

Truthful Reporting of Competence with Minimal Verification

Reshef Meir
 Technion—Israel Institute of
 Technology
 Haifa, Israel
 reshefm@technion.ac.il

Jonathan Wagner
 Technion—Israel Institute of
 Technology
 Haifa, Israel
 swagner@campus.technion.ac.il

Omer Ben-Porat
 Technion—Israel Institute of
 Technology
 Haifa, Israel
 omerbp@technion.ac.il

ABSTRACT

Suppose you run a home exam, where students should report their own scores but can cheat freely. You can, if needed, call a limited number of students to class and verify their actual performance against their reported score. We consider the class of mechanisms where truthful reporting is a dominant strategy, and truthful agents are never penalized—even off-equilibrium.

How many students do we need to verify, in expectation, if we want to minimize the bias, i.e., the difference between agents’ competence and their expected grade? When perfect verification is available, we characterize the best possible tradeoff between these requirements and provide a simple parametrized mechanism that is optimal in the class for any distribution of agents’ types. When verification is noisy, the task becomes much more challenging. We show how proper scoring rules can be leveraged in different ways to construct truthful mechanisms with a good (though not necessarily optimal) tradeoff.

KEYWORDS

Game theory; mechanism design; auditing

ACM Reference Format:


Reshef Meir, Jonathan Wagner, and Omer Ben-Porat. 2026. Truthful Reporting of Competence with Minimal Verification. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/CZFV1739>

1 INTRODUCTION

We consider a scenario in which self-interested agents need to report their own performance, or score in a given task, such as a home exam, earnings, recycling, etc.

More precisely, each agent reports her own performance or competence (=expected score), and should be graded or compensated accordingly. However, agents always prefer *higher grades* and thus have an incentive to misreport.

The principal has limited verification capabilities: it can select some of the agents and observe their score directly by retaking the exam under supervision, auditing their financial status, tracking their consumption or recycling behavior, etc. The final grades can be based on both *reported* and *observed* scores, and thus agents that score lower than they reported could be punished. This abstraction of verification mechanisms is common in the economic literature

 This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CZFV1739>

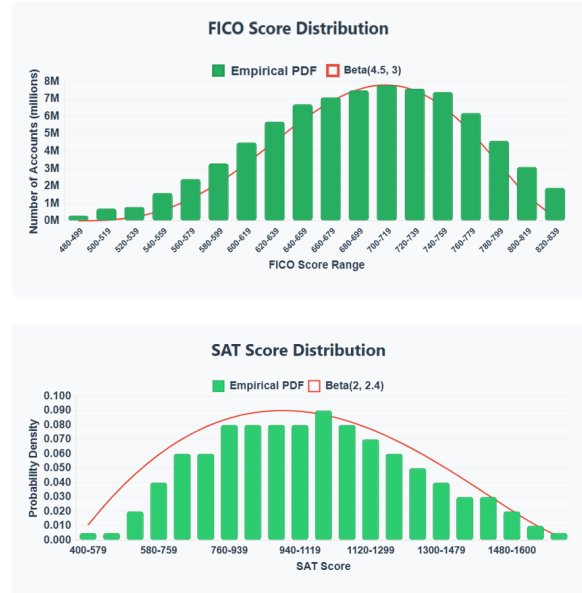


Figure 1: Top: Typical distribution of credit score [1]. Bottom: Typical distribution of SAT scores [31].

when the underlying problem varies from allocation of one or more goods [6, 10], to incentivizing certain actions [23], or revealing the value for the principal [20]. Intuitively, we would like a mechanism where truthful reports are a dominant strategy, as is standard both in the economics and mechanism design literature.

However, we focus on a design goal that has not been studied to the best of our knowledge, despite being very natural: specifically, the principal would like to assign the agent(s) their true competence as their grade. This alone could be trivially obtained by verifying everyone, by threatening cheaters with an excessive penalty, or by rewarding truth-tellers. However, we would like to audit as few agents as possible, and excessive penalties may not be allowed.

Moreover, some agents may misreport even in a truthful mechanism. This can happen if agents are irrational [17], have unintentional mistakes in their reports [30], cannot see why the mechanism is truthful [3, 19], cannot verify the behavior of the mechanism [7], etc. We therefore require that a truthful agent is never punished (i.e., graded lower than her competence) regardless of others’ reports.

In many situations, there is public information on the distribution of agents’ types, that the principal can use. See Fig. 1 and Fig. 3. Ideally, however, the properties of the mechanism (e.g., truthfulness) should not critically depend on access to the distribution, as it may be inaccurate.

1.1 Related Work

There is much work on incentivizing self-interested agents to report private information. We can distinguish these settings according to the main reason agents may want to lie.

In the first type of setting, agents do not really care how the principal uses the reported information but may lie because they do not have it or because obtaining the correct information requires effort. Common solutions in the single agent case include the use of scoring rules [28, 33] or more complicated contracts that assume some ability to verify the outcome [8]. In the multi-agent case, there are many variants of mechanisms built on cross-validation, such as Bayesian Truth Serum [27, 32].

In contrast, the mechanism design literature that is closer to our motivation deals with situations where the agents have a direct interest in the outcome, which may be a common alternative, assignment of resources, and so on. We can further split problems into those that allow monetary transfers (which often make it easy to align incentives, despite possible computational hardness [5]); and those that are not [25]. In the latter type, negative results are abundant and successful mechanisms are only possible when the structure of the problem can be exploited.

Our problem relates to the second type, and moreover, agents have very clear incentives to inflate their reported scores. The closest line of work within this framework is probably *strategic classification*, where agents report information and would like it to be classified in a certain way [15, 18, 21, 34]. For example, Hardt et al. [15] consider a designer (‘jury’) that sets up the classification algorithm. An agent (‘contestant’) can try and game the input, at a cost, in an attempt to improve her outcome (say, increase her chances for a loan). In our model, agents can lie freely, but might be caught and punished.

Some papers also combine payments/penalties and verification. E.g., in [9], partial verification is used to reduce payments; and other works consider payments/penalties contingent on some value realization or action [11, 26] (implying it can be freely verified). Mylovanov and Zapechelnuk [24] highlight the issue of *limited liability*, as the penalty is exogenously bounded to a fraction of the value of the allocated good.

Closer to our motivation, Akasiadis and Chalkiadakis [2] use scoring rules to induce truthful reports of electricity consumption shifting costs. However, in these works, all values are eventually revealed without the principal taking an explicit action.

Mechanism design with verification. In the classical principal-agent model with verification [14], the principal has no ability to reveal the true state, but can a-priori restrict the allowed responses of the agent, which is irrelevant in our case. In a recent follow-up work by Ball and Kattwinkel [4], they add *tests* that the principal can assign to the agent following her report, similarly to our noisy verification model. The test is costless, and the challenge is to match each report with the best test.

