{Cost}(f(\tilde{\mathbf{b}}),\tilde{\mathbf{b}})\\ &\leq \mathrm{Cost}(f(\mathbf{b}),\tilde{\mathbf{b}})\\ &= \mathrm{Cost}(f(\mathbf{b}),\mathbf{b})\end{aligned}$$

Hence, the two inequalities must hold with '$=$'. Since f uses a fixed tie breaking rule, we have $f(\mathbf{b}) = f(\tilde{\mathbf{b}})$. □

We are now in a position to prove Theorem 28:

PROOF OF THEOREM 28. By contradiction, assume that f is not n-resistant. That is, there exists a vector $\mathbf{u}$ which is n-different from $\mathbf{v}$ such that verification does not catch $\mathbf{u}$ misreported to $\mathbf{v}$ with $v_i^{\oplus} < u_i^{\oplus}$, for some i not selected in $f(\mathbf{v})$. Consider the following two vectors obtained from $\mathbf{u}$ and $\mathbf{v}$, respectively, by replacing their i's entry with the same value:

$$\begin{aligned}\tilde{\mathbf{u}} &:= (u_1,\ldots,u_{i-1},\tilde{u}_i,u_{i+1},\ldots,u_n)\\ \tilde{\mathbf{v}} &:= (v_1,\ldots,v_{i-1},\tilde{v}_i,v_{i+1},\ldots,v_n)\end{aligned}$$

with $\tilde{u}_i = \tilde{v}_i$ satisfying

$$v_i^{\oplus} < \tilde{v}_i = \tilde{u}_i < u_i^{\oplus}.$$

Next, we increase some of the entries of $\tilde{\mathbf{u}}$ in order to obtain $\hat{\mathbf{u}}$ such that for all j selected in $f(\hat{\mathbf{u}})$ it holds $\hat{u}_j \geq \tilde{v}_j$. This implies that

$$\mathrm{Cost}(f(\hat{\mathbf{u}}),\tilde{\mathbf{v}}) \leq \mathrm{Cost}(f(\hat{\mathbf{u}}),\hat{\mathbf{u}}). \tag{17}$$

Starting from $\hat{\mathbf{u}} = \tilde{\mathbf{u}}$, until there is some j selected in $f(\hat{\mathbf{u}})$ with $\hat{u}_j < \tilde{v}_j$, we replace $\hat{u}_j$ with $\tilde{v}_j$ in $\hat{\mathbf{u}}$. Notice that this process eventually stops since every j is considered at most once. By definition of $\hat{\mathbf{u}}$ and since f is threshold-resistant, we have $\tilde{u}_i^{\oplus} \leq \hat{u}_i^{\oplus}$. (This inequality holds at every step.) Since $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$ differ from $\mathbf{u}$ and $\mathbf{v}$ only in the i's entry, we have $\tilde{u}_i^{\oplus} = u_i^{\oplus}$ and $\tilde{v}_i^{\oplus} = v_i^{\oplus}$. In particular, since $\tilde{u}_i = \tilde{v}_i$, the construction of $\hat{\mathbf{u}}$ implies that $\hat{u}_i = \tilde{u}_i$. Putting all these things together we have

$$\tilde{v}_i^{\oplus} = v_i^{\oplus} < \tilde{v}_i = \tilde{u}_i = \hat{u}_i < u_i^{\oplus} = \tilde{u}_i^{\oplus} \leq \hat{u}_i^{\oplus}.$$

This and Definition 21 imply that $f(\tilde{\mathbf{v}}) \neq f(\hat{\mathbf{u}})$.

We conclude the proof by showing that $f(\tilde{\mathbf{v}}) = f(\hat{\mathbf{u}})$, thus a contradiction. We start by proving

$$\mathrm{Cost}(f(\tilde{\mathbf{v}}),\hat{\mathbf{u}}) \leq \mathrm{Cost}(f(\tilde{\mathbf{v}}),\tilde{\mathbf{v}}). \tag{18}$$

