Mechanism Design with Selective Verification

DIMITRIS FOTAKIS, National Technical University of Athens
CHRISTOS TZAMOS, Massachusetts Institute of Technology
MANOLIS ZAMPETAKIS, Massachusetts Institute of Technology

We introduce a general approach based on *selective verification* and obtain approximate mechanisms without money for maximizing the social welfare in the general domain of Utilitarian Voting. Having a good allocation in mind, a mechanism with verification selects few critical agents and detects, using a verification oracle, whether they have reported truthfully. If yes, the mechanism produces the desired allocation. Otherwise, the mechanism ignores any misreports and proceeds recursively with the remaining agents. We obtain randomized truthful (or almost truthful) mechanisms without money that verify only $O(\ln m/\varepsilon)$ agents, where m is the number of outcomes, independently of the total number of agents, and are $(1-\varepsilon)$-approximate for the social welfare. We also show that any truthful mechanism with a constant approximation ratio needs to verify $\Omega(\log m)$ agents. A remarkable property of our mechanisms is *immunity* (to agent misreports), namely that their outcome depends only on the reports of the truthful agents.

CCS Concepts:•**Theory of computation**→ **Algorithmic mechanism design;**

Additional Key Words and Phrases: Algorithmic Mechanism Design; Approximate Mechanism Design without Money; Mechanisms with Verification

1. INTRODUCTION

Let us consider a simple mechanism design setting where we place a facility on the line based on the preferred locations of n strategic agents. Each agent aims to minimize the distance of her preferred location to the facility and may misreport her location, if she finds it profitable. Our objective is to minimize the maximum distance of any agent to the facility and we insist that the facility allocation should be *truthful*, i.e., no agent should be able to improve her distance by misreporting her location. The optimal solution is to place the facility at the average of the two extreme locations. However, if we cannot incentivize truthfulness through monetary transfers (e.g., due to ethical or practical reasons), the optimal solution is not truthful. E.g., the leftmost agent may declare a location further on the left so that the facility moves closer to her preferred location. In fact, for the infinite real line, the optimal solution leads to no equilibrium declarations for the two extreme agents. The fact that in this simple setting, the optimal solution is not truthful was part of the motivation for the research agenda of *approximate mechanism design without money*, introduced by Procaccia and Tennenholtz [2013]. They proved that the best deterministic (resp. randomized) truthful mechanism achieves an approximation ratio of 2 (resp. $3/2$) for this problem.

Authors' addresses: D. Fotakis, Division of Computer Science, School of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece, fotakis@cs.ntua.gr . C. Tzamos, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, tzamos@mit.edu . M. Zampetakis, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, mzampet@mit.edu .

EC'16, July 24–28, 2016, Maastricht, The Netherlands. ACM 978-1-4503-3936-0/16/07 ...$15.00.

http://dx.doi.org/10.1145/2940716.2940743

Our work is motivated by the simple observation that the optimal facility allocation can be implemented truthfully if we inspect the declared locations of the two extreme agents and verify that they coincide with their preferred locations (e.g., for their home address, we may mail something there or visit them). Inspection of the two extreme locations happens before we place the facility. If both agents are truthful, we place the facility at their average. Otherwise, we ignore any false declarations and recurse on the remaining agents. This simple modification of the optimal solution is truthful, because non-extreme agents do not affect the facility allocation, while the two extreme agents cannot change the facility location in their favor, due to the verification step. Interestingly, this also applies to the optimal solution for k facilities on the line when we minimize the maximum agent-facility distance. Now, we need to inspect the locations of k extreme agent pairs, one for each facility. Moreover, the Greedy algorithm for k-Facility Location (see e.g., [Williamson and Shmoys 2011, Sec. 2.2]) becomes truthful if we verify the k agents allocated a facility and ignore any liars among them (see Section 4). Greedy is 2-approximate for minimizing the maximum agent-facility distance, in any metric space, while Fotakis and Tzamos [2014] show that there are no deterministic truthful mechanisms (without verification) that place $k \geq 2$ facilities in tree metrics and achieve any bounded (in terms of n and k) approximation ratio.

1.1. Selective Verification: Motivation and Justification

Verifying the declarations of most (or all) agents and imposing large penalties on liars should suffice for the truthful implementation of socially efficient solutions (see e.g., [Caragiannis et al. 2012]). But in our facility location examples, we truthfully implement the optimal (or an almost optimal) solution by verifying a small number of agents (independent of n) and by using a mild and reasonable penalty. Apparently, verification is successful in these examples because it is *selective*, in the sense that we verify only the critical agents for the facility allocation and fully trust the remaining agents.

Motivated by this observation, we investigate the power of *selective verification* in approximate mechanism design without money in general domains. We consider the general setting of Utilitarian Voting with m outcomes and n strategic agents, where each agent has a nonnegative utility for each outcome. We aim at truthful mechanisms that verify few critical agents and guarantee high utility for the agents. We do not include the verification cost in the social objective and refer to the total agent utility as the *social welfare*. The main reason is that in absence of money, it is not clear whether (and how) the agent utility and the verification cost can be expressed in the same unit. Instead, we analyze the verification cost as a separate efficiency criterion and require that it should be low. Our goal is to determine the best approximation guarantee for the social welfare achievable by truthful mechanisms with limited selective verification, so that we obtain a better understanding of the power of selective verification in mechanism design without money. Our main result is a smooth and essentially best possible tradeoff between the approximation ratio and the number of agents verified by randomized truthful (or almost truthful) mechanisms with selective verification.

Our general approach is to start from a (non-truthful) allocation rule f with a good approximation guarantee for the social welfare and to devise a mechanism F without money that incentivizes truthful reporting through selective verification. The mechanism F first selects an outcome o and an (ideally small) verification set of agents according to f (e.g., for facility location on the line, the allocation rule f is to take the average of the two extreme locations, the selected outcome o is the average of the two extremes in the particular instance and the verification set consists of the two extreme agents). Next, F detects, through the use of a *verification oracle*, whether the selected agents are truthful. If yes, the mechanism outputs o. Otherwise, F excludes any misreporting agents and is applied recursively to the remaining agents. We note that F

asks the verification oracle for a single bit of information about each agent verified: whether she has reported truthfully or not. F excludes misreporting agents from the allocation rule, so it does not need to know anything else about their true utilities.

Instead of imposing some explicit (i.e., monetary) penalty to the agents caught lying by verification, the mechanism F just ignores their reports, a reasonable reaction to their revealed attempt of manipulating the mechanism. We underline that liars still get utility from the selected outcome. It just happens that their preferences are not taken into account by the allocation rule. For these reasons, the penalty of exclusion from the mechanism is mild and compatible with the main assumption behind mechanisms without money (i.e., that in absence of monetary transfers, selecting an appropriate outcome is the only way in which the mechanism can affect the agent utilities).

Selective verification allows for an explicit quantification of the amount of verification and is applicable to essentially any domain. From a theoretical viewpoint, we believe that it can lead to a deep and delicate understanding of the power of limited verification in approximate mechanism design without money. From a practical viewpoint, the extent to which selective verification and the penalty of ignoring false declarations are natural depends on the particular domain / application. E.g., for applications of facility location, where utility is usually determined by the home address of each agent, public authorities have simple ways of verifying it. E.g., registration to a public service usually requires a certificate of address. Failure to provide such a certificate usually implies that the application is ignored, with no penalties attached.

1.2. Technical Approach and Results

A (randomized) mechanism with selective verification is *truthful* (in expectation) if no matter what the other agents report and whether they are truthful, truthful reporting maximizes the (expected) utility of each agent from the mechanism. Two nice features of our allocation rules (and mechanisms) is that they are *strongly anonymous* and *scale invariant*. The former means that the allocation only depends on the total agents' utility for each outcome (and not on each agent's contribution) and the latter means that multiplying all valuations by a positive factor does not change the allocation.

For mechanisms with selective verification, truthfulness is a consequence of two natural (and desirable) properties: immunity and voluntary participation. Immunity is a remarkable property made possible by selective verification. A mechanism with verification F is *immune* (to agent misreports) if F completely ignores any misreports and the resulting probability distribution is determined by the reports of truthful agents only. So, if F is immune, no misreporting agent can change the resulting allocation whatsoever. We achieve immunity through obliviousness of F to the declarations of misreporting agents not verified (see also [Fotakis and Tzamos 2013b, Sec. 5]). Specifically, a randomized mechanism F is *oblivious* if the probability distribution of F over the outcomes, conditional on the event that no misreporting agents are included in the verification set, is identical to the probability distribution of F if all misreporting agents are excluded from the mechanism. Namely, misreporting agents not verified do not affect the allocation of F. By induction on the number of agents, we show that obliviousness is a sufficient condition for immunity (Lemma 3.1). To the best of our knowledge, this is the first time that immunity (or a similar) property is considered in mechanism design. We defer the discussion about immunity to Section 11.

Immunity leaves each agent with essentially two options: either she reports truthfully and participates in the mechanism or she lies and is excluded from it. An allocation rule satisfies *voluntary participation* (or simply, *participation*) if each agent's utility when she is truthful is no less than her utility when she is excluded from the mechanism. Immunity and participation imply truthfulness (Lemma 3.2). We prove that strongly anonymous randomized allocation rules that satisfy participation are closely related to *maximal in distributional range* rules (see e.g., [Dobzinski and Dughmi 2013; Lavi and Swamy 2011]), i.e., allocation rules that maximize the expected social welfare over a subset of probability distributions over outcomes. Specifically, we show that maximizing the social welfare is sufficient for participation (Lemma 2.1), while for scale invariant and continuous allocation rules, it is also necessary (Lemma 2.2).

