Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Pandora's Problem with Combinatorial Cost

Ben Berger; , Tomer Ezra; , Michal Feldman; , Federico Fusco

Published online in Articles in Advance 18 Dec 2024

. https://doi.org/10.1287/moor.2023.0248

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–29
ISSN 0364-765X (print), ISSN 1526-5471 (online)
https://pubsonline.informs.org/journal/moor

Pandora's Problem with Combinatorial Cost

Ben Berger,[a] Tomer Ezra,[b] Michal Feldman,[c,d] Federico Fusco[e,*]

[a] Research Team, Offchain Labs, Inc., Clifton, New Jersey 07013; [b] The Center of Mathematical Sciences and Applications, Harvard University, Cambridge, Massachusetts 02138; [c] Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 6997801, Israel; [d] Microsoft Israel Development Center (ILDC), Microsoft Research, Herzliya 4672415, Israel; [e] Department of Computer, Control, and Management Engineering, Sapienza University of Rome, 00185 Rome, Italy

*Corresponding author

Contact: bberger@offchainlabs.com, https://orcid.org/0000-0002-9115-6160 (BB); tomer@cmsa.fas.harvard.edu, https://orcid.org/0000-0003-0626-4851 (TE); mfeldman@tauex.tau.ac.il, https://orcid.org/0000-0002-2915-8405 (MF); fuscof@diag.uniroma1.it, https://orcid.org/0000-0001-6250-945X (FF)

Received: August 11, 2023
Revised: May 31, 2024
Accepted: August 5, 2024
Published Online in Articles in Advance: December 18, 2024

MSC2020 Subject Classifications: Primary: 68W01; secondary: 90B35

https://doi.org/10.1287/moor.2023.0248

Abstract. 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 applications where search cost is nonadditive. Weitzman's celebrated algorithm (1979) demonstrates that for additive costs, the optimal search strategy is nonadaptive and computationally feasible. We inquire to which extent this structural and computational simplicity extends beyond additive costs. Our main result is that the class of submodular cost functions admits an optimal strategy that follows a fixed order, thus preserving the structural simplicity of additive costs. In contrast, for the more general class of subadditive (or even fractionally subadditive) cost functions, the optimal strategy may inevitably determine the search order adaptively. On the computational side, obtaining any approximation to the optimal utility requires superpolynomially many queries to the cost function, even for a strict subclass of submodular cost functions.

Funding: B. Berger and M. Feldman are partially supported by the HORIZON EUROPE European Research Council [Grant 866132] and the United States-Israel Binational Science Foundation [Grant 2020788]. T. Ezra is supported by the Harvard University Center of Mathematical Sciences and Applications. F. Fusco is partially supported by the National Recovery and Resilience Plan (PNRR) by Ministero dell'Università e della Ricerca [Project IR0000013-SoBigData.it] and the Future Artificial Intelligence Research project funded by the NextGenerationEU program within the National Recovery and Resilience Plan Extended Partnership for Artificial Intelligence (PNRR-PE-AI) scheme [Grant M4C2, investment 1.3, line on Artificial Intelligence]. Additionally, this work was supported by Ministero dell'Università e della Ricerca [Grant ALGADIMAR], the Israel Science Foundation [Grant 317/17], the Amazon Research Award, and the HORIZON EUROPE European Research Council [Grant 788893].

Keywords: Pandora's problem • optimal stopping • combinatorial cost functions

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 [30] 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.

Weitzman [30] showed that optimal strategies for this problem exhibit both structural simplicity and computational simplicity. In particular, there exists an optimal strategy that 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, this 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 [23], Singla [27]) or modifying the rules of exploration (e.g., Boodaghians et al. [8], Doval [11], Esfandiari et al. [13], Fu et al. [15]). 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, whereas 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 [30] 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 an optimal strategy may unavoidably be *adaptive*.[1] This is demonstrated in the following example.

