

Byzantine Game Theory: Sun Tzu's Boxes

Andrei Constantinescu

ETH Zürich
Zürich, Switzerland
aconstantine@ethz.ch

Roger Wattenhofer

ETH Zürich
Zürich, Switzerland
wattenhofer@ethz.ch

ABSTRACT

We introduce the Byzantine Selection Problem, living at the intersection of game theory and fault-tolerant distributed computing. Here, an event organizer is presented with a group of n agents, and wants to select $\ell < n$ of them to form a team. For these purposes, each agent i self-reports a positive skill value v_i , and a team's value is the sum of its members' skill values. Ideally, the value of the team should be as large as possible, which can be easily achieved by selecting agents with the highest ℓ skill values. However, an unknown subset of at most $t < n$ agents are byzantine and hence not to be trusted, rendering their true skill values as 0. In the spirit of the distributed computing literature, the identity of the byzantine agents is not random but instead chosen by an adversary aiming to minimize the value of the chosen team. Can we still select a team with good guarantees in this adversarial setting? As it turns out, deterministically, it remains optimal to select agents with the highest ℓ values. Yet, if $t \geq \ell$, the adversary can choose to make all selected agents byzantine, leading to a team of value zero. To provide meaningful guarantees, one hence needs to allow for randomization, in which case the expected value of the selected team needs to be maximized, assuming again that the adversary plays to minimize it. For this case, we provide linear-time randomized algorithms that maximize the expected value of the selected team.

KEYWORDS

Game Theory; Byzantine Fault Tolerance; Robust Optimization; Algorithms; Mechanism Design

ACM Reference Format:

Andrei Constantinescu and Roger Wattenhofer. 2025. Byzantine Game Theory: Sun Tzu's Boxes. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 10 pages.

1 INTRODUCTION

The Art of War is an ancient Chinese military treatise attributed to Sun Tzu. It is widely considered the first known origin of game theory. Sun Tzu advocates misleading the enemy through deception, discusses asymmetric information setups and even zero-sum games.

While Sun Tzu's discussions were qualitative only, around two millennia later, in the 16th and 17th centuries, the first foundations for analyzing games of chance were laid out. John von Neumann eventually turned game theory into a rigorous academic subject,

culminating with his celebrated 1944 book *Theory of Games and Economic Behavior* co-authored with Oskar Morgenstern.

Ever since, the agents participating in the game have been typically considered to be strategic, selfish, egoistic, or sometimes altruistic, but they would essentially always follow the rules of the game. In this paper, we take game theory back to Sun Tzu's origins of lies and deception. We believe this take on game theory opens up a whole class of exciting problems that have remained largely unexplored.

Let us introduce our paper with a concrete puzzle. A game show host puts four boxes in front of you, each with a value written on top: 8, 7, 5 and 4. Each box contains its advertised amount of value, but there is a catch: one box is a lie and contains nothing! You have to choose one box, and you will receive the money inside. What is your strategy to win the maximum possible amount of money?

You quickly understand that there is no way to guarantee winning any money: no matter which box you pick, it could be empty. However, what if you flipped a fair coin and chose the first box (with value 8) if the coin landed on heads and the second box (with value 7) if it landed on tails? At most one of those boxes can be empty, so with probability at least 50%, you do not choose the empty box, so in expectation you earn at least $\min(8, 7)/2 = 3.5$. But can you do even better? Before reading on, we would like to encourage you to think about this puzzle for a moment.

This problem can be generalized to n boxes promising monetary rewards $v_1, \dots, v_n$. At most $t < n$ boxes are a lie, and one has to choose $\ell < n$ boxes, winning the total inside these ℓ boxes. While we presented the problem as a game show puzzle, one can think of real-life situations where the same setup emerges: say you are running an auction for ℓ identical perishable items (e.g., food close to the expiration date, seat upgrades to first class on tomorrow's flight, leftover concert tickets). People place n bids $v_1, \dots, v_n$, and you want to maximize your revenue. For transparency reasons, you want to publish the auction mechanism prior to the auction, and you worry that doing so might lead to the emergence of malicious bidders that bid adversarially to win the auction and then walk away after winning an item, leaving you with unsold items. Provided no binding terms enforce honoring the winning bids, your best course of action is to design your mechanism to account for the presence of a certain threshold t on the number of malicious bidders.

Another example where the problem occurs is when we have to choose ℓ individuals out of n applicants, like when putting together a university chess team. Again, up to t applicants could be imposters, bragging about their chess strength (the number on their box) despite barely knowing the rules of chess (the box is empty).

Let us now return to our example with four boxes with values $\mathbf{v} = (8, 7, 5, 4)$. The first thought following the fair-coin strategy would be to use a biased coin to balance the expectations from the two cases: choose box 1 with probability p_1 and box 2 with probability



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

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

$p_2 = 1 - p_1$. Then, if the first box is empty, we get an expectation of $7 \cdot p_2$, while if the second box is empty, we get an expectation of $8 \cdot p_1$. Setting the two products to be equal gives us the solution $\mathbf{p} = (\frac{7}{15}, \frac{8}{15})$, giving in both cases an expectation of $\frac{7 \cdot 8}{15} \approx 3.73$, which is better than our previous expected value of 3.5. Going further, one would think of generalizing the idea to randomize between all four boxes: choose each box $i \in \{1, \dots, 4\}$ with probability p_i such that $\sum_{i=1}^4 p_i = 1$ and $v_1 \cdot p_1 = \dots = v_4 \cdot p_4$. Solving the ensuing equations yields the solution $\mathbf{p} = (\frac{35}{201}, \frac{40}{201}, \frac{56}{201}, \frac{70}{201})$. In all four cases (i.e., one for each possibility of which box is empty), the incurred expectation is $(4 - 1) \cdot \frac{280}{201} \approx 4.18$, which is better than our last solution — but is this the best possible? Intuitively, one could think that the more boxes we consider, the better our expectation will be. However, this turns out not to be the case. In particular, ignoring box 4 is a good idea: suppose we require that $\sum_{i=1}^3 p_i = 1$ and $v_1 \cdot p_1 = \dots = v_3 \cdot p_3$. Solving the equations yields $\mathbf{p} = (\frac{35}{131}, \frac{40}{131}, \frac{56}{131}, 0)$, incurring in all three cases (depending on which box out of the first three is empty) an expectation of $(3 - 1) \cdot \frac{280}{131} \approx 4.27$, which turns out to be the unique optimal solution.

The Problem. More formally, the boxes problem can be formulated as follows: n boxes indexed by $[n] := \{1, \dots, n\}$ are given with promised monetary amounts $\mathbf{v} = (v_1, \dots, v_n)$ written on them. Without loss of generality, assume $v_1 \geq \dots \geq v_n > 0$.¹ It is known that $t < n$ of them are empty,² termed *byzantine*, following tradition in the distributed computing literature; the other $n - t$ contain the advertised amounts. No prior over which boxes are byzantine is provided other than the fact that there are t of them. The goal is to select $\ell < n$ boxes to open so as to maximize the worst-case expected total amount of money in the selected boxes. Formally, a randomized algorithm is sought that samples from a probability distribution p over size- ℓ subsets of $[n]$ such that the quantity $\text{value}(p)$ is maximized, which is defined as the minimum over all possibilities for the size- t set $B \subseteq [n]$ of byzantine boxes of the expectation incurred by p assuming boxes in B are empty and the other boxes have the advertised values. See Section 3 for a fully formal definition and further discussion of the model.

Our Contribution. We give linear-time algorithms for sampling from a distribution p maximizing $\text{value}(p)$.³ We start by doing this for the case $\ell = 1$, for which we show that the approach we used to solve our four-boxes example is indeed correct in general: a $\text{value}(p)$ -maximizing distribution p always exists such that for some $t + 1 \leq i \leq n$ we have $v_1 \cdot p_1 = \dots = v_i \cdot p_i$ and $p_{i+1} = \dots = p_n = 0$. For a fixed i , there exists a single distribution, given by setting $p_1, \dots, p_i$ proportionally to $\frac{1}{v_1}, \dots, \frac{1}{v_i}$. Finding the i maximizing the expectation can then be easily achieved in linear time. To show that an optimal solution maximizing $\text{value}(p)$ has this shape, we combine an exchange argument with combinatorial reasoning about the dual of a linear program (LP). Afterwards, we move on to the general- ℓ case, for which such reasoning no longer suffices. Instead, we combine multiple techniques to recover a linear-time algorithm for this case as well. First, we use prior results on randomized rounding to show that it suffices to reason

about the marginals of p along the n boxes. Then, we use a succession of exchange arguments to prove increasingly more refined forms of an optimal solution. The final form can be elegantly interpreted through a visual metaphor: pouring ℓ units of water into n unit-volume vessels. Using this metaphor, we devise a linear-time two-pointer approach that iterates through possible solutions in a principled manner and returns an optimal one. A technical overview of our results and techniques can be found in Section 4.