Other works use costly verification in a very similar way to us, when the problem is to allocate item(s) to the agents with the highest private value [6, 10], or eligibility based on private features [12], where value/features are sampled independently from *known* individual distributions. Li studies a similar assignment problem, but

where the agents privately know the value of the item to the principal [20]. Our setting is most similar to [23], where agents’ utility is a function of their reported type and the penalties, but the principal’s goals are very different.

1.2 Our Contribution

In the next section, we define the class of elicitation mechanisms and our formal requirements. Section 3 contains our main result, which is the single-parameter subclass of Monotone-Cutoff Verification (MCV) mechanisms. MCV provides an optimal tradeoff between minimal verification and minimal bias, assuming deterministic verification. When the type distribution is known, one can choose the parameter to bound either the bias or the verification cost. When it is not known, we show how near-optimal bounds can still be guaranteed in a multiagent setting. In Section 4, we provide three mechanisms that only assume noisy verification. All mechanisms leverage *proper scoring rules*, albeit in different ways. The last section shows empirically that our mechanisms provide a good tradeoff of bias vs. verification on common distributions.

All missing proofs are included in the full version of the paper [22].

2 SETTING

There are N agents, each one with a type t_i , which is an element in the set \mathcal{T} , where $\mathcal{T} \subseteq [0, 1]$. The type, which is private information, measures the competence of agent i , where better agents have a higher type. Agents report their types to a central authority that we call the *principal*. We denote the report of agent i by $\hat{t}_i \in \mathcal{T}$. The principal wishes to grade agents, where the grades should be as close as possible to the real types. Agents are selfish and wish to maximize their grade, as explained in the introduction. Thus, agents will potentially lie to maximize their expected grade.

Verification. The principal can verify a particular agent’s type, e.g., by forcing them to re-take an exam in class, audit economic reports, or track individual consumption. We assume the principal has access to a *verification process*, which returns t_i in the deterministic case. In the noisy verification case, the process returns $s_i \sim T_i$ where T_i is an arbitrary distribution on \mathcal{T} with mean t_i . Other than the expectation, we make no assumptions on the distribution T_i . In particular, it may be correlated with other agents, and need not be known to the designer.

Mechanisms. A mechanism $\mathcal{M} = (\mathcal{M}_V, \mathcal{M}_G)$ consists of two different functions:

- (1) Verification selection function $\mathcal{M}_V : \mathcal{T} \rightarrow [0, 1]$. This function takes as input an agent’s reported types \hat{t}_i and possibly other available information (see below), and returns an *audit probability* $q_i \in [0, 1]$ for every $i \in N$: the probability that i will be selected for verification. Since all agents with same reported type t have the same q_i , we also denote q_t ;
- (2) Grading function $\mathcal{M}_G : \mathcal{T} \times (\mathcal{T} \cup \{\perp\}) \rightarrow \mathbb{R}$. This function takes as input a reported type \hat{t}_i and a realized performance sample s_i (or \perp if i is not verified) and sets a grade g_i . The grade is not necessarily in the set \mathcal{T} and could e.g. be negative unless specified otherwise.

Mechanism 1: A general elicitation mechanism

Require: $(\hat{t}_i)_{i \in N}, \mathcal{M}_V, \mathcal{M}_G, [\text{Optional: Information } I]$

- 1: **for all** $i \in N$ **do**
- 2: $q_i \leftarrow \mathcal{M}_V(\hat{t}_i)$ //audit probability
- 3: **for all** $i \in N$ **do**
- 4: with probability $1 - q_i$, set $g_i \leftarrow \mathcal{M}_G(\hat{t}_i, \perp)$;
- 5: with the remaining probability, obtain a verified score $s_i \sim T_i$ and set grade $g_i \leftarrow \mathcal{M}_G(\hat{t}_i, s_i)$.
- 6: **return** $(g_i)_{i \in N}$

The functions above may rely on external parameters or additional information, that the mechanism takes as input.

To specify a concrete mechanism, we only need to formally define \mathcal{M}_V and \mathcal{M}_G (see Mechanism 1). Three simple examples are:

VERIFY-ALL: $\mathcal{M}_V^{VA}(\hat{t}_i) = 1 \forall \hat{t}_i$; and grades such that

$$\mathcal{M}_G^{VA}(\hat{t}_i, s_i) = \hat{t}_i \text{ if } \hat{t}_i = s_i, \text{ and } 0 \text{ otherwise.}$$

PAY-ALL: \mathcal{M}^{PA} just pays 1 to all agents without verification, i.e.,

$$\mathcal{M}_V^{PA} \equiv 0, \mathcal{M}_G^{PA} \equiv 1.$$

HUGE-PENALTY: $\mathcal{M}_V^{HP}(\hat{t}_i) = \varepsilon \forall \hat{t}_i$ for some small $\varepsilon > 0$; and

$$\text{grades each such that } \mathcal{M}_G^{HP}(\hat{t}_i, s_i) := \begin{cases} \hat{t}_i & \text{if } \hat{t}_i = s_i \text{ (or } \perp) \\ -2/\varepsilon & \text{if } \hat{t}_i \neq s_i \end{cases}$$

Note that once we fix the mechanism \mathcal{M} and the profile \hat{t} , the expected grade of an agent can be determined. Let

$$\bar{g}_t(\hat{t}) := \mathcal{M}_V(\hat{t}) E_{s_i \sim T_i}[\mathcal{M}_G(\hat{t}_i, s_i)] + (1 - \mathcal{M}_V(\hat{t})) \mathcal{M}_G(\hat{t}_i, \perp) \quad (1)$$

be the expected grade of a type t agent reporting \hat{t} , and denote by $\bar{g}_t := \bar{g}_t(\hat{t})$ the expected grade of a truthful type- t agent.

Hard requirements.

HR1. Truth-telling must be a weakly dominant-strategy;

HR2. Truth-tellers are never punished. I.e., $g_i \geq t_i$ for all $t_i \in \mathcal{T}$ even if others lie;

HR3. Realized grades must be at least $-\xi$, where $\xi \geq 0$ is a parameter.¹

The special case of $\xi = 0$ is called *limited liability* and means we cannot ‘punish’ cheaters, beyond disqualifying their reward.

In the *deterministic verification* case, a mechanism satisfying **HR1-HR3** is called *valid* (we later discuss relaxations for the noisy setting). Clearly if we only care about validity, it could be obtained trivially, by either one of the **VERIFY-ALL** or **PAY-ALL** mechanisms. **HUGE-PENALTY** may or may not be valid, depending on ξ .

Soft requirements and tradeoffs. Ideally, we would like to minimize both the amount of verified agents, and the grading error $|g_i - t_i|$. Since the grade may be affected by random factors, a more modest goal would be to minimize the *bias* $|\bar{g}_t - t|$ of every type t . By **HR2** the bias of a valid mechanism must be non-negative.

The metrics we would like to optimize may also depend on the type distribution $p = (p(t))_{t \in \mathcal{T}}$,² where $t_{\min} \neq t_{\max}$ are the minimal and maximal types with positive support. All soft requirements refer to a truthful profile:

¹We could similarly add a requirement about the *maximal* realized grade (e.g., a course grade may be capped at ‘100’ by the system). For simplicity, we avoid this restriction, but note that most of our mechanisms indeed keep realized grades at most 1.