Since $\tilde{v}_i > \tilde{v}_i^{\oplus}$, we have i not selected in $f(\tilde{\mathbf{v}})$. Recall that by our initial hypothesis i is not selected in $f(\mathbf{v})$. From Lemma 29 we obtain $f(\mathbf{v}) = f(\tilde{\mathbf{v}})$. Now assume, by contradiction, that (18) does not hold, that is, $\mathrm{Cost}(f(\mathbf{v}),\hat{\mathbf{u}}) > \mathrm{Cost}(f(\mathbf{v}),\tilde{\mathbf{v}})$. Then, there is some j selected in $f(\mathbf{v})$ with $\hat{u}_j > \tilde{v}_j$. Hence, it must be $j \neq i$ and, by definition of $\hat{\mathbf{u}}$, we have $\hat{u}_j = \tilde{u}_j$. (Indeed, since $\hat{u}_j > \tilde{v}_j$, then in the construction of $\hat{\mathbf{u}}$ the j-th coordinate is not updated.) Putting things together we obtain that:

$$u_j = \tilde{u}_j = \hat{u}_j > \tilde{v}_j = v_j.$$

The inequality $u_j > v_j$ and the hypothesis j selected in $f(\mathbf{v})$ imply that verification catches $\mathbf{u}$ misreported to $\mathbf{v}$, which is a contradiction. We conclude that (18) must hold. Then the following inequalities hold too:

$$
\begin{aligned}
\mathrm{Cost}(f(\hat{\mathbf{u}}), \hat{\mathbf{u}}) &\leq && \text{(since } f \text{ is MIR)} && (19)\\
\mathrm{Cost}(f(\tilde{\mathbf{v}}), \hat{\mathbf{u}}) &\leq && \text{(by Equation 18)}\\
\mathrm{Cost}(f(\tilde{\mathbf{v}}), \tilde{\mathbf{v}}) &\leq && \text{(since } f \text{ is MIR)}\\
\mathrm{Cost}(f(\hat{\mathbf{u}}), \tilde{\mathbf{v}}) &\leq && \text{(by Equation 17)} && (20)\\
\mathrm{Cost}(f(\hat{\mathbf{u}}), \hat{\mathbf{u}}) &.
\end{aligned}
$$

Hence, these inequalities must hold with '='. Since f uses a fixed tie breaking rule, from (19) and (20) we obtain $f(\hat{\mathbf{u}}) \leq f(\tilde{\mathbf{v}})$ and $f(\tilde{\mathbf{v}}) \leq f(\hat{\mathbf{u}})$, respectively. This implies $f(\hat{\mathbf{u}}) = f(\tilde{\mathbf{v}})$, which is the desired contradiction. □

The class of threshold-resistant strongly MIR social choice functions turns out to be fundamental for the results of the next section where we consider again the arbitrary bounded domains of Section 3, but this time for *arbitrary* cost functions.

C General Optimizer (GO-)Mechanisms

In this section we consider *arbitrary cost functions* $\mathrm{Cost}()$ together with *arbitrary bounded domains* with $\ell_{\min} \geq 0$. We derive truthful mechanisms for any general optimizer over any arbitrary bounded domain. The main idea is to regard each agent as a *coalition* of (virtual) single-parameter agents.

C.1 Arbitrary domains as coalitions of single-parameter agents

We call every agent whose domain is an arbitrary bounded domain a *multidimensional* agent. Since there are a alternative outcomes, and n multidimensional agents, we simply consider n coalitions $C_1, \ldots, C_n$, where each coalition C_i consists of a (virtual) single-parameter agents that corresponds to the multidimensional agent i. (These are actually "known" coalitions which we use only for the purpose of defining the payments and analyzing the resulting mechanism.) This "new game" has $N = n \cdot a$ single-parameter agents and a alternative outcomes. For each outcome x, we have a single single-parameter agent per coalition being selected: denoted by $1^{(i)}, \ldots, a^{(i)}$ the agents in coalition C_i, we have agent $x^{(i)}$ being selected in x, and every other agent in C_i being not selected in x; this holds for all coalitions above. We choose the parameter of the (virtual) single-parameter agents in the coalition C_i so that the cost for an agent $x^{(i)}$ selected is equal to the cost of the multidimensional agent i when outcome x is selected. That is, for all i and all outcomes x, the parameter $t_{x^{(i)}}$ of agent $x^{(i)}$ is equal to $t^i(x)$.