2 RELATED WORK

Our paper brings together multiple fields often studied separately. In the following, we outline numerous connections between byzantine elements (and, more specifically, our selection problem) and other areas studying related notions and problems. Outlining these connections can lead to fruitful interdisciplinary future insights.

Fault-Tolerant Distributed Computing. The fault-tolerant distributed computing literature often considers a setup with n parties, out of which at most t are corrupted and may exhibit arbitrary deviations from the intended behavior. Following the tradition of the field, such parties are called *byzantine*, while the other (at least) $n - t$ parties are referred to as *honest* and will follow the intended behavior. One of the most prominently-studied problems is binary *Byzantine Agreement* (BA), where each party i has an input $x_i \in \{0, 1\}$, and must eventually provide an output $y_i \in \{0, 1\}$, such that any two honest parties give the same output (the *agreement* condition) and this output is moreover the input of some honest party (the *validity* condition). Depending on additional assumptions (e.g., synchronous communication, cryptographic setup, randomization), the difficulty of designing a correct BA protocol varies, with protocols resilient against $t < n/2$ corruptions existing in certain settings. The core challenge here lies in the global lack of trust: nobody can be sure who are the corrupted parties, and hence protocols resilient against even a small number of corruptions can be highly non-trivial. In contrast, assuming the existence of a *trusted third party* (TTP) that is guaranteed to be honest makes BA essentially trivial: all parties send their x_i 's to the TTP, which takes a majority vote and then sends back the outcome to all parties. Consequently, efforts in the literature boil down to simulating the behavior of the TTP in an untrusted setting. The keen reader might have already noted that the simple TTP protocol above fails to work past the $t < n/2$ corruption threshold (if too many parties are byzantine, they can outnumber the honest votes), and, in fact, no TTP protocol can surpass this bound, implying that $t < n/2$ is best possible in the distributed setting too. Hence, from a fresh point of view, understanding the difficulty of problems in fault tolerance begins with understanding their difficulty in the TTP setting. While this was mostly straightforward for binary BA (and is hence rarely considered), the story becomes interesting again for agreement problems over other domains:

When the inputs are real numbers, *Honest-Range Validity* [25] requires that the agreed-upon value is between the smallest and largest honest inputs. A stronger requirement stands in variations of *Median Validity* [18, 47, 56], asking that the output is close to the median of the honest inputs. An orthogonal generalization of *Honest-Range Validity* is when the inputs are D -dimensional real vectors, where *Convex Validity* [48, 57] requires the output vector

¹Boxes i with $v_i = 0$ can be safely ignored. This removes edge cases in the analysis.

²Alternatively, at most t are empty. This does not change the analysis.

³The underlying distribution may have exponentially-large support, so outputting it explicitly may not be possible. We also show that, at the cost of an extra factor of n in the time complexity, an explicit p (with linearly-sized support) can be computed.

to be in the convex hull of the honest input vectors. This has been further generalized to abstract convexity spaces, including graphs and lattices [19, 52]. When the inputs are linear orders over a set of alternatives, also known as *rankings* in social-choice theoretic terms, [46] considers *Pareto Validity*, requiring that whenever all honest inputs rank a above b , so does the output ranking. The same paper also studies a second problem, requiring an output ranking that is close to the *Kemeny median* of the honest rankings. Another notion inspired by social-choice is that of *Voting Validity* [59].

Note how all these problems are non-trivial even in the TTP setting: how should the TTP aggregate n votes to yield a valid outcome when t of the votes might be corrupted and should not be considered? The problem we study is precisely of this flavor: the TTP receives the box values $v_1, \dots, v_n$ and must output a set of boxes to open to minimize the loss from empty boxes⁴ (possibly using randomization). In general, recent results indicate that the difficulty of solving the “centralized” TTP problem often matches that of its decentralized version, at least when considering BA with non-trivial validity conditions [13, 14, 16]. Hence, it is imperative to understand the difficulty of the centralized problem in more settings, and our paper sets out to do just this for our selection problem.

Malice in Game Theory. While most of game theory considers rational (selfish) agents, a few works consider the presence of malicious agents, who derive their utility from making the system perform poorly or from others’ disutility. This has been particularly investigated in distributed settings, where considering malicious actors is the norm. An excellent survey of work at the interface of rational and malicious behavior can be found in [50], which also introduces the *Price of Malice*, quantifying the degradation of the performance of a system of rational agents with the introduction of malicious actors.⁵ The same paper analyzes the price of malice in a virus inoculation game, and subsequent papers study it and other notions of malice in other classes of games [5, 11, 53]. Particularly relevant to the distributed computing field is the concept of *BAR fault tolerance* [2], which requires distributed protocols to withstand Byzantine, Altruistic, and Rational (BAR) behavior. Specifically, they must tolerate a constant fraction of Byzantine agents, as is standard, while ensuring that the remaining agents — who are rational — have sufficient incentives to follow the protocol. Subsequent work adopted the notion and applied it in a variety of distributed settings, spanning theory and practice [15, 41, 42]. A number of works define fault-tolerant solution concepts, e.g., fault-tolerant Nash Equilibria, and apply them to study the fault tolerance of various games [26, 35]. Also relevant are the appealing results in [31], which find that “large” games are naturally fault-tolerant. Malice can also manifest as *spite*. In this context, [9, 49] examine auctions with spiteful agents, i.e., agents who derive utility from others’ disutility. Conversely, several works consider altruistic agents (which, in the previous context, referred to agents who follow the protocol regardless of incentives, but this is not a strict requirement), modeled as agents for whom larger utilities of others translate to larger utilities for themselves [34, 45].

⁴Technically, maximize the win from non-empty boxes (the sum of the two is non-constant, so there is a distinction).

⁵The notion is inspired by the celebrated *Price of Anarchy*, which instead targets the difference between selfish and collaborating agents.

Stackelberg and Security Games. Our problem can be seen as a *zero-sum Stackelberg (maxi-min) game* where the leader commits to a (potentially randomized) strategy of picking ℓ boxes, and then the follower (i.e., the adversary), knowing the leader’s choice, chooses t boxes to nullify. The leader (follower) attempts to maximize (minimize) the expected sum of the ℓ selected boxes post-nullification. Related is the class of *security games*, which have received extensive attention in recent years [55]. In one variant [39], there are n potential targets, an attacker, and a defender. The defender moves first, chooses t targets to protect (using a potentially randomized strategy), and then the attacker chooses ℓ targets to attack, knowing the strategy of the defender. The utilities of the two players additively depend solely on attacked targets, with the goal of the attacker being to attack undefended targets and that of the defender being to protect attacked targets, but the game is not necessarily zero-sum. The existing literature largely concerns the non-zero-sum case with $\ell = 1$ attacked targets. On the computational front, security games are largely amenable to techniques from combinatorial optimization [58], often linear programming [38]. Results become scarcer when seeking more efficient algorithms: [37] for the $\ell = 1$ case, and [39] for the general- ℓ case if instead of Stackelberg Equilibria we require Nash Equilibria, the former becoming computationally demanding. Our problem corresponds to a zero-sum variant of the previously described security game with the order the attacker and defender play in reversed. We note that this does not fundamentally change the game since for zero-sum games, the maxi-min and mini-max values coincide by von Neumann’s theorem [51]. Stackelberg and Nash Equilibria are also closely tied in the zero-sum case, so the algorithm of [39] can be used for our setting, but it is arguably more complicated than our approach (and quadratic instead of linear).

Robust Combinatorial Optimization. An active area of operations research concerns optimization under uncertainty [6]: when there is uncertainty in the constraints or objective. For our purposes, let us restrict ourselves to a problem template where only the objective is uncertain: the (not necessarily continuous) feasible region is known, denoted by $\mathcal{F} \subseteq \mathbb{R}^n$, and we are interested in $\max_{x \in \mathcal{F}} c^T x$. Instead of knowing c , we only know an uncertainty set $\mathcal{U}$ such that $c \in \mathcal{U}$. No probability distribution over $\mathcal{U}$ is supplied: we seek a solution x that maximizes $c^T x$ in the worst-case, making for the *maxi-min objective* $\max_{x \in \mathcal{F}} \min_{c \in \mathcal{U}} c^T x$. This setup is very flexible, as $\mathcal{F}$ can range from a polytope in the continuous case to the set of s - t paths or spanning trees of a graph in the discrete one. The maxi-min objective can be replaced with regret-inspired variants. The uncertainty set $\mathcal{U}$ can take various shapes, with the two most prominent ones being *discrete uncertainty*: $\mathcal{U}_D = \{c_1, \dots, c_k\}$ and *interval uncertainty*: $\mathcal{U}_I = [a_1, b_1] \times \dots \times [a_n, b_n]$. Interval uncertainty admits a variant in the spirit of byzantine fault-tolerance, introduced in [7]: given a threshold Γ define Γ -*interval uncertainty* $\mathcal{U}_I^\Gamma$ such that $c \in \mathcal{U}_I^\Gamma$ if $c \in \mathcal{U}_I$ and $|\{i : c_i \neq b_i\}| \leq \Gamma$. See [1, 10, 29, 30, 36] for excellent surveys of results and techniques. One of the basic cases considered in the robust optimization literature concerns $\mathcal{F} = \{x \in \{0, 1\}^n : \sum_{i=1}^n x_i = \ell\}$, the so-called *Selection Problem*: see [36] for a compilation of results under various objective and uncertainty set assumptions. Of particular interest to us is the case with the normal maxi-min objective and $\Gamma = t$ -interval uncertainty for $\mathcal{U}_I = [0, v_1] \times \dots \times [0, v_n]$: this corresponds

exactly with choosing to open ℓ boxes, out of which the adversary can nullify t . Note, however, that this only models the deterministic part of our paper (which is straightforward): it cannot model committing to a randomized strategy of which ℓ boxes to open to which the adversary replies by nullifying t so as to minimize the expectation. One of the few papers considering such randomized strategies is [44], where the authors show that for discrete and interval uncertainty sets, it is possible to optimize the regret objective in polynomial time as long as (non-robust) optimization over $\mathcal{F}$ is feasible polynomially (which trivially holds for the Selection Problem). Note that their result does not target Γ -interval uncertainty and is for the regret objective, hence not applicable to us.