²For ease of presentation we treat p as a discrete distribution but all definitions and results (except in Section 4.2) naturally extend to continuous distributions.

- Minimize the expected proportion of agents we verify $\text{ver}(M, p) := E_{t \sim p}[q_t] = \sum_{t \in \mathcal{T}} p(t) q_t$.
- Minimize the **maximal** bias (always nonnegative by **HR2**) $\mathbf{B}(M) := \max_{t \in \mathcal{T}} (\bar{g}_t - t)$.
- Minimize the **expected** bias $\text{bias}(M, p) := E_{t \sim p}[\bar{g}_t - t] = \sum_{t \in \mathcal{T}} p(t) (\bar{g}_t - t)$.

While the soft requirements necessarily depend on the type distribution p , our hard requirements including truthfulness should hold even if p is unknown (in contrast to most previous work [6, 12]).

3 DETERMINISTIC VERIFICATION

In this section we assume verification is deterministic, i.e., whenever i is audited, we get $s_i = t_i$ w.p. 1. Even so, the three soft requirements capture design goals that are generically contradictory: the lower the audit probability, the higher the required threat to maintain truth-telling as an equilibrium. The **VERIFY-ALL** and **HUGE-PENALTY** mechanisms represent the two extremes. It is thus natural to ask what is the efficient compromise between accuracy and small sample, in terms of a theoretical bound and, moreover, whether we can find mechanisms that implement it.

Definition 3.1 (Domination). We say that a mechanism \mathcal{M} (weakly) *dominates* some other mechanism \mathcal{M}' if for any type distribution p , \mathcal{M} outperforms \mathcal{M}' in all soft requirements. If in addition \mathcal{M} strictly outperforms \mathcal{M}' on one of those requirements, we say that \mathcal{M} *strictly dominates* \mathcal{M}' .

In other words, \mathcal{M} dominates \mathcal{M}' if \mathcal{M} verifies fewer agents, and yet obtains lower maximal bias and expected bias. Furthermore,

Definition 3.2 (Efficiency). A mechanism is *efficient* if there is no other mechanism that strictly dominates it.

3.1 The Monotone-cutoff Verification Class

Next, we introduce the class of *Monotone-Cutoff Verification* mechanism (hereinafter MCV). The class is parameterized with $\gamma \in [0, 1]$, which governs the audit probabilities. We emphasize that for an MCV mechanism and the results in this subsection, the number of agents is irrelevant, and we could have stated them for a single agent. However, our extensions will make use of the entire population.

The Monotone-Cutoff Verification (MCV) mechanism. The mechanism $\mathcal{M}^{\gamma, \xi}$ takes the parameters $\gamma \in [0, 1]$ and $\xi \geq 0$ as input. Since ξ is fixed and derived from **HR3** we will omit it to simplify the notation, writing \mathcal{M}^γ rather than $\mathcal{M}^{\gamma, \xi}$.

- Verification: agent i with report \hat{t}_i is verified by \mathcal{M}_V^γ with a probability of $\mathcal{M}_V^\gamma(\hat{t}_i) = \begin{cases} \frac{\hat{t}_i - \gamma}{\hat{t}_i + \xi} & \hat{t}_i > \gamma \\ 0 & \hat{t}_i \leq \gamma \end{cases}$
- Grading: if agent i was not verified, then she receives a score of $\mathcal{M}_G^\gamma(\hat{t}_i, \perp) = \begin{cases} \hat{t}_i & \hat{t}_i > \gamma \\ \gamma & \hat{t}_i \leq \gamma \end{cases}$.
Otherwise, if the verification yielded a result s_i , then $\mathcal{M}_G^\gamma(\hat{t}_i, s_i) = \begin{cases} \hat{t}_i & \hat{t}_i = s_i \\ -\xi & \hat{t}_i \neq s_i \end{cases}$.




Figure 2: Illustration of an MCV mechanism. The horizontal line represents the report $\hat{t}_i = t$, and the vertical axis plays the role of both the grade and the audit probabilities. The audit probability (red double line) is zero for reports below γ , and any such report gets γ (green horizontal segment). The dotted diagonal line represents the identity. When $\hat{t}_i > \gamma$, audit probability is increasing, both the default (green) and truthful (dashed orange) grades match \hat{t}_i , and there is a fixed maximal penalty of $-\xi$ (blue dash-dots).

We visualize an example of a MCV mechanism in Figure 2. The red and gray shaded areas demonstrate the overall audit probability and bias, respectively (for a uniform type distribution).

Notice that \mathcal{M}^Y effectively lets agents report their type as “not greater than γ ” without being verified, and $\bar{g}_t = \max\{t, \gamma\}$.

For $\xi = 0$ and the extreme values $\gamma = 0$ and $\gamma = 1$, we get the VERIFY-ALL and PAY-ALL mechanisms, respectively. Setting $\gamma = 0$, $\xi = \frac{1}{\epsilon}$ almost coincides with the HUGE-PENALTY mechanism. Thus, the MCV class of mechanisms aims to optimally balance between the solutions, where every pick of γ, ξ corresponds to a different point on the Pareto frontier.

PROPOSITION 3.3. *All MCV mechanisms are valid.*

PROOF. **HR2** and **HR3** are immediate by definition. For truthfulness, note that in \mathcal{M}^Y truth telling guarantees at least γ for all agents, whereas the maximum expected payoff for an untruthful agent,

$$\max_{\hat{t}_i \neq t_i} (1 - \mathcal{M}'_V(\hat{t}_i)) \mathcal{M}'_G(\hat{t}_i, \perp) + \mathcal{M}'_V(\hat{t}_i) \cdot \mathcal{M}'_G(\hat{t}_i, t_i),$$

is exactly γ . If $\hat{t}_i \leq \gamma$, the agent receives a score of γ , without being verified. For $\hat{t}_i > \gamma$:

$$\begin{aligned} & (1 - \mathcal{M}'_V(\hat{t}_i)) \mathcal{M}'_G(\hat{t}_i, \perp) + \mathcal{M}'_V(\hat{t}_i) \cdot \mathcal{M}'_G(\hat{t}_i, t_i) \\ &= \left(1 - \frac{\hat{t}_i - \gamma}{\hat{t}_i + \xi}\right) \hat{t}_i + \frac{\hat{t}_i - \gamma}{\hat{t}_i + \xi} (-\xi) = \gamma. \quad \square \end{aligned}$$

As we mentioned above, setting $\gamma = 0$ and $\gamma = 1$ (under limited liability) results in the VERIFY-ALL and PAY-ALL mechanisms, respectively. Next, we show that the class $\Gamma := \{\mathcal{M}^Y : \gamma \in [0, 1]\}$ constitutes the efficiency boundary of all valid mechanisms.

