Pandora's Problem with Combinatorial Cost

BEN BERGER, Tel Aviv University, Israel
TOMER EZRA, Sapienza Università di Roma, Italy
MICHAL FELDMAN, Tel Aviv University, Israel
FEDERICO FUSCO, Sapienza Università di Roma, Italy

Pandora's problem is a fundamental model in economics that studies optimal search strategies under costly inspection. In this paper we initiate the study of Pandora's problem with *combinatorial costs*, capturing many real-life scenarios where search cost is non-additive. Weitzman's celebrated algorithm [1979] establishes the remarkable result that, for additive costs, the optimal search strategy is non-adaptive and computationally feasible.

We inquire to which extent this structural and computational simplicity extends beyond additive cost functions. Our main result is that the class of submodular cost functions admits an optimal strategy that follows a fixed, non-adaptive order, thus preserving the structural simplicity of additive cost functions. In contrast, for the more general class of subadditive (or even XOS) cost functions the optimal strategy may already need to determine the search order adaptively. On the computational side, obtaining any approximation to the optimal utility requires super polynomially many queries to the cost function, even for a strict subclass of submodular cost functions.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; **Computational pricing and auctions**.

Additional Key Words and Phrases: Pandora's Problem; Pandora's Box Problem

ACM Reference Format:
Ben Berger, Tomer Ezra, Michal Feldman, and Federico Fusco. 2023. Pandora's Problem with Combinatorial Cost. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23), July 9–12, 2023, London, United Kingdom.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3580507.3597699

1 INTRODUCTION

Pandora's problem captures the challenge of searching for a good alternative among multiple options, under costly evaluation. This problem was introduced in the seminal paper of Weitzman [1979], as a stochastic search problem over n boxes, each associated with an independent hidden stochastic value, and an exploration cost. At every point in time, the decision maker chooses which box (if any) to open. Upon opening a box, the decision maker incurs its exploration cost, and observes its realized value. Then, the decision maker can either decide to open an additional box or halt and obtain the maximum value observed so far. The goal is to maximize the expected maximum value over the set of opened boxes minus the sum of their exploration costs.

This setting captures many real-life scenarios, such as hiring employees or searching for an apartment, where there is an inherent tension between the desire to explore many options in an attempt to find one with high reward, and the desire to minimize the total exploration cost.

The full version of the paper is available at https://arxiv.org/abs/2303.01078.

EC '23, July 9–12, 2023, London, United Kingdom

ACM ISBN 979-8-4007-0104-7/23/07…$15.00
https://doi.org/10.1145/3580507.3597699

Weitzman [1979] showed that the optimal strategy for this problem exhibits both structural simplicity and computational simplicity. In particular, it opens the boxes according to a fixed-order, determined at the outset; only the stopping time is determined online, depending on the observed values. Moreover, the entire optimal strategy can be computed efficiently.

The last few years have seen a renewed interest in Pandora's problem, leading to a line of work that studies several extensions of the original model. Most studies focus on extending one of two features of the original problem: either considering a different notion of value derived from the set of opened boxes [e.g., Olszewski and Weber 2015; Singla 2018]; or modifying the rules of exploration [e.g., Boodaghians et al. 2020; Doval 2018; Esfandiari et al. 2019; Fu et al. 2018]. However, all of them share one fundamental assumption, namely that each box is associated with an individual cost, and these costs accumulate additively just like in the original model.

However, in many real-life scenarios, exploring one alternative may affect the exploration cost of other alternatives. For instance, when recruiting a new employee, there is a fixed cost for setting up the hiring process, while evaluating each additional candidate induces a small marginal cost. As another example, when searching for an apartment, each individual visit incurs a cost, but visiting multiple apartments in the same neighborhood is clearly less expensive than the sum of the costs of visiting them separately.

In this paper, we initiate the study of Pandora's problem with combinatorial cost functions, namely, a cost function that assigns a real value to every set of boxes. In this model, a decision maker who opens an extra box, given a set S of opened boxes, incurs its *marginal cost* given S. We inquire to which extent the structural and computational simplicity of Weitzman [1979] extends beyond additive cost functions. As it turns out, the structural simplicity of the original problem does not carry over to general cost functions. In particular, the exploration order in the optimal strategy may unavoidably be *adaptive*. This is demonstrated in the following example.

Example 1.1. Consider an instance with 3 boxes. The value in box 1 is 10 with probability $\frac{1}{2}$ and 0 otherwise. The value in box 2 is 12 with probability $\frac{1}{2}$ and 0 otherwise. The value in box 3 is 10 with probability 1. The total cost of exploring a set of boxes from the collection $\{\varnothing, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}\}$ is 0, and the total cost of exploring a set of boxes from the collection $\{\{2, 3\}, \{1, 2, 3\}\}$ is 20. It is not too difficult to observe that opening both boxes 2 and 3 is too expensive for any reasonable strategy. In fact, it can be shown (see the full version of the paper for details) that the (unique) optimal strategy for this instance is the following: open box 1. If its value is 10, then open box 2, otherwise (*i.e.*, the value in box 1 is 0), open box 3.

In the example above, boxes 2 and 3 exhibit strong complementarity in their cost; namely, the cost of opening both of them is (much) greater than the sum of their individual costs (which is 0). Many real-life scenarios, however, exhibit the opposite phenomenon, where the cost of the whole is smaller than the sum of the costs of its parts. This structure is captured by the class of subadditive cost functions, where $c(S \cup T) \leq c(S) + c(T)$ for any sets of boxes S and T, also known as complement-free functions.

A widely-encountered subclass of subadditive functions is the class of submodular functions, defined by *decreasing marginal contribution.* Indeed, many real-life exploration tasks exhibit this structure; e.g., where some fixed cost is incurred, followed by smaller individual costs. A hierarchy of complement-free functions has been provided by Lehmann et al. [2006], including the prominent classes of additive, submodular, and subadditive functions, as well as fractionally-subadditive functions (also known as XOS), where additive $\subset$ submodular $\subset$ XOS $\subset$ subadditive.

Given the prevalence of complement-free cost functions in real-life exploration scenarios, it is natural to study the structure of optimal strategies in these scenarios, and the corresponding computational problem. These are the main problems that drive us in this work. In particular, we ask whether Pandora's problem under different classes of complement-free cost functions preserves the structural and computational simplicity of the original problem with additive costs.

1.1 Our Results

As mentioned above, Example 1.1 shows an example of a *general* cost function, where an adaptive exploration order is inevitable. We first show that this phenomenon is not unique to cost functions that exhibit complementarities. Indeed, there exist instances with XOS cost functions for which an adaptive exploration order is inevitable[1].

Theorem 1: There exists an instance of the Pandora's problem with an XOS cost function that admits no optimal strategy with non-adaptive exploration order.

On the face of it, the above theorem seems to be unrelated to Example 1.1, where the cost-function exhibits strong complementarity. However, we identify a close connection between the two results. In particular, we show that every instance with a (monotone and normalized) cost function over n boxes induces an "equivalent" instance with an XOS cost function over $n + 1$ boxes, that inherits the adaptive exploration order of its source instance. With this result, the necessity of an adaptive order under XOS cost functions can be derived from Example 1.1. We refer the reader to the full version of the paper for more details.

A key property of XOS functions that enables this construction is that a marginal function of an XOS function c (namely, for some fixed T, $c'(S) := c(S \mid T) := c(S \cup T) - c(S)$) is unrestricted, and in particular can exhibit complementarities.

In stark contrast, the class of submodular functions is closed under marginal value; namely, if the cost function c is submodular, then so is the function $c(\cdot \mid T)$ for any fixed set T. In particular, the scenario depicted in Example 1.1, where the combined cost of opening boxes 2 and 3 is excessive, while opening each of them separately is cheap, cannot be replicated in an example utilizing a submodular cost function, even with the addition of more boxes.

A natural question is then whether instances of Pandora's problem with submodular cost functions preserve the structural simplicity of additive costs. That is, we ask whether these instances admit optimal strategies that open the boxes according to a fixed, non-adaptive order. Our first main result answers this question in the affirmative (see Sections 3 and 4).

Theorem 2 (see Theorem 4.4): Every instance of Pandora's problem with a submodular cost function admits an optimal strategy with non-adaptive exploration order.

Our second main result shows that, while the structural simplicity is preserved under submodular cost functions, the computational simplicity is not preserved. In particular, in Section 5 we prove the following stronger result.

Theorem 3 (see Theorem 5.3): The problem of deciding whether a given instance of Pandora's problem with a submodular cost function admits a strategy that attains strictly positive utility requires super-polynomially many queries to the cost function.

Notably, this theorem implies that no approximation to the optimal utility can be obtained with polynomially-many cost queries.

[1] Notably, for the larger class of subadditive cost functions, we find that an example demonstrating the necessity of adaptive order can be induced by a (seemingly unrelated) example that has been given in a completely different model of Pandora's box under constrained exploration order [Boodaghians et al. 2020] (see the full version of the paper for details).

1.2 Our Techniques

The main technical tool to solve Pandora's problem is the notion of reservation value of a box [e.g., Boodaghians et al. 2020; Esfandiari et al. 2019; Kleinberg and Kleinberg 2018; Singla 2018; Weitzman 1979]. This is the maximum value, presumably among those observed in previously opened boxes, for which opening the box achieves the same marginal utility as not opening it. Formally, the reservation value of a box with random reward V and (additive) cost c is the solution z of the following equation: $\mathbb{E}\left[(V-z)^+\right] = c$. Weitzman's optimal strategy opens the boxes in decreasing order of reservation value, halting when the current maximum observed reward exceeds the reservation value of the next unopened box. Since they are also easy to compute, reservation values simultaneously establish structural and computational simplicity for the problem. In the combinatorial setting that we study, however, this approach may yield an arbitrarily bad performance.

Example 1.2. Consider an instance with 2 identical boxes, each with a random reward of 2 with probability $\frac{1}{3}$ (and 0 otherwise), and a symmetric unit-demand cost function with a cost of 1 (i.e., $c(\{1\}) = c(\{2\}) = c(\{1,2\}) = 1$, and $c(\varnothing) = 0$). The reservation value of the two boxes is negative, thus Weitzman's strategy would not open any one of them. However, the best strategy for this instance opens both boxes, achieving an expected utility of $2 \cdot \frac{5}{9} - 1 > 0$.

The example illustrates why the reservation value is not suitable in the presence of combinatorial costs: the intrinsic importance of a box in the exploration is not solely determined by its random reward or its current marginal cost, but also by its influence on the marginal cost of all the (exponentially many) possible subsets of boxes that may be opened in the future.

In what follows we describe our techniques for our structural and computation results. We first present our techniques for the main structural result for Bernoulli instances, and then show how to extend it from Bernoulli to general instances. Finally, we present our techniques for our computational impossibility result.