Further Related Work. In election bribery [27], the goal is to alter the voting outcome by bribing at most t voters to change their votes, akin to a byzantine adversary. However, the focus is typically on understanding the impact under a given voting rule rather than designing fault-tolerant rules. Our byzantine selection problem shares similarities with various settings studied in online decision-making, including adversarial bandit and expert learning [43, 60], the byzantine secretary problem [8], and, more broadly, prophet inequalities [21]. Byzantine elements have also been explored in statistical learning with adversarial noise [12, 22–24, 40]. In the full version of our paper [20], we provide a more detailed discussion of the connections outlined in this paragraph.

3 PRELIMINARIES

Sets and Distributions. Given a non-negative integer k , write $[k] := \{1, 2, \dots, k\}$. Given a set S and a non-negative integer k , we write $S^{(k)} := \{S' \subseteq S \mid |S'| = k\}$ for the set of k -element subsets of S . For finite S , we write $\Delta(S) := \{p : S \rightarrow [0, 1] \mid \sum_{s \in S} p_s = 1\}$ for the set of probability distributions over S . For technical reasons, given a non-negative number x , we also define $\Delta_x(S) := \{p : S \rightarrow [0, 1] \mid \sum_{s \in S} p_s = x\}$, the set of what we will call *pseudo-distributions* of sum x over S . Note that $\Delta_1(S) = \Delta(S)$.

Setup. We consider a setting with n boxes indexed by the set $[n]$. Each box $i \in [n]$ has a real number $v_i > 0$ written on it. For ease of presentation, except where stated otherwise, we will assume that $v_1 \geq \dots \geq v_n$. Each box i may either be *honest*, in which case it contains v_i units of money, or *byzantine*, in which case it contains no money, without loss of generality. Assume $B \subseteq [n]$ is the set of byzantine boxes, then $[n] \setminus B$ is the set of *honest* boxes. Given the set of byzantine boxes B , the payoff incurred by opening box i is $\text{pay}(i, B) = v_i$ if $i \notin B$ and $\text{pay}(i, B) = 0$ otherwise. The total payoff incurred by opening a set of boxes $S \subseteq [n]$ is defined additively as $\text{pay}(S, B) := \sum_{i \in S} \text{pay}(i, B)$.

Problem Definition. We are interested in designing a (for now deterministic) mechanism M that, given n , the mapping $v : [n] \rightarrow \mathbb{R}^+$ and two numbers $1 \leq t, \ell < n$ select a size- ℓ subset of boxes $M(n, v, t, \ell) = S \subseteq [n]$ to open such that the worst-case payoff is maximized with respect to all options for the size- t subset $B \subseteq [n]$ of byzantine boxes.⁶ Formally, we want to design a M such that:

$$M(n, v, t, \ell) \in \operatorname{argmax}_{S \subseteq [n]^{(\ell)}} \min_{B \in [n]^{(t)}} \text{pay}(S, B) \quad (1)$$

⁶Note that requiring $|S| \leq \ell$ and $|B| \leq t$ instead of $|S| = \ell$ and $|B| = t$ would not change the problem.

The right-hand side of Eq. (1) can be seen as a game between the mechanism and an adversary: the mechanism picks the set S of boxes to open, then the adversary (knowing S) selects the set B of byzantine boxes. This is a zero-sum game where the mechanism aims to maximize the payoff, while the adversary to minimize it.⁷

Let us define $\text{value}(S) := \min_{B \in [n]^{(t)}} \text{pay}(S, B)$, in which case Eq. (1) asks that the selected S maximizes $\text{value}(S)$. Our problem, as stated so far, is easy to solve: from any set S of selected boxes, the adversary will choose to nullify the largest $\min\{t, \ell\}$ values by including the corresponding boxes in B . Hence, to maximize payoff, the mechanism should select boxes with the ℓ largest values, from which the largest $\min\{t, \ell\}$ will be nullified by the adversary.

Theorem 1. Among *deterministic* mechanisms, selecting $S^* = \{1, 2, \dots, \ell\}$ achieves the highest possible payoff:

$$\max_{S \in [n]^{(\ell)}} \min_{B \in [n]^{(t)}} \text{pay}(S, B) = \text{value}(S^*) = \begin{cases} \sum_{i=t+1}^{\ell} v_i & t < \ell \\ 0 & t \geq \ell \end{cases}$$

As an immediate consequence, deterministic mechanisms fail to provide any meaningful guarantees in the simplest case $t = \ell = 1$. Naturally, the next step is to allow for randomization, in which case we need to discuss the power of the adversary, i.e., how much they are allowed to know about the random decisions of the mechanism prior to selecting B . For instance, a *strong adversary*⁸ knows both the mechanism and the randomness before picking which boxes are byzantine. Since knowing the randomness renders any randomized mechanism deterministic, randomization does not help against a strong adversary. On the other hand, an *oblivious/weak adversary* has access to the mechanism, but not to the random bits used. Formally, the game played against an oblivious adversary proceeds as follows: the mechanism outputs a probability distribution over sets S of ℓ boxes each, say $p \in \Delta([n]^{(\ell)})$, then the adversary (knowing p) selects the set B of at most t byzantine boxes,⁹ and finally a set $S \sim p$ is sampled, incurring a payoff of $\text{pay}(S, B)$. Before the final set S is sampled, the outcome of the game is a probability distribution over possible payoffs, allowing for the formulation of various optimization goals. Most prominently, one can optimize for the expected payoff, but this is not the only option: a risk-averse mechanism user might prefer an expected payoff of 99 that guarantees a payoff of at least 50 in all realizations of the randomness to an expected payoff of 100 that leaves a positive probability of getting payoff 1. In the paper, we will assume the expectation objective and an oblivious adversary, unless stated otherwise. With these assumptions we seek a randomized mechanism M such that:

$$M(n, v, t, \ell) \in \operatorname{argmax}_{p \in \Delta([n]^{(\ell)})} \min_{B \in [n]^{(t)}} \mathbb{E}_{S \sim p}[\text{pay}(S, B)] \quad (2)$$

As before, we define $\text{value}(p) := \min_{B \in [n]^{(t)}} \mathbb{E}_{S \sim p}[\text{pay}(S, B)]$, in which case Eq. (2) asks that the selected p maximizes $\text{value}(p)$.

4 RESULTS AND TECHNICAL OVERVIEW

The case $\ell = 1$ allows for simpler arguments and characterizations of optimal solutions (and generally less machinery). To solve it, we

⁷In fact, Eq. (1) corresponds to the game-theoretic maxi-min solution concept.

⁸Also known as an *adaptive offline adversary* in the context of online algorithms.

⁹Since the adversary plays second, it does not help them to randomize their strategy.

will use a natural exchange argument to show that an optimal solution satisfying $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$ always exists. Restricting our search to such distributions p is particularly useful because we then know that the worst-case set of byzantine boxes is $\{1, \dots, t\}$. These observations let us cast the problem as a simple linear program (LP). Reasoning combinatorially about its dual then shows that an optimal solution taking a very elegant form always exists. In particular, this solution satisfies $v_1 \cdot p_1 = \dots = v_i \cdot p_i$ and $p_{i+1} = \dots = p_n = 0$, for some $t+1 \leq i \leq n$. In other words, it is optimal to select a prefix of boxes $1, \dots, i$ of length at least $t+1$ and put all the probability mass on it such that the expected values $v_j \cdot p_j$ are equal on the prefix. For a fixed i , the resulting distribution is unique and given by setting $(p_1, \dots, p_n)$ proportionally to $(\frac{1}{v_1}, \dots, \frac{1}{v_i}, 0, \dots, 0)$, in which case one can check that $\text{value}(p) = \frac{i-t}{\sum_{j=1}^i \frac{1}{v_j}}$. It is then easy to compute this value in linear time for all $t+1 \leq i \leq n$ and take the maximum. We prove the required claims in section Section 5.