THEOREM 3.4. *For any valid mechanism M , there exists $\mathcal{M}^Y \in \Gamma$ that weakly dominates M .*

Meaning, if minimum possible grade is $-\xi$ and we demand **HR1** and **HR2**, we only need to choose $\gamma \in [0, 1]$ for which the corresponding $\mathcal{M}^Y \in \Gamma$ best suits our ideal trade-off between accuracy and verification (where these depend also on the prior distribution of types p). In particular, as the proof shows, $\text{ver}(\mathcal{M}^Y, p)$, $\mathbf{B}(\mathcal{M}^Y)$ and $\text{bias}(\mathcal{M}^Y, p)$ are easy to compute for any given type dist. p .

PROOF. Denote, for any type distribution p and $\gamma \in [0, 1]$ the following measures that are independent of the mechanism:

$$\text{BIAS}(\gamma, p) := \sum_{t \leq \gamma} p(t)(\gamma - t); \quad (2)$$

$$\text{VER}(\gamma, p) := \sum_{t > \gamma} p(t) \frac{t - \gamma}{t + \xi} \quad (3)$$

We first argue that any mechanism \mathcal{M}^Y meets these measures exactly, i.e., $\forall \gamma \in [0, 1]$ and $p \in \Delta([0, 1])$, $\text{bias}(\mathcal{M}^Y, p) = \text{BIAS}(\gamma, p)$ and $\text{ver}(\mathcal{M}^Y, p) = \text{VER}(\gamma, p)$.

Indeed, these equalities follow directly from the definitions of \mathcal{M}'_G and \mathcal{M}'_V , respectively. Also note that $\mathbf{B}(\mathcal{M}^Y) = \gamma - t_{\min}$.

Now take M as described in the theorem. Let M' be a mechanism identical to M , except $\mathcal{M}'_G(\hat{t}_i, s_i) := -\xi$ for $\hat{t}_i \neq s_i$. We argue that M' is also valid, and dominates M . First, since penalty for cheating is weakly increasing by **HR3**, being truthful remains a weakly dominant strategy. Second, grades of all truthful agents are unchanged, and thus no change in bias. Next, let

$$\begin{aligned} \tilde{\gamma} &:= \max_{\hat{t}_i \neq t_i} (1 - \mathcal{M}'_V(\hat{t}_i)) \mathcal{M}'_G(\hat{t}_i, \perp) + \mathcal{M}'_V(\hat{t}_i) \cdot \mathcal{M}'_G(\hat{t}_i, t_i) \\ &= \max_{\hat{t}_i} (1 - \mathcal{M}'_V(\hat{t}_i)) \mathcal{M}'_G(\hat{t}_i, \perp) + \mathcal{M}'_V(\hat{t}_i) \cdot (-\xi). \end{aligned}$$

$\tilde{\gamma}$ the maximum gain for untruthful agents under M' , regardless of their true type. We will see that $\mathcal{M}^{\tilde{\gamma}} \in \Gamma$ dominates M' (and thus M).

If a grade of $\tilde{\gamma}$ is achievable by lying, then by **HR1** all truthful agents are graded at least $\tilde{\gamma}$ in equilibrium, therefore

$$\text{bias}(M', p) \geq \sum_{t \leq \tilde{\gamma}} p(t)(\tilde{\gamma} - t) = \text{bias}(\mathcal{M}^{\tilde{\gamma}})$$

and, in particular, a type-zero agent has at least $g_0 = \tilde{\gamma}$, thus $\mathbf{B}(M') \geq \tilde{\gamma} - t_{\min} = \mathbf{B}(\mathcal{M}^{\tilde{\gamma}})$. Moreover, for all $\hat{t}_i \in [0, 1]$ we have $\mathcal{M}'_G(\hat{t}_i, \perp) \geq \hat{t}_i$ by **HR2**, and thus by definition of $\tilde{\gamma}$,

$$\tilde{\gamma} \geq (1 - \mathcal{M}'_V(\hat{t}_i)) \hat{t}_i - \mathcal{M}'_V(\hat{t}_i) \xi \implies \mathcal{M}'_V(\hat{t}_i) \geq \frac{\hat{t}_i - \tilde{\gamma}}{\hat{t}_i + \xi}$$

Summing over all types $\tilde{\gamma} < \hat{t}_i \leq 1$ we get

$$\text{ver}(M') \geq \sum_{t > \tilde{\gamma}} p(t) \frac{t - \tilde{\gamma}}{t + \xi} = \text{ver}(\mathcal{M}^{\tilde{\gamma}}). \quad \square$$

COROLLARY 3.5. *Every MCV mechanism is efficient.*

PROOF. If \mathcal{M}^Y is strictly dominated by any other mechanism, the theorem gives that it is also strictly dominated by some $\mathcal{M}^{Y'} \in \Gamma$. But this is impossible since for any distribution p and for any two $0 \leq \gamma < \gamma' \leq 1$ (where $p(\gamma') > 0$), $\text{ver}(\mathcal{M}^{\gamma'}, p) < \text{ver}(\mathcal{M}^{\gamma}, p)$ while $\text{bias}(\mathcal{M}^{\gamma'}, p) > \text{bias}(\mathcal{M}^{\gamma}, p)$. \square




Figure 3: US income distribution by tax brackets in 2013 [16].

Optimizing the cutoff point. Theorem 3.4 provides us with a way to explicitly bound the tradeoff on bias/verification.

For $\beta \in [0, 1]$ let $\gamma(\beta, p) := \max\{\gamma \in [0, 1] : \text{BIAS}(\gamma, p) \leq \beta\}$.

COROLLARY 3.6. For every $\beta \in [0, 1]$, type dist. p , $\text{bias}(\mathcal{M}^{\gamma(\beta, p)}) = \beta$; and $\text{ver}(\mathcal{M}^{\gamma(\beta, p)}) = \text{VER}(\gamma(\beta, p), p)$.

As we saw in the proof of the theorem, no mechanism with lower bias can verify less.

Limited liability. In particular, if $\xi = 0$, the following corollary is implied, showing that *some* bias is inescapable unless we verify each and every agent.

COROLLARY 3.7. Suppose $p(0) > 0$. Under limited liability, no valid mechanism exists such that $\text{bias}(M, p) = 0$ and $\text{ver}(M, p) < 1$.

Example. Suppose we apply the MCV mechanism to decide on tax audits on the population from Fig. 3, where a grade of ‘1’ means full tax exemption (0% tax). Note that there are no ‘grade 0’ (100% tax) agents so Corollary 3.7 does not apply. In fact by setting $\gamma = 0.6$ and $\xi = 0$ (i.e., tax full income of cheaters) we can get 0 bias by auditing 30% of the population.

In contrast, if we maintain limited liability and set $\gamma = 0.75$ (cap max tax bracket at 25%) then we will have an expected bias of 0.003, auditing only 14.6% of the people.