Bernoulli instances. In Section 3 we prove Theorem 2 for Bernoulli instances, *i.e.*, instances where each box i has value v_i with probability p_i and value 0 otherwise.

A key notion in our analysis is that of an *impulsive strategy*. Such a strategy is determined by an ordered subset of boxes, and proceeds by opening them in the given order and halting upon the first time that a non-zero value of a box is observed (or if all boxes of the strategy have been opened). We show that every Bernoulli instance admits an optimal strategy that takes the form of an impulsive strategy. To establish this result, we follow the following steps.

We first show that we may assume the existence of an optimal strategy π^* that takes the following form: It starts by opening an arbitrary box r. If its non-zero value is realized, then it executes some impulsive sub-strategy π^Y, and if its realized value is 0, then it executes another impulsive sub-strategy π^N. This is proved by induction, using the fact that the marignal cost of a submodular function is also submodular.

Under this assumption, we proceed as follows: Assume towards contradiction that there is no optimal strategy which is impulsive. If all boxes of π^N appear also in π^Y, then it is straightforward to argue that the impulsive strategy that first executes π^Y, and then opens r if no non-zero value was observed, is an impulsive strategy that yields at least the same utility as π^*, and we are done.

Therefore it remains to handle the case where there exists a box in π^N that does not appear in π^Y. In this case, we show that there exists a subset of $\pi^N \setminus \pi^Y$ that can be concatenated to π^Y to improve the overall utility and thus obtain a contradiction.

The main tool we use to this end is the notion of an *impulsive strategy with dummies*. This is a randomized strategy which is determined by a (deterministic) impulsive strategy π and a subset A

of its boxes, denoted π_A, and proceeds as follows: For a box $i \in A$, it proceeds as usual (open the box, observe its value, incur its marginal cost and halt if the observed value is non-zero). For a box $i \notin A$, instead of opening i, it halts with probability p_i and otherwise continues to the next box. In particular π_A only opens boxes from A.

Such a strategy is appealing, since it restricts the set of boxes that might be opened while retaining some of the properties of the original strategy. For example, the contribution of any box $i \in A$ to the expected reward is the same in π_A as in π. Furthermore, every impulsive strategy with dummies is a probability distribution over deterministic impulsive strategies. Thus, any lower bound on its utility applies also to the utility of the best impulsive strategy in its support.

We use the notion of impulsive strategies with dummies to identify a strategy that can be concatenated to π^Y which has a positive marginal utility, thus reaching a contradiction. In particular, we prove that given an impulsive strategy π and any partition $A \uplus B$ of its boxes, the utility attained by π is at most the utility of π_B plus the *marginal* utility of π_A when executed after the boxes in B have been opened. The submodularity of the cost function is crucial to obtain this technical property. The desired strategy that can be concatenated to π^Y can now be identified, by applying this lemma with $\pi := \pi^N$, $A = \pi^N \setminus \pi^Y$, $B = \pi^Y \cap \pi^N$. In particular, we prove that there exists such a strategy in the support of π^N_A.

From Bernoulli to arbitrary instances. In Section 4, we show how to extend Theorem 2 to hold for arbitrary distributions. We do so using the following steps: We devise a transformation that, given an arbitrary instance $\mathcal{I}$ creates a Bernoulli instance $\mathcal{I}'$, which maintains submodularity of the cost function as well as other properties. First, the transformation discretizes the (possibly) continuous and unbounded distributions to have finite supports, and then it "Bernoullifies" each box by associating it with a set of Bernoulli boxes.

We then show a correspondence between strategies for the two instances in which an impulsive strategy for $\mathcal{I}'$ is associated with a fixed order strategy for $\mathcal{I}$. The correspondence preserves the utility up to an arbitrarily small precision. We conclude that if there is an instance that admits a gap between the best fixed-order strategy and the best arbitrary strategy, then it implies that there is a Bernoulli instance that admits a gap between the best impulsive strategy and best arbitrary strategy, contradicting the main result of Section 3. The instance-transformation we use might be of independent interest and find applications in other stochastic settings (such as prophet setting).

Computational Hardness. In Section 5 we prove Theorem 3 even for a very simple subclass of submodular functions (*i.e*, matroid rank functions). To this end, we follow the construction of Svitkina and Fleischer [2011], and design two instances of Pandora's box problem whose cost functions are "indistinguishable" using polynomially many cost queries, but only one of them admits a strategy that yields positive utility. Since no algorithm can distinguish between them efficiently, we conclude that the problem of deciding whether a given instance admits a strategy that attains positive utility is unsolvable with polynomially many cost queries. Moreover, this implies that no approximation can be obtained by an efficient algorithm.

1.3 Related Work

Pandora's Problem originated in economics but has suscitated a keen interest in the computer science community. Weitzman's optimal solution is based on the clever idea of reservation value, a quantity that captures the intrinsic value of a box in the exploration process. The reservation value has a deep connection with the notion of Gittins index [Weber et al. 1992]; actually, Dumitriu et al. [2003] showed that it is possible to rephrase Pandora's problem as a Markov game whose Gittins index coincides with the reservation value. Recently, a simpler proof of the optimality of

Weitzman's rule was also given by Kleinberg et al. [2016]. Following these papers, many interesting modifications of Pandora's Problem have been considered.

Singla [2018] used an adaptivity gap approach to approximately solve Pandora's problem under various combinatorial models, while Olszewski and Weber [2015] studied to which extent a threshold strategy like Weitzman's is optimal when the definition of the reward of the exploration goes beyond the max function.

A successful line of work has also focused on Pandora's Problem with non-obligatory inspection. Here, at the end of the exploration, the decision maker can decide to select an unopened box without having to open it (and thus without paying its cost). Doval [2018] introduced this model, highlighting the surprising property that there are instances where the optimal strategy is adaptive in the order of boxes it chooses. A sequence of papers then closed this problem from the computational perspective [Beyhaghi and Cai 2022; Beyhaghi and Kleinberg 2019; Fu et al. 2022]: Pandora's problem with non-obligatory inspection is NP-hard to solve but a PTAS exists for it. Interestingly enough, this minor tweak in the exploration rule (i.e., giving the possibility of getting a single box "for free" without inspection) hindered both the computational and structural simplicity of the original setting.

Constraints on the order in which the boxes can be opened have also been studied. Esfandiari et al. [2019] considered the case where the boxes have to be opened consistently with a total ordering of the boxes (possibly skipping some). In contrast, Boodaghians et al. [2020] investigated partial orderings on the boxes modeled by precedence graphs. In that work, the authors investigated to which extent the simplicity of the original Pandora's problem extends under order constraints: when the partial ordering on the boxes is represented by a tree, then there exists an optimal strategy that is fixed order and can be computed efficiently; however, under general partial ordering the problem becomes NP-hard to solve, and there are instances where adaptivity is needed to achieve optimality. We further elaborate on the relations with our work in the full version of the paper.

Fu et al. [2018]; Segev and Singla [2021] studied Pandora's Problem with commitment, when, similarly to what happens in online selection problems like secretary or prophet inequalities, only the reward in the last opened box can be collected. Chawla et al. [2021, 2020] investigated what happens when the assumption on the independence of the random rewards in the boxes is dropped, Alaei et al. [2021] introduced the revenue maximization version of the problem, while Bechtel et al. [2022] considered a delegated version of Pandora's problem. Finally, Pandora's problem has also been studied from the learning perspective, both in the sample complexity framework [Guo et al. 2021], and in online learning [Gatmiry et al. 2022; Gergatsouli and Tzamos 2022].

2 PRELIMINARIES

In Pandora's problem there are n boxes, containing hidden values V_i which are distributed according to the independent non-negative distributions D_i. We denote by supp the union of the supports of these distributions. The cost of inspecting a set of boxes is given by a combinatorial cost function $c : 2^{[n]} \to \mathbb{R}_{\geq 0}$, where $[n]$ denotes the set $\{1, \dots, n\}$. We assume that c is always normalized and monotone, *i.e*, $c(\emptyset) = 0$ and $S \subseteq T$ implies $c(S) \leq c(T)$. We also use $(x)^+$ to denote $\max(x, 0)$ for any number $x \in \mathbb{R}$.

We denote an *instance* of the problem by $\mathcal{I} = (D_1, \dots, D_n, c)$. Given an instance $\mathcal{I}$, a *strategy* π for $\mathcal{I}$ inspects the boxes in a sequential manner where each inspection of box i reveals its hidden (random) value V_i. At each round the strategy may choose any uninspected box to inspect next, or it may halt and attain as utility the difference between the largest observed value and the cost of the set of opened boxes. The decisions are based on the given instance and the sequence of opened boxes and realized values so far. Given a strategy π for $\mathcal{I}$, we use the following notation:

- $S(\pi)$ - the (random) ordered set of boxes opened by π.
- $V(\pi) := \max_{i \in S(\pi)} V_i$ - the maximum value observed by π. We also refer to this as the reward obtained by π.
- $\mathrm{u}(\pi) := \mathbb{E}\left[V(\pi)\right] - \mathbb{E}\left[c(S(\pi))\right]$ - the expected utility (i.e., value minus cost) achieved by π.

Note that the quantities defined above depend on the given instance. When not clear from the context, we shall use $\mathrm{u}\left(\mathcal{I};\pi\right)$ to denote the expected utility of strategy π for instance $\mathcal{I}$.

A *randomized* strategy can toss coins before every decision point. Note that these coins can be tossed a priori before the first box is inspected. Therefore every randomized strategy is a distribution over deterministic strategies. In particular, for every randomized strategy there is a deterministic strategy that achieves at least the same utility (the one with the highest utility in the support of the distribution).

Given an instance $\mathcal{I}$, we denote by Π the set of all strategies for $\mathcal{I}$. An *optimal* strategy π^* is a strategy that maximizes the utility, *i.e.*, $\pi^* = \arg\max_{\pi\in\Pi} \mathrm{u}(\pi)$. By the paragraph above, we can assume without loss of generality that π^* is deterministic. Note also that if there is some i for which $\mathbb{E}\left[V_i\right] = \infty$, then the strategy that opens i and halts achieves infinite utility. We thus assume that all distributions D_i have finite expectations.

A *fixed order strategy* π is a strategy in which the order of inspection is non-adaptive. Formally, such a strategy is characterized by a permutation $\sigma : [n] \to [n]$ such that at every round i, the strategy either opens the box $\sigma(i)$, or halts. A strategy π is called a *fixed order strategy with thresholds* $t_1, \ldots, t_n \in \mathbb{R}$ if it is fixed order, and at every round i, π halts if and only if the maximum value inspected so far is at least t_i. The proof of the following observation is deferred to the full version of the paper.

Observation 2.1. For every permutation σ, the optimal strategy with fixed order σ is a fixed order strategy with thresholds.