As a proof of concept[1], we apply selective verification to k-Facility Location problems (Section 4), which have served as benchmarks in approximate mechanism design without money (see e.g., [Procaccia and Tennenholtz 2013; Alon et al. 2010; Lu et al. 2010; Fotakis and Tzamos 2013a] and the references therein). We show that the Greedy allocation ([Williamson and Shmoys 2011, Section 2.2]) and the Proportional allocation [Lu et al. 2010; Arthur and Vassilvitskii 2007] satisfy participation and are immune and truthful, if we verify the k agents allocated the facilities (Theorems 4.1 and 4.2).

For the general domain of Utilitarian Voting, we aim at strongly anonymous randomized allocation rules that are maximal in distributional range, so that they satisfy participation, and oblivious, so that they achieve immunity. In Section 5, we present the *Power mechanism*, which selects each outcome o with probability proportional to the ℓ-th power of the total utility for o, where $\ell \geq 0$ is a parameter. Power provides a smooth transition from the (immune and truthful) uniform allocation, where each outcome is selected with probability $1/m$, for $\ell = 0$, to the optimal solution, for $\ell \to \infty$. Power is scale invariant, approximately maximizes the social welfare and approximately satisfies participation. Exploiting the proportional nature of its probability distribution, we make Power oblivious and immune by verifying at most ℓ agents. Using $\ell = \ln m/\varepsilon$, we obtain that for any $\varepsilon > 0$, Power with selective verification of $\ln m/\varepsilon$ agents, is immune, ε-truthful and $(1-\varepsilon)$-approximate for the social welfare (Theorem 5.5).

To quantify the improvement on the approximation ratio due to selective verification, we show that without verification, in the general setting of Utilitarian Voting, the best possible approximation ratio of any randomized truthful mechanism is $1/m$. In a more restricted setting with injective valuations [Filos-Ratsikas and Miltersen 2014], the best known randomized truthful mechanism achieves an approximation ratio of $\Theta(m^{-3/4})$ and the best possible approximation ratio is $O(m^{-2/3})$.

In Section 7, we characterize the class of scale invariant and strongly anonymous truthful mechanisms that verify $o(n)$ agents and achieve *full allocation*, i.e., result in some outcome with probability 1. We prove that such mechanisms employ a constant allocation rule, i.e., a probability distribution that does not depend on the agent reports. Therefore, they cannot achieve nontrivial approximation guarantees. Our characterization reveals an interesting connection between continuity (which is necessary for low verification), full allocation and maximal in distributional range mechanisms.

Relaxing some of the properties in the characterization, we obtain truthful mechanisms with low verification. Relaxing full allocation, we obtain the *Partial Power mechanism* (Section 8), and relaxing scale invariance, we obtain the *Exponential mechanism* (Section 9). Both are immune and truthful. For any $\varepsilon > 0$, they verify $O(\ln m/\varepsilon^2)$ agents in the worst-case and $\ln m/\varepsilon$ agents in expectation, respectively. Partial Power is $(1-\varepsilon)$-approximate, while Exponential has an additive error of εn. The amount of verification is essentially best possible, since any truthful mechanism with constant approximation ratio needs to verify $\Omega(\log m)$ agents (Theorem 6.1). We match this lower bound, that applies to all mechanisms, by strongly anonymous and scale invariant mechanisms. All our mechanisms can be implemented in time polynomial in n and m.

[1]One may verify that the optimal solution when we place k facilities on the line and minimize the maximum distance of any agent to the nearest facility is scale invariant, not strongly anonymous, oblivious (and thus, immune) and satisfies participation. Immunity and participation imply truthfulness.

Power	Partial Power	Exponential
ε-truthful	truthful	truthful
full allocation	**partial allocation**	full allocation
scale invariant	scale invariant	**not scale invariant**
immune	immune	immune
$(1-\varepsilon)$-approximation	$(1-\varepsilon)$-approximation	additive error εn
verification $\ln m/\varepsilon$	verification $O(\ln m/\varepsilon^2)$	expected verification $\ln m/\varepsilon$

Fig. 1. The main properties of our mechanisms. Partial allocation means that the mechanism may result in an artificial outcome of valuation 0 for all agents (e.g., we may refuse to allocate anything, for private goods, or to provide the service, for public goods). We depict in bold the property whose relaxation allows the mechanism to escape the characterization of Section 7.

The properties of our mechanisms are summarized in Fig. 1. In all cases, we achieve a smooth tradeoff between the number of agents verified and the quality of approximation. Rather surprisingly, the verification depends on m, the number of outcomes, but not on n, the number n of agents. Our results imply that especially for public good allocation, where m is typically independent of n, selective verification of few agents leads to truthful (or almost truthful) mechanisms that are almost as efficient as VCG, but do not require any monetary transfers. As a concrete example in this direction, we discuss, in Section 10, an application to the Combinatorial Public Project problem (see e.g., [Schapira and Singer 2008; Papadimitriou et al. 2008]).

Due to the space limitations, some proofs and technical claims are omitted from this extended abstract. The full version of this work is available at [Fotakis et al. 2015b].

1.3. Related Previous Work

Previous work [Archer and Kleinberg 2008; Caragiannis et al. 2012; Fotakis and Zampetakis 2015] shows that partial verification is essentially useless in truthful mechanism design. Hence, verification should be exact, i.e., it should forbid even negligible deviations from the truth, at least for some misreports. So, most research has focused on the power of exact verification schemes with either limited or costly verification.

Caragiannis et al. [2012] introduced *probabilistic verification* as a general framework. The idea is that any deviation from the truth is detectable with a probability depending on the distance of the misreport to the true type. They proved that virtually any allocation rule can be implemented with money and probabilistic verification if (i) the detection probability is positive for all agents and for any deviation from the truth; and (ii) each liar incurs a sufficiently large penalty. Instead, we use selective verification and the reasonable penalty of ignoring misreports, and we verify only a small subset of agents instead of almost all of them.

Selective verification is similar to *costly verification*, introduced by Townsend [1979]. However, the model of [Townsend 1979] (and other similar models) involve a single agent and allow for monetary transfers between the agent and the mechanism. As for mechanisms without monetary transfers, Glazer and Rubinstein [2004] consider mechanisms that decide on a request based on some claims about an agent's type. The mechanism can verify few of these claims and aims to maximize the probability that the request is accepted iff it is justified by the agent's true type. [Ben-Porath et al. 2014] and [Erlanson and Kleiner 2015] also consider mechanisms without money and with costly verification, where the objective is to maximize the total agent utility minus the verification cost. Ben-Porath et al. [2014] study truthful allocations of an indivisible good without money, and Erlanson and Kleiner [2015] study Bayesian incentive compatible mechanisms without money that choose among two outcomes. Although their approach is conceptually similar to ours, our setting and our mechanisms are much more general, we resort to approximate mechanisms (rather than exact ones) and treat the verification cost as a separate efficiency criterion (instead of including it in the social objective). Similar to the mechanisms presented in this work have been used by [Fotakis et al. 2015a] to extend the setting of Hartline and Roughgarden [2008], where they aim to maximize the social welfare minus the payments charged for truthful single-unit and multi-unit auctions.

A significant amount of previous work on mechanism design with verification either characterizes the optimal mechanism (see e.g., [Sher and Vohra 2014]) or shows that mechanisms with money and verification achieve better approximation guarantees than mechanisms without verification (see e.g., [Auletta et al. 2009; Krysta and Ventre 2015]). To the best of our knowledge, our work is the first where truthful mechanisms with selective verification (instead of "one-sided" verification applied to all agents with positive utility, see e.g., [Fotakis et al. 2014; Fotakis and Tzamos 2013b; Pountourakis and Schäfer 2014]) are shown to achieve best possible approximation guarantees for the general domain of Utilitarian Voting and for Combinatorial Public Project.

From a technical viewpoint, the idea of partial allocation in approximate mechanism design without money has been applied successfully in [Cole et al. 2013]. However, this idea alone (i.e., without verification) cannot achieve any strong approximation guarantees for the social welfare in general domains, such as Utilitarian Voting and Combinatorial Public Project. Moreover, our motivation for using the exponential mechanism with selective verification came from [Nissim et al. 2012; Huang and Kannan 2012], due their tradeoffs between the approximation guarantee and the probability of the gap mechanism (resp. amount of payments) required for truthfulness.

2. NOTATION AND PRELIMINARIES

For any integer $m \geq 1$, we let $[m] \equiv \{1, \dots, m\}$. For an event E, $\mathbb{P}\mathrm{r}[E]$ denotes the probability of E. For a random variable X, $\mathbb{E}[X]$ is the expectation of X. For a finite set S, $\Delta(S)$ is the unit simplex that includes all probability distributions over S. For a vector $\vec{x} = (x_1, \dots, x_m)$ and some $j \in [m]$, $\vec{x}_{-j}$ is $\vec{x}$ without x_j. For a nonempty $S \subseteq [m]$, $\vec{x}_S = (x_j)_{j \in S}$ is the projection of $\vec{x}$ to S. For vectors $\vec{x}$ and $\vec{y}$, $\vec{x}+\vec{y} = (x_1+y_1, \dots, x_m+y_m)$ denotes their coordinate-wise sum. For a vector $\vec{x}$ and an $\ell \geq 0$, $\vec{x}^\ell = (x_1^\ell, \dots, x_m^\ell)$ is the coordinate-wise power of $\vec{x}$ and $\|\vec{x}\|_\ell = (\sum_{j=1}^m |x_j|^\ell)^{1/\ell}$ is the ℓ-norm of $\vec{x}$. For brevity, we let $|\vec{x}| \equiv \|\vec{x}\|_1$. Moreover, $\|\vec{x}\|_\infty = \max_{j \in [m]}\{x_j\}$ is the infinity norm of $\vec{x}$.