Theorem 2. Assume $\ell = 1$ and define p^i for $t+1 \leq i \leq n$ to be the unique distribution such that $(p_1^i, \dots, p_n^i)$ is proportional to $(\frac{1}{v_1}, \dots, \frac{1}{v_i}, 0, \dots, 0)$, in which case $\text{value}(p^i) = \frac{i-t}{\sum_{j=1}^i \frac{1}{v_j}}$. Then, among distributions $p \in \Delta([n])$, the maximum $\text{value}(p)$ is attained at one of $p^{t+1}, \dots, p^n$, and we can determine which one in linear time.

For the harder general case, it is no longer the case that a value-maximizing distribution of such an attractive shape exists. First and foremost, this is because we are now dealing with distributions over size- ℓ subsets of $[n]$. To overcome this first obstacle, linearity of expectation gives us that $\text{value}(p)$ for some $p \in \Delta([n]^{(\ell)})$ is uniquely determined by the marginals $p'_i := \mathbb{P}_{S \sim p'}(i \in S)$ for $i \in [n]$. Naturally, the marginals are between 0 and 1 and sum up to ℓ since the sampled set S satisfies $|S| = \ell$, from which $p' \in \Delta_\ell([n])$. Note p' is *not* just a distribution scaled up by a factor of ℓ because of the constraint that $p'_i \in [0, 1]$ for $i \in [n]$. Still, it would be convenient to optimize directly over $p' \in \Delta_\ell([n])$, a much lower-dimensional object. Perhaps surprisingly, this is something that we will be able to do by invoking results on randomized rounding: for any $p' \in \Delta_\ell([n])$, there exists $p \in \Delta([n]^{(\ell)})$ such that p' gives the marginals of p , i.e., $p'_i = \mathbb{P}_{S \sim p}(i \in S)$ for $i \in [n]$, and we can sample from such a p in linear time. Hence, our new goal will be to find a $p' \in \Delta_\ell([n])$ maximizing the appropriately redefined $\text{value}(p')$.

Having retaken the problem into the realm of combinatorial tractability, we then use a similar exchange argument to show that a value-maximizing p' exists such that $v_1 \cdot p'_1 \geq \dots \geq v_n \cdot p'_n$. Past this point, unfortunately, optimal solutions no longer exhibit the elegant form that we could prove by considering the dual in the $\ell = 1$ case.¹⁰ To pass this hurdle, we take a different approach, proving a series of more refined forms of the previous inequality. The last of these can be attractively visualized through a physical metaphor involving pouring ℓ units of liquid into a histogram consisting of n rectangular vessels where box i corresponds to a $v_i \times (1/v_i)$ vessel. The amount of water poured into each vessel i corresponds to p'_i , justifying why the rectangles were chosen to have area 1. The pseudo-distributions

¹⁰We can still cast the problem as an LP and take its dual, but this requires adding the constraints $p'_i \leq 1$ to the primal since they no longer follow from the sum constraint, leading to a dual that is harder to analyze and no longer implies the required property.

p' that we will need to consider will be uniquely determined by the water level in the first container. This way, we will simulate in linear time (in a two-pointer fashion) the process of continuously decreasing the water level in the first container and find the moment in time where the value of the corresponding p' is maximized. We note that this algorithm can, of course, also be used when $\ell = 1$, but the details are considerably trickier and do not fully exploit the structure present in that case. All these considerations are carried out in Section 6.

Theorem 3. Assume $\ell \geq 1$. We can sample from a distribution $p \in \Delta([n]^{(\ell)})$ that maximizes $\text{value}(p)$ in linear time. If required, an explicit such p can be computed in $O(n^2)$ time.

5 THE CASE $\ell = 1$

As a warm-up, to build intuition and avoid some notational burdens of the general case, we begin with the case $\ell = 1$, where the goal is to select a single “winner” box with the highest value possible. In the presence of t byzantine boxes, for randomized mechanisms, this means outputting a distribution $p \in \Delta([n])$ maximizing

$$\text{value}(p) = \min_{B \in [n]^{(t)}} \mathbb{E}_{i \sim p}[\text{pay}(i, B)] = \min_{B \in [n]^{(t)}} \sum_{i \notin B} v_i \cdot p_i = \sum_{i=t+1}^n x_i \quad (3)$$

where $x_1, \dots, x_n$ are the values $v_1 \cdot p_1, \dots, v_n \cdot p_n$ ordered such that $x_1 \geq \dots \geq x_n$. The first equality follows by linearity of expectation, while the second holds because the best strategy for the adversary is to choose the byzantine boxes have the highest t values $v_i \cdot p_i$.

Lemma 4. There exists a distribution $p \in \Delta([n])$ maximizing $\text{value}(p)$ that satisfies $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$.

PROOF. Among distributions p maximizing $\text{value}(p)$, consider one such that for any $i < j$ with $v_i = v_j$ we have $p_i \geq p_j$. We show that this p has the required property. Assume the contrary, then for some $i < j$ we have $v_i \cdot p_i < v_j \cdot p_j$. Recall that $v_i \geq v_j$, and moreover, since in case $v_i = v_j$ the condition $v_i \cdot p_i < v_j \cdot p_j$ reduces to $p_i < p_j$, which can not hold by our choice of p , we have $v_i > v_j$. Let us now construct $p' : [n] \rightarrow [0, 1]$ as follows:

$$p'_k = \begin{cases} \frac{v_j \cdot p_j}{v_i} & k = i \\ \frac{v_i \cdot p_i}{v_j} & k = j \\ p_k & k \notin \{i, j\} \end{cases}$$

First, note that the values $v_1 \cdot p_1, \dots, v_n \cdot p_n$ and $v_1 \cdot p'_1, \dots, v_n \cdot p'_n$ coincide except for the i 'th and j 'th entries being swapped, so $\text{value}(p) = \text{value}(p')$. Furthermore, it turns out that p' does not use the whole available probability mass:

$$\sum_{k=1}^n p'_k = p'_i + p'_j + \sum_{\substack{k=1 \\ k \notin \{i, j\}}}^n p_k = p'_i + p'_j + 1 - p_i - p_j = \frac{v_j \cdot p_j}{v_i} + \frac{v_i \cdot p_i}{v_j} + 1 - p_i - p_j = 1 - (v_i - v_j) \left(\frac{v_j \cdot p_j - v_i \cdot p_i}{v_i \cdot v_j} \right) < 1$$

Where the last inequality holds because $v_i > v_j$ and $v_i \cdot p_i < v_j \cdot p_j$. Let $\alpha > 1$ be such that $\sum_{k=1}^n p'_k = \frac{1}{\alpha}$. Then, the distribution $p'' \in \Delta([n])$ given by $p''_k = \alpha \cdot p'_k$ satisfies $\text{value}(p'') = \alpha \cdot \text{value}(p) > \text{value}(p)$, contradicting the optimality of p . $\square$

With this lemma in place, let us restrict ourselves to distributions p satisfying $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$. For this case Eq. (3) simply becomes $\text{value}(p) = \sum_{i=t+1}^n v_i \cdot p_i$. With this observation, we now prove the following lemma, implying Theorem 2:

Lemma 5. *There exists a distribution $p \in \Delta([n])$ maximizing $\text{value}(p)$ such that for some $i \geq t+1$ we have $v_1 \cdot p_1 = \dots = v_i \cdot p_i$ and $p_{i+1} = \dots = p_n = 0$.*

PROOF. By Lemma 4, requiring that $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$ does not change the maximum achievable $\text{value}(p)$, so let us restrict ourselves to this case. As a result, we know that $\text{value}(p) = \sum_{i=t+1}^n v_i \cdot p_i$, so a maximizing p is a solution to the following linear program with n variables and $2n$ constraints:

$$\begin{aligned} & \text{maximize} && \sum_{i=t+1}^n v_i \cdot p_i \\ & \text{subject to} && v_{i+1} \cdot p_{i+1} - v_i \cdot p_i \leq 0, \quad i = 1, \dots, n-1 \\ & && p_i \geq 0, \quad i = 1, \dots, n \\ & && \sum_{i=1}^n p_i = 1 \end{aligned}$$

Let us take its dual. For this purpose, introduce non-negative variables $y_1, \dots, y_{n-1}$ corresponding to constraints of the first kind and a variable $T \in \mathbb{R}$ corresponding to the last constraint. For convenience, we define $y_0 = y_n = 0$ and $v'_1, \dots, v'_n$ such that $v'_i = v_i$ if $i \geq t+1$ and $v'_i = 0$ otherwise. Then, the dual is:

$$\begin{aligned} & \text{minimize} && T \\ & \text{subject to} && T - v_i \cdot y_i + v_i \cdot y_{i-1} \geq v'_i, \quad i = 1, \dots, n \\ & && y_i \geq 0, \quad i = 1, \dots, n-1 \\ & && y_i = 0 \quad i \in \{0, n\} \end{aligned}$$

We want to find the optimum value of the dual. This amounts to understanding for which values of T the dual is feasible. Let us fix a value T and study the feasibility of the dual.