A *Bernoulli instance* is an instance where all distributions D_i are weighted Bernoulli distributions, *e.g.*, with probability 0.7 $V_i = 18$ and otherwise $V_i = 0$. An *impulsive* strategy for a Bernoulli instance is a fixed order strategy that immediately halts if the value of the currently inspected box is non-zero (the strategy can also halt if the currently observed value is zero). An example of such a strategy is: inspect box 1 and halt if its value is non-zero. Otherwise, inspect box 2 and halt if its value is non-zero. Otherwise, inspect box 7 and halt (regardless of the findings). An example of a non-impulsive strategy is: inspect box 1. If its value is non-zero, inspect box 2 and halt. Otherwise, inspect box 3 and halt. Note that an impulsive strategy is a fixed order strategy, where each threshold equals the weight of its corresponding Bernoulli box (except for the threshold corresponding to the last box, which equals 0). We also remark that the empty strategy which halts immediately without inspecting any boxes is considered an impulsive strategy.

Combinatorial functions. In this paper we study combinatorial cost functions. In particular, given a base set X of elements, we say that a function $c : 2^X \to \mathbb{R}_{\geq 0}$ is

- submodular if $c(x \mid B) \leq c(x \mid A)$ for all $A \subseteq B \subseteq X$, $x \in X \setminus B$, where $c(x \mid S) := c(S \cup \{x\}) - c(S)$ denotes the marginal contribution of element x to set S.
- fractionally subadditive (XOS) if there exists a family of linear function $\{c_i\}$ such that $c(A) = \max_i c_i(A)$, for all $A \subseteq X$.
- subadditive if $c(A \cup B) \leq c(A) + c(B)$ for all $A, B \subseteq X$.

It is known that submodular $\subset$ XOS $\subset$ subadditive, with strict inclusions [Lehmann et al. 2006].

Computational setting. The computational problem we consider is the following (Pandora's) decision problem: given an instance $\mathcal{I}$, decide whether there exists a strategy π for $\mathcal{I}$ that achieves positive utility, *i.e*, $\mathrm{u}(\mathcal{I};\pi) > 0$. An algorithm for this problem gets access to the given cost function via cost queries (analogous to value queries for a combinatorial valuation function); namely, given a set S of elements, a cost query returns $c(S)$.

3 IMPULSIVE OPTIMAL STRATEGIES FOR BERNOULLI INSTANCES

In this section, we prove our main structural result for the special case of Bernoulli instances.

THEOREM 3.1. *For every Bernoulli instance with a submodular cost function there exists an optimal strategy that is impulsive.*

The crux of the proof of Theorem 3.1 is captured by the following Lemma, which is the main technical result of the paper.

LEMMA 3.2. *Let $\mathcal{I}$ be a Bernoulli instance with a submodular cost function. If there exists an optimal strategy for $\mathcal{I}$ of the following form:*

- *Inspect some first box, denoted r, that follows the distribution $V_r = v_r > 0$ with probability $p_r > 0$.*
- *If $V_r = v_r$, execute an impulsive sub-strategy π^Y.*
- *If $V_r = 0$, execute an impulsive sub-strategy π^N.*

Then, there exists an optimal strategy for $\mathcal{I}$ which is impulsive.

Before proving Lemma 3.2, we show how it implies Theorem 3.1.

PROOF OF THEOREM 3.1. We prove this by induction on the number of boxes n. For $n = 1$, the claim is trivially true, since every strategy is an impulsive strategy. Assume by induction that for any Bernoulli instance $\mathcal{I}'$ on $n-1$ boxes with a submodular cost function there exists a deterministic optimal strategy that is impulsive. Let $\mathcal{I} = (D_1, \dots, D_n, c)$ be a Bernoulli instance with n boxes whose cost function is submodular, and let π^* be a deterministic optimal strategy for $\mathcal{I}$. Since the strategy π^* is deterministic, it either does not open any box (and thus π^* is an impulsive strategy), or there exists a box $i \in [n]$ that it inspects first. Note that if $V_i = 0$ with probability 1, then π'can be weakly improved by skipping i and proceeding to the next box: the value obtained by this new strategy is the same for any realization of the boxes, but the incurred cost is weakly improved (by monotonicity of the cost function). Thus we can assume without loss of generality that $V_i = v_i > 0$ with some probability $p_i > 0$, and $V_i = 0$ otherwise.

For each of the two possible realizations of box i, the instance remaining after opening box i is either $\mathcal{I}^N = (D_1, \dots, D_{i-1}, D_{i+1}, \dots, D_n, c')$ if $V_i = 0$, or $\mathcal{I}^Y = (D'_1, \dots, D'_{i-1}, D'_{i+1}, \dots, D'_n, c')$ if $V_i = v_i$, where $c' : [n] \setminus \{i\} \to \mathbb{R}_{\geq 0}$ is the cost function $c'(S) = c(S \cup \{i\}) - c(\{i\}) = c(S \mid \{i\})$, and D'_j is the weighted Bernoulli distribution of $(v_j - v_i)^+$ with probability p_j where D_j is the Bernoulli distribution of having a value of v_j with probability p_j. Note that since c' is the marginal function of c given $\{i\}$ and since c is submodular, then c' is a submodular function. By the induction hypothesis (since $\mathcal{I}^Y, \mathcal{I}^N$ have submodular cost functions and $n-1$ boxes), there exist two optimal strategies π^Y, π^N for $\mathcal{I}^Y, \mathcal{I}^N$, respectively, that are impulsive. Thus there exists an optimal strategy that opens box i, if its value is non-zero executes the sub-strategy π^Y, and otherwise it execute the sub-strategy π^N. By applying Lemma 3.2, we establish that there exists an optimal strategy for $\mathcal{I}$ that is impulsive. □

The remainder of this section is dedicated to the proof of Lemma 3.2. In Section 3.1 we make the required preparation, and in Section 3.2 we provide the full proof of the lemma.

3.1 Setup for Lemma 3.2

In this section we introduce the notation and constructs that we shall need for the proof of Lemma 3.2. Let $\mathcal{I} = (D_r, D_1, \dots, D_n, c)$ be a Bernoulli instance where c is a submodular cost function. For every $i \in \{r\} \cup [n]$, the random value V_i in box i is set to v_i with probability p_i, and to 0 otherwise (with probability $q_i = 1 - p_i$). We can assume without loss of generality that $v_i > 0$ and $p_i > 0$ for every box $i \in \{r\} \cup [n]$, since otherwise $V_i = 0$ with probability 1, in which case any strategy that does open i can be weakly improved by skipping i and proceeding as if its value 0 was observed. If a box i satisfies $p_i = 1$ then we say it is a *deterministic* box.

An impulsive (sub-)strategy π is given by a tuple of box indices (with no repetitions), *e.g*, $(1, 2, 7)$ stands for the impulsive strategy that first inspects box 1 and halts if $V_1 = v_1$, otherwise it proceeds to inspect box 2 and halts if $V_2 = v_2$, and otherwise it proceeds to inspect box 7 and halts. An impulsive strategy can also be given by a tuple of impulsive sub-strategies $(\pi_1, \dots, \pi_k)$, *e.g.*, $((1, 2), (7))$ stands for the strategy $(1, 2, 7)$. The empty strategy that does not inspect any box is also considered an impulsive strategy and is denoted by the tuple $(\varnothing)$. We shall occasionally abuse notation and identify an impulsive strategy π with the set of boxes that form π, *e.g.*, $i \in (1, 2, 7)$ stands for $i \in \{1, 2, 7\}$, and $\pi \subseteq \{1, 2, 3, 4\}$ means that all boxes outside of $\{1, 2, 3, 4\}$ are never inspected by π.

Let π^* be a deterministic optimal strategy for $\mathcal{I}$ in the form given by the statement of Lemma 3.2. Thus, π^* first inspects box r; if it observes that $V_r = v_r$ then it executes the impulsive sub-strategy π^Y and otherwise it executes the impulsive sub-strategy π^N. Note that π^Y and π^N both inspect boxes with indices from $[n]$. We can assume without loss of generality that for every $i \in \pi^Y$, we have $v_i \geq v_r$: Otherwise π^* can be weakly improved by removing i from π^Y — note that the reward obtained in the end of the process is unaffected by the realized value of V_i in this case, and therefore continuing to the suffix of π^Y after i is also optimal. We also assume that each of π^Y and π^N contains at most one deterministic box, in which case it is the last one in the tuple. This too is without loss of generality since impulsive strategies always halt after inspecting a deterministic box. Note that if π^Y is the empty strategy (*i.e* it halts immediately without opening any boxes), then π^* is an impulsive strategy by itself, and we are done. Thus we assume that π^Y is not empty, *i.e.*, $|\pi^Y| \geq 1$. Finally, out of all optimal strategies that satisfy the assumptions above, we also assume that π^* maximizes $|\pi^Y| + |\pi^N|$.

Assume towards contradiction that there is no impulsive strategy for $\mathcal{I}$ that achieves the same utility as π^*. We show that in this case we can replace either π^Y or π^N by impulsive sub-strategies of bigger size, without losing utility. This would constitute a contradiction to the definition of π^*.

Given an impulsive strategy π, We denote by $p_{(\pi)}$ the probability that one of the boxes inspected by π has a non-zero value, *i.e.*, the probability that there is some $i \in S(\pi)$ such that $V_i = v_i$. We denote by $q_{(\pi)} := 1 - p_{(\pi)}$ the probability that $V_i = 0$ for every $i \in \pi$. For the empty strategy we define $p_{(\varnothing)} = 0$ (or equivalently $q_{(\varnothing)} = 1$). Note that by our assumption that $p_i > 0$ for every i, we have $p_{(\pi)} > 0$ for every non-empty impulsive strategy, and $p_{(\pi)} = 1$ if and only if π contains a deterministic box.

Observation 3.3. Let $\pi = (i_1, \dots, i_k) \subseteq [n]$ be an impulsive strategy. Then $p_{(\pi)} = \sum_{j=1}^{k} q_{(i_1, \dots, i_{j-1})} \cdot p_{i_j}$ and $q_{(\pi)} = \prod_{j=1}^{k} q_{i_j} = \prod_{j=1}^{k} (1 - p_{i_j})$.

Note that Observation 3.3 also holds when the coordinates i_j are by themselves impulsive sub-strategies which are not singletons. Also observe that if π^1, π^2 are impulsive strategies such that $\pi^1 \subseteq \pi^2$, then $p_{(\pi^1)} \leq p_{(\pi^2)}$.

We now introduce notation for the marginal utility achieved by an impulsive (sub) strategy executed at some point after inspecting box r. Note that this quantity depends on whether the observed value V_r equals v_r or 0. We thus introduce notation for both cases, and it shall be useful to define these utilities conditioned on already having inspected some set of boxes T. We also introduce a third "non-lower-bounded utility" that we shall need.