Agent Valuations. We consider a set N of n strategic agents with private preferences over a set O of outcomes. We focus on combinatorial problems, assume that O is finite and let $m \equiv |O|$ be the number of different outcomes. The preferences of each agent i are given by a *valuation function* or *type* $\vec{x}_i : O \to \mathbb{R}_{\geq 0}$ that i seeks to maximize. The set of possible valuations is the *domain* $D = \mathbb{R}^m_{\geq 0}$. We usually regard each valuation as a vector $\vec{x}_i = (x_i(j))_{j \in [m]}$, where $x_i(j)$ is i's valuation for outcome j. A valuation profile is a tuple $\vec{x} = (\vec{x}_1, \dots, \vec{x}_n)$ consisting of the agent valuations. Given a valuation profile $\vec{x}$, $\vec{w}(\vec{x}) = \vec{x}_1 + \cdots + \vec{x}_n$ is the vector of the total valuation, or simply, of the *weight*, for each outcome. We usually write $\vec{w}$, instead of $\vec{w}(\vec{x})$, when $\vec{x}$ is clear from the context.

Allocation Rules. A (randomized) *allocation rule* $f : D^n \to \Delta(O)$ maps each valuation profile to a probability distribution over O. To allow for the exclusion of some agents from f, we assume that f is well defined for any number of agents $n' \leq n$. We regard the probability distribution of f on input $\vec{x}$ as a vector $f(\vec{x}) = (f_j(\vec{x}))_{j \in [m]}$, where $f_j(\vec{x})$ is the probability of outcome j. Then, the expected utility of agent i from $f(\vec{x})$ is equal to the dot product $\vec{x}_i \cdot f(\vec{x})$. An allocation rule is *constant* if for all valuation profiles $\vec{x}$ and $\vec{y}$, $f(\vec{x}) = f(\vec{y})$, i.e., the allocation of f is independent of the valuation profile. E.g., the uniform allocation, which selects each outcome with probability $1/m$, is constant.

An allocation rule f achieves *full allocation* if for all $\vec{x}$, $|f(\vec{x})| = 1$, and *partial allocation* if $|f(\vec{x})| < 1$, for some $\vec{x}$. A full allocation rule always outputs an outcome $o \in O$, while a partial allocation rule may also output an artificial (or *null*) outcome not in O. We assume that all agents have valuation 0 for the null outcome.

Two desirable properties of our allocation rules are strong anonymity and scale invariance. An allocation rule f is *scale invariant* if for any valuation profile $\vec{x}$ and any $\alpha \in \mathbb{R}_{>0}$, $f(\alpha\vec{x}) = f(\vec{x})$, i.e., scaling all valuations by α does not change the probability distribution of f. An allocation rule f is *strongly anonymous* if $f(\vec{x})$ depends only on the vector $\vec{w}(\vec{x})$ with outcome weights. Formally, for all valuation profiles $\vec{x}$ and $\vec{y}$ (possibly with a different number of agents) with $\vec{w}(\vec{x}) = \vec{w}(\vec{y})$, $f(\vec{x}) = f(\vec{y})$. Hence, a strongly anonymous rule can be regarded as a one-agent allocation rule $f : D \to \Delta(O)$. Next, all allocation rules (and mechanisms) are strongly anonymous, unless stated otherwise.

Approximation Guarantee. The social efficiency of an allocation rule f is evaluated by an objective function $g : D^n \times O \to \mathbb{R}_{\geq 0}$. We mostly consider the objective of *social welfare*, where we seek to maximize $\sum_{i=1}^{n} \vec{x}_i \cdot f(\vec{x})$. The optimal social welfare of a valuation profile $\vec{x}$ is $\| \sum_{i=1}^{n} \vec{x}_i \|_\infty$. An allocation rule f has *approximation ratio* $\rho \in (0, 1]$ (resp. *additive error* $\delta > 0$) if for all valuation profiles $\vec{x}$, $\sum_{i=1}^{n} \vec{x}_i \cdot f(\vec{x}) \geq \rho \, \|\sum_{i=1}^{n} \vec{x}_i\|_\infty$ (resp. $\sum_{i=1}^{n} \vec{x}_i \cdot f(\vec{x}) \geq \|\sum_{i=1}^{n} \vec{x}_i\|_\infty - \delta$).

Voluntary Participation and MIDR. An allocation rule f satisfies *voluntary participation* (or simply, *participation*) if for any agent i and any valuation profile $\vec{x}$, $\vec{x}_i \cdot f(\vec{x}) \geq \vec{x}_i \cdot f(\vec{x}_{-i})$, i.e., i's utility does not decrease if she participates in the mechanism. For some small $\varepsilon \in (0, 1]$, f satisfies *ε-participation* if for any agent i and any valuation profile $\vec{x}$, $\vec{x}_i \cdot f(\vec{x}) \geq \varepsilon \, \vec{x}_i \cdot f(\vec{x}_{-i})$.

An allocation rule f is *maximal in distributional range* (MIDR) if there exist a range $Z \subseteq \{\vec{z} \in \mathbb{R}^m_{\geq 0} : |\vec{z}| \leq 1\}$ of (possibly partial) allocations and a function $h : Z \to \mathbb{R}$ such that for all valuation profiles $\vec{x}$, $f(\vec{x}) = \arg\max_{\vec{z} \in Z} \sum_{i=1}^{n} \vec{x}_i \cdot \vec{z} + h(\vec{z})$ (see e.g., [Lavi and Swamy 2011]). We first show that MIDR is a sufficient condition for participation. We can also show that for scale invariant and strongly anonymous continuous allocation rules, MIDR is a necessary condition for participation.

LEMMA 2.1. *Let f be any MIDR allocation rule. Then, f satisfies participation.*

PROOF. Let i be any agent. Since the allocation rule f is MIDR, we obtain that

$$\sum_{j=1}^{n} \vec{x}_j \cdot f(\vec{x}) + h(f(\vec{x})) \geq \sum_{j=1}^{n} \vec{x}_j \cdot f(\vec{x}_{-i}) + h(f(\vec{x}_{-i}))$$

$$\sum_{j \neq i} \vec{x}_j \cdot f(\vec{x}_{-i}) + h(f(\vec{x}_{-i})) \geq \sum_{j \neq i} \vec{x}_j \cdot f(\vec{x}) + h(f(\vec{x}))$$

We apply the MIDR condition first to $\vec{x}$ and next to $\vec{x}_{-i}$. Summing up the two inequalities, we obtain that $\vec{x}_i \cdot f(\vec{x}) \geq \vec{x}_i \cdot f(\vec{x}_{-i})$, i.e., the participation condition. □

LEMMA 2.2. *For any scale invariant and strongly anonymous continuous allocation rule f that satisfies participation, there is a range Z of (possibly partial) allocations such that $f(\vec{x}) = \arg\max_{\vec{z} \in Z} \vec{x} \cdot \vec{z}$.*

3. MECHANISMS WITH SELECTIVE VERIFICATION AND BASIC PROPERTIES

A *mechanism with selective verification* F takes as input a reported valuation profile $\vec{y}$ and has *oracle access* to a binary *verification vector* $\vec{s} \in \{0, 1\}^n$, with $s_i = 1$ if agent i has reported truthfully $\vec{y}_i = \vec{x}_i$, and $s_i = 0$ otherwise. We assume that F verifies an agent i through a *verification oracle* ver that on input i, returns $\mathrm{ver}(i) = s_i$. So, we regard a mechanism with verification as a function $F : D^n \times \{0, 1\}^n \to \Delta(O)$. We highlight that although the entire vector $\vec{s}$ appears as a parameter of F, for notational convenience, the outcome of F actually depends on few selected coordinates of $\vec{s}$. We let $V(\vec{y}) \subseteq N$, or simply V, be the *verification set*, i.e., the set of agents verified by F on input $\vec{y}$. As for allocation rules, we treat the probability distribution of F over outcomes as an m-dimensional vector and assume that F is well defined for any number of agents $n' \leq n$.

We start from an allocation rule f and devise a mechanism F that motivates truthful reporting by selective verification. We say that a mechanism F with selective verification is *recursive* if there is an allocation rule f such that F operates as follows: on a valuation profile $\vec{y}$, F selects an outcome o, with probability $f_o(\vec{y})$, and a verification set $V(\vec{y})$, and computes the set $L = \{i \in V(\vec{y}) : \mathrm{ver}(i) = 0\}$ of misreporting agents in $V(\vec{y})$. If $L = \emptyset$, F returns o. Otherwise, F recurses on $\vec{y}_{-L}$. Our mechanisms are recursive, except for Partial Power (Section 8), which adopts a slightly different reaction to $L \neq \emptyset$.

Given an allocation rule f, we say that a mechanism with verification F is an *extension* of f if for all valuation profiles $\vec{x}$, $F(\vec{x}, \vec{1}) = f(\vec{x})$. Namely, F behaves exactly as f given that all agents report truthfully. For the converse, given a mechanism F, we say that F *induces* an allocation rule f if for all $\vec{x}$, $f(\vec{x}) = F(\vec{x}, \vec{1})$. For clarity, we refer to mechanisms with selective verification simply as *mechanisms*, and denote them by uppercase letters, and to allocation rules simply as *rules* or *algorithms*, and denote them by lowercase letters. A mechanism F has a property of an allocation rule (e.g., scale invariance, partial or full allocation, participation, approximation ratio) iff the induced rule f has this property.