Note that constraints of the first kind in the dual can be conveniently rewritten as:

$$y_i - y_{i-1} \leq \frac{T - v'_i}{v_i} = \frac{T}{v_i} - [i \geq t+1]$$

where $[i \geq t+1] := 1$ if $i \geq t+1$ and 0 otherwise. Denote the right-hand side of the inequality with c_i^T , which is a constant in terms of the fixed T . Hence, we want to check the feasibility of:¹¹

$$\begin{aligned} & y_i - y_{i-1} \leq c_i^T, \quad i = 1, \dots, n \\ & y_i \geq 0, \quad i = 1, \dots, n-1 \\ & y_i = 0 \quad i \in \{0, n\} \end{aligned}$$

It is natural to reinterpret this in terms of $\delta_i := y_i - y_{i-1}$:

$$\begin{aligned} & \delta_i \leq c_i^T, \quad i = 1, \dots, n \\ & \sum_{j=1}^i \delta_j \geq 0, \quad i = 1, \dots, n-1 \\ & \sum_{j=1}^n \delta_j = 0 \end{aligned}$$

If we relax the last constraint to $\sum_{j=1}^n \delta_j \geq 0$, this is clearly feasible if and only if setting $\delta_i = c_i^T$ satisfies the non-negativity constraints. Doing so might however lead to $\sum_{j=1}^n \delta_j > 0$. If this is the case, it suffices to decrease δ_n to make the total equal 0, so this is not a problem.

¹¹The accustomed reader will notice the similarity with the linear program for distances in the graph $n \rightarrow \dots \rightarrow 0$.

As a result, a value of T is achievable for the dual if and only if $\sum_{j=1}^i c_j^T \geq 0$ for all $1 \leq i \leq n$. For a fixed i , this is equivalent to:

$$\begin{aligned} \sum_{j=1}^i \frac{T}{v_j} - \sum_{j=1}^i [j \geq t+1] &= T \cdot \sum_{j=1}^i \frac{1}{v_j} - \max\{0, i-t\} \geq 0 \\ \iff T &\geq \frac{\max\{0, i-t\}}{\sum_{j=1}^i \frac{1}{v_j}} \end{aligned}$$

For $i \leq t$, this simplifies to $T \geq 0$, and otherwise it simplifies to a stricter inequality, so the case $i \leq t$ can be discarded from the condition. We hence get that the optimum value of the dual is:

$$T^* = \max_{t+1 \leq i \leq n} \frac{i-t}{\sum_{j=1}^i \frac{1}{v_j}}$$

We will now give a feasible solution p^* to the primal with the property required in the statement of the lemma. We will show that p^* achieves value T^* implying, by weak duality, that it is optimal for the primal, concluding the proof.

The distribution p^* is constructed as follows: for $j > i^*$, we set $p_j^* = 0$, while for $j \leq i^*$ we set $p_j^* = \frac{1}{v_j \cdot C}$, where $C := \sum_{k=1}^{i^*} \frac{1}{v_k}$ is a normalizing factor. First, p^* is clearly a well-defined distribution, as its entries sum up to 1 and are non-negative. Moreover, $v_j \cdot p_j$ is $\frac{1}{C}$ for $j \leq i^*$ and 0 for $j > i^*$, implying that it is a feasible solution for the primal with the property required in the lemma statement. It remains to show that p^* achieves a value of T^* in the primal to get the optimality of p^* , completing the proof. This amounts to a simple algebraic verification. $\square$

6 THE GENERAL- ℓ CASE

The general case concerns value-maximizing distributions $p \in \Delta([n]^\ell)$, i.e., over size- ℓ subsets of $[n]$. For $\ell > 1$, these are combinatorially difficult to reason about. However, by linearity of expectation, we can write $\mathbb{E}_{S \sim p}[\text{pay}(S, B)] = \sum_{i \in B} v_i \cdot \mathbb{P}_{S \sim p}(i \in S)$, so much of the information contained in p is redundant: the expectation only depends on the *marginals* $p'_i := \mathbb{P}_{S \sim p}(i \in S)$ for all $i \in [n]$. In particular, $\mathbb{E}_{S \sim p}[\text{pay}(S, B)] = \sum_{i=1}^n v_i \cdot p'_i$. Note that the marginals define a pseudo-distribution $p' \in \Delta_\ell([n])$: they are between 0 and 1 and sum up to ℓ .¹² Trying to optimize for p' directly seems particularly attractive given its lower-dimensional nature. In particular, instead of optimizing for $\text{value}(p)$ among $p \in \Delta([n]^\ell)$, we can optimize for $\text{value}(p') := \min_{B \in [n]^\ell} \sum_{i=1}^n v_i \cdot p'_i$ among $p' \in \Delta_\ell([n])$. While attractive, this approach could be prone to a serious pitfall: it could be that value-maximizing pseudo-distributions p' cannot occur as the marginals of a true-distribution p . However, perhaps surprisingly, this is never the case: any pseudo-distribution p' can be (efficiently) implemented through a true distribution p :

Lemma 6. *Let $\ell \geq 1$ and $p' \in \Delta_\ell([n])$ be arbitrary. Then, there exists a distribution $p \in \Delta([n]^\ell)$ such that $\mathbb{P}_{S \sim p}(i \in S) = p'_i$ for $i \in [n]$. Moreover, such a p exists with support of size at most n and can be computed in $O(n^2)$ time given p' . If computing p explicitly is not required, we can sample from such a p in $O(n)$ time per sample.*

PROOF. The linear time per sample part follows from applying dependent randomized rounding [28] to a star graph with n

¹²By linearity of expectation: $\sum_{i=1}^n p'_i = \sum_{i=1}^n \mathbb{P}_{S \sim p}(i \in S) = \mathbb{E}_{S \sim p}[|S|] = \ell$.

edges of weights $p'_1, \dots, p'_n$ (even ignoring the dependent part). The quadratic explicit construction follows by applying the *AllocationFromShares* algorithm in [3]. See also [32] and the informative discussion in [4].

We note that other existential proofs have appeared (sometimes implicitly) in the literature, using tools like the Birkhoff-von Neumann theorem and Carathéodory's theorem [38, 44]. Given existence, a distribution can be computed [44] by setting up an LP with exponentially-many variables and polynomially-many constraints and solving its dual using the ellipsoid method with a poly-time separation oracle [33, 54] (however, not in strongly polynomial time). $\square$

As a result, our new goal is to give a linear-time algorithm outputting a pseudo-distribution $p' \in \Delta_\ell([n])$ that maximizes $\text{value}(p')$. The rest of the paper is dedicated to this task. Knowing p' , we can then apply Lemma 6 to get our main result: Theorem 3.

6.1 Value-Maximizing Pseudo-Distributions

In order not to overburden the notation, henceforth, we drop the apostrophe and ask for a pseudo-distribution $p \in \Delta_\ell([n])$ that maximizes $\text{value}(p)$. It will also be convenient to, without loss of generality, allow the elements in p to sum to *at most* ℓ instead of exactly ℓ : we write $p \in \Delta_{\leq \ell}([n])$ from now on to signal this fact.

To achieve our goal, we would now like to show similar structural results to Lemmas 4 and 5. For Lemma 4, a similar argument shows that it suffices to consider pseudo-distributions satisfying $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$. However, an immediate analog of Lemma 5 is not possible. Instead, by going through an intermediary lemma, we show that it suffices to restrict ourselves to what we call *nice pseudo-distributions*, which admit a nice geometric interpretation through a water-in-vessels analogy. We achieve our final goal by a careful, efficient simulation of continuously decreasing the water level in the first vessel. We begin with the following lemma, whose proof is very similar to that of Lemma 4 and hence deferred to the full version of our paper [20].

Lemma 7. *There exists a pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ maximizing $\text{value}(p)$ that satisfies $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$.*

Lemma 8. *There exists a pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ maximizing $\text{value}(p)$ satisfying $v_1 \cdot p_1 = \dots = v_{t+1} \cdot p_{t+1} \geq \dots \geq v_n \cdot p_n$.*


PROOF. By Lemma 7, let p be a maximizer of $\text{value}(p)$ such that $v_1 \cdot p_1 \geq \dots \geq v_n \cdot p_n$. If p satisfies the required property, then we are done. Otherwise, construct $p' : [n] \rightarrow [0, 1]$ as follows:

$$p'_j = \begin{cases} \frac{v_{t+1} \cdot p_{t+1}}{v_j} & j \leq t \\ p_j & j > t \end{cases}$$


Note that for $j \leq t$ we have $p'_j = \frac{v_{t+1} \cdot p_{t+1}}{v_j} \leq \frac{v_j \cdot p_j}{v_j} = p_j$, so $\sum_{j=1}^n p'_j \leq \sum_{j=1}^n p_j \leq \ell$, from which $p' \in \Delta_{\leq \ell}([n])$. Moreover, by construction, we have $\text{value}(p') = \text{value}(p)$ and that p' satisfies $v_1 \cdot p'_1 = \dots = v_{t+1} \cdot p'_{t+1} \geq \dots \geq v_n \cdot p'_n$. $\square$

Definition 9. A pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ is (E, i) -nice for some $E \geq 0$ and $i \geq t+1$ if the following conditions are satisfied:

- (1) $v_k \cdot p_k = E$ for all $1 \leq k \leq t+1$;



(a) Example instance for $n = 7, t = 1, \ell = 5$.