Example 1. Consider an instance with three 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 $\{\emptyset, \{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}\}$ is 0, and the total cost of exploring a set of boxes from the collection $\{\{2,3\}, \{1,2,3\}\}$ is 20. Opening both boxes 2 and 3 is too expensive for any reasonable strategy. In fact, it can be shown (see Claim 1) 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 zero), 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 zero). 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. [22], 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 as illustrated in Figure 1.

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 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.[2]

Theorem 1. *There exists an instance of Pandora's problem with an XOS cost function that admits no optimal strategy with nonadaptive exploration order.*

Figure 1. (Color online) The complement-free hierarchy of combinatorial functions. In the upper part (lower part, respectively) of the thickly shaded classes, the structural (computational, respectively) simplicity is preserved, and in the lightly shaded classes, it is not. Thm., theorem.

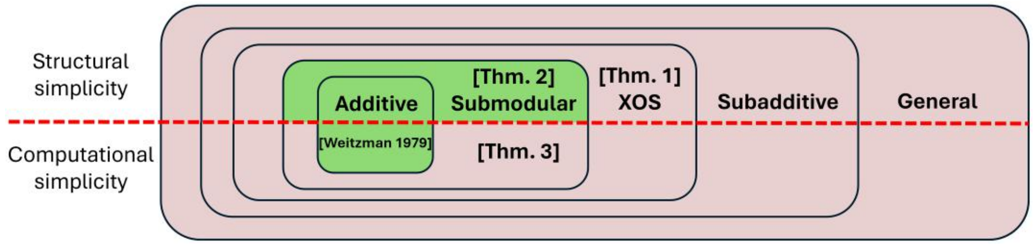

On the face of it, the above theorem seems to be unrelated to Example 1, where the cost function exhibits strong complementarity. However, we identify a close connection between the two results. In particular, we show that every instance with a (monotone and normalized) cost function over n boxes induces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

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, $c'(S) := c(S|T) := c(S \cup T) - c(S)$) 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 $c(\cdot|T)$ for any fixed set T. In particular, the scenario depicted in Example 1, where the combined cost of opening boxes 2 and 3 is excessive, whereas 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, nonadaptive order. Our first main result answers this question in the affirmative (see Sections 4 and 5).

Theorem 2 (Informal). *Every instance of Pandora's problem with a submodular cost function admits an optimal strategy with nonadaptive exploration order.*

Our second main result shows that although the structural simplicity is preserved under submodular cost functions, the computational simplicity is not preserved. In particular, in Section 6, we prove the following stronger result.

Theorem 3 (Informal). *The problem of deciding whether a given instance of Pandora's problem with a submodular* (*or even gross-substitutes*) *cost function admits a strategy that attains strictly positive utility requires superpolynomially 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.

1.2. 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. [8], Esfandiari et al. [13], Kleinberg and Kleinberg [20], Singla [27], Weitzman [30]). 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: $\mathbb{E}[(V-z)^+] = c$. Weitzman's optimal strategy opens the boxes in decreasing order of reservation value (breaking ties arbitrarily), halting when the current maximum observed reward exceeds the reservation value of the next unopened box. Because 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 2. Consider an instance with two identical boxes, each with a random reward of two with probability $\frac{1}{3}$ (and zero otherwise) and a symmetric unit-demand cost function with a cost of one (i.e., $c(\{1\}) = c(\{2\}) = c(\{1,2\}) = 1$ and $c(\emptyset) = 0$). 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, is determined 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.

1.2.1. Bernoulli Instances. In Section 4, we prove Theorem 2 for Bernoulli instances (i.e., instances where each box i has value v_i with probability p_i and value of zero 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 nonzero 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 proceed along 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 nonzero value is realized, then it executes some impulsive substrategy π^Y, and if its realized value is zero, then it executes another impulsive substrategy π^N. This is proved by induction using the fact that the marginal cost of a submodular function is also submodular.