A mechanism F is ε*-truthful*, for some $\varepsilon \in (0, 1]$, if for any agent i, for all valuation pairs $\vec{x}_i$ and $\vec{y}_i$ and for any reported valuation $\vec{y}_{-i}$ and verification vector $\vec{s}_{-i}$,

$$\vec{x}_i \cdot F((\vec{y}_{-i}, \vec{x}_i), (\vec{s}_{-i}, 1)) \geq \varepsilon\, \vec{x}_i \cdot F((\vec{y}_{-i}, \vec{y}_i), (\vec{s}_{-i}, 0))$$

A mechanism F is *truthful* if it is 1-truthful. Namely, no matter the reported valuations of the other agents and whether they report truthfully or not, the expected utility of agent i is maximized if she reports truthfully.

Immunity and Obliviousness. A remarkable property of our mechanisms is *immunity*, namely that they ignore any misreporting agents and let their outcome depend on the valuations of truthful agents only. Formally, a mechanism F is *immune* if for all reported valuations $\vec{y}$ and verification vectors $\vec{s}$, $F(\vec{y}, \vec{s}) = F(\vec{y}_{T(\vec{s})}, (1, \ldots, 1))$, with the equality referring to the probability distribution of F, where $T(\vec{s}) = \{i \in N : s_i = 1\}$ is the set of truthful agents in $\vec{y}$. Next, we simply use T, instead of $T(\vec{s})$.

A mechanism with selective verification F is *oblivious* (to the declarations of misreporting agents not verified) if for all valuation profiles $\vec{y}$ and verification vectors $\vec{s}$, with $L = N \setminus T(\vec{s})$, and any outcome o,

$$\mathbb{P}\mathrm{r}[F(\vec{y}, \vec{s}) = o \,|\, V(\vec{y}) \cap L = \emptyset] = \mathbb{P}\mathrm{r}[F(\vec{y}_{-L}, \vec{1}) = o] \qquad (1)$$

I.e., if the misreporting agents are not included in the verification set, they do not affect the probability distribution of F (see also [Fotakis and Tzamos 2013b]). By induction on the number of agents, we show that obliviousness is sufficient for immunity.

LEMMA 3.1. *Let F be any oblivious recursive mechanism with selective verification. Then, F is immune.*

PROOF. We fix a valuation profile $\vec{y}$ and a verification vector $\vec{s}$. We show that for any outcome $o \in O$, $\mathbb{P}\mathrm{r}[F(\vec{y}, \vec{s}) = o] = \mathbb{P}\mathrm{r}[F(\vec{y}_T, \vec{1}) = o]$, where $T \subseteq N$ is the set of truthful agents in $\vec{y}$. The proof is by induction on the number of agents N.

If $N = \emptyset$, the statement is obvious. So, we assume inductively that the statement holds for every proper subset of N. Let $L = N \setminus T$ be the set of misreporting agents in

$\vec{y}$ and let V be the verification set of F on input $\vec{y}$. Then,

$$\mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o] = \sum_{L' \subseteq L} \mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o \,|\, V \cap L = L'] \, \mathbb{P}\mathrm{r}[V \cap L = L'] \tag{2}$$

We have that $\mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o \,|\, V \cap L = \emptyset] = \mathbb{P}\mathrm{r}[F(\vec{y}_T, \vec{1}) = o]$, by (1), since F is oblivious. If V includes a non-empty set $L' = V \cap L$, since F is recursive, it ignores their declarations and recurses on $\vec{y}_{-L'}$. Therefore, for all $\emptyset \neq L' \subseteq L$,

$$\mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o \,|\, V \cap L = L'] = \mathbb{P}\mathrm{r}[F(\vec{y}_{-L'}, \vec{s}_{-L'}) = o] = \mathbb{P}\mathrm{r}[F(\vec{y}_T, \vec{1}) = o]\,,$$

where the last equality follows from the induction hypothesis, because the set of agents in $\vec{y}_{-L'}$ is a proper subset of the set of agents in $\vec{y}$. Therefore, using that $\mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o \,|\, V \cap L = L'] = \mathbb{P}\mathrm{r}[F(\vec{y}_T, \vec{1}) = o]$, for all $L' \subseteq L$, in (2), we obtain that $\mathbb{P}\mathrm{r}[F(\vec{y},\vec{s}) = o] = \mathbb{P}\mathrm{r}[F(\vec{y}_T, \vec{1}) = o]$, i.e., that F is immune. □

Immunity, Participation and Truthfulness. We next show that immunity and participation imply truthfulness (note that the converse may not be true, since a truthful mechanism with verification does not need to be immune). Then, by Lemma 2.1 and Lemma 3.1, we can focus on MIDR allocation rules for which the outcome and the verification set can be selected in an oblivious way.

LEMMA 3.2. *For any $\varepsilon \in (0,1]$, if a mechanism with selective verification F is immune and satisfies ε-participation, then F is ε-truthful.*

PROOF. Since F is immune, for any agent i, for any valuation pair $\vec{x}_i$ and $\vec{y}_i$ and for all reported valuations $\vec{y}_{-i}$ and verification vectors $\vec{s}_{-i}$,

$$F((\vec{y}_{-i}, \vec{x}_i), (\vec{s}_{-i}, 1)) = F(((\vec{y}_T)_{-i}, \vec{x}_i), (1, \ldots, 1)) \quad \text{and}$$
$$F((\vec{y}_{-i}, \vec{y}_i), (\vec{s}_{-i}, 0)) = F((\vec{y}_T)_{-i}, (1, \ldots, 1))$$

We assume here that $\vec{x}_i$ is i's true type and $\vec{y}_i \neq \vec{x}_i$ is a misreport. Moreover, using that f (i.e., the rule induced by F on truthful reports) satisfies ε-participation, we get that:

$$\vec{x}_i \cdot F(((\vec{y}_T)_{-i}, \vec{x}_i), (1, \ldots, 1)) \geq \varepsilon\, \vec{x}_i \cdot F((\vec{y}_T)_{-i}, (1, \ldots, 1))$$

Combining the three equations above, we conclude the proof of the lemma. □

Quantifying Verification. Focusing on truthful mechanisms with verification, where the agents do not have any incentive to misreport, we bound the amount of verification when the agents are truthful (similarly to the definition of the approximation ratio of F as the approximation ratio of the induced allocation rule f). For a truthful mechanism F, this is exactly the amount of verification required so that F motivates truthfulness.

Given a mechanism with selective verification F, its *worst-case verification* is $\mathrm{Ver}(F) \equiv \max_{\vec{x} \in D^n} |V(\vec{x})|$, i.e., the maximum number of agents verified by F in any truthful valuation profile. If F is randomized, its *expected verification* is $\mathbb{E}[\mathrm{Ver}(F)] \equiv \max_{\vec{x} \in D^n} \mathbb{E}[|V(\vec{x})|]$, where expectation is over all coin tosses of the mechanism.

4. FACILITY LOCATION MECHANISMS WITH SELECTIVE VERIFICATION

As a proof of concept, we apply mechanisms with verification to k-Facility Location. In Facility Location problems, we are given some metric space (M, d), where M is a finite set of points and d is a metric distance function. The possible outcomes are all subsets of k locations in M. Each agent i has a preferred location $t_i \in M$ and her "valuation" for outcome C is $\vec{x}_i(C) = -d(t_i, C)$, i.e., minus the distance of her preferred location to the nearest facility in C. So, each agent i aims at minimizing $d(t_i, C)$. The mechanism F takes a profile $\vec{z} = (z_1, \ldots, z_n)$ of reported locations as an input. Using access to a verification oracle, F maps $\vec{z}$ to a set C of k facility locations.

MECHANISM 1: The Power Mechanism $\mathrm{Pow}^\ell(\vec{x}, \vec{s})$

let N be the set of the remaining agents and let $L \leftarrow \emptyset$
pick an outcome $j \in O$ and a tuple $\vec{t} \in N^\ell$ with probability proportional to
  the value of the term $x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j)$
for each agent $i \in \vec{t}$ **do**
  if $\mathrm{ver}(i) \neq 1$ **then** $L \leftarrow L \cup \{i\}$
end
if $L \neq \emptyset$ **then return** $\mathrm{Pow}^\ell(\vec{x}_{-L}, \vec{s}_{-L})$
else return outcome j

Maximum Cost. To minimize $\max_{i\in N}\{d(t_i, F(\vec{t}, \vec{1}))\}$, i.e, the maximum distance of any agent to the nearest facility, we use the 2-approximate Greedy algorithm for k-Center (see e.g., [Williamson and Shmoys 2011, Sec. 2.2]). On input $\vec{z}$, Greedy first allocates a facility to an arbitrary agent. As long as $|C| < k$, the next facility is allocated to the agent i maximizing $d(z_i, C)$. Verification inspects the reported location z_i of each agent i allocated a facility. If all these agents are truthful, we place the k facilities at their locations. Otherwise, we exclude any liars among them and recurse on the remaining agents. The following theorem summarizes the properties of Greedy with selective verification. Recall that there are no deterministic truthful mechanisms (without verification) that place $k \geq 2$ facilities in tree metrics and achieve a bounded (in terms of n and k) approximation ratio (see [Fotakis and Tzamos 2014]).