(b) The maximal 7-nice pseudo-distribution p .

Figure 1: Consider an example with $n = 7, t = 1, \ell = 5$ and $v = (12, 8, 8, 6, 4, 3, 2)$. This is depicted in Fig. 1a by rectangles with heights given by v , each of area 1. One can understand pseudo-distributions p using a water-filling metaphor: each rectangle corresponds to a container comprising one unit of volume and each value $p_i \in [0, 1]$ corresponds to pouring p_i units of water into the i -th container. By the choice of widths, pouring p_i units of water into the i -th container makes the water rise to height $h_i := v_i \cdot p_i$ inside the container. Given (t, ℓ) , a pseudo-distribution is (E, i) -nice if: (0) it uses at most ℓ units of water; (1) the water rises to level E in the first $t+1$ containers; (2) each container k among the first i is *saturated* for E , i.e., either the water rises to height E in k or $E < v_k$ and k is full; (3) container $i+1$ (if it exists) is not saturated for E ; (4) all subsequent containers are empty. Given E , there can exist at most one maximal E -nice pseudo-distribution: fill in the first $t+1$ containers to level E , if this exceeds the water budget ℓ , then no solution exists, otherwise continue in order through the next containers, saturating them until there is not enough water left to saturate current container. This is demonstrated for $E = 7$ in Fig. 1b: the first $3 \geq t+1$ containers rise to level 7, the next two containers have $p_4 = p_5 = 1$ (hence $i = 5$), the following container is not saturated: $p_6 = \frac{2}{3}$, and the last container is empty: $p_7 = 0$. Overall, $p = (\frac{7}{12}, \frac{7}{8}, \frac{7}{8}, 1, 1, \frac{2}{3}, 0)$, whose entries sum up to $\ell = 5$.

- (2) $p_k = \min\{1, \frac{E}{v_k}\}$ for all $1 \leq k \leq i$;
- (3) $p_{i+1} < \min\{1, \frac{E}{v_{i+1}}\}$;¹³
- (4) $p_k = 0$ for all $i+2 \leq k \leq n$.

Moreover, p is E -nice if it is (E, i) -nice for some $i \geq t+1$ and it is nice if it is E -nice for some $E \geq 0$.

¹³Provided $i+1 \leq n$, otherwise we consider the last two conditions vacuously true.

An E -nice pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ is *maximal* (for E) if there is no E -nice pseudo-distribution $p' \in \Delta_{\leq \ell}([n])$ such that $p'_i \geq p_i$ for all $1 \leq i \leq n$ and $p'_i > p_i$ for at least one i .

The notion of nice pseudo-distributions might appear complex at first. However, it has an elegant geometric interpretation by a water-filling metaphor presented in Fig. 1 and its caption.

Lemma 10. *There exists a pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ maximizing $\text{value}(p)$ that is nice.*

PROOF. By Lemma 8, let p be a maximizer of $\text{value}(p)$ such that $v_1 \cdot p_1 = \dots = v_{t+1} \cdot p_{t+1} \geq \dots \geq v_n \cdot p_n$. Define $E := v_1 \cdot p_1$ and let $t+1 \leq i \leq n$ be the largest index such that $p_i = \min\{1, \frac{E}{v_i}\}$. Among such pseudo-distributions p , further assume that we selected one maximizing the pair (i, p_{i+1}) lexicographically, where we consider $p_{n+1} := 0$ in order for this to always be well-defined. If $i = n$, then p is (E, i) -nice, so we are done, otherwise, assume $i+1 \leq n$, implying $p_{i+1} < \min\{1, \frac{E}{v_{i+1}}\}$, so p satisfies the first three conditions for being (E, i) -nice. If the fourth condition is also satisfied, then we are done. Otherwise, write $x := \sum_{k=i+2}^n p_k$, and observe that $x > 0$. Furthermore, note that $p_i \cdot v_i > p_{i+1} \cdot v_{i+1}$. To see this, assume the contrary, i.e., $p_i \cdot v_i = p_{i+1} \cdot v_{i+1}$, from which $p_{i+1} = \frac{v_i \cdot p_i}{v_{i+1}} = \frac{v_i}{v_{i+1}} \min\{1, \frac{E}{v_i}\} = \min\{\frac{v_i}{v_{i+1}}, \frac{E}{v_{i+1}}\} \geq \min\{1, \frac{E}{v_{i+1}}\}$, a contradiction. Now, define $\varepsilon := \min\{x, \frac{v_i \cdot p_i}{v_{i+1}} - p_{i+1}\}$ and note that $\varepsilon > 0$.¹⁴ Construct $p' : [n] \rightarrow [0, 1]$ as follows:

$$p'_k = \begin{cases} p_k & k \leq i \\ p_k + \varepsilon & k = i+1 \\ p_k \cdot \frac{x-\varepsilon}{x} & k \geq i+2 \end{cases}$$

Let us first show that $p' \in \Delta_{\leq \ell}([n])$. For this, we need to show that all entries are non-negative and sum up to at most ℓ . Non-negativity follows since $\frac{x-\varepsilon}{x} \geq 0$, and the sum is a simple algebraic calculation.

Let us now show that p' satisfies $v_1 \cdot p'_1 = \dots = v_{t+1} \cdot p'_{t+1} \geq \dots \geq v'_n \cdot p'_n$. The equalities hold by construction of p' from p and because $i \geq t+1$. It then remains to show that for all $1 \leq k \leq n-1$ we have $v_k \cdot p'_k \geq v_{k+1} \cdot p'_{k+1}$. The cases $k < i$ and $k \geq i+2$ are immediate by construction. The case $k = i+1$ is also straightforward: $v_{i+1} \cdot p'_{i+1} \geq v_{i+1} \cdot p_{i+1} \geq v_{i+2} \cdot p_{i+2} \geq v_{i+2} \cdot p'_{i+2}$. The final case $i = k$ amounts to $v_i \cdot p_i \geq v_{i+1} \cdot (p_{i+1} + \varepsilon)$, which holds by ε 's definition.

Finally, let us show that the existence of p' is a contradiction. First, $\text{value}(p') \geq \text{value}(p)$. This is because p' is constructed from p by moving ε probability mass from positions $i+2, \dots, n$ to position $i+1$, and v_{i+1} is no smaller than $v_{i+2}, \dots, v_n$ and hence $\text{value}(p') = \sum_{k=t+1}^n v_k \cdot p'_k \geq \sum_{k=t+1}^n v_k \cdot p_k = \text{value}(p)$. If $\text{value}(p') > \text{value}(p)$ then this is already a contradiction, otherwise, $\text{value}(p') = \text{value}(p)$ holds, but even in that case p' has by construction a lexicographically strictly higher (i, p'_{i+1}) pair, again a contradiction. $\square$

Before proceeding further, we emphasize the observation made in the caption of Fig. 1: Given E , there exists at most one maximal E -nice pseudo-distribution, obtained by the presented left-to-right water-filling argument. In fact, something even stronger holds: E -nice pseudo-distributions are linearly ordered by the amount of water used: the only way to create non-maximal E -nice pseudo-distributions is to follow the same water-filling argument but stop

it early before it had a chance to use all possible water. For a fixed E , pseudo-distributions with more water can not lead to a worse value, so we can augment Lemma 10 to get that it suffices to look at *maximal* pseudo-distributions, which are uniquely determined by E whenever they exist. Let us give the existence conditions in light of the water-filling argument: (i) $E \leq v_{t+1}$, as otherwise $p_k = \frac{E}{v_k} > 1$ would hold for some $1 \leq k \leq t+1$, in particular for $k = t+1$ and (ii) $\sum_{k=1}^{t+1} \frac{E}{v_k} \leq \ell \iff E \leq \ell \cdot (\sum_{k=1}^{t+1} \frac{1}{v_k})^{-1}$, since otherwise the ℓ units of water we have at our disposal are insufficient to make the water rise to level E in the first $t+1$ containers. Hence, (maximal) E -nice pseudo-distributions exist for $0 \leq E \leq E_{\max} := \min\{v_{t+1}, \ell \cdot (\sum_{k=1}^{t+1} \frac{1}{v_k})^{-1}\}$. It remains to identify $E \in [0, E_{\max}]$ yielding a maximum-value E -nice pseudo-distribution. The main idea to do this in linear time is to note that the value of the maximal E -nice distribution is piece-wise linear in E . We prove this fact and give a linear-time algorithm that produces and iterates through the breakpoints in order from $E_{\max}$ down to 0, at the same time computing the value of the corresponding pseudo-distributions and outputting the best one at the end (the maximum has to happen at a break-point due to piece-wise linearity). The details can be found in the proof of the lemma below, which is rather involved and hence deferred to the full version of our paper [20]. To give a different view, our algorithm simulates the continuous process of decreasing the water level in the first container (i.e., E , starting at $E_{\max}$), outputting the best achievable value along the way in overall linear time.