3.2 Unknown Prior

The definition of the MCV class, its strategy-proofness and efficiency do not require an a-priori assumed type distribution. However, running \mathcal{M}^γ for some arbitrary $\gamma \in [0, 1]$ offers no guarantee in terms of average bias and expected sample size.

In what comes next, we show that we can, in fact, approximate pretty well any desired goals on the Pareto curve, even without a prior, provided there are enough agents from the same distribution p . Notably, the agents’ types need not be independent.

The basic principle of the mechanism we construct is that for every agent i , we pick a slightly different MCV mechanism, based on the observed distribution of all agents excluding her. As we will see, this scheme will approximate the desired bias goal by an additive $O(1/n)$ factor, with no additional verification.

Definition 3.8. For all $i \in N$, let $p_{-i} : [0, 1] \rightarrow [0, 1]$ be the observed (reported) type distribution aggregated over $\hat{t}_j, j \neq i$. Since p_{-i} is identical for every two agents of the same type, we also use the notation p_t .

Agent-Specific MCV (AS-MCV) mechanism: For all $\beta \in [0, 1]$, define \mathcal{AM}^β as follows. For every agent i , let $\gamma_i := \gamma(\frac{n}{n-1}\beta, p_{-i})$; then apply \mathcal{M}^{γ_i} to agent i .

We first observe that \mathcal{AM}^β is valid, since any \mathcal{M}^{γ_i} is valid, and since the report of i cannot influence which mechanism she will face. β is our desired bound on the expected bias, and we wish to minimize $\text{ver}(M, p)$ without knowing p .

THEOREM 3.9. For every $\beta \in [0, 1]$, and any type dist. p ,

- (1) $\text{bias}(\mathcal{AM}^\beta, p) \leq \beta + \frac{1}{n}$
- (2) $\text{ver}(\mathcal{AM}^\beta, p) \leq \text{VER}(\gamma(\beta, p), p)$.

Compare these bounds to Cor. 3.6: the theorem shows that we only lose an additive factor of $\frac{1}{n}$ on the average bias by not knowing p in advance, without additional verification.

Intuitively, the challenge is that there is no longer a single cutoff. But the maximal and minimal γ_{-i} can be used to consider all relevant types for bias and ver, respectively. Since p_{-i} is close to p the increase in bias is bounded. Then we show that all agents that affect $\text{ver}()$ use the same mechanism $\mathcal{M}^{\gamma(\beta, p)}$.

It is not hard to see how one can similarly minimize the bias subject to bounded verification capacity.

4 NOISY VERIFICATION

The MCV mechanism depends on the fact that verification is deterministic, i.e., that liars are always caught, if verified. In this section we relax the assumption of deterministic verification, meaning that T_i is not guaranteed to return t_i , but rather some random variable $s_i \sim T_i$ with expected value t_i .

To see why MCV (and any mechanism with penalties) fails under noisy verification, note that even a truthful agent might get $s_i < \hat{t}_i$ and be punished, in violation of **HR2**. We therefore have to relax **HR2** to **HR2’**: we still require unverified agents to get at least \hat{t} , and truthful agents should get at least \hat{t} only in expectation (i.e., $\bar{g}_t \geq t$). In other words, we require that the bias is nonnegative.

All mechanisms we construct in this section make use of *proper scoring rules*.

Scoring rules. Based on [13], a *proper scoring rule* is a function $R : [0, 1]^2 \rightarrow \mathbb{R}$ s.t. $E_{s \sim T}[R(x, s)]$ is maximized for $x = E_{s \sim T}[s]$ for any distribution T .³

A common example of a proper scoring rule is the *continuous semi power scoring rule*

$$R^\alpha(x, s) := \alpha \cdot s \cdot x^{\alpha-1} - (\alpha - 1)x^\alpha. \tag{4}$$

This is an asymmetric variant of the power scoring rules studied in [29]. For $\alpha = 2$ we get the semi-quadratic scoring rule $R^2(x, s) = 2sx - x^2$. It is known from [13] that a function can be implemented by a proper scoring rule if and only if it is strictly convex, and this will be useful for us later on. In particular, $E_{s \sim T}[R^\alpha(t, s)] = t^\alpha$ for any distribution T with mean $t \in [0, 1]$, and any $\alpha > 1$.

Naïve mechanisms. In the deterministic verification case, there were simple valid mechanisms that minimized either the bias or the audit probability (while performing poorly on the other). Indeed, the PAX-ALL mechanism still applies in the noisy setting. A parallel of

³More generally, scoring rules allow the agent to report her full distribution T rather than just its expectation, but this simplified form will suffice for our needs.

VERIFY-ALL is also possible, by applying the scoring rule $R^{1+\varepsilon}(\hat{t}, s)$ to all agents, for some arbitrarily small $\varepsilon > 0$.⁴

However, designing a mechanism with nontrivial bounds both on the bias and the verification cost is much more challenging than in the deterministic case, since we can never be certain that a particular agent has cheated. E.g., it is not a-priori clear how to implement a mechanism similar to HUGE-PENALTY (even ignoring HR3), as in the noisy setting large penalties may violate both HR1 and HR2'. We will later return to this point. We also show in the full version that applying a proper scoring rule only to some agents is *not* truthful.

We therefore provide two partial solutions in this section. The first maintains all of our previous assumptions, implementing mechanisms that behave like ‘scoring-rules-in-expectation’. The second mechanism exploits the size of the population and assumes access to the accurate distribution of the true types, and has a unique equilibrium with no verification and arbitrarily small bias.

Deliberate failures. One may worry about situations where agents deliberately under-perform when getting verified. We highlight that although we made no explicit requirement regarding this type of manipulation, the grade in all of our mechanisms is monotonic in the realization s , and thus agents cannot gain from such behavior.

4.1 Scoring Rules ‘in Expectation’

The idea is not to apply a scoring rule only on verified agents, but to set up the mechanism such that every agent is subject to a scoring rule *in expectation*, and thus has an incentive to be truthful.

We demonstrate this first with a simple, parameter-free mechanism with linear audit probability.

Linear Verification (LV) mechanism:

- $\mathcal{LVM}_V(\hat{t}_i) := \hat{t}_i$.
- $\mathcal{LVM}_G(\hat{t}_i, \perp) := \frac{1}{4} + \hat{t}_i$.
- $\mathcal{LVM}_G(\hat{t}_i, s_i) := 2s_i - \frac{3}{4}$.

LEMMA 4.1. $\bar{g}_t(\hat{t}) = \frac{1}{4} + R^2(\hat{t}, t)$, where grade expectation is taken over both the random verification decision of the mechanism, and the performance of the verified agent.

COROLLARY 4.2. The LV mechanism satisfies HR1, HR2', and HR3 with $\xi = \frac{3}{4}$.

PROOF. Truthfulness (HR1) follows immediately from Lemma 4.1 since the expected grade is an affine transformation of R^2 , and thus also a proper scoring rule. As for HR2', note that the expected grade of a truthful agent of type t is $\frac{1}{4} + t^2$. It is easy to check that this is at least t . HR3 is immediate from the definition. \square

Note that the constant boost of the grade cannot be lower than $\frac{1}{4}$, since for $t^* = \frac{1}{2}$ the expected grade is exactly $\frac{1}{2}$.