THEOREM 4.1. *The Greedy mechanism with verification for k-Facility Location is truthful and immune, is 2-approximate for the maximum cost and verifies k agents.*

Social Cost. To minimize $\sum_{i=1}^n d(t_i, F(\vec{t}, \vec{1}))$, i.e, the total cost of the agents, we use the Proportional mechanism [Lu et al. 2010], which is $\Theta(\ln k)$-approximate [Arthur and Vassilvitskii 2007]. Proportional first allocates a facility to an agent chosen uniformly at random. As long as $|C| < k$, agent i is allocated the next facility with probability proportional to $d(z_i, C)$. Verifying the reported location of every agent that is allocated a facility, we obtain that:

THEOREM 4.2. *Proportional with verification for k-Facility Location is truthful and immune, is $\Theta(\ln k)$-approximate for the social cost and verifies k agents.*

5. THE POWER MECHANISM WITH SELECTIVE VERIFICATION

In this section, we present the Power mechanism, a recursive mechanism with verification that approximates the social welfare in the general domain of Utilitarian Voting.

Power with parameter $\ell \geq 0$ (or Pow^ℓ, in short, see also Mechanism 1) is based on a strongly anonymous and scale invariant allocation rule that assigns probability proportional to the weight of each outcome raised to ℓ. Hence, for each valuation profile $\vec{x}$, the outcome of Pow^ℓ is determined by the weight vector $\vec{w} = \sum_{i=1}^n \vec{x}_i$. Assuming that all agents are truthful, Pow^ℓ results in outcome j with probability $w_j^\ell / \sum_{q=1}^m w_q^\ell$, i.e., proportional to w_j^ℓ (note that for $\ell = 0$, we get the uniform allocation, while for $\ell = \infty$, the outcome of maximum weight gets probability 1). To implement this allocation rule with low verification, we observe that each term w_j^ℓ can be expanded in n^ℓ terms:

$$w_j^\ell = \left(\sum_{i\in N} x_i(j)\right)^\ell = \sum_{\vec{t}\in N^\ell} x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j) \tag{3}$$

Hence, choosing an outcome j and a tuple $\vec{t} \in N^\ell$ with probability proportional to $x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j)$, we end up with outcome j with probability proportional to w_j^ℓ.

Example 5.1. Let $n = 3$, $\ell = 2$ and $w_j = x_1 + x_2 + x_3$ (in this example, we omit j from x's, for clarity). In (3), we expand w_j^2 in $3^2 = 9$ terms as follows $w_j^2 = (x_1 + x_2 + x_3)^2 = x_1x_1 + x_1x_2 + x_1x_3 + x_2x_1 + x_2x_2 + x_2x_3 + x_3x_1 + x_3x_2 + x_3x_3$. Hence, $\vec{t} \in \{1, 2, 3\} \times \{1, 2, 3\}$. Given that outcome j is chosen, each of these terms (and the corresponding tuple $\vec{t}$) is selected with probability proportional to its value. E.g., x_1x_2 and $\vec{t} = (1, 2)$ are selected with probability x_1x_2/w_j^2 and x_3x_3 and $\vec{t} = (3, 3)$ are selected with probability x_3^2/w_j^2.

In Mechanism 1, to select an outcome j and a tuple $\vec{t} \in N^\ell$ efficiently, we sample from (3) in $O(m + n\ell)$ steps. Specifically, we first select outcome j with probability $w_j^\ell/|\vec{w}^\ell|$. Then, conditional on j, we select an agent i in each position of $\vec{t}$ independently with probability $x_i(j)/w_j$. Each tuple $\vec{t}$ is picked with probability $x_{t_1}(j)\cdots x_{t_\ell}(j)/|\vec{w}^\ell|$.

The verification set of Pow^ℓ consists of the agents in $\vec{t}$, i.e., of at most ℓ agents. Moreover, due to its proportional nature, Pow^ℓ is oblivious to the declarations of any misreporting agents not verified, and thus, immune. The following formalizes this intuition.

LEMMA 5.2. *For any $\ell \geq 0$, Pow^ℓ is immune and verifies at most ℓ agents.*

PROOF. If all agents are truthful, Pow^ℓ picks an outcome j and a tuple $\vec{t} \in N^\ell$ and returns j after verifying all agents in $\vec{t}$. Since there are at most ℓ different agent indices in $\vec{t}$, the verification of Pow^ℓ is at most ℓ.

We next show that Pow^ℓ is oblivious to the declarations of any misreporting agents not verified. Specifically, we show that for all valuation profiles $\vec{x}$ and verification vectors $\vec{s}$, with $L = N \setminus T(\vec{s})$, and all outcomes j,

$$\mathbb{P}\mathrm{r}[\mathrm{Pow}^\ell(\vec{x}, \vec{s}) = j \,|\, V(\vec{x}) \cap L = \emptyset] = \mathbb{P}\mathrm{r}[\mathrm{Pow}^\ell(\vec{x}_{-L}, \vec{1}) = j] \tag{4}$$

To prove (4), we observe that for any outcome j, the condition $V(\vec{x}) \cap L = \emptyset$ implies that $\mathrm{Pow}^\ell(\vec{x}, \vec{s})$ selects only terms $x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j)$ with truthful agents in $T(\vec{s})$. Every term with some valuation $x_{t'}(j)$ of a misreporting agent $t' \in L$ is excluded, since $V(\vec{x}) \cap L = \emptyset$ implies that $\vec{t} \in T(\vec{s})^\ell$. Therefore, for any outcome j, $\mathrm{Pow}^\ell(\vec{x}, \vec{s})$, conditional on $V(\vec{x}) \cap L = \emptyset$, and $\mathrm{Pow}^\ell(\vec{x}_{-L}, \vec{1})$ have exactly the same set of "allowable" terms from which they select $x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j)$ and $\vec{t}$. In both, every such term is selected with probability proportional to its value, i.e., with identical probability. Taking all outcomes into account, we obtain that the distribution of $\mathrm{Pow}^\ell(\vec{x}, \vec{s})$, conditional on $V(\vec{x}) \cap L = \emptyset$, and the distribution of $\mathrm{Pow}^\ell(\vec{x}_{-L}, \vec{1})$ are identical.

Therefore, Pow^ℓ is oblivious. Since it is also recursive, Lemma 3.1 implies that Pow^ℓ is immune. □

We next establish the approximation ratio of Pow^ℓ for the objective of social welfare. The intuition is that as ℓ increases from 0 to ∞, the probability distribution of Pow^ℓ sharpens from the uniform allocation, where each outcome is selected with probability $1/m$, to the optimal allocation. The rate of this transition determines the approximation ratio and is quantified in the following lemma.

LEMMA 5.3. *For any $\ell \geq 0$, Pow^ℓ is $m^{-1/(\ell+1)}$-approximate for the social welfare.*

PROOF. Let us fix any valuation profile $\vec{x}$ and let $\vec{w} \equiv \vec{w}(\vec{x})$ be the outcome weights in $\vec{x}$. For the approximation ratio, we can assume that all agents are truthful. So, we let $\mathrm{Pow}^\ell(\vec{w}) \equiv \mathrm{Pow}^\ell(\vec{x}, \vec{1})$, for convenience.

The optimal social welfare is $\|\vec{w}\|_\infty$. The expected social welfare of the mechanism is $\vec{w} \cdot \mathrm{Pow}^\ell(\vec{w}) = |\vec{w}^{\ell+1}|/|\vec{w}^\ell|$. So the approximation ratio of Pow^ℓ is equal to:

$$\frac{|\vec{w}^{\ell+1}|}{|\vec{w}^\ell|\,\|\vec{w}\|_\infty} = \frac{(\|\vec{w}\|_{\ell+1})^{\ell+1}}{(\|\vec{w}\|_\ell)^\ell\,\|\vec{w}\|_\infty} = \left(\frac{\|\vec{w}\|_{\ell+1}}{\|\vec{w}\|_\ell}\right)^\ell \frac{\|\vec{w}\|_{\ell+1}}{\|\vec{w}\|_\infty}$$

Using that $\|\vec{w}\|_\infty \leq \|\vec{w}\|_{\ell+1}$ and that $\|\vec{w}\|_\ell \leq m^{\left(\frac{1}{\ell}-\frac{1}{\ell+1}\right)}\|\vec{w}\|_{\ell+1} = m^{1/\ell(\ell+1)}\|\vec{w}\|_{\ell+1}$, we obtain that the approximation ratio of Pow^ℓ is at least $m^{-1/(\ell+1)}$. □

Unfortunately, Power does not satisfy participation. For a simple example, we consider $m = 2$ outcomes and $n = 2$ agents with valuations $\vec{x}_1 = (1, 0)$ and $\vec{x}_2 = (3/4, 1/4)$. Then, agent 2 prefers outcome 1, but her participation decreases its probability from 1, when agent 1 is alone, to something less than 1, when both agents participate.

However, Power satisfies participation approximately. This follows from the fact that the Partial Power allocation is MIDR, by definition, and essentially a smoothed version of Power. Using that the probabilities that each outcome is selected in Partial Power and in Power are close to each other and the fact that Partial Power satisfies participation, we obtain the following.

LEMMA 5.4. *For any $\ell \geq 0$, Pow^ℓ satisfies $m^{-1/(\ell+1)}$-participation.*

Since Pow^ℓ is immune and satisfies $m^{-1/(\ell+1)}$-participation, Lemma 3.2 implies that Pow^ℓ is $m^{-1/(\ell+1)}$-truthful. Using $\ell = \ln m/\varepsilon$ in the three lemmas above, we obtain that:

THEOREM 5.5. *For any $\varepsilon > 0$, Pow^ℓ with $\ell = \ln m/\varepsilon$ is immune and $(1-\varepsilon)$-truthful, has worst-case verification $\ln m/\varepsilon$, and achieves an approximation ratio of $(1-\varepsilon)$ for the objective of social welfare.*

6. LOGARITHMIC VERIFICATION IS BEST POSSIBLE

In this section, we describe a random family of instances where truthfulness requires a logarithmic expected verification. Thus, we show that the verification bound of Theorem 5.5 is essentially best possible.