Lemma 11. *A pseudo-distribution $p \in \Delta_{\leq \ell}([n])$ maximizing the quantity $\text{value}(p)$ can be computed in time $O(n)$.*

7 CONCLUSION AND FUTURE WORK

We introduced the Byzantine Selection Problem and gave attractive linear-time algorithms sampling from a value-maximizing distribution. It would also be interesting to study the problem for other notions of optimality inspired by algorithm design under uncertainty, such as minimizing regret or maximizing the competitive ratio. Introducing an online element, where boxes arrive one at a time, like in the secretary problem, could also prove fruitful. Finally, our model assumed a known threshold t on the number of empty boxes (exactly t purely for analysis). If such a threshold is unavailable, our algorithms can still be used with $t = n-1$, though this may be overly conservative. One can choose to work with a lower t depending on their risk aversion. Studying how the optimum varies with t for an instance or assuming a distribution on the number of corruptions could be promising avenues for future work.

More broadly, our paper argues for an exciting yet relatively unexplored research avenue: introducing byzantine elements into classic non-distributed settings, such as the centralized settings emerging in social choice theory. The recent work in [17] brings this element to stable matchings. While their paper targets a distributed setting, one of their contributions is the development of a fault-tolerant notion of stable matchings. It would be appealing to conduct similar investigations in other application domains, such as preference aggregation or fair division. The paper [26] targets so-called *fault-tolerant implementations*, and could also serve as a good starting point for further investigation in this direction when incentives are to be considered as well.

¹⁴This is because $x > 0$ and $p_i \cdot v_i > p_{i+1} \cdot v_{i+1}$.

REFERENCES

- [1] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. 2009. Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research* 197, 2 (2009), 427–438. <https://doi.org/10.1016/j.ejor.2008.09.012>
- [2] Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Mike Dahlin, Jean-Philippe Martin, and Carl Porth. 2005. BAR fault tolerance for cooperative services. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles* (Brighton, United Kingdom) (SOSP'05). Association for Computing Machinery, New York, NY, USA, 45–58. <https://doi.org/10.1145/1095810.1095816>
- [3] Haris Aziz, Omer Lev, Nicholas Mattei, Jeffrey S. Rosenschein, and Toby Walsh. 2019. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artificial Intelligence* 275 (2019), 295–309. <https://doi.org/10.1016/j.artint.2019.06.004>
- [4] Haris Aziz, Xinhua Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. 2024. Best-of-Both-Worlds Fairness in Committee Voting. In *Web and Internet Economics*, Jugal Garg, Max Klimm, and Yuqing Kong (Eds.). Springer Nature Switzerland, Cham, 676. Full version arXiv:2303.03642 [cs.GT].
- [5] Moshe Babaioff, Robert Kleinberg, and Christos H. Papadimitriou. 2009. Congestion games with malicious players. *Games and Economic Behavior* 67, 1 (2009), 22–35. <https://doi.org/10.1016/j.geb.2009.04.017> Special Section of Games and Economic Behavior Dedicated to the 8th ACM Conference on Electronic Commerce.
- [6] D. Bertsimas and D. den Hertog. 2022. *Robust and Adaptive Optimization*. Dynamic Ideas LLC. https://books.google.ch/books?id=V_RPzwEACAAJ
- [7] Dimitris Bertsimas and Melvyn Sim. 2004. The Price of Robustness. *Oper. Res.* 52, 1 (Jan. 2004), 35–53. <https://doi.org/10.1287/opre.1030.0065>
- [8] Domagoj Bradac, Anupam Gupta, Sahil Singla, and Goran Zuzic. 2020. Robust Algorithms for the Secretary Problem. In *11th Innovations in Theoretical Computer Science Conference (ITCS) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 151)*, Thomas Vidick (Ed.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 32:1–32:26. <https://doi.org/10.4230/LIPIcs.ITCS.2020.32>
- [9] Felix Brandt, Tuomas Sandholm, and Yoav Shoham. 2007. Spiteful bidding in sealed-bid auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (Hyderabad, India) (IJCAI'07)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1207001214.
- [10] Christoph Buchheim and Jannis Kurtz. 2018. Robust combinatorial optimization under convex and discrete cost uncertainty. *EURO Journal on Computational Optimization* 6, 3 (2018), 211–238. <https://doi.org/10.1007/s13675-018-0103-0>
- [11] Deeparnab Chakrabarty, Chinmay Karande, and Ashish Sangwan. 2009. The Effect of Malice on the Social Optimum in Linear Load Balancing Games. arXiv:0910.2655 [cs.GT] <https://arxiv.org/abs/0910.2655>
- [12] Moses Charikar, Jacob Steinhardt, and Gregory Valiant. 2017. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing* (Montreal, Canada) (STOC 2017). Association for Computing Machinery, New York, NY, USA, 47–60. <https://doi.org/10.1145/3055399.3055491>
- [13] Pierre Civi, Seth Gilbert, Rachid Guerraoui, Jovan Komatovic, Anton Paramonov, and Manuel Vidigueira. 2024. All Byzantine Agreement Problems Are Expensive. In *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing* (Nantes, France) (PODC'24). Association for Computing Machinery, New York, NY, USA, 157–169. <https://doi.org/10.1145/3662158.3662780>
- [14] Pierre Civi, Seth Gilbert, Rachid Guerraoui, Jovan Komatovic, and Manuel Vidigueira. 2023. On the Validity of Consensus. In *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing*. 332–343.
- [15] Allen Clement, Harry C. Li, Jeff Napper, Jean-Philippe Martin, Lorenzo Alvisi, and Michael Dahlin. 2008. BAR primer. In *The 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2008, June 24-27, 2008, Anchorage, Alaska, USA, Proceedings*. IEEE Computer Society, 287–296. <https://doi.org/10.1109/DSN.2008.4630097>
- [16] Andrei Constantinescu, Marc Dufay, Diana Ghinea, and Roger Wattenhofer. 2024. Validity in Network-Agnostic Byzantine Agreement. arXiv:2410.19721 [cs.DC]
- [17] Andrei Constantinescu, Marc Dufay, Diana Ghinea, and Roger Wattenhofer. 2025. Byzantine Stable Matching. arXiv:2502.05889 [cs.DC] <https://arxiv.org/abs/2502.05889>
- [18] Andrei Constantinescu, Diana Ghinea, Lioba Heimbach, Zilin Wang, and Roger Wattenhofer. 2024. A Fair and Resilient Decentralized Clock Network for Transaction Ordering. In *27th International Conference on Principles of Distributed Systems (OPODIS 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 286)*, Alysson Bessani, Xavier Défago, Junya Nakamura, Koichi Wada, and Yukiko Yamauchi (Eds.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 8:1–8:20. <https://doi.org/10.4230/LIPIcs.OPODIS.2023.8>
- [19] Andrei Constantinescu, Diana Ghinea, Roger Wattenhofer, and Floris Westermann. 2024. Convex Consensus with Asynchronous Fallback. In *38th International Symposium on Distributed Computing (DISC 2024) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 319)*, Dan Alistarh (Ed.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 15:1–15:23. <https://doi.org/10.4230/LIPIcs.DISC.2024.15>
- [20] Andrei Constantinescu and Roger Wattenhofer. 2025. Byzantine Game Theory: Sun Tzu's Boxes. arXiv:2502.14812 [cs.GT] <https://arxiv.org/abs/2502.14812>
- [21] Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2019. Recent developments in prophet inequalities. *SIGecom Exch.* 17, 1 (May 2019), 61–70. <https://doi.org/10.1145/3331033.3331039>
- [22] Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. 2019. Robust Estimators in High-Dimensions Without the Computational Intractability. *SIAM J. Comput.* 48, 2 (2019), 742–864. <https://doi.org/10.1137/17M1126680>
- [23] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. 2018. Robustly learning a gaussian: getting optimal error, efficiently. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, Louisiana) (SODA'18). Society for Industrial and Applied Mathematics, USA, 2683–2702.
- [24] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. 2018. Learning geometric concepts with noisy noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (Los Angeles, CA, USA) (STOC 2018). Association for Computing Machinery, New York, NY, USA, 1061–1073. <https://doi.org/10.1145/3188745.3188754>
- [25] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. 1986. Reaching approximate agreement in the presence of faults. *J. ACM* 33, 3 (May 1986), 499–516. <https://doi.org/10.1145/5925.5931>
- [26] Kfir Eliaz. 2002. Fault Tolerant Implementation. *The Review of Economic Studies* 69, 3 (2002), 589–610. <http://www.jstor.org/stable/1556711>
- [27] Piotr Faliszewski and Jörg Rothe. 2016. Control and Bribery in Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D.Editors Procaccia (Eds.). Cambridge University Press, 146–168.
- [28] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. 2006. Dependent rounding and its applications to approximation algorithms. *J. ACM* 53, 3 (may 2006), 324–360. <https://doi.org/10.1145/1147954.1147956>
- [29] Marc Goerigk and Anita Schöbel. 2016. Algorithm Engineering in Robust Optimization. In *Algorithm Engineering: Selected Results and Surveys*, Lasse Kliemann and Peter Sanders (Eds.). Springer International Publishing, Cham, 245–279. https://doi.org/10.1007/978-3-319-49487-6_8
- [30] Bram L. Gorissen, İhsan Yanikoğlu, and Dick den Hertog. 2015. A practical guide to robust optimization. *Omega* 53 (2015), 124–137. <https://doi.org/10.1016/j.omega.2014.12.006>
- [31] Ronen Gradwohl and Omer Reingold. 2014. Fault tolerance in large games. *Games and Economic Behavior* 86 (2014), 438–457. <https://doi.org/10.1016/j.geb.2013.06.007>
- [32] Geoffrey Grimmett. 2004. Stochastic Apportionment. *The American Mathematical Monthly* 111, 4 (2004), 299–307. <http://www.jstor.org/stable/4145239>
- [33] Martin Grötschel, László Lovász, and Alexander Schrijver. 1993. *Geometric Algorithms and Combinatorial Optimization*. Springer Berlin Heidelberg, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-78240-4>
- [34] Martin Hoefer and Alexander Skopalik. 2013. Altruism in Atomic Congestion Games. *ACM Trans. Econ. Comput.* 1, 4, Article 21 (Dec. 2013), 21 pages. <https://doi.org/10.1145/2542174.2542177>
- [35] George Karakostas and Anastasios Viglas. 2007. Equilibria for networks with malicious users. *Math. Program.* 110, 3 (Sept. 2007), 591–613.
- [36] Adam Kasperski and Paweł Zieliński. 2016. Robust Discrete Optimization Under Discrete and Interval Uncertainty: A Survey. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, Michael Doumpos, Constantin Zopounidis, and Evangelos Grigoroudis (Eds.). Springer International Publishing, Cham, 113–143. https://doi.org/10.1007/978-3-319-33121-8_6
- [37] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (Budapest, Hungary) (AAMAS'09)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 689–696.
- [38] Dmytro Korzhuk, Vincent Conitzer, and Ronald Parr. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (Atlanta, Georgia) (AAAI'10). AAAI Press, 805–810.
- [39] Dmytro Korzhuk, Vincent Conitzer, and Ronald Parr. 2011. Security games with multiple attacker resources. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One (Barcelona, Catalonia, Spain) (IJCAI'11)*. AAAI Press, 273–279.
- [40] K. A. Lai, A. B. Rao, and S. Vempala. 2016. Agnostic Estimation of Mean and Covariance. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, Los Alamitos, CA, USA, 665–674. <https://doi.org/10.1109/FOCS.2016.76>
- [41] Harry C. Li, Allen Clement, Mirco Marchetti, Manos Kapritsos, Luke Robison, Lorenzo Alvisi, and Mike Dahlin. 2008. FlightPath: obedience vs. choice in cooperative services. In *Proceedings of the 8th USENIX Conference on Operating*