Under this assumption, we proceed as follows. Assume toward contradiction that there is no optimal strategy, which is impulsive. If all boxes of π^N also appear in π^Y, then it is straightforward to argue that the impulsive strategy that first executes π^Y and then opens r if no nonzero 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 that 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 nonzero). 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 because 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 (because the randomness deciding whether to stop at box $i \notin A$ can be drawn at the beginning of the process). Thus, any lower bound on its utility also applies 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 $A\dot{\cup}B$ 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.

1.2.2. From Bernoulli to Arbitrary Instances. In Section 5, 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 4. The instance transformation that we use might be of independent interest and find applications in other stochastic settings (such as prophet setting).

1.2.3. Computational Hardness. In Section 6, we prove Theorem 3 even for a very simple subclass of submodular functions (i.e., matroid rank functions (MRFs)). To this end, we follow the construction of Svitkina and Fleischer [28]; we 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. Because 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 [29]); actually, Dumitriu et al. [12] 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. [21]. Following these papers, many interesting modifications of Pandora's problem have been considered.

Singla [27] used an adaptivity gap approach to approximately solve Pandora's problem under various combinatorial models, whereas Olszewski and Weber [23] 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 nonobligatory 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 [11] 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 [5], Beyhaghi and Kleinberg [6], Fu et al. [14]); Pandora's problem with nonobligatory inspection is NP hard to solve, but a polynomial-time approximation scheme 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. [13] considered the case where the boxes have to be opened consistently with a total ordering of the boxes (possibly skipping some); in a related setting, Berger et al. [4] studied Pandora's problem in presence of deadline constraints, where boxes are available only in certain time slots. In contrast, Boodaghians et al. [8] 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 Appendix B.

Fu et al. [15] and Segev and Singla [26] 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. [9], Chawla et al. [10], and Gergatsouli and Tzamos [18] investigated what happens when the assumption on the independence of the random rewards in the boxes is dropped. Alaei et al. [1] introduced the revenue maximization version of the problem, whereas Bechtel et al. [3] 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. [19]) and in online learning (Gatmiry et al. [16], Gergatsouli and Tzamos [17]).

2. Preliminaries

In Pandora's problem, there are n boxes containing hidden values V_i, which are distributed according to the independent nonnegative 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 $\{1, \ldots, n\}$. 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 $\mathcal{I} = (D_1, \ldots, D_n, c)$. 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. More formally, if we denote by supp the union of the supports of the distributions D_i, then a deterministic strategy for $\mathcal{I}$ can be defined[3] by a function $\pi : \cup_{i=0}^{n-1} ([n]^i \times \mathsf{supp}^i) \to [n] \cup \{\perp\}$.

A strategy might use randomization that is independent of the realizations of the boxes, in which case we call it a *randomized strategy*. There are two ways to visualize the behavior of a randomized strategy; either a deterministic strategy is randomly drawn before the first box is inspected, or the strategy randomizes each decision, possibly as a function of the previous history. Because the random values are independent of the strategy's decisions, these two perspectives are equivalent so that 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 a subset of boxes S and a value v, let $\mathcal{I}(S,v)$ denote the subinstance where the unopened boxes are S, and the highest value observed so far is v. The subinstance $\mathcal{I}(S,v)$ is equivalent[4] to an instance $((D_i{}')_{i\in S}, c')$, where $D_i{}'$ is the distribution of $\max(X_i - v, 0)$ where $X_i \sim D_i$ and $c' : 2^S \to \mathbb{R}_{\geq 0}$ is such that $c'(T) = c(T \cup ([n] \setminus S)) - c([n] \setminus S)$. We use the term substrategy for a description of a strategy on a subinstance or on the equivalent instance of the mentioned subinstance.

Given a strategy π for $\mathcal{I}$, we use the following notation.

- $S(\pi)$—the (random) ordered set of boxes opened by π.
- $\mathsf{u}(\pi) := \mathbb{E}[\max_{i\in S(\pi)} V_i - c(S(\pi))]$—the expected utility achieved by π. The first[5] term represents the reward obtained by π (i.e., the maximum value observed by the strategy). The second term represents the exploration cost induced by π.