THEOREM 6.1. *Let F be randomized truthful mechanism that achieves a constant approximation ratio for any number of agents n and any number of outcomes m. Then, F needs expected verification $\Omega(\log m)$.*

PROOF SKETCH. We consider m outcomes and m disjoint groups of agents. Each group has a large number ν of agents. An agent in group j has valuation either 1 or δ for outcome j, where $\delta > 0$ is tiny, and valuation 0 for any other outcome. In each group j, the probability that k agents, $0 \leq k \leq \nu$, have valuation 1 for outcome j is $2^{-(k+1)}$. The expected maximum social welfare of such instances is $\Theta(\log m)$.

We next focus on a group j of agents and fix $\vec{x}_{-j}$, i.e., the agent declarations in all other groups. We can show that it is wlog. to assume that the probability of outcome j depends only on the number of agents in group j that declare 1 for j. Thus, the mechanism induces a sequence of probabilities $p_0, p_1, \ldots, p_k, \cdots$, where p_k is the probability of outcome j, given that the number of agents that declare 1 for j is k. By truthfulness, if k agents declare 1 for outcome j, we need to verify each of them with probability at least $p_k - p_{k-1}$. Otherwise, an agent with valuation δ can declare 1 and improve her expected utility. Therefore, for any fixed $\vec{x}_{-j}$, when k agents declare 1 for outcome j, we need an expected verification of at least $k(p_k - p_{k-1})$ for agents in group j.

Assuming truthful reporting and taking the expectation over the number of agents in group j with valuation 1, we find that the expected verification for agents in group j is at least half the mechanism's expected welfare from group j, conditional on $\vec{x}_{-j}$, minus half the probability of outcome j, conditional on $\vec{x}_{-j}$. Removing the conditioning on $\vec{x}_{-j}$ and summing up over all groups j, we conclude that the expected verification is at least half the mechanism's expected welfare minus $1/2$. Since the mechanism is $O(1)$-approximate, there are instances where the expected verification is $\Omega(\log m)$. □

7. CHARACTERIZATION OF STRONGLY ANONYMOUS MECHANISMS

Next, we characterize the class of scale invariant and strongly anonymous truthful mechanisms that verify $o(n)$ agents. The characterization is technically involved and consists of four main steps. We first prove that these rules are continuous.

LEMMA 7.1. *Let f be any scale invariant and strongly anonymous allocation rule. If f is discontinuous, every truthful extension F of f needs to verify $\Omega(n)$ agents in expectation, for arbitrarily large n.*

PROOF SKETCH. First, we prove that if f has a discontinuity, there are $\Omega(n)$ agents that have a very small valuation $\delta > 0$ and can change the allocation by a constant factor, independent of n and δ. Next, we focus on any truthful extension F of f and show that for every agent i that has the ability to change the allocation by a constant factor, the probability that F verifies i should be at least a constant, say ζ, due to truthfulness. Therefore, the expected verification of F is at least $\zeta \times \Omega(n) = \Omega(n)$. □

Therefore, if a truthful mechanism F verifies $o(n)$ agents and induces a scale invariant and strongly anonymous allocation rule f, then f needs to be continuous. We can prove that such an allocation f satisfies participation. Then, by Lemma 2.2, we obtain that such an allocation rule f is MIDR. Moreover, we can show that any full allocation and MIDR rule f is either constant, i.e., its probability distribution does not depend on the input $\vec{x}$, or has a discontinuity at $\vec{1}$. Thus, we obtain the following:

THEOREM 7.2. *Let F be any truthful mechanism that verifies $o(n)$ agents, is scale invariant and strongly anonymous and achieves full allocation. Then, F induces a constant allocation rule.*

8. THE PARTIAL POWER MECHANISM WITH SELECTIVE VERIFICATION

The Power mechanism, in Section 5, escapes the characterization of Theorem 7.2 by relaxing participation (and thus, truthfulness). Next, we present Partial Power which escapes the characterization by relaxing full allocation. Thus, Partial Power results in some outcome in O with probability less than 1, and with the remaining probability, it results in an artificial *null* outcome for which all agents have valuation 0.

Lemma 2.2 implies that social welfare maximization is essentially necessary for participation. The proof of Theorem 7.2 implies that maximizing the social welfare over $\Delta(O)$ results in discontinuous mechanisms that need $\Omega(n)$ verification (e.g., let $m = 2$ and consider welfare maximization for weight vectors $(1, 1+\epsilon)$ and $(1, 1-\epsilon)$). Hence, for Partial Power, we optimize over a smooth surface that is close to $\Delta(O)$, but slightly curved towards the corners, so that the resulting welfare maximizers are continuous.

More formally, aiming at a strongly anonymous allocation rule f that satisfies $f(\vec{w}) = \arg\max_{\vec{z}\in Z} \vec{w}\cdot\vec{z}$, for all weight vectors $\vec{w} \in \mathbb{R}^m_{\geq 0}$, we consider welfare maximization over the family of sets

$$Z_{\ell,r} = \left\{ \vec{z} \in \mathbb{R}^m_{\geq 0} : \|\vec{z}\|_{1+1/\ell} \leq (1-1/r)m^{-1/(\ell+1)} \right\}, \text{ for all integers } \ell, r \geq 1 .$$

Using the fact that the range $Z_{\ell,r}$ is smooth and strictly convex, we can show that for any $\ell, r \geq 1$, the allocation rule $f^{(\ell,r)}(\vec{w})$ obtained as the solution to the optimization problem $\max_{\vec{z}\in Z_{\ell,r}} \vec{w}\cdot\vec{z}$ is:

MECHANISM 2: The Partial Power Mechanism $\mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s})$

1 pick r tuples $\vec{t}^{(1)},...,\vec{t}^{(r)} \in N^{\ell+1}$ with probability proportional to
2 the value of the term $\sum_{j\in O} x_{t_1}(j)x_{t_2}(j)\cdots x_{t_{\ell+1}}(j)$
3 **for each** $k \in \{1,...,r\}$ and agent $i \in \vec{t}^{(k)}$ **do**
4 | **if** $\mathrm{ver}(i) \neq 1$ **then return** $\perp$
5 **end**
6 with probability $1-\sum_j f_j^{(\ell,r)}(\vec{w})$ **return** null
7 pick an outcome $j \in O$ and a tuple $\vec{t} \in N^{\ell}$ with probability proportional to
8 the value of the term $x_{t_1}(j)x_{t_2}(j)\cdots x_{t_\ell}(j)$
9 **for each** agent $i \in \vec{t}$ **do**
10 | **if** $\mathrm{ver}(i) \neq 1$ **then return** $\perp$
11 **end**
12 **return** outcome j

$$f^{(\ell,r)}(\vec{w}) = \frac{(1-1/r)}{m^{1/(\ell+1)}} \cdot \frac{\vec{w}^{\ell}}{\|\vec{w}^{\ell}\|_{1+1/\ell}}$$

Essentially by definition, $f^{(\ell,r)}(\vec{w})$ is a continuous allocation, MIDR and satisfies participation. Moreover, we can show that for any $\ell \geq 1$, the partial allocation $f^{(\ell,r)}$ has approximation ratio $(1-1/r)m^{-1/(\ell+1)}$ for the social welfare.

We next show that there exists an immune extension $\mathrm{PartPow}^{\ell,r}$ of the allocation rule $f^{(\ell,r)}$ that uses reasonable verification. Thus, we establish that $\mathrm{PartPow}^{\ell,r}$ is truthful. To this end, we introduce Mechanism 2. Since $f^{(\ell,r)}$ is strongly anonymous, we consider below the weight vector $\vec{w} \equiv \vec{w}(\vec{x})$ instead of the valuation profile $\vec{x}$. If all agents are truthful, $\mathrm{PartPow}^{\ell,r}$ samples exactly from $f^{(\ell,r)}(\vec{w})$. In particular, assuming truthful reporting, steps 1-5 never result in $\perp$, step 6 outputs null with probability $1-|f^{(\ell,r)}(\vec{w})|$, and steps 7-12 work identically to Pow^{ℓ}, since given that the null outcome is not selected, each outcome j is chosen with probability proportional to w_j^{ℓ}.

The most interesting case is when some agents misreport their valuations. To achieve immunity, we need to ensure that the probability distribution is identical to the case where all misreporting agents are excluded from the mechanism. Similarly to Pow^{ℓ}, misreporting agents cannot affect the relative probabilities of each outcome. In $\mathrm{PartPow}^{\ell,r}$ however, they may affect the probability of the null outcome. Thus, $\mathrm{PartPow}^{\ell,r}$ is not oblivious and we cannot establish immunity through Lemma 3.1.