Definition 3.4. Given an impulsive strategy $\pi \subseteq [n]$ and a set of boxes $T \subseteq [n]$ such that $T \cap \pi = \varnothing$, we define

- $\mathrm{u}_Y(\pi \mid T) := \mathbb{E}\left[\max_{i \in S(\pi)}(V_i - v_r)^+\right] - \mathbb{E}\left[c\left(S(\pi) \mid \{r\} \cup T\right)\right]$, the marginal utility of π, given that $V_r = v_r$ and that the boxes in T were already opened.
- $\mathrm{u}_N(\pi \mid T) := \mathbb{E}\left[\max_{i \in S(\pi)}(V_i)\right] - \mathbb{E}\left[c\left(S(\pi) \mid \{r\} \cup T\right)\right]$, the marginal utility of π, given that $V_r = 0$ and that the boxes in T were already opened.
- $\mathrm{u}_M(\pi \mid T) := p_{(\pi)} \cdot \mathbb{E}\left[\max_{i \in S(\pi)}(V_i - v_r) \mid \exists i \in S(\pi) \text{ s.t. } V_i = v_i\right] - \mathbb{E}\left[c\left(S(\pi) \mid \{r\} \cup T\right)\right]$.

We write $\mathrm{u}_Y(\pi), \mathrm{u}_N(\pi), \mathrm{u}_M(\pi)$ instead of $\mathrm{u}_Y(\pi \mid \varnothing), \mathrm{u}_N(\pi \mid \varnothing), \mathrm{u}_M(\pi \mid \varnothing)$, respectively. We observe that since c is submodular, then for any sets of boxes $T_1 \subseteq T_2$ that do not intersect π, we have $\mathrm{u}_Y(\pi \mid T_1) \le \mathrm{u}_Y(\pi \mid T_2), \mathrm{u}_N(\pi \mid T_1) \le \mathrm{u}_N(\pi \mid T_2)$ and $\mathrm{u}_M(\pi \mid T_1) \le \mathrm{u}_M(\pi \mid T_2)$.

Observation 3.5. Let $\pi = (i_1, \ldots, i_k) \subseteq [n]$ be an impulsive strategy. Then:

$$\mathrm{u}_N(\pi) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \mathrm{u}_N(i_j \mid \{i_\ell\}_{\ell \in [j-1]}) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \left(p_{i_j} \cdot v_{i_j} - c\left(i_j \mid \{r\} \cup \{i_\ell\}_{\ell \in [j-1]}\right)\right).$$

$$\mathrm{u}_Y(\pi) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \mathrm{u}_Y(i_j \mid \{i_\ell\}_{\ell \in [j-1]}) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \left(p_{i_j} \cdot \left(v_{i_j} - v_r\right)^+ - c\left(i_j \mid \{r\} \cup \{i_\ell\}_{\ell \in [j-1]}\right)\right).$$

$$\mathrm{u}_M(\pi) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \mathrm{u}_M(i_j \mid \{i_\ell\}_{\ell \in [j-1]}) = \sum_{j=1}^{k} q_{(i_1,\ldots,i_{j-1})} \cdot \left(p_{i_j} \cdot \left(v_{i_j} - v_r\right) - c\left(i_j \mid \{r\} \cup \{i_\ell\}_{\ell \in [j-1]}\right)\right).$$

The corresponding expressions $\mathrm{u}_N(\pi \mid T), \mathrm{u}_Y(\pi \mid T), \mathrm{u}_M(\pi \mid T)$ for a set $T \subseteq [n]$ such that $T \cap \pi = \varnothing$ follow the same equations above with the addition of a "T" term after every "|" symbol. As a concrete example, for the strategy $\pi = (1, 2, 7)$ and set of boxes $T = \{4, 5\}$, we have

$$\mathrm{u}_N(\pi \mid T) = p_1 v_1 - c(1 \mid \{r, 4, 5\}) + q_1\left(p_2 v_2 - c(2 \mid \{r, 4, 5, 1\})\right) + q_{(1,2)}\left(p_7 v_7 - c(7 \mid \{r, 4, 5, 1, 2\})\right).$$

Furthermore, observe that

$$\mathrm{u}(\pi^*) = p_r \cdot (v_r + \mathrm{u}_Y(\pi^Y)) - c(r) + q_r \cdot \mathrm{u}_N(\pi^N).$$

Our goal is to replace either π^Y or π^N with a strategy π that achieves at least as much marginal utility, but (potentially) inspects more boxes. This will constitute a contradiction to the assumption that π^* maximizes $\left|\pi^Y\right| + \left|\pi^N\right|$.

The proof of the following straightforward observation can be found in the full version of the paper.

Observation 3.6. Let $\pi \subseteq [n]$ be an impulsive strategy. Then for any set of boxes $T \subseteq [n]$ such that $T \cap \pi = \varnothing$, we have:

- $\mathrm{u}_M(\pi \mid T) \le \mathrm{u}_Y(\pi \mid T) \le \mathrm{u}_N(\pi \mid T)$.
- $\mathrm{u}_M(\pi \mid T) = \mathrm{u}_N(\pi \mid T) - p_{(\pi)} \cdot v_r$.
- If $\pi \subseteq \pi^Y$, then $\mathrm{u}_M(\pi \mid T) = \mathrm{u}_Y(\pi \mid T)$.

Impulsive Strategies with Dummies. Our proof makes use of a particular family of strategies that are distributions over impulsive strategies: an *impulsive strategy with dummies* is given by a (regular) impulsive strategy π, and a subset of boxes $P \subseteq \pi$. The strategy is denoted π_P, and proceeds exactly as π would, with the following single difference: when considering index $i \in \pi$, if it is also the case that $i \notin P$ (*i.e.*, $i \in \pi \setminus P$), then instead of inspecting box i the strategy rather only halts with probability p_i and otherwise proceeds to the next coordinate of the tuple. We refer to the boxes in $\pi \setminus P$ as *dummy* boxes. As an example, the strategy $(2, 1, 4, 7)_{\{1,7\}}$ first halts with probability p_2, then, if it did not halt it proceeds to inspect box 1 and halts if $V_1 = v_1$, otherwise it halts with probability p_4, and then, if it did not halt it proceeds to inspect box 7 and halts. Observe that such a strategy is a distribution over deterministic impulsive strategies. For example, $(2, 1, 4, 7)_{\{1,7\}}$ equals the empty strategy with probability p_2, the strategy (1) with probability $q_2 \cdot p_4$, and the strategy $(1, 7)$ with probability $q_{(2,4)} = q_2 \cdot q_4$. The marginal utility quantities in Definition 3.4 carry over to impulsive strategies with dummies. For example, given the strategy $\pi = (2, 1, 4, 7)$, subset of boxes $P = \{1, 7\}$ and another set of boxes $T = \{4, 5\}$ that has presumably already been opened, we have

$$\mathrm{u}_M(\pi_P \mid T) = q_2 \left(p_1(v_1 - v_r) - c(1 \mid \{r, 4, 5\})\right) + q_{(2,1,4)} \left(p_7(v_7 - v_r) - c(7 \mid \{1\} \cup \{r, 4, 5\})\right).$$

Note that the value and cost terms corresponding to boxes 1 and 7 are multiplied by the factors q_2 and $q_{(2,1,4)}$, respectively, and that these are the same factors these terms are multiplied by in the expression for the utility of π. Also note that $T \cap \pi \neq \varnothing$ in this example, but we allow this since the strategy we are computing the utility for, π_P, never inspects boxes from T. Furthermore, the expression $p_{(\pi_P)}$ — the probability that one of the boxes inspected by π_P has a non-zero value — is also well defined. *E.g*, in the example above this probability equals $q_2 p_1 + q_{(2,1,4)} p_7$.

Observation 3.6 also carries over to impulsive strategies with dummies, where the third bullet there holds for any such strategy π_P where $P \subseteq \pi^Y$. Finally, observe that for $P = \varnothing$ we have that π_P is the empty strategy, and that for $P = \pi$ we have that π_P coincides with π.

3.2 Proof of Lemma 3.2

The first step of the proof of Lemma 3.2 is the following inequality.

Lemma 3.7. *It holds that* $p_{(\pi^Y)} < p_{(\pi^N)}$.

Proof. Assume towards contradiction that $p_{(\pi^Y)} \geq p_{(\pi^N)}$. This implies

$$\mathrm{u}_N\left(\pi^Y\right) = \mathrm{u}_Y\left(\pi^Y\right) + p_{(\pi^Y)} v_r \geq \mathrm{u}_Y\left(\pi^N\right) + p_{(\pi^N)} v_r \geq \mathrm{u}_N\left(\pi^N\right) \geq \mathrm{u}_N\left(\pi^Y\right),$$

where the equality and the second inequality hold by Observation 3.6, the first inequality holds by the optimality of π^Y for the scenario that $V_r = v_r$, and the last inequality holds by the optimality of π^N for the scenario that $V_r = 0$.

Thus all expressions in the above chain are equal and in particular we have $\mathrm{u}_N\left(\pi^Y\right) = \mathrm{u}_N\left(\pi^N\right)$. This implies that the strategy π' that first inspects r and then executes π^Y regardless of the realization of V_r is also optimal. Now consider the impulsive strategy $\pi'' = \left(\pi^Y, r\right)$. Since $v_i \geq v_r$ for any $i \in \pi^Y$, then the maximum value observed by π'' coincides with that of π' for any realization of the boxes. On the other hand the cost incurred by π'' is weakly less then that of π', again for any realization of the boxes. Thus the impulsive strategy π'' is optimal as well, a contradiction. □

The following lemma is the main technical tool needed for the rest of the proof.

Lemma 3.8. *Let $\pi \subseteq [n]$ be any impulsive strategy, and let $\pi = A \uplus B$ be a partition of the set of boxes corresponding to π. Then we have* $\mathrm{u}_N(\pi) \leq \mathrm{u}_N(\pi_A \mid B) + \mathrm{u}_N(\pi_B)$.

Proof. Let π, A, B be as in the lemma statement. Recall that the expressions in the inequality are each made up of (expected) value terms and (expected) cost terms. We first show that the value terms cancel out. Explicitly, we show that $\mathbb{E}\left[\max_{i \in S(\pi)}(V_i)\right] = \mathbb{E}\left[\max_{i \in S(\pi_A)}(V_i)\right] + \mathbb{E}\left[\max_{i \in S(\pi_B)}(V_i)\right]$.

To see this, denote π without loss of generality as $\pi = (1, \ldots, k)$. Then, for any $i \in [k]$, the value term corresponding to i when expanding $\mathbb{E}\left[\max_{i \in S(\pi)}(V_i)\right]$ is $q_{(1,\ldots,i-1)} p_i v_i$. Furthermore, regardless of whether $i \in A$ or $i \in B$, this would also be the value term corresponding to i when expanding the right-hand side of the equation — if $i \in A$ then this would appear in the expansion of $\mathbb{E}\left[\max_{i \in S(\pi_A)}(V_i)\right]$ and if $i \in B$ then this would appear in the expansion of $\mathbb{E}\left[\max_{i \in S(\pi_B)}(V_i)\right]$. In fact, this last discussion also shows:

Observation 3.9. For any impulsive strategy $\pi \subseteq [n]$ and for any partition $\pi = A \cup B$ of the set of boxes corresponding to π, we have $p_{(\pi)} = p_{(\pi_A)} + p_{(\pi_B)}$.