As for the soft requirements, it is immediate that $\mathbf{B}(\mathcal{LVM}) = \frac{1}{4}$. The expected verification [resp., bias] depend on the distribution p , but are easy to compute by integrating t [resp., $\frac{1}{4} + t^2 - t$] over p . E.g. for a uniform type distribution we would get $\text{ver}(\mathcal{LVM}, p) = \frac{1}{2}$ and $\text{bias}(\mathcal{LVM}, p) = \frac{1}{12}$.

⁴This introduces a small *negative* bias, which is not allowed, but we can fix this by adding a small constant to the grade, see below.

Reducing the bias. While increasing the parameter θ from 0 towards 1 clearly increases the amount of verification, we saw that the effect on the expected bias is non-monotone. Increasing it beyond 1 is unlikely to reduce the bias and therefore we need a different mechanism.

Polynomial Verification (PV) mechanism: The mechanism takes two parameters:

- an integer parameter $\kappa \geq 1$;
- a real-valued parameter $\theta \in [0, 1]$;

and is defined as follows.

- $\mathcal{PVM}_V^{\theta, \kappa}(\hat{t}_i) := \theta \cdot (\hat{t}_i)^{\frac{1}{\kappa}}$.
- $\mathcal{PVM}_G^{\theta, \kappa}(\hat{t}_i, \perp) := \frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + \frac{1}{\kappa \cdot \theta^{\kappa+1}} \sum_{\ell=1}^{\kappa} (\theta \cdot \hat{t}_i)^{\ell/\kappa}$.
- $\mathcal{PVM}_G^{\theta, \kappa}(\hat{t}_i, s_i) := \frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + \frac{1}{\theta} (1 + \frac{1}{\kappa}) s_i - \frac{1}{\kappa \cdot \theta^{\kappa+1}}$.

Note that for $\kappa = 1, \theta = 1$ the PV mechanism coincides with the LV mechanism.⁵ Just like the LV mechanism, any PV mechanism implements a (semi-polynomial) scoring rule. See the dash-dotted purple line in Fig. 4.

LEMMA 4.3. $\bar{g}_t(\hat{t}) = \frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + R^{1+\frac{1}{\kappa}}(\hat{t}, t)$;

LEMMA 4.4. $\frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + t^{1+\frac{1}{\kappa}} \in [t, t + \frac{1}{\kappa}]$.

PROOF. Clearly, $\frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}}$ is positive and upper bounded by $\frac{1}{\kappa+1}$. Now consider $t - t^{1+\frac{1}{\kappa}}$. This term is nonnegative and concave. By derivation, the maximum is obtained at $(\frac{\kappa}{\kappa+1})^\kappa$. Thus

$$\begin{aligned} t - t^{1+\frac{1}{\kappa}} &\leq \left(\frac{\kappa}{\kappa+1}\right)^\kappa - \left(\frac{\kappa}{\kappa+1}\right)^{\kappa(1+\frac{1}{\kappa})} \\ &= \left(\frac{\kappa}{\kappa+1}\right)^\kappa - \left(\frac{\kappa}{\kappa+1}\right)^{\kappa+1} = \frac{\kappa^\kappa(\kappa+1)}{(\kappa+1)^{\kappa+1}} - \frac{\kappa^{\kappa+1}}{(\kappa+1)^{\kappa+1}} \\ &= \frac{\kappa^{\kappa+1} + \kappa^\kappa - \kappa^{\kappa+1}}{(\kappa+1)^{\kappa+1}} = \frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}}. \quad \square \end{aligned}$$

PROPOSITION 4.5. The PV mechanism satisfies HR1 and HR2'. It satisfies HR3 iff $\frac{1}{\kappa \cdot \theta^{\kappa+1}} - \frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} \leq \xi$.

PROOF. Truthfulness (HR1) follows immediately from Lemma 4.3 since the expected grade is an affine transformation of $R^{1+\frac{1}{\kappa}}$, and thus also a proper scoring rule. As for HR2', note that the expected grade of a truthful agent of type t is $\frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + t^{1+\frac{1}{\kappa}}$. By Lemma 4.4 this is at least t . HR3 is respected by our restriction on the parameters. \square

As for the soft requirements, it is easy to see that the audit probability is monotonically increasing in both parameters. As $\kappa \rightarrow \infty$, it approaches a constant function that verifies every agent (unless $\hat{t}_i = 0$) w.p. θ .

We observe that the bias does not depend on θ at all, which may at first seem surprising. The bias for a truthful type- t agent is $\frac{\kappa^\kappa}{(\kappa+1)^{\kappa+1}} + t^{1+\frac{1}{\kappa}} - t$. By Lemma 4.4 this is at most $\frac{1}{\kappa+1}$ and thus both the maximal and the expected bias go to 0 as $\kappa \rightarrow \infty$.

Therefore, by selecting sufficiently large κ and sufficiently small θ , we approach a ‘perfect’ mechanism with no verification and no bias. However, we must not forget that the parameters are restricted

⁵We can also see that when $\theta = 1, \kappa \rightarrow \infty$ the mechanism approaches VERIFY-ALL. We get PAY-ALL when $\theta = \kappa = 0$ although we did not formally allow $\kappa = 0$.




Figure 4: Visualization of the Polynomial Verification mechanism with $\kappa = 2, \theta = 0.5$. The purple dash-dot line shows the expected grade of a truthful agent. As in Fig. 2, the shaded areas show the overall verification and bias under a uniform type distribution.

by the maximal allowed penalty ξ . Indeed, ignoring this restriction provides us with the noisy counterpart of the HUGE-PENALTY mechanism.

From two parameters to one. By the above discussion, once ξ is fixed and κ was selected, we should choose the minimal θ that satisfies HR3. Denote $\theta^*(\xi, \kappa) := \left((1 + 1/\kappa)^{-(\kappa+1)} + \kappa \cdot \xi \right)^{-1/(\kappa+1)}$.

COROLLARY 4.6. *For any $\xi \geq 0$ and $\kappa \in \mathbb{Z}_+$, any valid $\mathcal{PVM}^{\theta, \kappa}$ mechanism is dominated by $\mathcal{PVM}^{\theta^*(\xi, \kappa), \kappa}$.*

This is simply since mechanisms with $\theta < \theta^*(\kappa, \xi)$ violate HR3 (thus not valid), and mechanisms with $\theta > \theta^*(\kappa, \xi)$ have higher verification rate and same bias.

Thus, as in the deterministic verification case, we get a three-way tradeoff: fixing the maximal penalty ξ , the (integer) parameter κ balances verification vs. bias. However, in the PM class, allowing negative grades is essential. We conjecture that under limited liability, no valid mechanism is better than a convex combination of VERIFY-ALL and PAY-ALL.