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

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., $\pi^* = \arg\max_{\pi\in\Pi} \mathsf{u}(\pi)$). 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 $\mathbb{E}[V_i] = \infty$, 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 nonadaptive. 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 Appendix C.

Observation 1. For every permutation σ, there exists an optimal strategy with fixed order σ that is a fixed-order strategy with thresholds.

A *Bernoulli instance* is an instance where all distributions D_i are weighted Bernoulli distributions[6] (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 nonzero (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 nonzero. Otherwise, inspect box 2, and halt if its value is nonzero. Otherwise, inspect box 7, and halt (regardless of the findings). An example of a nonimpulsive strategy is inspect box 1. If its value is nonzero, 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 zero). We also remark that the empty strategy that halts immediately without inspecting any boxes is considered an impulsive strategy.

2.1. 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 $c(x|B) \leq c(x|A)$ for all $A \subseteq B \subseteq X$, $x \in X \setminus B$, where $c(x|S) := c(S \cup \{x\}) - c(S)$ 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 $A \subseteq X$; or
- subadditive if $c(A \cup B) \leq c(A) + c(B)$ for all $A, B \subseteq X$.

It is known that submodular ⊂ XOS ⊂ subadditive, with strict inclusions (Lehmann et al. [22]).

3. Adaptive Order Is Necessary Beyond Submodular Cost Functions

In this section, we expand our discussion on Example 1 and use it to prove Theorem 1. We start describing the (only) optimal strategy for that instance.

Claim 1. The only optimal strategy for the instance of Example 1 is adaptive.

Proof. Consider the following strategy. First, open box 1, and if the value of box 1 is 10, then open box 2; if the value of box 1 is zero, then open box 3. We prove that this adaptive strategy is the only optimal one. As mentioned in Section 2, every randomized strategy is a distribution over deterministic strategies, and therefore, it is sufficient to prove that this strategy is optimal among deterministic strategies. We first claim that an optimal strategy will never open both boxes 2 and 3. This holds because after opening one of them, then opening the other leads to a cost of 20, but it can only increase the value by at most 12. If the first box the strategy opens is 1, then if its value is 10, then the only box that has positive marginal utility is 2; therefore, the strategy should open it, and if its value is zero, then because box 3 has higher marginal utility (and we never open both boxes 2 and 3), then the strategy must open box 3. If box 2 is opened first, then if its value is 12, all other boxes have nonpositive marginal utility, and if its value is 0, then box 1 is the only one with positive utility; so, the strategy should open it. This leads to a lower utility than the one that starts with opening box 1. If box 3 is opened first, then all other boxes have nonpositive marginal utility, and this leads to a lower utility than the one that starts with opening box 1. □

The function described in Example 1 is not fractionally subadditive. For this reason, as a second step in our analysis, we argue that starting from any monotone and normalized cost function (like the one in Example 1), it is possible to construct a suitable XOS function.

Claim 2. For every monotone and normalized function $f: 2^X \to \mathbb{R}_{\geq 0}$, the function $g: 2^{X \cup \{0\}} \to \mathbb{R}_{\geq 0}$ defined as follows is XOS:

$$g(S) = n \cdot f(X) \cdot \mathbb{I}_{\{S \neq \emptyset\}} + f(S \cap X).$$

Moreover, the marginal of g with respect to the set $\{0\}$ is f (i.e., for every $S \subseteq X$, it holds that $f(S) = g(S \cup \{0\}) - g(\{0\})$).

Proof. To show that g is XOS, we need to present additive functions such that the maximum over them is g. For every $S \subseteq X$, we define an additive function a^S over the elements in $X \cup \{0\}$ as follows. If $S = \emptyset$, then $a^S(0) = n \cdot f(X)$, and $a^S(i) = 0$ for every $i \in [n]$. If $S \neq \emptyset$, then $a^S(i) = 0$ for $i \in (X \setminus S) \cup \{0\}$, and $a^S(i) = \frac{n \cdot f(X) + f(S)}{|S|}$ for $i \in S$. Observe that under this definition, for every $T \subseteq X \cup \{0\}$, we have

$$a^{T \cap X}(T) = \begin{cases} \emptyset & T = \emptyset \\ n \cdot f(X) + f(T \cap X) & \text{otherwise.} \end{cases}$$