It remains to handle the cost terms. For ease of exposition we omit the "$\{r\}$" terms inside the conditional cost terms. This has no effect on the proof. Thus, in the remainder of the proof we establish the following inequality:

$$\mathbb{E}\left[c\left(S(\pi_A) \mid B\right)\right] + \mathbb{E}\left[c\left(S(\pi_B)\right)\right] - \mathbb{E}\left[c\left(S(\pi)\right)\right] \leq 0. \tag{1}$$

We prove inequality (1) by induction on $|A| + |B|$, and we start with the base case $|A| + |B| = 0$. In this case, π is the empty strategy implying that all three summands in inequality (1) equal 0, and the inequality follows.

We now assume that $|A| + |B| \geq 1$. Denote π again without loss of generality as $\pi = (1, \ldots, k)$, where $k = |A| + |B|$. We expand each of the expressions in inequality (1):

$$\begin{aligned}
\mathbb{E}\left[c\left(S(\pi)\right)\right] &= \sum_{i=1}^{k} q_{(1,\ldots,i-1)} \cdot c\left(i \mid \{1, \ldots, i-1\}\right) \\
\mathbb{E}\left[c\left(S(\pi_B)\right)\right] &= \sum_{i \in B} q_{(1,\ldots,i-1)} \cdot c\left(i \mid \{1, \ldots, i-1\} \cap B\right) \\
\mathbb{E}\left[c\left(S(\pi_A) \mid B\right)\right] &= \sum_{i \in A} q_{(1,\ldots,i-1)} \cdot c\left(i \mid \{1, \ldots, i-1\} \cup B\right)
\end{aligned}$$

By plugging these into inequality (1) and taking out common "q" factors, we get the equivalent inequality

$$\begin{aligned}
&\sum_{i \in A} q_{(1,\ldots,i-1)} \cdot \left[c\left(i \mid \{1, \ldots, i-1\} \cup B\right) - c\left(i \mid \{1, \ldots, i-1\}\right)\right] && (2) \\
&+ \sum_{i \in B} q_{(1,\ldots,i-1)} \cdot \left[c\left(i \mid \{1, \ldots, i-1\} \cap B\right) - c\left(i \mid \{1, \ldots, i-1\}\right)\right] && (3) \\
&\leq 0
\end{aligned}$$

We now split to two cases. In the first (easy) case we assume that $k \in A$, *i.e.*, the last box potentially to be inspected by π is a box from A. Note that in this case we have $\{1, \ldots, k-1\} \cup B = \{1, \ldots, k-1\}$, and thus the summand in line (2) corresponding to $i = k$ cancels out and equals 0. Therefore, if we denote $\pi^{(k)} := (1, \ldots, k-1)$, then inequality (1) is equivalent to

$$\mathbb{E}\left[c\left(S\left(\pi^{(k)}_{A \setminus \{k\}}\right) \mid B\right)\right] + \mathbb{E}\left[c\left(S\left(\pi^{(k)}_B\right)\right)\right] - \mathbb{E}\left[c\left(S\left(\pi^{(k)}\right)\right)\right] \leq 0,$$

which holds by the induction hypothesis.

We now handle the case $k \in B$. Note that if $A = \varnothing$ and $B = [k]$, then inequality (1) holds trivially — the summands in line (2) do not exist, and the summands in line (3) cancel out. We thus assume that $|A| \geq 1$.

The rest of the proof involves a systematic manipulation of the inequality. Mostly, we shall make repeated use of the following observation, which we term the "cancellation lemma". and whose simple proof can be found in the full version of the paper. We use colors in the lemma statement so that it will be easier to see how we apply it in the rest of the proof.

Lemma 3.10. (Cancellation Lemma) *For every cost function $c : 2^{[n]} \to \mathbb{R}_{\geq 0}$, subset $T \subseteq [n]$ and elements $h, \ell \in [n] \setminus T$, we have*

$$c\left(h \mid T \cup \{\ell\}\right) - c\left(\ell \mid T \cup \{h\}\right) = c\left(h \mid T\right) - c\left(\ell \mid T\right).$$

Denote $A = \{a_1, \ldots, a_w\}$, where $w = |A| \geq 1$ and where the order $a_1, \ldots, a_w$ is consistent with the relative ordering of A in π, *i.e.*, $a_1 < \cdots < a_w$. Thus we can rewrite line (2) as follows:

$$\sum_{i=1}^{w} q_{(1,\ldots,a_i-1)} \cdot \left[c\left(a_i \mid \{a_1, \ldots, a_{i-1}\} \cup B\right) - c\left(a_i \mid \{1, \ldots, a_i - 1\}\right)\right] \tag{4}$$

For each $i = 1, \ldots, w$ we shall refer to the corresponding summand in the above sum as the "a_i-summand". We can also rewrite the summand in line (3) corresponding to $i = k$ (recall that we are in the case that $k \in B$) as

$$q_{(1,\ldots,k-1)} \cdot \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid (B \setminus \{k\}) \cup \{a_1, \ldots, a_{w-1}, a_w\}\right)\right].$$

and we shall refer to it as the "k-summand". Note the coloring of k and a_w, highlighting their roles in the (first upcoming) application of the cancellation lemma. Consider the a_w-summand:

$$q_{(1,\ldots,a_w-1)} \cdot \left[c\left(a_w \mid \{a_1, \ldots, a_{w-1}\} \cup (B \setminus \{k\}) \cup \{k\}\right) - c\left(a_w \mid \{1, \ldots, a_w - 1\}\right)\right]$$

We cannot directly apply the lemma on the k-summand and the a_w-summand because of the different "q" factors. To get around this issue, denote the difference inside the square parentheses in the k-summand by

$$D := \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid (B \setminus \{k\}) \cup \{a_1, \ldots, a_{w-1}, a_w\}\right)\right],$$

and note that $D \geq 0$ since c is submodular. Furthermore, note that $q_{(1,\ldots,a_w-1)} \geq q_{(1,\ldots,k-1)}$. Thus we can (weakly) increase the k-summand as follows:

$$\begin{aligned} q_{(1,\ldots,k-1)} \cdot D &\leq q_{(1,\ldots,a_w-1)} \cdot D \\ &= q_{(1,\ldots,a_w-1)} \cdot \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid (B \setminus \{k\}) \cup \{a_1, \ldots, a_{w-1}, a_w\}\right)\right] \end{aligned}$$

Now the "q" factors are the same and we can apply the cancellation lemma. Thus we remove the $\{k\}$ term from the a_w-summand, which now becomes

$$q_{(1,\ldots,a_w-1)} \cdot \left[c\left(a_w \mid \{a_1, \ldots, a_{w-1}\} \cup (B \setminus \{k\})\right) - c\left(a_w \mid \{1, \ldots, a_w - 1\}\right)\right]$$

We also remove the a_w term from the k-summand, which now becomes

$$q_{(1,\ldots,a_w-1)} \cdot \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid (B \setminus \{k\}) \cup \{a_1, \ldots, a_{w-2}, a_{w-1}\}\right)\right].$$

Note the coloring of k and a_{w-1} highlighting the next application of the cancellation lemma. Consider now the a_{w-1}-summand:

$$q_{(1,\ldots,a_{w-1}-1)} \cdot \left[c\left(a_{w-1} \mid \{a_1, \ldots, a_{w-2}\} \cup (B \setminus \{k\}) \cup \{k\}\right) - c\left(a_{w-1} \mid \{1, \ldots, a_{w-1} - 1\}\right)\right]$$

As before, we cannot directly apply the cancellation lemma due to the different "q" factors. As before, we get around this by using the fact that $q_{(1,\ldots,a_{w-1}-1)} \geq q_{(1,\ldots,a_w-1)}$ and the fact that c is submodular in order to replace the $q_{(1,\ldots,a_w-1)}$ factor in the k-summand by the factor $q_{(1,\ldots,a_{w-1}-1)}$, making the k-summand (weakly) larger by doing so.

After the application of the cancellation lemma, we remove the $\{k\}$ term from the a_{w-1}-summand. We also remove the a_{w-1} term from the k-summand, which becomes

$$q_{(1,\ldots,a_{w-1}-1)} \cdot \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid (B \setminus \{k\}) \cup \{a_1, \ldots, a_{w-3}, a_{w-2}\}\right)\right],$$

and again note the coloring of k and a_{w-2} highlighting the next application of the cancellation lemma. We continue this way, applying the cancellation lemma to the summands corresponding to the pairs $(k, a_{w-2}), (k, a_{w-3}), \ldots, (k, a_1)$.

After the last application, the k-summand becomes

$$q_{(1,\ldots,a_1-1)} \cdot \left[c\left(k \mid B \setminus \{k\}\right) - c\left(k \mid B \setminus \{k\}\right)\right] = 0,$$

and the sum of the a_i-summands (Line (4)) is modified by replacing "B" with "$B \setminus \{k\}$". Thus, recalling the notation $\pi^{(k)} := (1, \ldots, k-1)$, we have shown in the above process that

$$\mathbb{E}\left[c\left(S(\pi_A) \mid B\right)\right] + \mathbb{E}\left[c\left(S(\pi_B)\right)\right] - \mathbb{E}\left[c\left(S(\pi)\right)\right] \leq$$
$$\mathbb{E}\left[c\left(S\left(\pi_A^{(k)}\right) \mid B \setminus \{k\}\right)\right] + \mathbb{E}\left[c\left(S\left(\pi_{B\setminus\{k\}}^{(k)}\right)\right)\right] - \mathbb{E}\left[c\left(S\left(\pi^{(k)}\right)\right)\right],$$

and the bottom expression is upper-bounded by 0, by the induction hypothesis. This concludes the proof of Lemma 3.8. □

The remainder of the proof of Lemma 3.2 proceeds as follows. By Lemma 3.7 we have $\pi^N \setminus \pi^Y \neq \emptyset$, since otherwise $\pi^N \subseteq \pi^Y$ which implies $p_{(\pi^N)} \leq p_{(\pi^Y)}$. Lemma 3.7 also implies that $p_{(\pi^Y)} < 1$,*i.e.*, π^Y does not contain a deterministic box. To prove Lemma 3.2 we show that there exists a non-empty impulsive sub-strategy made from boxes in $\pi^N \setminus \pi^Y$ that we can concatenate to π^Y without decreasing utility. This would constitute a contradiction to the definition of π^*.