In contrast to the deterministic case, we have no guarantee that the PM class provides the best possible tradeoff. However, the class provides a strong baseline for mechanisms that balance between verification and bias.

4.2 Alternative Approach: Type Histogram

The MCV mechanism can work without any information on the type distribution, but can exploit such information to better optimize the bias.

In this last section, we will assume the set of possible types is discrete with $|t - t'| > \epsilon$ for any $t \neq t'$. We further assume we have access to the *exact type histogram* (i.e., we know exactly the number of agents of each type, denoted $H(t) := |\{i \in N : t_i = t\}|$), and

show a ‘toy mechanism’ that exploits this information to provide a ‘perfect’ outcome: a unique Nash equilibrium in which *no agent is verified*, and the *maximal bias is arbitrarily small*, even when verification is noisy.

The equilibrium is not in dominant strategies, and thus the mechanism does not belong in the same class as the previous mechanisms, relaxing HR1 to HR1’: there is a unique Nash equilibrium, which is truthful. The mechanism also holds HR2’ and limited liability.⁶

The Histogram Mechanism: The mechanism \mathcal{HM}^α takes a parameter $\alpha \in (1, 1 + \epsilon)$ as input, as well as the histogram H of all agents’ types (as “additional information”). Define $\kappa := \lceil \frac{2}{\alpha-1} \rceil > \frac{2}{\epsilon}$.

- Set $\hat{H}(t) := |\{i : \hat{t}_i = t\}|$;
- $t_V \leftarrow \max\{t \in \mathcal{T} \text{ s.t. } \hat{H}(t) > H(t)\}$; // (maximal violating type if exists)
- Set $\mathcal{HM}_V^\alpha(\hat{t}_i) = 1$ if $\hat{t}_i = t_V$ and 0 otherwise;
- $\mathcal{HM}_G^\alpha(\hat{t}_i, \perp) = \epsilon + \hat{t}_i$;
- $\mathcal{HM}_G^\alpha(\hat{t}_i, s) = \epsilon + R^{1+\frac{1}{\kappa}}(\hat{t}_i, s)$.

PROPOSITION 4.7. *The histogram mechanism \mathcal{HM}^α respects limited liability, and has a unique Nash equilibrium, in which all agents are truthful and none are verified. Moreover, for any truthful type t agent, $\bar{g}_t \in [t, t + \epsilon]$ even off-equilibrium.*

Intuitively, the proof relies on two observations: first, in every profile where agents only lie upwards, there is at least one cheater that is verified with certainty, and thus experiences the scoring rule fully. Second, by Lemma 4.4 the proper scoring rule $R^{1+\frac{1}{\kappa}}$ implements a near-identity function when α is close to 1 (and thus κ is large), thus approximating agents’ true types to arbitrary precision.

5 EMPIRICAL DEMONSTRATION

In this section, we demonstrate the triple tradeoff between maximal penalty, average bias and verification for the MCV and PV classes. Since the tradeoff relies on an actual type distribution we employ both synthetic and real-world data.

Data. As real-world datasets, we used (a) SAT 2022 test scores: a standardized test widely used for college admissions in the United States, with statistics taken from [31]; (b) FICO scores statistics from 2018 [1].

Additionally, we use the synthetic Beta distribution with various parameters. $B(1, 1)$ is the Uniform distribution, and higher parameters correspond to lower variance, where MCV is expected to perform better. Both of our real-world datasets are shown in Fig. 3, together with the best-fit Beta distributions.

Hardware. We used a standard PC with intel Core i7-9700k CPU and 16GB RAM for running the simulations. The entire execution takes a couple of seconds.

Experimental setup. We set the value of the maximal penalty parameter ξ . Note that both mechanism classes essentially have a single parameter: $\gamma \in [0, 1]$ for MCV, and $\kappa \in \mathbb{Z}_+$ for PV. For each type distribution p , and $M \in \{\mathcal{M}^{\gamma, \xi}, \mathcal{PVM}^{\theta(\xi, \kappa), \kappa}\}$ we compute

⁶In fact, the mechanism holds a stronger version of HR2’, since truthful agents get at least their type in expectation whether they are verified or not.

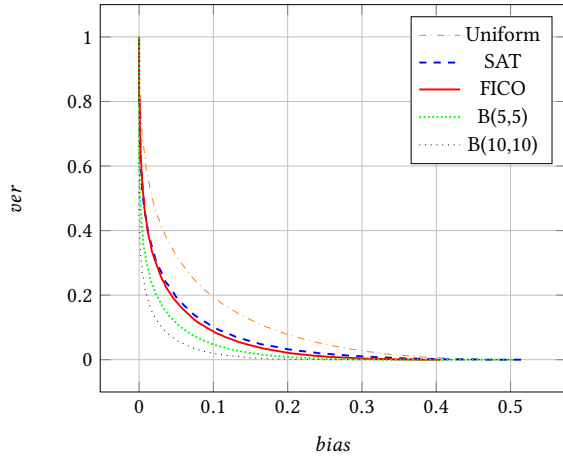


Figure 5: Efficiency curves for MCV mechanism under *Limited Liability* on several synthetic and real type distributions. The γ parameter ranges from 0 to 1, with higher values to the right.




Figure 6: Efficiency curves of MCV and PV mechanisms on the SAT type distribution. Lines are shown for $\xi = 1, 2, 4$ (and also $\xi = 0$ for MCV). The gray line shows the convex combination of VERIFY-ALL and PAY-ALL.

the bias(\mathcal{M}, p) and ver(\mathcal{M}, p) for different values of the parameter in range. We select γ as multiples of 0.01, and $\kappa \in \{1, \dots, 15, 20, 50\}$.

We compute bias() and ver() as *expected values*, by integrating over the type distribution p . We do not sample the types, the verification decision, or (in the noisy setting) the performance of the verified agent.⁷

5.1 Results

Deterministic verification. Figure 5(top) shows Efficiency curves of the MCV mechanism (i.e., under deterministic verification) for

⁷In fact, we do not even need to define the performance distribution T_i , since we only use its expectation t_i .

the various distributions, using $\xi = 0$ (limited liability). Higher γ values correspond to points towards the bottom right, i.e., with less verification and a higher bias.

We can see that MCV performs better as the variance of p decreases, and is doing much better on SAT and FICO than on the Uniform distribution.

For example, in SAT, we could guarantee an expected bias of less than 5% of the grade, with verifying only 20% of the students. Note that in the noiseless setting, the bias of the MCV mechanism also equals the error in grading (as the grade is deterministic).

Triple tradeoff. Fig. 6 compares the performance we get under deterministic verification (using MCV) and noisy verification (using PV), for maximal penalty $\xi \in \{1, 2, 4\}$. The figure shows results on the SAT distribution, results on other distributions are similar.

We also show results for MCV with limited liability, which still does better than PV with a high maximal penalty. Note that the best we can do in the noisy setting with limited liability is the baseline.