That is, $a^{T \cap X}(T) = g(T)$. Thus, in order to establish that $g = \max_{S \subseteq X} a^S$, it remains to show that for every $T \subseteq X \cup \{0\}$ and for every $S \subseteq X$, we have $a^S(T) \leq g(T)$. Let T and S be as described above. If $T = \emptyset$, then the claim clearly holds (the inequality becomes $0 \leq 0$). We thus assume that $T \neq \emptyset$. Now, if $S = \emptyset$, then it is straightforward to see that

$$a^{\emptyset}(T) = \begin{cases} n \cdot f(X) & 0 \in T \\ 0 & \text{otherwise} \end{cases} \leq n \cdot f(X) \cdot \mathbb{I}_{\{T \neq \emptyset\}} + f(T \cap X) = g(T).$$

If $S \neq \emptyset$, then $a^S(T) = a^S(T \cap S)$. If we also have $S \subseteq T$, then $a^S(T) = a^S(S) \leq g(T)$ by monotonicity of f. If, on the other hand, we have $S \setminus T \neq \emptyset$, then

$$a^S(T) = \frac{|T \cap S|}{|S|}(n \cdot f(X) + f(S)) \leq n \cdot f(X) \leq g(T).$$

This concludes the proof. □

As a final ingredient for Theorem 1, we prove that starting from any instance (possibly with non-XOS costs) that only admits optimal adaptive strategy, it is possible to construct an instance with XOS costs for which the optimal strategy must be adaptive as well.

Lemma 1. *Given an instance $\mathcal{I} = (D_1, \ldots, D_n, c)$ where every optimal strategy must use adaptive order, there exists a distribution D_0 and an XOS cost function $c': 2^{[n] \cup \{0\}} \to \mathbb{R}_{\geq 0}$, where every optimal strategy for instance $\mathcal{I}' = (D_0, D_1, \ldots, D_n, c')$ must use an adaptive order.*

Proof. Consider an instance with an additional box (which we will refer to as box 0). The cost function c' is the cost function described in Claim 2. The random variable V_0 for box 0 is defined as follows. We draw an independent sample s_i from each D_i, and we also draw a Bernoulli random variable b, which equals one with probability $1/2$ and equals zero otherwise. Then, V_0 is defined to be $2 \cdot b \cdot (1 + n \cdot c([n]) + \max_i s_i)$.

Assume toward contradiction that there exists a deterministic optimal strategy $\pi^*_{\mathcal{I}'}$ for instance $\mathcal{I}'$ that uses a fixed order, and let $\sigma : [n+1] \to [n] \cup \{0\}$ be that order. Let $\tau_1, \dots, \tau_{n+1}$ be the optimal thresholds that $\pi^*_{\mathcal{I}'}$ uses (i.e., $\pi^*_{\mathcal{I}'}$ opens $\sigma(i)$ if up to round $i-1$, the maximum realized value so far is strictly less than τ_i).

We first claim that $\tau_1 > 0$ (i.e., the strategy $\pi^*_{\mathcal{I}'}$ opens at least one box). If $\pi^*_{\mathcal{I}'}$ does not open any box, then $\pi^*_{\mathcal{I}'}$ can be strictly improved by opening box 0 because $\mathbb{E}[V_0] = 2 \cdot \mathbb{E}[b] \cdot \mathbb{E}[1 + n \cdot c([n]) + \max_i s_i)] > c(\{0\})$.