Let A and B be the sets of boxes defined by $A = \pi^N \setminus \pi^Y$, $B = \pi^Y \cap \pi^N \subseteq \pi^Y$. Note that $\pi^N = A \uplus B$ and that $A \neq \emptyset$. We can write π^N as a concatenation of contiguous sub-strategies made up of boxes from A or B as follows: $\pi^N = \left(B^{\text{pre}}, A^1, B^1, \ldots, A^k, B^k, A^{\text{suff}}\right)$, where $A = \left(\uplus_{i=1}^k A^k\right) \uplus A^{\text{suff}}$, $B = B^{\text{pre}} \uplus \left(\uplus_{i=1}^k B^k\right)$. The only sub-strategies that we allow to be empty in this presentation are B^{pre}, for the case that π^N starts with a box from A, and A^{suff}, for the case that π^N ends with a box from B (in the latter case we must have $k \geq 1$ as otherwise $A = \emptyset$ and we get a contradiction).

We define the strategies $\pi^A = \left(A^1, A^2, \ldots, A^k, A^{\text{suff}}\right)$, $\pi^B = \left(B^{\text{pre}}, B^1, B^2, \ldots, B^k\right)$, and note the difference between π^A, π^B, and π_A^N, π_B^N. The former are deterministic strategies, whereas the latter are strategies with dummies.

Claim 3.11. We have $\mathsf{u}_M\left(\pi^A \mid \pi^Y\right) < 0$.

Proof. Assume towards contradiction that $\mathsf{u}_M\left(\pi^A \mid \pi^Y\right) \geq 0$. Then by Observation 3.6 we also have $\mathsf{u}_Y\left(\pi^A \mid \pi^Y\right) \geq 0$. Thus we can concatenate π^A to π^Y to obtain a new optimal strategy that contradicts the definition of π^* as the maximizer of $\left|\pi^Y\right| + \left|\pi^N\right|$. Formally, consider the strategy obtained from π^* by replacing the strategy π^Y with $\left(\pi^Y, \pi^A\right)$. Then the utility obtained does not decrease, since

$$\mathsf{u}_Y\left(\left(\pi^Y, \pi^A\right)\right) = \mathsf{u}_Y\left(\pi^Y\right) + q_{(\pi^Y)}\mathsf{u}_Y\left(\pi^A \mid \pi^Y\right) \geq \mathsf{u}_Y\left(\pi^Y\right).$$

Thus, the new strategy is optimal as well, and as discussed above we get a contradiction. □

Observe that

$$\mathsf{u}_N\left[\left(\pi^Y, \pi^A\right)\right] \leq \mathsf{u}_N\left(\pi^N\right) \leq \mathsf{u}_N(\pi_A^N \mid B) + \mathsf{u}_N(\pi_B^N) \leq \mathsf{u}_N(\pi_A^N \mid \pi^Y) + \mathsf{u}_N(\pi_B^N),$$

where the first inequality holds by the optimality of π^N for the scenario where $V_r = 0$, the second inequality holds by Lemma 3.8, and the third holds by submodularity of the cost function c since $B \subseteq \pi^Y$. Now, since the strategy $\left(\pi^Y, \pi^A\right)$ is a superset of π^N, then in particular we have

$$p_{(\pi^Y, \pi^A)} \geq p_{(\pi^N)} = p_{(\pi_A^N)} + p_{(\pi_B^N)},$$

where the second equality holds by Observation 3.9. Therefore, the chain of inequalities above implies:

$$\begin{aligned}
\mathsf{u}_M\left[\left(\pi^Y,\pi^A\right)\right] &= \mathsf{u}_N\left[\left(\pi^Y,\pi^A\right)\right] - p_{(\pi^Y,\pi^A)}v_r \\
&\le \mathsf{u}_N(\pi_A^N \mid \pi^Y) + \mathsf{u}_N(\pi_B^N) - \left(p_{(\pi_A^N)} + p_{(\pi_B^N)}\right)v_r \\
&= \left(\mathsf{u}_N(\pi_A^N \mid \pi^Y) - p_{(\pi_A^N)}v_r\right) + \left(\mathsf{u}_N(\pi_B^N) - p_{(\pi_B^N)}v_r\right) \\
&= \mathsf{u}_M(\pi_A^N \mid \pi^Y) + \mathsf{u}_M(\pi_B^N),
\end{aligned}$$

where the first and last equalities hold by Observation 3.6. Since $\mathsf{u}_M\left[\left(\pi^Y,\pi^A\right)\right] = \mathsf{u}_M\left(\pi^Y\right) + q_{(\pi^Y)}\mathsf{u}_M\left(\pi^A \mid \pi^Y\right)$, then the above inequality implies

$$\mathsf{u}_M\left(\pi^Y\right) - \mathsf{u}_M(\pi_B^N) \le \mathsf{u}_M(\pi_A^N \mid \pi^Y) - q_{(\pi^Y)}\mathsf{u}_M\left(\pi^A \mid \pi^Y\right). \tag{5}$$

In the following claim we rule out the case that $k = 0$, *i.e*, that in π^N all boxes from B are inspected before all boxes from A. The proof can be found in the full veresion of the paper.

Claim 3.12. There exist boxes $a \in A, b \in B$ such that π^N inspects b only after inspecting A, *i.e.*, $k \ge 1$.

In the remainder we show that for some $j \in [k]$, we can concatenate the (non-empty) strategy $\left(A^1,\ldots,A^j\right)$ to π^Y without losing utility. This would constitute a contradiction to the assumption that π^* maximizes $\left|\pi^Y\right| + \left|\pi^N\right|$. To this end we analyze both sides of inequality (5). First, the left hand side satisfies

$$0 \le \mathsf{u}_Y\left(\pi^Y\right) - \mathsf{u}_Y(\pi_B^N) = \mathsf{u}_M\left(\pi^Y\right) - \mathsf{u}_M(\pi_B^N) \tag{6}$$

where the equality holds by Observation 3.6 (recall that $B \subseteq \pi^Y$), and the inequality holds since π^Y is the optimal sub-strategy for the scenario where $V_r = v_r$. For the right hand side we have the following claim which is derived through a careful algebraic manipulation that mostly applies Observation 3.5. Its proof can be found in the full version of the paper.

Claim 3.13. The right hand side of inequality (5) satisfies

$$\mathsf{u}_M(\pi_A^N \mid \pi^Y) - q_{(\pi^Y)}\mathsf{u}_M\left(\pi^A \mid \pi^Y\right) \le \sum_{j=1}^{k} q_{(B^{\mathrm{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}\mathsf{u}_M\left(\left(A^1,\ldots,A^j\right) \mid \pi^Y\right).$$

We plug the inequality in Claim 3.13 and inequality (6) into inequality (5), to get

$$0 \le \sum_{j=1}^{k} q_{(B^{\mathrm{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}\mathsf{u}_M\left(\left(A^1,\ldots,A^j\right) \mid \pi^Y\right) \tag{7}$$

Note that all the factors $q_{(B^{\mathrm{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}$ are strictly positive since π^Y does not have a deterministic box and B is a subset of π^Y. This implies that at least one of the expressions $\mathsf{u}_M\left(\left(A^1,\ldots,A^j\right) \mid \pi^Y\right)$, for $j \in [k]$, is non-negative. Choose some j that satisfies this. To conclude the proof we would like to say that we can concatenate $\left(A^1,\ldots,A^j\right)$ to π^Y without decreasing the utility, thus obtaining the desired contradiction, analogously to what we did in the proof of Claim 3.11. The (small) problem is that there might be boxes $i \in \left(A^1,\ldots,A^j\right)$ for which $v_i < v_r$.

To get around this, we note that it cannot be the case that all boxes $i \in \left(A^1,\ldots,A^j\right)$ satisfy $v_i < v_r$, since the contribution of these boxes to $\mathsf{u}_M\left(\left(A^1,\ldots,A^j\right) \mid \pi^Y\right)$ is strictly negative. Consider then the impulsive strategy with dummies $\pi^{\mathsf{rand}} = \left(A^1,\ldots,A^j\right)_{\{i \mid v_i \ge v_r\} \cap (A^1,\ldots,A^j)}$ which is obtained from $\left(A^1,\ldots,A^j\right)$ by replacing the inspection of every box i for which $v_i < v_r$ with a decision to halt with probability p_i and otherwise continue to the next box. Then we have $\mathsf{u}_M\left(\pi^{\text{rand}} \mid \pi^Y\right) \geq 0$. Furthermore, by the observation above this strategy has non-empty deterministic strategies in its support (recall that an impulsive strategy with dummies is a distribution over deterministic impulsive strategies). Thus, there exists one such strategy, denoted π^{diff}, for which $\mathsf{u}_M\left(\pi^{\text{diff}} \mid \pi^Y\right) \geq 0$, and which satisfies $v_i \geq v_r$ for every $i \in \pi^{\text{diff}}$. This in turn implies $\mathsf{u}_Y\left(\pi^{\text{diff}} \mid \pi^Y\right) \geq 0$, by Observation 3.6. We now concatenate π^{diff} to π^Y without decreasing the utility, analogously to what we did in the proof of Claim 3.11, and get a contradiction to the definition of π^*. This concludes the proof of Lemma 3.2.

4 REDUCTION TO BERNOULLI INSTANCES

In this section, we show how Theorem 3.1 implies that for any instance with arbitrary distributions and a submodular cost function there is an optimal strategy with a fixed-order. We do so by transforming an instance with arbitrary distributions, to an instance with Bernoulli distributions. We first discretize the support of the distributions using a discretization parameter ϵ and by capping the values by a sufficiently large number (that depends on the distributions and on ϵ). This leads to a modified instance with finite support. Then, we replace each box with a finite number of boxes with weighted Bernoulli distributions. Both transformations maintain several key properties of the instance. The goal of these transformations is to modify the instance to have only a finite number of weighted Bernoulli boxes, for which we can apply Theorem 3.1.

Transformation 1: Transformation $\mathcal{T}^\epsilon$, defined by a parameter $\epsilon > 0$, proceeds as follows: given an instance $\mathcal{I} = (D_1, \ldots, D_n, c)$, let $\kappa_\epsilon := \min\{\kappa \geq 0 \mid \sum_{i=1}^n \mathbb{E}\left[(V_i - \kappa)^+\right] \leq \epsilon\}$. Such a constant κ_ϵ is well defined for every $\epsilon > 0$ since $\sum_{i=1}^n \mathbb{E}\left[(V_i - \kappa)^+\right]$ is a monotone continuous decreasing function in κ, and the limit as κ approaches infinity is 0 (see full version of the paper for details). Using κ_ϵ, for every i, D_i^ϵ is defined to be the distribution of the random variable $\bar{V}_i = \epsilon \cdot \lfloor \frac{\min(V_i, \kappa_\epsilon)}{\epsilon} \rfloor$. We remark that since κ_ϵ is finite, then the support of the new set of distributions is finite. Finally, the output of $\mathcal{T}^\epsilon$ is $\mathcal{T}^\epsilon(D_1, \ldots, D_n, c) = (D_1^\epsilon, \ldots, D_n^\epsilon, c)$. The proof of the following proposition can be found in the full version.