Immunity of $\mathrm{PartPow}^{\ell,r}$ is obtained through the special action $\perp$, triggered when verification reveals some misreporting agents. Then, $\mathrm{PartPow}^{\ell,r}$ needs to allocate appropriate probabilities to each outcome j and to the null outcome so that the unconditional probability distribution of $\mathrm{PartPow}^{\ell,r}$ is identical to $f^{(\ell,r)}(\vec{w}_T)$, where T is the set of truthful agents. Hence, when $\mathrm{PartPow}^{\ell,r}$ returns $\perp$, we verify all agents, compute the weight vector $\vec{w}_T$ for the truthful agents, and return each outcome j with probability:

$$p_j = \frac{f_j^{(\ell,r)}(\vec{w}_T) - \mathbb{P}\mathrm{r}[\mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s}) = j \mid \mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s}) \neq \perp]\,\mathbb{P}\mathrm{r}[\mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s}) \neq \perp]}{\mathbb{P}\mathrm{r}[\mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s}) = \perp]}$$

The null outcome is returned with probability $1-\sum_j p_j$. We highlight that these probabilities are chosen so that we cancel the effect of misreporting agents in the unconditional probability distribution of $\mathrm{PartPow}^{\ell,r}$ and achieve exactly the probability distribution $f^{(\ell,r)}(\vec{w}_T)$. Moreover, if the mechanism returns $\perp$, we verify all agents. So, it is always possible to compute there probabilities correctly.

The crucial step is to show that p_j's are always non-negative and their sum is at most 1. We employ steps 1-5 for this reason. These steps implement additional verification and ensure that $\mathbb{P}\mathrm{r}[\mathrm{PartPow}^{\ell,r}(\vec{x},\vec{s}) = \perp]$ is large enough for this property to hold.

THEOREM 8.1. *For every $\varepsilon > 0$, there exist integers $\ell, r \geq 1$, such that Partial Power is truthful, immune, $(1-\varepsilon)$-approximate for the objective of social welfare and verifies at most $O(\ln m/\varepsilon^2)$ agents in the worst case.*

9. THE EXPONENTIAL MECHANISM WITH SELECTIVE VERIFICATION

Next, we consider the Exponential mechanism (or Expo, for brevity) and show that it escapes the characterization of Section 7 by relaxing scale invariance. Expo is strongly anonymous and assigns a probability proportional to the exponential of the weight of each outcome. Specifically, for any profile $\vec{x}$, the outcome of Expo depends on $\vec{w} \equiv \sum_{i=1}^{n} \vec{x}_i$. If all agents are truthful, $\mathrm{Expo}^{\alpha}(\vec{w})$ results in outcome j with probability $e^{w_j/\alpha}/\sum_{q=1}^{m} e^{w_q/\alpha}$, i.e., proportional to $e^{w_j/\alpha}$, where $\alpha > 0$ is a parameter. As in Section 5, we expand every term $e^{w_j/\alpha}$ and verify only the agents in the tuple $\vec{t}$ corresponding to each term in the expansion (the sampling is implemented as in Section 5):

$$e^{w_j/\alpha} = \sum_{\ell=0}^{\infty} \frac{(w_j/\alpha)^{\ell}}{\ell!} = \sum_{\ell=0}^{\infty} \frac{\alpha^{-\ell}}{\ell!} \sum_{\vec{t} \in N^{\ell}} x_{t_1}(j) x_{t_2}(j) \cdots x_{t_\ell}(j) \tag{5}$$

The detailed description of Expo^{α} is similar to Mechanism 1, with the only difference that, in the second step, we pick an outcome $j \in O$, an integer $\ell \geq 0$ and a tuple $\vec{t} \in N^{\ell}$ with probability proportional to the value of the term $x_{t_1}(j) x_{t_2}(j) \cdots x_{t_\ell}(j)/(\alpha^{\ell} \ell!)$. The following summarizes the properties of Expo.

THEOREM 9.1. *For any $\alpha > 0$, $\mathrm{Expo}^{\alpha}(\vec{w})$ is immune and truthful, achieves an additive error of $\alpha \ln m$ for the social welfare and has expected verification $\|\vec{w}\|_\infty/\alpha$.*

PROOF SKETCH. Similarly to the proof of Lemma 5.2, we can show that Expo^{α} is oblivious (note that the allocation of Pow^{ℓ} is obtained from the allocation of Expo^{α} if we condition on a particular exponent ℓ). Then, immunity follows from Lemma 3.1, because Expo^{α} is a recursive mechanism. As for participation, the Exponential allocation is known to be MIDR with range $Z = \Delta(O)$ and function $h(\vec{z}) = -\alpha \sum_j z_j \ln z_j$ (see e.g., [Huang and Kannan 2012]). Therefore, by Lemma 2.1, Expo^{α} satisfies participation. Since it is also immune, Lemma 3.2 implies that Expo^{α} is truthful.

For the verification, (5) implies that when all agents are truthful, the number of agents verified, given that the selected outcome is j, follows a Poisson distribution with parameter $w_j/\alpha \leq \|\vec{w}\|_\infty/\alpha$. Therefore, the expected verification is at most $\|\vec{w}\|_\infty/\alpha$.

As for the approximation guarantee, the optimal social welfare $\|\vec{w}\|_\infty$ and the objective maximized by Expo^{α} differ by α times the entropy of the allocation, which is at most $\alpha \ln m$. □

In many settings, we know (or can obtain in a truthful way, e.g., by random sampling) an estimation E of $\|\vec{w}\|_\infty$ with $E \geq \|\vec{w}\|_\infty \geq \rho E$, for some $\rho \in (0,1)$. Then, we can choose $\alpha = \varepsilon \rho E/\ln m$ and obtain an approximation ratio of $1-\varepsilon$ with expected verification $\ln m/(\rho\varepsilon)$, for any $\varepsilon > 0$. E.g., if for all agents i, $|\vec{x}_i| = 1$, $n \geq \|\vec{w}\|_\infty \geq n/m$. Then, using $\alpha = n\varepsilon/\ln m$, we have an additive error of εn with verification $\ln m/\varepsilon$. Moreover, with $\alpha = n\varepsilon/(m \ln m)$, we have approximation ratio $1-\varepsilon$ with verification $m \ln m/\varepsilon$. Finally, note that, since the number of agents verified follows a Poisson distribution, by Chernoff bounds, the verification bounds also hold with high probability.

10. AN APPLICATION TO COMBINATORIAL PUBLIC PROJECT

The *Combinatorial Public Project Problem* (CPPP) was introduced in [Schapira and Singer 2008; Papadimitriou et al. 2008] and has received considerable attention since then. An instance consists of a set R with r resources, a parameter k, $1 \leq k \leq r$, and n agents. Each agent i has a function $\vec{x}_i : 2^R \to \mathbb{R}_{\geq 0}$ that assigns a non-negative valuation $\vec{x}_i(S)$ to each resource subset $S \subseteq R$. We want to compute a set C of k resources that maximizes $\sum_i \vec{x}_i(C)$, i.e., the social welfare from C. We assume that all valuations $\vec{x}_i$ are *normalized*, i.e., $\vec{x}_i(\emptyset) = 0$, and *monotone*, i.e., $\vec{x}_i(S_1) \leq \vec{x}_i(S_2)$ for all $S_1 \subseteq S_2$.

The valuation functions $\vec{x}_i$ are implicitly represented through a value oracle that returns $\vec{x}_i(S)$, for any resource subset S, in $O(1)$ time. CPPP is $\mathbf{NP}$-hard and practically inapproximable in polynomial time, under standard computational complexity assumptions [Schapira and Singer 2008]. If the valuation functions $\vec{x}_i$ are *submodular*, i.e., each $\vec{x}_i$ satisfies $\vec{x}_i(S_1 \cup S_2) + \vec{x}_i(S_1 \cap S_2) \leq \vec{x}_i(S_1) + \vec{x}_i(S_2)$, for all $S_1, S_2 \subseteq R$, CPPP can be approximated in polynomial time within a factor of $1 - 1/e$. If the valuations $\vec{x}_i$ are *subadditive*, i.e., each $\vec{x}_i$ satisfies $\vec{x}_i(S_1 \cup S_2) \leq \vec{x}_i(S_1) + \vec{x}_i(S_2)$, for all $S_1, S_2 \subseteq R$, CPPP can be approximated in polynomial time within a factor of $r^{-1/2}$, while approximating it within any factor better than $r^{-1/4+\varepsilon}$, for any constant $\varepsilon > 0$, requires exponential communication. Papadimitriou et al. [2008] proved that CPPP with submodular valuations cannot be approximated in polynomial time (or with polynomial communication) by deterministic truthful mechanisms (with money) within $r^{-1/2+\varepsilon}$, for any constant $\varepsilon > 0$. Dobzinski [2011] proved a similar communication complexity lower bound for randomized truthful in expectation mechanisms with money.

CPPP can be naturally cast to our Utilitarian Voting framework. The outcome set O consists of all resource subsets S with $|S| = k$ (hence, $m \leq r^k$). So, CPPP with general valuations can be approximated as follows by mechanisms with selective verification:

— For any $\varepsilon > 0$, *Power* allocates a set of k resources, is immune, ε-truthful, achieves an approximation ratio of $1 - \varepsilon$ and verifies at most $k \log r/\varepsilon$ agents.
— For any $\varepsilon > 0$, *Partial Power* allocates a set of k resources with probability $1 - O(\varepsilon)$, is immune, truthful, achieves an approximation ratio of $1 - \varepsilon$ and verifies $O(k \log r/\varepsilon^2)$ agents. In this case, the empty set corresponds to the null outcome.
— For Exponential, we need to assume that for any agent i, $\max_{S \subseteq R, |S| \leq k} \vec{x}_i(S) \leq 1$. Then, for any $\varepsilon > 0$, *Exponential* allocates a set of k resources, is immune, truthful, and achieves an additive error of εn with verification of $O(k \log r/\varepsilon)$ agents, or achieves an approximation ratio of $1 - \varepsilon$ with verification of $O(kr^k \log r/\varepsilon)$ agents.

These guarantees are very strong and rather surprising, especially if the number of agents n is significantly larger than $k \log r$, which is true in many practical settings. We almost reach the optimal social welfare using truthful mechanisms without money that verify a small number of agents independent of n.