We next show that there is a fixed-order strategy with the same utility as $\pi^*_{\mathcal{I}'}$ that opens box 0 first. Let $i_0 = \sigma^{-1}(0)$. If $i_0 = 1$, then $\pi^*_{\mathcal{I}'}$ opens box 0 first, and we are done. Else, consider now the following fixed-order strategy π with order $\sigma' = (0, \sigma(1), \dots, \sigma(i_0 - 1), \sigma(i_0 + 1), \dots, \sigma(n+1))$. The strategy opens box 0; then, it always opens box $\sigma(1)$. For $i = 2, \dots, i_0 - 1, i_0 + 1, \dots, n+1$, it opens box $\sigma(i)$ if the maximum among the values of boxes $\sigma(1), \dots, \sigma(i-1)$ is strictly less than τ_i. Note that for $i < i_0$, the strategy ignores the value of box 0 for the decision of whether to halt, whereas for $i > i_0$, the strategy acts exactly as in $\pi^*_{\mathcal{I}'}$. For every realization of the values of the boxes, the costs of π and $\pi^*_{\mathcal{I}'}$ are the same, and the value obtained by π is at least as large as the value of $\pi^*_{\mathcal{I}'}$. Thus, π is an optimal strategy.

Under the case that the value of box $V_0 = 0$ (which happens with a probability of at least 0.5), the marginal instance that π is facing is exactly $\mathcal{I}$, and π must also be optimal for this case (because it happens with a nonzero probability). But, because π is a fixed-order strategy, its marginal strategy for this case is also a fixed-order strategy, and this contradicts that no fixed-order strategy is optimal for $\mathcal{I}$. This concludes the proof. □

The previous claim together with Example 1 yields Theorem 1.

Theorem 1. *There exists an instance of Pandora's problem with an XOS cost function that admits no optimal strategy with nonadaptive exploration order.*

4. Impulsive Optimal Strategies for Bernoulli Instances

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

Theorem 4. *For every Bernoulli instance with a submodular cost function, there exists an optimal strategy that is impulsive.*

The crux of the proof of Theorem 4 is captured by the following lemma, which is the main technical result of the paper.

Lemma 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 substrategy π^Y.*
- *If $V_r = 0$, execute an impulsive substrategy π^N.*

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

Before proving Lemma 2, we show how it implies Theorem 4.

Proof of Theorem 4. We prove this by induction on the number of boxes n. For $n = 1$, the claim is trivially true because 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 $\mathcal{I}$. Because the strategy π^* is deterministic, either it 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 the 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 subinstance remaining after opening box i is equivalent either to $\mathcal{I}^N = (D_1, \dots, D_{i-1}, D_{i+1}, \dots, D_n, c')$ if $V_i = 0$ or to $\mathcal{I}^Y = (D'_1, \dots, D'_{i-1}, D'_{i+1}, \dots, D'_n, c')$ if $V_i = v_i$, where $c' : [n] \setminus \{i\} \to \mathbb{R}_{\geq 0}$ is the cost function $c'(S) = c(S \cup \{i\}) - c(\{i\}) = c(S | \{i\})$ 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 because c' is the marginal function of c given $\{i\}$ and c is submodular, then c' is a submodular function. By the induction hypothesis ($\mathcal{I}^Y, \mathcal{I}^N$ 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 nonzero and executes the substrategy π^Y, and otherwise, it executes the substrategy π^N. By applying Lemma 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 2. In Section 4.1, we make the required preparation, and in Section 4.2, we provide the full proof of the lemma.

4.1. Setup for Lemma 2

In this section, we introduce the notation and constructs that we shall need for the proof of Lemma 2. Let $\mathcal{I} = (D_r, D_1, \ldots, D_n, c)$ 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 zero 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 $i \in \{r\} \cup [n]$ because 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 of zero was observed. If a box i satisfies $p_i = 1$, then we say that it is a *deterministic* box.

An impulsive (sub-)strategy π is given by a tuple of box indices (with no repetitions). For example, (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 substrategies $(\pi_1, \ldots, \pi_k)$ (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 $(\emptyset)$. 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 2. Thus, π^* first inspects box r; if it observes that $V_r = v_r$, then it executes the impulsive substrategy π^Y, and otherwise, it executes the impulsive substrategy π^N. Note that π^Y and π^N both inspect boxes with indices from $[n]$.

We make the following assumptions.