- Systems Design and Implementation* (San Diego, California) (OSDI'08). USENIX Association, USA, 355–368.
- [42] Harry C. Li, Allen Clement, Edmund L. Wong, Jeff Napper, Indrajit Roy, Lorenzo Alvisi, and Michael Dahlin. 2006. BAR gossip. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation* (Seattle, Washington) (OSDI'06). USENIX Association, USA, 191–204.
 - [43] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. 2018. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (Los Angeles, CA, USA) (STOC 2018). Association for Computing Machinery, New York, NY, USA, 114–122. <https://doi.org/10.1145/3188745.3188918>
 - [44] Andrew Mastin, Patrick Jaillet, and Sang Chin. 2015. Randomized Minmax Regret for Combinatorial Optimization Under Uncertainty. In *Algorithms and Computation*, Khaled Elbassioni and Kazuhisa Makino (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 491–501.
 - [45] Dominic Meier, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer. 2008. On the windfall of friendship: inoculation strategies on social networks. In *Proceedings of the 9th ACM Conference on Electronic Commerce* (Chicago, IL, USA) (EC'08). Association for Computing Machinery, New York, NY, USA, 294–301. <https://doi.org/10.1145/1386790.1386836>
 - [46] Darya Melnyk, Yuyi Wang, and Roger Wattenhofer. 2018. Byzantine Preferential Voting. In *Web and Internet Economics*, George Christodoulou and Tobias Harks (Eds.). Springer International Publishing, Cham, 327–340.
 - [47] Darya Melnyk and Roger Wattenhofer. 2018. Byzantine Agreement with Interval Validity. In *2018 IEEE 37th Symposium on Reliable Distributed Systems (SRDS)*. 251–260. <https://doi.org/10.1109/SRDS.2018.00036>
 - [48] Hammurabi Mendes, Maurice Herlihy, Nitin Vaidya, and Vijay K. Garg. 2015. Multidimensional agreement in Byzantine systems. *Distrib. Comput.* 28, 6 (Dec. 2015), 423–441. <https://doi.org/10.1007/s00446-014-0240-5>
 - [49] John Morgan, Ken Steiglitz, and George Reis. 2003. The Spite Motive and Equilibrium Behavior in Auctions. *Contributions in Economic Analysis & Policy* 2, 1 (2003), 1–25. <https://doi.org/10.2202/1538-0645.1102>
 - [50] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. 2009. The Price of Malice: A Game-Theoretic Framework for Malicious Behavior. *Internet Mathematics* 6, 2 (2009), 125–156.
 - [51] J. von Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100 (1928), 295–320. <http://eudml.org/doc/159291>
 - [52] Thomas Nowak and Joel Rybicki. 2019. Byzantine Approximate Agreement on Graphs. In *33rd International Symposium on Distributed Computing (DISC 2019)* (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 146), Jukka Suomela (Ed.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 29:1–29:17. <https://doi.org/10.4230/LIPIcs.DISC.2019.29>
 - [53] Aaron Roth. 2008. The Price of Malice in Linear Congestion Games. In *Proceedings of the 4th International Workshop on Internet and Network Economics* (Shanghai, China) (WINE'08). Springer-Verlag, Berlin, Heidelberg, 118–125. https://doi.org/10.1007/978-3-540-92185-1_20
 - [54] Alexander Schrijver. 1999. *Theory of linear and integer programming*. Wiley, Wiley-Interscience series in discrete mathematics and optimization.
 - [55] Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. 2018. Stackelberg Security Games: Looking Beyond a Decade of Success. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 5494–5501. <https://doi.org/10.24963/ijcai.2018/775>
 - [56] David Stolz and Roger Wattenhofer. 2016. Byzantine Agreement with Median Validity. In *19th International Conference on Principles of Distributed Systems (OPODIS 2015)* (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 46), Emmanuelle Anceaume, Christian Cachin, and Maria Potop-Butucaru (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 22:1–22:14. <https://doi.org/10.4230/LIPIcs.OPODIS.2015.22>
 - [57] Nitin H. Vaidya and Vijay K. Garg. 2013. Byzantine vector consensus in complete graphs. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing* (Montréal, Québec, Canada) (PODC'13). Association for Computing Machinery, New York, NY, USA, 65–73. <https://doi.org/10.1145/2484239.2484256>
 - [58] Haifeng Xu. 2016. The Mysteries of Security Games: Equilibrium Computation Becomes Combinatorial Algorithm Design. In *Proceedings of the 2016 ACM Conference on Economics and Computation* (Maastricht, The Netherlands) (EC'16). Association for Computing Machinery, New York, NY, USA, 497–514. <https://doi.org/10.1145/2940716.2940796>
 - [59] Zhangchen Xu, Yuetai Li, Chenglin Feng, and Lei Zhang. 2023. Exact Fault-Tolerant Consensus with Voting Validity. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 842–852. <https://doi.org/10.1109/IPDPS54959.2023.00089>
 - [60] Lin Yang, Mohammad H. Hajiesmaili, M. Sadegh Talebi, John C.S. Lui, and Wing S. Wong. 2020. Adversarial bandits with corruptions. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (NIPS'20). Curran Associates Inc., Red Hook, NY, USA, Article 1674, 10 pages.