The mechanisms above run in time polynomial in r^k and n. So, if valuations are represented by value oracles, they are not computationally efficient. However, we still need to resort to approximate solutions, because, in absence of money, the optimal solution is not truthful. We should highlight that computational inefficiency is unavoidable, since our approximation ratio of $1 - \varepsilon$, for any constant $\varepsilon > 0$, is dramatically better than known lower bounds on the polynomial time approximability of CPPP.

If we seek computationally efficient mechanisms without money for CPPP, we can combine our mechanisms with existing Maximal-in-Range mechanisms. E.g., for CPPP with subadditive valuation functions, we can use the maximal-in-range mechanism of [Schapira and Singer 2008, Sec. 3.2] and obtain randomized polynomial-time truthful mechanisms without money that achieve an approximation ratio of $O(\min\{k, \sqrt{r}\})$ for the social welfare with selective verification of $O(k \log r)$ agents.

11. CONCLUSIONS AND DISCUSSION

In this work, we introduce a general approach to approximate mechanism design without money and with selective verification, and apply it to the general domain of Utilitarian Voting, to Combinatorial Public Project and to k-Facility Location. We present (mostly) randomized mechanisms that are truthful (or almost truthful) and achieve essentially best possible approximation guarantees by verifying only few agents.

A remarkable property of our mechanisms is immunity to agent misreports. To the best of our knowledge, this is the first time that immunity (or a similar) property is considered in mechanism design. Immunity is a strong property that is possible due to selective verification. Actually, with the exception of constant mechanisms, whose probability distribution over outcomes is independent of the agent declarations, a mechanism can be immune only if it uses exact verification. For a comparison against truthfulness, immunity means that provided that an agent misreports, her lie cannot change the final allocation *whatsoever*, while truthfulness means that a liar cannot change the allocation *in her favor*. Hence, truthfulness assumes a utility function that the agents maximize by truthful reporting. Immunity, on the other hand, does not refer to the agent utilities or incentives. It just ensures that the allocation depends only on the declarations admitted as truthful ones by the verification oracle.

We believe that immunity can be very useful when the agents do not explicitly declare their utility functions to the mechanism, but instead have (and declare) some observable types (e.g., address, age, income), and the mechanism translates them into utility functions. Translating observable types to utility functions may introduce some error (with respect to the actual agent utilities), which could affect several properties of the mechanism. E.g., inaccuracy in the utility functions may affect the approximation ratio, which should be computed with respect to the actual utilities, not those assumed by the mechanism. But assuming reasonable error bounds, the asymptotics of the approximation ratio should not change (i.e., a constant approximation ratio should remain constant, but its value may increase). On the other hand, the validity of binary properties, such as participation and truthfulness, crucially depends on the accuracy of utility functions. Interestingly, immunity, also a binary property, is not affected by what the mechanism assumes about agent utilities, because it is only related to the verification of the agent declarations. E.g., in Facility Location, the agents declare their preferred locations to the mechanism. Then, the mechanism assumes that each agent wants a facility close to her declared location and that her cost increases linearly with the distance (see also [Fotakis and Tzamos 2013a]). Approximation ratio, participation and truthfulness depend crucially on this assumption. Immunity however only depends on whether each agent declares her true location (e.g., her true home address) to the mechanism, not on how agent costs depend on the distance.

ACKNOWLEDGMENTS

This research was partially supported by the project Handling Uncertainty in Data Intensive Applications (DELUGE), co-financed by the European Union (European Social Fund) and Greek national funds, through the Research Funding Program THALES, by NSF Award CCF-1101491, and by an award for Graduate Students in Theoretical Computer Science by the Simons Foundation.

REFERENCES

N. Alon, M. Feldman, A.D. Procaccia, and M. Tennenholtz. 2010. Strategyproof Approximation of the Minimax on Networks. *Mathematics of Operations Research* 35, 3 (2010), 513–526.

A. Archer and R. Kleinberg. 2008. Truthful Germs are Contagious: A local-to-global characterization of truthfulness. In *Proc. of the 9th ACM Conference on Electronic Commerce (EC '08)*. 21–30.

D. Arthur and S. Vassilvitskii. 2007. k-means++: the advantages of careful seeding. In *Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*. 1027–1035.

V. Auletta, R. De Prisco, P. Penna, and G. Persiano. 2009. The Power of Verification for One-Parameter Agents. *Journal of Computer and System Sciences* 75 (2009), 190–211.

E. Ben-Porath, E. Dekel, and B.L. Lipman. 2014. Optimal Allocation with Costly Verification. *American Economic Review* 104, 12 (2014), 3779–3813.

I. Caragiannis, E. Elkind, M. Szegedy, and L. Yu. 2012. Mechanism design: from partial to probabilistic verification. In *Proc. of the 13th ACM Conference on Electronic Commerce (EC '12)*. 266–283.

R. Cole, V. Gkatzelis, and G. Goel. 2013. Mechanism design for fair division: allocating divisible items without payments. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC '13)*. 251–268.

S. Dobzinski. 2011. An Impossibility Result for Truthful Combinatorial Auctions with Submodular Valuations. In *Proc. of the 43rd ACM Symposium on Theory of Computing (STOC '11)*. 139–148.

S. Dobzinski and S. Dughmi. 2013. On the Power of Randomization in Algorithmic Mechanism Design. *SIAM Journal on Computing* 42, 6 (2013), 2287–2304.

A. Erlanson and A. Kleiner. 2015. Costly Verification in Collective Decisions. *Working Paper* (2015).

A. Filos-Ratsikas and P.B. Miltersen. 2014. Truthful Approximations to Range Voting. In *Proc. of the 10th Workshop on Internet and Network Economics (WINE '14)*, (LNCS) 8877. 175–188.

D. Fotakis, P. Krysta, and C. Ventre. 2014. Combinatorial Auctions without Money. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '14)*. 1029–1036.

D. Fotakis, D. Tsipras, C. Tzamos, and E. Zampetakis. 2015a. Efficient Money Burning in General Domains. In *Proc. of the 8th Symposium on Algorithmic Game Theory (SAGT '15)*, LNCS 9347. 85–97.

D. Fotakis and C. Tzamos. 2013a. Strategyproof Facility Location with Concave Costs. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC '13)*. 435–452.

D. Fotakis and C. Tzamos. 2013b. Winner-Imposing Strategyproof Mechanisms for Multiple Facility Location Games. *Theoretical Computer Science* 472 (2013), 90–103.

D. Fotakis and C. Tzamos. 2014. On the Power of Deterministic Mechanisms for Facility Location Games. *ACM Transactions on Economics and Computation* 2, 4 (2014), 15:1–15:37.

D. Fotakis, C. Tzamos, and E. Zampetakis. 2015b. Who to Trust for Truthfully Maximizing Welfare? *CoRR* abs/1507.02301 (2015).

D. Fotakis and E. Zampetakis. 2015. Truthfulness Flooded Domains and the Power of Verification for Mechanism Design. *ACM Transactions on Economics and Computation* 3, 4 (2015), 20.

J. Glazer and A. Rubinstein. 2004. On Optimal Rules of Persuasion. *Econometrica* 72 (2004), 1715-1736.

J.D. Hartline and T. Roughgarden. 2008. Optimal mechanism design and money burning. In *Proc. of the 40th ACM Symposium on Theory of Computing (STOC '08)*. 75–84.

Z. Huang and S. Kannan. 2012. The Exponential Mechanism for Social Welfare: Private, Truthful, and Nearly Optimal. In *Proc. of the 53rd IEEE Symposium on Foundations of Computer Science (FOCS '12)*. 140–149.

P. Krysta and C. Ventre. 2015. Combinatorial Auctions with Verification Are Tractable. *Theoretical Computer Science* 571 (2015), 21–35.

R. Lavi and C. Swamy. 2011. Truthful and Near-Optimal Mechanism Design via Linear Programming. *Journal of the ACM* 58, 6 (2011), 25.

P. Lu, X. Sun, Y. Wang, and Z.A. Zhu. 2010. Asymptotically Optimal Strategy-Proof Mechanisms for Two-Facility Games. In *Proc. of the 11th ACM Conference on Electronic Commerce (EC '10)*. 315–324.

K. Nissim, R. Smorodinsky, and M. Tennenholtz. 2012. Approximately Optimal Mechanism Design via Differential Privacy. In *Proc. of the 3rd Conference on Innovations in Theoretical Computer Science (ITCS '12)*. 203–213.

C.H. Papadimitriou, M. Schapira, and Y. Singer. 2008. On the Hardness of Being Truthful. In *Proc. of the 49th IEEE Symposium on Foundations of Computer Science (FOCS '08)*. 250–259.

E. Pountourakis and G. Schäfer. 2014. Mechanisms for Hiring a Matroid Base without Money. In *Proc. of the 7th International Symposium on Algorithmic Game Theory (SAGT '14)*, LNCS 8768. 255–266.

A.D. Procaccia and M. Tennenholtz. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation* 1, 4 (2013), 18.

M. Schapira and Y. Singer. 2008. Inapproximability of Combinatorial Public Projects. In *Proc. of the 4th Workshop on Internet and Network Economics (WINE '08)*, LNCS 5385. 351–361.

I. Sher and R. Vohra. 2014. Price Discrimination through Communication. *Theoretical Economics* 10 (2014), 597–648.

R.M. Townsend. 1979. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21 (1979), 265–293.

D.P. Williamson and D.B. Shmoys. 2011. *The Design of Approximation Algorithms*. Cambridge Univ. Press.