- 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 because 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., $|\pi^Y| \geq 1$).
- Of all optimal strategies that satisfy the assumptions above, we also assume that π^* maximizes $|\pi^Y| + |\pi^N|$.
- Lastly, we assume toward 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 substrategies 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 nonzero 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 $p_{(\emptyset)} = 0$ (or equivalently, $q_{(\emptyset)} = 1$). Note that by our assumption that $p_i > 0$ for every i, we have $p_{(\pi)} > 0$ for every nonempty impulsive strategy and $p_{(\pi)} = 1$ if and only if π contains a deterministic box.

Observation 2. Let $\pi = (i_1, \ldots, i_k) \subseteq [n]$ be an impulsive strategy. Then,

- $p_{(\pi)} = \sum_{j=1}^{k} q_{(i_1, \ldots, i_{j-1})} \cdot p_{i_j}$.
- $q_{(\pi)} = \prod_{j=1}^{k} q_{i_j} = \prod_{j=1}^{k} (1 - p_{i_j})$.

Note that Observation 2 also holds when the coordinates i_j are by themselves impulsive substrategies, 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 zero. 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 "nonlower-bounded utility" that we shall need.

Definition 1. Given an impulsive strategy $\pi \subseteq [n]$ and a set of boxes $T \subseteq [n]$ such that $T \cap \pi = \emptyset$, we define

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

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

Observation 3. Let $\pi = (i_1, \dots, i_k) \subseteq [n]$ be an impulsive strategy. Then,

$$\begin{aligned}
\mathsf{u}_N(\pi) &= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot \mathsf{u}_N(i_j|\{i_\ell\}_{\ell\in[j-1]}) \\
&= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot (p_{i_j} \cdot v_{i_j} - c(i_j|\{r\}\cup\{i_\ell\}_{\ell\in[j-1]})). \\
\mathsf{u}_Y(\pi) &= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot \mathsf{u}_Y(i_j|\{i_\ell\}_{\ell\in[j-1]}) \\
&= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot (p_{i_j} \cdot (v_{i_j} - v_r)^+ - c(i_j|\{r\}\cup\{i_\ell\}_{\ell\in[j-1]})). \\
\mathsf{u}_M(\pi) &= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot \mathsf{u}_M(i_j|\{i_\ell\}_{\ell\in[j-1]}) \\
&= \sum_{j=1}^{k} q_{(i_1,\dots,i_{j-1})} \cdot (p_{i_j} \cdot (v_{i_j} - v_r) - c(i_j|\{r\}\cup\{i_\ell\}_{\ell\in[j-1]})).
\end{aligned}$$

The corresponding expressions $\mathsf{u}_N(\pi|T), \mathsf{u}_Y(\pi|T), \mathsf{u}_M(\pi|T)$ for a set $T \subseteq [n]$ such that $T \cap \pi = \emptyset$ 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

$$\mathsf{u}_N(\pi|T) = p_1v_1 - c(1|\{r,4,5\}) + q_1(p_2v_2 - c(2|\{r,4,5,1\})) + q_{(1,2)}(p_7v_7 - c(7|\{r,4,5,1,2\})).$$

Furthermore, observe that

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

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 $|\pi^Y| + |\pi^N|$.

The proof of the following claim is deferred to Appendix D.

Claim 3. Let $\pi \subseteq [n]$ be an impulsive strategy. Then, for any set of boxes $T \subseteq [n]$ such that $T \cap \pi = \emptyset$, we have
- $\mathsf{u}_M(\pi|T) \le \mathsf{u}_Y(\pi|T) \le \mathsf{u}_N(\pi|T)$.
- $\mathsf{u}_M(\pi|T) = \mathsf{u}_N(\pi|T) - p_{(\pi)} \cdot v_r$.
- If $\pi \subseteq \pi^Y$, then $\mathsf{u}_M(\pi|T) = \mathsf{u}_Y(\pi|T)$.