Proposition 4.1. *For every instance $\mathcal{I}$ and every $\epsilon > 0$, let $\mathcal{I}^\epsilon = \mathcal{T}^\epsilon(\mathcal{I})$. Then:*

1. *For every strategy π on instance $\mathcal{I}$ there exists a strategy π' on instance $\mathcal{I}^\epsilon$ such that* $\mathsf{u}\left(\mathcal{I}; \pi\right) \leq \mathsf{u}\left(\mathcal{I}^\epsilon; \pi'\right) + 2\epsilon$.
2. *For every strategy π' on instance $\mathcal{I}^\epsilon$ there exists a strategy π on instance $\mathcal{I}$ such that* $\mathsf{u}\left(\mathcal{I}; \pi\right) \geq \mathsf{u}\left(\mathcal{I}^\epsilon; \pi'\right)$. *Furthermore, if π' is a fixed-order strategy, then there exists such a fixed-order strategy π.*

Transformation 2: Transformation $\mathcal{T}$ receives an instance $\mathcal{I} = (D_1, \ldots, D_n, c)$ with distributions with finite supports, and returns a Bernoulli instance by the following process: We can assume without loss of generality that 0 is in the union of the supports supp, then, we can rename the elements of the union of the supports in an increasing order supp $= \{v_1, \ldots, v_m\}$, where $0 = v_1 < v_2 < \ldots < v_m$. For every $i \in [n]$ and $j \in [m]$, let $D_{i,j}$ be the weighted Bernoulli distribution that returns the value v_j with probability $\frac{\mathbb{P}(V_i = v_j)}{\mathbb{P}(V_i \leq v_j)}$, and 0 otherwise (where $\frac{0}{0}$ is interpreted as 0). Let $c' : 2^{[n]\times[m]} \to \mathbb{R}_{\geq 0}$ be the cost function where for every $S \subseteq [n] \times [m]$,

$$c'(S) := c(\{i \mid \exists j \in [m] \text{ such that } (i, j) \in S\}).$$

Then $\mathcal{T}(\mathcal{I}) = (D_{1,1}, \ldots, D_{n,m}, c')$. One can easily verify that $\mathcal{T}$ maintains monotonicity and normalization of the cost function. The following claim shows that it also maintains submodularity of the cost function, and its proof is deferred to the full version of the paper.

Proposition 4.2. *If c is submodular, then c' obtained by transformation $\mathcal{T}$ is also submodular.*

In the full version of the paper we show that $\mathcal{T}$ maintains also MRF, GS, coverage, XOS and subadditivity of the cost function, but not budget additive.

The next proposition shows that the new instance $\mathcal{I}' = \mathcal{T}(\mathcal{I})$ is in some sense equivalent to $\mathcal{I}$. The proof of the following proposition is deferred to the full version of the paper.

Proposition 4.3. *For every instance $\mathcal{I}$, let $\mathcal{I}' = \mathcal{T}(\mathcal{I})$. Then:*

(1) *For every strategy π on instance $\mathcal{I}$, there exists a strategy π' on instance $\mathcal{I}'$ such that* $\mathsf{u}(\mathcal{I};\pi) \leq \mathsf{u}(\mathcal{I}';\pi')$.
(2) *For every strategy π' on instance $\mathcal{I}'$, there exists a strategy π on instance $\mathcal{I}$ such that* $\mathsf{u}(\mathcal{I};\pi) \geq \mathsf{u}(\mathcal{I}';\pi')$. *Furthermore, if π' is impulsive, then there exists such a fixed-order strategy π.*

In Section 3 we showed that for weighted Bernoulli instances with submodular costs, there exists an optimal strategy that is impulsive. We next show that this implies our main theorem:

Theorem 4.4. *For every instance $\mathcal{I} = (D_1, \ldots, D_n, c)$ where c is submodular, there exists an optimal strategy that is a fixed order strategy with thresholds.*

Proof. Let π^* be an optimal strategy for $\mathcal{I}$ and let π be the optimal strategy for $\mathcal{I}$ among the strategies with a fixed order. Assume towards contradiction that $\mathsf{u}(\mathcal{I};\pi^*) > \mathsf{u}(\mathcal{I};\pi)$. Let $\epsilon = \frac{\mathsf{u}(\mathcal{I};\pi^*)-\mathsf{u}(\mathcal{I};\pi)}{4}$, and let $\mathcal{I}^\epsilon = \mathcal{T}^\epsilon(\mathcal{I})$. By Proposition 4.1, there exists π_1 such that $\mathsf{u}(\mathcal{I};\pi^*) \leq \mathsf{u}(\mathcal{I}^\epsilon;\pi_1) + 2\epsilon$. Let $\mathcal{I}' = \mathcal{T}(\mathcal{I}^\epsilon)$. Then, by proposition 4.3 there exists π_2 such that $\mathsf{u}(\mathcal{I}^\epsilon;\pi_1) \leq \mathsf{u}(\mathcal{I}';\pi_2)$. By Theorem 3.1 there exists an impulsive strategy π_3 such that $\mathsf{u}(\mathcal{I}';\pi_2) \leq \mathsf{u}(\mathcal{I}';\pi_3)$. By Proposition 4.3 there exists a fixed-order π_4 such that $\mathsf{u}(\mathcal{I}^\epsilon;\pi_4) \geq \mathsf{u}(\mathcal{I}';\pi_3)$, and by Proposition 4.1 bthere exists a fixed-order π_5 such that $\mathsf{u}(\mathcal{I};\pi_5) \geq \mathsf{u}(\mathcal{I}^\epsilon;\pi_4)$. All together we have:

$$\mathsf{u}(\mathcal{I};\pi_5) \geq \mathsf{u}(\mathcal{I}^\epsilon;\pi_4) \geq \mathsf{u}\left(\mathcal{I}';\pi_3\right) \geq \mathsf{u}\left(\mathcal{I}';\pi_2\right) \geq \mathsf{u}(\mathcal{I}^\epsilon;\pi_1) \geq \mathsf{u}(\mathcal{I};\pi^*) - 2\epsilon > \mathsf{u}(\mathcal{I};\pi),$$

which contradicts the assumption that π is the optimal fixed-order strategy. □

5 COMPUTATIONAL RESULTS

In this section we show that the task of finding an optimal strategy for Pandora's problem with submodular costs does not admit a polynomial time algorithm. In fact, we show a stronger result, namely that there exists no algorithm for the Pandora's *decision problem* that uses a polynomial number of cost queries. An algorithm for the Pandora's decision problem takes as input an instance of the Pandora's problem with a combinatorial cost function and outputs whether there exists a strategy yielding strictly positive utility on that instance.[2]

5.1 Distinguishing Submodular Functions

To formalize our argument we use the notion of distinguishability of submodular functions [Svitkina and Fleischer 2011]. We say that an algorithm distinguishes between two cost functions c_1 and c_2 if it produces different outputs when given oracle access to c_1 versus oracle access to c_2. Here, we construct a family of cost functions and a baseline cost function that are hard to distinguish using polynomially many cost queries, similarly to the construction of Svitkina and Fleischer [2011]. Let X be a set of n boxes, and let $\alpha = \left\lceil \ln n \cdot \frac{\sqrt{n}}{5} \right\rceil$ and $\beta = \left\lceil \frac{\ln^2 n}{5} \right\rceil$. On this set of boxes we define a

[2]We remark that even a demand oracle to the cost function, in the sense of Blumrosen and Nisan [2007], would not allow us to solve the decision problem with polynomially many queries. The reason is that our impossibility result already holds for matroid rank functions, a strict subclass of gross substitutes, for which a demand query can be simulated by polynomially many cost queries; see the full version of the paper for definitions of gross substitutes and matroid rank functions.

"baseline" cost function $c_0(S) = \min\{|S|, \alpha\}$. Then, for any subset $R \subseteq X$ of boxes with $|R| = \alpha$, we define the cost function c_R:

$$c_R(S) = \min\{|S|, \alpha, \beta + |S \cap R^C|\}. \tag{8}$$

It is immediate to see that c_0 and c_R are submodular and differ on sets S such that $\beta + |S \cap R^C|$ is strictly smaller than $\min\{\alpha, |S|\}$. Consider now a random set $\mathcal{R}$ that is drawn uniformly at random from all the subsets of X of cardinality α. It is possible to show that no deterministic algorithm can distinguish c_R (for a random set $R \sim \mathcal{R}$) from c_0, with high probability. We formalize this result in the following theorem whose proof is deferred to the full version.

Theorem 5.1. *Let $\mathcal{A}$ be any deterministic algorithm that has a cost oracle access to a submodular function $c_R \sim c_{\mathcal{R}}$ over a set X of n elements, which outputs a set $S \subseteq X$ using polynomially many cost queries. Then, for any sufficiently large n we have: $\mathbb{P}(\mathcal{A}$ distinguishes $c_{\mathcal{R}}$ from $c_0) \leq n^{-1}$.*

Note that, by the proof of Theorem 5.1, n^{-1} as a bound on the probability that $\mathcal{A}$ distinguishes $c_{\mathcal{R}}$ from c_0 can be replaced by n^{-b} for any constant b. The corresponding "sufficiently large n" condition would then be $n > \frac{a+2+b}{-\ln(0.851)}$.

5.2 A Family of Difficult Instances

As we show next, the family of submodular cost functions introduced above induces a family of instances of Pandora's problem such that (i) the baseline instance admits no strategy that gives positive utility, and (ii) every other instance in the family admits a strategy obtaining positive utility.

Formally, fix any large enough n and consider the following class of instances of Pandora's Problem with submodular cost functions: there is a set X of n boxes with i.i.d. values distributed according to the following weighted Bernoulli distribution: the value of every box in X is $M = 5\beta > 0$ with probability $p = \frac{1}{\alpha}$, and 0 otherwise. For each $R \subseteq X$, with $|R| = \alpha$, we define the instance $\mathcal{I}_R$ with the above random values and the cost function c_R that is given in Equation (8). Moreover, we construct the baseline instance $\mathcal{I}_0$ using the same random variables, but with cost function c_0. There is a crucial difference between $\mathcal{I}_0$ and $\mathcal{I}_R$: With c_R it is possible to find a subset of α boxes such that only the first β of them have non-zero marginal cost, while this is impossible under c_0. With our choice of M and p it is possible to leverage this property and show the following Lemma, whose proof is deferred to the full version.