Observe that any type b^i in the domain of the multidimensional agent i can be seen as a vector

$$\mathbf{b}^i := (b^i_1, \ldots, b^i_a),$$

with $b^i_x = b^i(x)$ for every alternative outcome x. In particular, $\mathbf{b}^i$ is the vector of the parameters of the a agents in C_i, that is, b^i_x is the parameter of agent $x^{(i)}$. Consider a strongly MIR social choice function g over the multidimensional agents, and fix the bids $\mathbf{b}^{-i}$ of all agents but i. Then the resulting *single player function* $g(b^i, \mathbf{b}^{-i})$ can be seen as the social choice function $f(\mathbf{b}^i)$ over the domain given by the a single-parameter agents in C_i. In the sequel, we will show that every such single player function is a-resistant (see Lemma 32 below). Based on this fact, we can apply the techniques developed for single-parameter agents:

Definition 30 (GO-mechanism) *For any strongly MIR social choice function g we consider its* single player function*, depending on $\mathbf{b}^{-i}$, as $f(\mathbf{b}^i) := g(b^i, \mathbf{b}^{-i})$. In this case, the single player function f has its set C_i of* virtual *single-parameter agents. We define payment functions $q^i(b^i, \mathbf{b}^{-i}) := \sum_{j \in C_i} p^j(\mathbf{b}^i)$ where each $p^i()$ is the payment function of Theorem 25 when applied to f above and to the single-parameter agents in C_i. The resulting mechanism with verification (g, q) is called a* strongly MIR social choice function mechanism *(in short* GO-mechanism*).*

In the sequel we prove that every GO-mechanism is truthful *for the multidimensional agents*. In order to prove this result, we first observe that the GO-mechanism needs only be resistant to the "known" coalitions defined above:

Lemma 31 *If every single player function f of g is a-resistant with respect to its virtual single-parameter agents, then the GO-mechanism is truthful for the multidimensional agents.*

PROOF SKETCH. We observe that the utility of a multidimensional agent i is the sum of the utilities of all single-parameter agents in the coalition C_i. Therefore, if (g, q) was not truthful, then the mechanism (f, p) would not be a-truthful, thus contradicting Theorem 25. $\square$

So, in order to prove that GO-mechanisms are truthful for the multidimensional agents, it suffices to show that every strongly MIR social choice function satisfies the hypothesis of Lemma 31. Next lemma accomplishes such a task. We stress that, in its proof, we make use of the sufficient condition for n-truthfulness for single-parameter agents given by Theorem 27.

Lemma 32 *Every strongly MIR social choice function satisfies the hypothesis of Lemma 31.*

PROOF. Fix an agent i and an arbitrary bid vector $\mathbf{b}^{-i}$ for all other agents. Consider the corresponding single player function $f(\mathbf{u}^i) := g(u^i, \mathbf{b}^{-i})$, with u^i in the domain of the multidimensional agent i. We shall prove that f is threshold-resistant with respect to its virtual single-parameter agents.

First of all, we have to show that f is monotone. For any value α in $[0, \ell]$, we let

$$(\alpha, \mathbf{u}^i_{-j}) := (u^i_1, \ldots, u^i_{j-1}, \alpha, u^i_{j+1}, \ldots, u^i_a).$$

Recall that, for every outcome x, the single-parameter agent $j \in C_i$ is selected in x if and only if $j = x$. Consider any two vectors $\mathbf{u}^\uparrow$ and $\mathbf{u}^\downarrow$ that are obtained from $\mathbf{u}^i$ by changing u^i_j into $u^\uparrow_j$ and $u^\downarrow_j < u^\uparrow_j$, respectively:

$$\begin{aligned} \mathbf{u}^\uparrow &:= (u^\uparrow_j, \mathbf{u}^i_{-j}), \\ \mathbf{u}^\downarrow &:= (u^\downarrow_j, \mathbf{u}^i_{-j}). \end{aligned}$$

We shall prove that

$$j \text{ not selected in } f(\mathbf{u}^\downarrow) \;\Rightarrow\; f(\mathbf{u}^\uparrow) = f(\mathbf{u}^\downarrow), \tag{21}$$

$$j \text{ selected in } f(\mathbf{u}^\uparrow) \;\Rightarrow\; j \text{ selected in } f(\mathbf{u}^\downarrow). \tag{22}$$

Notice that (21) implies the following: if j is not selected in $f(\mathbf{u}^\downarrow)$ then j is not selected in $f(\mathbf{u}^\uparrow)$ as well. Therefore, (22) follows directly from (21). In Section C.1.1 we prove (21). We have to show the existence, for every agent j and for every vector $\mathbf{u}^i$ as above, of a threshold value $u^\oplus_j$ which satisfies Definition 21.