One thing that is evident from the figure is that we can significantly reduce the bias in both settings, but it is harder to control the verification amount in the noisy setting, as we currently have no solution between the Linear Verification mechanism ($\kappa = 1$) and the PAY-ALL mechanism.

6 CONCLUSION

In this paper, we considered possible ways to incentivize agents to truly report their type, when a high type leads to better utility (as in the case of grading an exam). While agents can lie freely about their type, the principal has certain verification power, that she prefers to use scarcely.

Our main result characterizes the optimal tradeoff between required verification and maximal/expected bias, when verification is guaranteed to reveal the true type. We note that our mechanism can also be used when agents have multi-dimensional types as in [12], as long as allocation probability is only determined by the agent’s own type (thus can be treated as ‘grade’), i.e., when there is no competition.

Noisy verification makes the task much more challenging, but still possible if we allow ‘negative grades’, and at the cost of larger bias and/or more verification. We further showed how having access to the exact type histogram leads to a mechanism that essentially eliminates both the bias and the required verification. While having such access is unlikely, we hope the ideas used in the proof could be useful to better exploit information on the type distribution.

The main remaining challenge is to prove a tight approximation bound for the noisy verification setting.

ACKNOWLEDGMENTS

The work of O. Ben-Porat was partially supported by the Israel Science Foundation (ISF; Grant No. 3079/24). We thank the anonymous reviewers for their valuable feedback.

REFERENCES

- [1] About-Rankings. 2018. Blog Post. <https://aboutranking.com/2018/01/21/how-a-credit-score-is-calculated-and-how-objective-it-is/>.
- [2] Charilaos Akasiadis and Georgios Chalkiadakis. 2017. Cooperative electricity consumption shifting. *Sustainable Energy, Grids and Networks* 9 (2017), 38–58.

- [3] Itai Ashlagi and Yannai A Gonczarowski. 2018. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory* 177 (2018), 405–425.
- [4] Ian Ball and Deniz Kattwinkel. 2025. Probabilistic Verification in Mechanism Design. arXiv:1908.05556 [econ.TH] <https://arxiv.org/abs/1908.05556>
- [5] Omer Ben-Porat, Yishay Mansour, Michal Moshkovitz, and Boaz Taitler. 2024. Principal-agent reward shaping in mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9502–9510.
- [6] Elchanan Ben-Porath, Eddie Dekel, and Barton L Lipman. 2014. Optimal allocation with costly verification. *American Economic Review* 104, 12 (2014), 3779–3813.
- [7] Simina Brânzei and Ariel D Procaccia. 2015. Verifiably truthful mechanisms. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*. 297–306.
- [8] Gabriel Carroll. 2019. Robust incentives for information acquisition. *Journal of Economic Theory* 181 (2019), 382–420.
- [9] Sofia Ceppi, Ian Kash, and Rafael Frongillo. 2019. Partial verification as a substitute for money. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1837–1844.
- [10] Geoffrey A Chua, Gaoji Hu, and Fang Liu. 2023. Optimal multi-unit allocation with costly verification. *Social Choice and Welfare* 61, 3 (2023), 455–488.
- [11] Markos Epitropou and Rakesh Vohra. 2019. Optimal on-line allocation rules with verification. In *Algorithmic Game Theory: 12th International Symposium, SAGT 2019, Athens, Greece, September 30–October 3, 2019, Proceedings 12*. Springer, 3–17.
- [12] Andrew Estornell, Sanmay Das, and Yevgeniy Vorobeychik. 2021. Incentivizing truthfulness through audits in strategic classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5347–5354.
- [13] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association* 102, 477 (2007), 359–378.
- [14] Jerry R Green and Jean-Jacques Laffont. 1986. Partially verifiable information and mechanism design. *The Review of Economic Studies* 53, 3 (1986), 447–456.
- [15] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. 2016. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*. 111–122.
- [16] IRS. 2013. SOI Tax Stats. <https://www.irs.gov/statistics/soi-tax-stats-individual-statistical-tables-by-tax-rate-and-income-percentile>.
- [17] Daniel Kahneman. 2003. Maps of bounded rationality: Psychology for behavioral economics. *American economic review* 93, 5 (2003), 1449–1475.
- [18] Sagi Levanon and Nir Rosenfeld. 2021. Strategic classification made practical. In *International Conference on Machine Learning*. PMLR, 6243–6253.
- [19] Shengwu Li. 2017. Obviously strategy-proof mechanisms. *American Economic Review* 107, 11 (2017), 3257–3287.
- [20] Yunan Li. 2020. Mechanism design with costly verification and limited punishments. *Journal of Economic Theory* 186 (2020), 105000.
- [21] Reshef Meir, Ariel D Procaccia, and Jeffrey S Rosenschein. 2012. Algorithms for strategyproof classification. *Artificial Intelligence* 186 (2012), 123–156.
- [22] Reshef Meir, Jonathan Wagner, and Omer Ben-Porat. 2026. Truthful Reporting of Competence with Minimal Verification. <https://doi.org/10.48550/arXiv.2602.14076> arXiv:2602.14076 [cs.GT] Full version of a paper accepted to AAMAS 2026.
- [23] Dilip Mookherjee and Ivan Png. 1989. Optimal auditing, insurance, and redistribution. *The Quarterly Journal of Economics* 104, 2 (1989), 399–415.
- [24] Tymofiy Mylovanov and Andriy Zapechelnyuk. 2017. Optimal allocation with ex post verification and limited penalties. *American Economic Review* 107, 9 (2017), 2666–2694.
- [25] Noam Nisan et al. 2007. Introduction to mechanism design (for computer scientists). *Algorithmic game theory* 9 (2007), 209–242.
- [26] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. 2008. Fault tolerant mechanism design. *Artificial Intelligence* 172, 15 (2008), 1783–1799.
- [27] Drazen Prelec. 2004. A Bayesian truth serum for subjective data. *science* 306, 5695 (2004), 462–466.
- [28] Leonard J Savage. 1971. Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* 66, 336 (1971), 783–801.
- [29] Reinhard Selten. 1998. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics* 1 (1998), 43–61.
- [30] Reinhard Selten and R Selten Bielefeld. 1988. *Reexamination of the perfectness concept for equilibrium points in extensive games*. Springer.
- [31] Wikipedia contributors. 2025. SAT – Wikipedia, The Free Encyclopedia. <https://en.wikipedia.org/wiki/SAT>. <https://en.wikipedia.org/wiki/SAT> Accessed: 2025-07-28.
- [32] Jens Witkowski and David Parkes. 2012. A robust bayesian truth serum for small populations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 26. 1492–1498.
- [33] Aviv Zohar and Jeffrey S Rosenschein. 2008. Mechanisms for information elicitation. *Artificial Intelligence* 172, 16-17 (2008), 1917–1939.
- [34] Tijana Zrnic, Eric Mazumdar, Shankar Sastry, and Michael Jordan. 2021. Who leads and who follows in strategic classification? *Advances in Neural Information Processing Systems* 34 (2021), 15257–15269.