Lemma 5.2. *For any sufficiently large n, no strategy extracts positive utility from $\mathcal{I}_0$, while for any R there exists a strategy that extracts positive utility from $\mathcal{I}_R$.*

5.3 The Computational Impossibility Result

We are ready for the main theorem of the section: since it is not possible to distinguish in polynomial time between c_R and the baseline c_0, then it is not possible to assess, in polynomial time, whether an instance of Pandora's problem can yield positive utility (as $\mathcal{I}_R$) or not (as the baseline instance $\mathcal{I}_0$). This immediately implies that no computationally efficient approximation for Pandora's problem with Submodular cost exists.

To formalize our result we introduce the concept of positivity oracle: a (possibly randomized) algorithm $\mathcal{O}$ is a positivity oracle for Pandora's problem with Submodular cost if it takes in input an instance $\mathcal{I}$ of the problem (i.e. knowledge of the distributions of the random rewards and cost oracle access to the cost function) and outputs an answer to the question whether it exists or not a strategy π such that $\mathrm{u}(\mathcal{I}; \pi) > 0$. We say that $\mathcal{O}$ is correct on instance $\mathcal{I}$ with a certain probability p if it outputs the correct answer to Pandora's decision problem on that instance with probability at least p, where the probability is with respect to the internal randomization of $\mathcal{O}$. In other words, a positivity oracle is an algorithm for Pandora's decision problem. We say that a positivity oracle $\mathcal{O}$ is efficient if there exists a constant a (that depends on $\mathcal{O}$ but not the specific instance) such that $\mathcal{O}$ issues at most n^a cost queries with probability 1 on all instances.

Theorem 5.3. *Fix any efficient positivity oracle $\mathcal{O}$ and positive constant $\varepsilon > 0$. Then there exists an instance $\mathcal{I}$ on n boxes (for n sufficiently large) such that $\mathcal{O}$ outputs the correct answer on $\mathcal{I}$ with probability at most $0.5 + \varepsilon$.*

Proof. The possibly randomized positivity oracle $\mathcal{O}$ is just a distribution over deterministic ones, so for any $R \subseteq X$ of cardinality α and any deterministic positivity oracle O, we denote with $\mathcal{E}(O, R)$ the event that O gives a different output when receiving $\mathcal{I}_R$ and $\mathcal{I}_0$ as input. Recall that $\mathcal{R}$ is a set of cardinality α drawn uniformly at random. Denote with $\mathbf{O}$ (respectively, $\mathbf{O}_d$) the set of all the randomized (resp., deterministic) efficient positivity oracles. Yao's principle gives the following:

$$\min_R \mathbb{P}\left(\mathcal{E}(\mathcal{O}, R)\right) \leq \min_R \max_{\mathcal{O}^* \in \mathbf{O}} \mathbb{P}\left(\mathcal{E}(\mathcal{O}^*, R)\right) \leq \max_{O \in \mathbf{O}_d} \mathbb{P}\left(\mathcal{E}(O, \mathcal{R})\right). \tag{9}$$

Consider the rightmost term; each deterministic positivity oracle O is an algorithm with cost oracle access to the underlying submodular cost function, which gives different outputs on $\mathcal{I}_R$ and $\mathcal{I}_0$ if it distinguishes $c_{\mathcal{R}}$ from c_0 (see definition of distinguishability), given that the rest of the input is identical. From Equation (9) we have then:

$$\min_R \mathbb{P}\left(\mathcal{E}(\mathcal{O}, R)\right) \leq \max_{O \in \mathbf{O}_d} \mathbb{P}\left(\mathcal{E}(O, \mathcal{R})\right) \leq \frac{1}{n} \leq \varepsilon, \tag{10}$$

where the second inequality follows from Theorem 5.1, for any n sufficiently large.

What we have shown so far is that there exists a set R such that $\mathcal{O}$ gives the same output on both $\mathcal{I}_0$ and $\mathcal{I}_R$ with probability at least $1 - \varepsilon$ even though the correct answer to Pandora's decision problem on the two instances is different. Let now $\mathcal{G}_0$, respectively $\mathcal{G}_R$, be the event that $\mathcal{O}$ is correct on input $\mathcal{I}_0$, respectively $\mathcal{I}_R$. If the probability of $\mathcal{G}_0$ is smaller than $0.5 + \varepsilon$ then there is nothing else to prove, as we can choose $\mathcal{I} = \mathcal{I}_0$; otherwise, we have the following:

$$\mathbb{P}\left(\mathcal{G}_R\right) = \mathbb{P}\left(\mathcal{G}_R \cap \mathcal{E}(\mathcal{O}, R)\right) + \mathbb{P}\left(\mathcal{G}_R \setminus \mathcal{E}(\mathcal{O}, R)\right) \leq \mathbb{P}\left(\mathcal{E}(\mathcal{O}, R)\right) + \mathbb{P}\left(\mathcal{G}_0^C\right) \leq \varepsilon + (0.5 - \varepsilon) = 0.5.$$

To see why the previous formula holds we study separately the two summands. The event $\mathcal{G}_R \cap \mathcal{E}(\mathcal{O}, R)$ is clearly contained in $\mathcal{E}(\mathcal{O}, R)$, and we know that its probability is smaller than ε by Equation (10). The event $\mathcal{G}_R \setminus \mathcal{E}(\mathcal{O}, R)$, on the other hand, is disjoint from $\mathcal{G}_0$; in fact, we know that if $\mathcal{O}$ gives the same output for $\mathcal{I}_0$ and $\mathcal{I}_R$, at most one of the two instances receives the correct answer to its decision problem. Finally, we are under the assumption that $\mathbb{P}\left(\mathcal{G}_0\right) \geq 0.5 + \varepsilon$, thus its complementary has at most a probability $0.5 - \varepsilon$ to realize. □

The previous result directly implies that no approximation result can be achieved for Pandora's problem with submodular cost functions using polynomially many cost queries: assume by contradiction that such an algorithm exists, then it would be easy to construct a positivity oracle that violates the previous theorem, e.g. by repeatedly simulating the algorithm and using concentration.

ACKNOWLEDGMENTS

Michal Feldman and Ben Berger are partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), by an Amazon Research Award, and by the NSF-BSF (grant number 2020788). Federico Fusco and Tomer Ezra are partially supported by ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets" and MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets". The work of Federico Fusco is also partially supported by PNRR MUR project PE0000013-FAIR" and PNRR MUR project IR0000013-SoBigData.it.

REFERENCES

Saeed Alaei, Ali Makhdoumi, and Azarakhsh Malekian. 2021. Revenue Maximization Under Unknown Private Values With Non-Obligatory Inspection. In *EC*. ACM, 27–28.

Curtis Bechtel, Shaddin Dughmi, and Neel Patel. 2022. Delegated Pandora's Box. In *EC*. ACM, 666–693.

Hedyeh Beyhaghi and Linda Cai. 2022. Pandora's Problem with Nonobligatory Inspection: Optimal Structure and a PTAS. *To appear in STOC 2023, preprint on the arXiv* abs/2212.01524 (2022).

Hedyeh Beyhaghi and Robert Kleinberg. 2019. Pandora's Problem with Nonobligatory Inspection. In *EC*. ACM, 131–132.

Liad Blumrosen and Noam Nisan. 2007. *Algorithmic game theory.* Cambridge University Press, New York, USA.

Shant Boodaghians, Federico Fusco, Philip Lazos, and Stefano Leonardi. 2020. Pandora's Box Problem with Order Constraints. In *EC*. ACM, 439–458.

Shuchi Chawla, Evangelia Gergatsouli, Jeremy McMahan, and Christos Tzamos. 2021. Approximating Pandora's Box with Correlations. *CoRR* abs/2108.12976 (2021).

Shuchi Chawla, Evangelia Gergatsouli, Yifeng Teng, Christos Tzamos, and Ruimin Zhang. 2020. Pandora's Box with Correlations: Learning and Approximation. In *FOCS*. IEEE, 1214–1225.

Laura Doval. 2018. Whether or not to open Pandora's box. *Journal of Economic Theory* 175 (2018), 127–158.

Ioana Dumitriu, Prasad Tetali, and Peter Winkler. 2003. On Playing Golf with Two Balls. *SIAM J. Discret. Math.* 16, 4 (2003), 604–615.

Hossein Esfandiari, Mohammad Taghi Hajiaghayi, Brendan Lucier, and Michael Mitzenmacher. 2019. Online Pandora's Boxes and Bandits. In *AAAI*. AAAI Press, 1885–1892.

Hu Fu, Jiawei Li, and Daogao Liu. 2022. Pandora Box Problem with Nonobligatory Inspection: Hardness and Approximation Scheme. *To appear in STOC 2023, preprint on the arXiv* abs/2207.09545v2 (2022).

Hao Fu, Jian Li, and Pan Xu. 2018. A PTAS for a Class of Stochastic Dynamic Programs. In *ICALP (LIPIcs, Vol. 107)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 56:1–56:14.

Khashayar Gatmiry, Thomas Kesselheim, Sahil Singla, and Yifan Wang. 2022. Bandit Algorithms for Prophet Inequality and Pandora's Box. *CoRR* abs/2211.08586 (2022).

Evangelia Gergatsouli and Christos Tzamos. 2022. Online Learning for Min Sum Set Cover and Pandora's Box. In *ICML (Proceedings of Machine Learning Research, Vol. 162)*. PMLR, 7382–7403.

Chenghao Guo, Zhiyi Huang, Zhihao Gavin Tang, and Xinzhi Zhang. 2021. Generalizing Complex Hypotheses on Product Distributions: Auctions, Prophet Inequalities, and Pandora's Problem. In *COLT (Proceedings of Machine Learning Research, Vol. 134)*. PMLR, 2248–2288.

Jon M. Kleinberg and Robert Kleinberg. 2018. Delegated Search Approximates Efficient Search. In *EC*. ACM, 287–302.

Robert D. Kleinberg, Bo Waggoner, and E. Glen Weyl. 2016. Descending Price Optimally Coordinates Search. In *EC*. ACM, 23–24.

Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games Econ. Behav.* 55, 2 (2006), 270–296.

Wojciech Olszewski and Richard Weber. 2015. A more general Pandora rule? *Journal of Economic Theory* 160 (2015), 429–437.

Danny Segev and Sahil Singla. 2021. Efficient Approximation Schemes for Stochastic Probing and Prophet Problems. In *EC*. ACM, 793–794.

Sahil Singla. 2018. The Price of Information in Combinatorial Optimization. In *SODA*. SIAM, 2523–2532.

Zoya Svitkina and Lisa Fleischer. 2011. Submodular Approximation: Sampling-based Algorithms and Lower Bounds. *SIAM J. Comput.* 40, 6 (2011), 1715–1737.

Richard Weber et al. 1992. On the Gittins index for multiarmed bandits. *The Ann. Appl. Probab.* 2, 4 (1992), 1024–1033.

Martin L Weitzman. 1979. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society* (1979), 641–654.