First, if j not selected in $f(\mathbf{u}^i)$ for all values $u_j \in [0, \ell]$, then we can simply set $u_j^\oplus := 0$. Otherwise, we consider

$$u_j^\oplus := \sup\{\alpha \in [0, \ell] | \; j \text{ selected in } f(\alpha, \mathbf{u}_{-j}^i)\}.$$

By definition and from (21-22), this value satisfies Definition 21. So, f is monotone w.r.t. its virtual single-parameter agents in C_i.

Next we show that the single player function f is threshold-resistant. We proceed by way of contradiction and assume that there exist two vectors $\mathbf{v}^i$ and $\mathbf{w}^i$ such that the following happens. The vector $\mathbf{w}^i$ is obtained from $\mathbf{v}^i$ by increasing only the value of v_j^i, that is, $w_j^i > v_j^i$ and

$$\begin{aligned} \mathbf{v}^i &:= (v_1^i, \ldots, v_{j-1}^i, v_j^i, v_{j+1}^i, \ldots, v_a^i), \\ \mathbf{w}^i &:= (v_1^i, \ldots, v_{j-1}^i, w_j^i, v_{j+1}^i, \ldots, v_a^i). \end{aligned}$$

Moreover, there exists an agent k whose threshold values $v_k^\oplus$ and $w_k^\oplus$ associated to $\mathbf{v}^i$ and $\mathbf{w}^i$ satisfy $w_k^\oplus < v_k^\oplus$. Observe that it must be $k \neq j$ since $\mathbf{v}^i$ and $\mathbf{w}^i$ are identical apart from the j's entry and thus $w_j^\oplus = v_j^\oplus$. Then we consider $\tilde{\mathbf{v}}^i$ and $\tilde{\mathbf{w}}^i$ obtained from $\mathbf{v}^i$ and $\mathbf{w}^i$, respectively, by replacing v_k^i and w_k^i with $\tilde{v}_k^i$ and $\tilde{w}_k^i$ such that $w_k^\oplus < \tilde{v}_k^i = \tilde{w}_k^i < v_k^\oplus$. Since f is monotone, we have k selected in $f(\tilde{\mathbf{v}}^i)$, k not selected in $f(\tilde{\mathbf{w}}^i)$, and because of (21), $f(\tilde{\mathbf{w}}^i) = f(\mathbf{w}^i) = y$. Since k is not selected in y we have $k \neq y^{(i)}$. From our "mapping" between the type w^i and the vector $\mathbf{w}^i$, we have that $w_y^i = \tilde{w}_y^i$, that is, $w^i(y) = \tilde{w}^i(y)$. This implies the following identity:

$$\begin{aligned} \mathrm{Cost}(y, (\tilde{w}^i, \mathbf{b}^{-i})) &= \mathrm{Cost}(b^1(y), \ldots, b^{i-1}(y), \tilde{w}^i(y), b^{i+1}(y), \ldots, b^n(y)) \\ &= \mathrm{Cost}(b^1(y), \ldots, b^{i-1}(y), w^i(y), b^{i+1}(y), \ldots, b^n(y)) \\ &= \mathrm{Cost}(y, (w^i, \mathbf{b}^{-i})). \end{aligned} \tag{23}$$

Putting things together we have

$$\begin{aligned} \mathrm{Cost}(g(\tilde{v}^i, \mathbf{b}^{-i}), (\tilde{v}^i, \mathbf{b}^{-i})) &\leq && \text{(since } g \text{ is MIR)} && (24) \\ \mathrm{Cost}(g(v^i, \mathbf{b}^{-i}), (\tilde{v}^i, \mathbf{b}^{-i})) &= && \text{(since } j \neq k \text{ and thus } \tilde{v}_k = v_k) \\ \mathrm{Cost}(g(v^i, \mathbf{b}^{-i}), (v^i, \mathbf{b}^{-i})) &\leq && \text{(since } g \text{ is MIR)} && (25) \\ \mathrm{Cost}(y, (v^i, \mathbf{b}^{-i})) &\leq && \text{(since } v_l \leq w_l \text{ for all } l) \\ \mathrm{Cost}(y, (w^i, \mathbf{b}^{-i})) &= && \text{(from Equation 23 above)} \\ \mathrm{Cost}(y, (\tilde{w}^i, \mathbf{b}^{-i})) &\leq && \text{(since } g \text{ is MIR)} && (26) \\ \mathrm{Cost}(g(\tilde{v}^i, \mathbf{b}^{-i}), (\tilde{w}^i, \mathbf{b}^{-i})) &= && \text{(since } k = f(\tilde{v}^i) = g(\tilde{v}^i, \mathbf{b}^{-i}) \text{ and } \tilde{w}_k = \tilde{v}_k) \\ \mathrm{Cost}(g(\tilde{v}^i, \mathbf{b}^{-i}), (\tilde{v}^i, \mathbf{b}^{-i})) &. \end{aligned}$$

Hence, these inequalities must hold with '='. Since g uses a fixed tie breaking rule, then (24), (25), and (26) imply $g(\tilde{\mathbf{v}}, \mathbf{b}^{-i}) \leq g(\mathbf{v}, \mathbf{b}^{-i})$, $g(\mathbf{v}, \mathbf{b}^{-i}) \leq y$, and $y \leq g(\tilde{\mathbf{v}}, \mathbf{b}^{-i})$, respectively. This implies $g(\tilde{\mathbf{v}}, \mathbf{b}^{-i}) = y = g(\tilde{\mathbf{w}}, \mathbf{b}^{-i})$. This is a contradiction to the fact that k is selected only in one of these two outcomes. We conclude that f is threshold-resistant for its a virtual single-parameter agents. By Theorem 27 the pair (f, p) is a-truthful. □

Hence, the main result of this section – already stated in Section 2 – follows:

Theorem 6 *Every strongly MIR social choice function admits an optimal truthful mechanism with verification over an arbitrary bounded domain (with $\ell_{\min} \geq 0$).*

Results analogous to Corollary 13 hold for the single-parameter mechanisms in Section B and for GO-mechanisms, if one assumes finite domains. The idea is to perform a binary search to determine the threshold values of Definition 21. For GO-mechanisms the running time is polynomial in the size of the input $\mathbf{t}$, where each t^i is a vector of a values, one for each outcome. Such an "explicit" representation of the input is in general necessary, as implied by computational complexity lower bounds for combinatorial auctions [NS06].

Corollary 33 *Every polynomial-time strongly MIR social choice function over an arbitrary finite domain admits a polynomial-time truthful mechanism with verification. For finite single-parameter domains, every polynomial-time c-resistant strongly MIR social choice function admits a polynomial-time c-truthful mechanism with verification.*

C.1.1 Proof of Equation 21

PROOF. From j not selected in $f(\mathbf{u}^{\downarrow})$ we have that $\mathrm{Cost}(f(\mathbf{u}^{\downarrow}),(w^i,\mathbf{b}^{-i}))$ does not depend on the value w^i_j, where

$$\mathbf{w}^i = (w^i_1, \ldots, w^i_j, \ldots, w^i_a)$$

is the vector associated to the type w^i. In particular, $\mathrm{Cost}(f(\mathbf{u}^{\downarrow}),\mathbf{u}^{\uparrow}) = \mathrm{Cost}(f(\mathbf{u}^{\downarrow}),\mathbf{u}^{\downarrow})$. Since g is MIR, then f is MIR too. We thus obtain

$$\begin{aligned}\mathrm{Cost}(f(\mathbf{u}^{\uparrow}),(\mathbf{u}^{\uparrow},\mathbf{b}^{-i})) &\leq \mathrm{Cost}(f(\mathbf{u}^{\downarrow}),(\mathbf{u}^{\uparrow},\mathbf{b}^{-i})) = \mathrm{Cost}(f(\mathbf{u}^{\downarrow}),(\mathbf{u}^{\downarrow},\mathbf{b}^{-i}))\\ &\leq \mathrm{Cost}(f(\mathbf{u}^{\uparrow}),(\mathbf{u}^{\downarrow},\mathbf{b}^{-i})) \leq \mathrm{Cost}(f(\mathbf{u}^{\uparrow}),(\mathbf{u}^{\uparrow},\mathbf{b}^{-i})),\end{aligned}$$

where the last inequality follows from the monotonicity of $\mathrm{Cost}()$ and from the fact that, by definition, $u^{\downarrow}_k \leq u^{\uparrow}_k$ for all k. (Indeed, we have $u^{\downarrow}_j < u^{\uparrow}_j$ while $u^{\downarrow}_k = u^{\uparrow}_k$, for $k \neq j$.) Hence the above inequalities hold with '='. Since f uses a fixed tie breaking rule, we have $f(\mathbf{u}^{\downarrow}) = f(\mathbf{u}^{\uparrow})$. This concludes the proof. □