4.1.1. 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 1 carry over to impulsive strategies with dummies. For example, given the strategy $\pi = (2,1,4,7)$, a subset of boxes $P = \{1,7\}$, and another set of boxes $T = \{4,5\}$ that has presumably already been opened, we have

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

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 that these terms are multiplied by in the expression for the utility of π. Also, note that $T \cap \pi \neq \emptyset$ in this example, but we allow this because the strategy that 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 nonzero value—is also well defined (e.g., in the example above, this probability equals $q_2 p_1 + q_{(2,1,4)} p_7$).

Claim 3 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 $P = \emptyset$, we have that π_P is the empty strategy and that for $P = \pi$, we have that π_P coincides with π.

4.2. Proof of Lemma 2

4.2.1. Proof Overview. Our goal is to show that we can contradict one of the assumptions discussed in the former section. Our focus is on showing how to extend one of the substrategies π^Y, π^N with boxes from the other substrategy without harming the utility. This contradicts the assumption that $|\pi^Y| + |\pi^N|$ maximizes the sum of the length of the substrategies. To do so, we first show in Lemma 3 that there exists a box in π^N, which is not in π^Y, and that π^Y does not contain a deterministic box. In Lemma 4, we establish a technical tool for handling the utility function u_N. In Claim 4, we show that concatenating all the boxes from the strategy π^N (when ignoring boxes from π^Y) to π^Y does not improve utility with respect to utility function u_M. In Claim 5, we show that in substrategy π^N, the part contained in π^Y and the part not contained in π^Y are interleaved. In Claim 6, we use the tools that we developed to devise a linear combination (with nonnegative coefficients) over nonempty substrategies of π^N that can be concatenated to π^Y with nonnegative utility. Thus, at least one of them has a nonnegative utility, which contradicts the assumption of maximality of the optimal strategy.

The first step of the proof of Lemma 2 is the following inequality.

Lemma 3. *It holds that* $p_{(\pi^Y)} < p_{(\pi^N)}$.

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

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

where the equality and the second inequality hold by Claim 3, 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 $\mathsf{u}_N(\pi^Y) = \mathsf{u}_N(\pi^N)$. 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 $\pi'' = (\pi^Y, r)$. Because $v_i \geq v_r$ 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 than 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. Here and in the rest of the paper, we may use $\dot\cup$ instead of $\cup$ to underline that the union is performed between two disjoint sets.

Lemma 4. *Let* $\pi \subseteq [n]$ *be any impulsive strategy, and let* $\pi = A \dot\cup B$ *be a partition of the set of boxes corresponding to* π. *Then, we have* $\mathsf{u}_N(\pi) \leq \mathsf{u}_N(\pi_A|B) + \mathsf{u}_N(\pi_B)$.

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 $\mathbb{E}[\max_{i \in S(\pi)}(V_i)] = \mathbb{E}[\max_{i \in S(\pi_A)}(V_i)] + \mathbb{E}[\max_{i \in S(\pi_B)}(V_i)]$.

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 $\mathbb{E}[\max_{i \in S(\pi)}(V_i)]$ 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 $\mathbb{E}[\max_{i \in S(\pi_A)}(V_i)]$, and if $i \in B$, then this would appear in the expansion of $\mathbb{E}[\max_{i \in S(\pi_B)}(V_i)]$. In fact, this last discussion also shows Observation 4.

Observation 4. For any impulsive strategy $\pi \subseteq [n]$ and for any partition $\pi = A \dot\cup B$ 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:

$$\mathbb{E}[c(S(\pi_A)|B)] + \mathbb{E}[c(S(\pi_B))] - \mathbb{E}[c(S(\pi))] \le 0. \tag{1}$$

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 zero, and the inequality follows.

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

$$\mathbb{E}[c(S(\pi))] = \sum_{i=1}^{k} q_{(1,\dots,i-1)} \cdot c(i|\{1,\dots,i-1\})$$
$$\mathbb{E}[c(S(\pi_B))] = \sum_{i\in B} q_{(1,\dots,i-1)} \cdot c(i|\{1,\dots,i-1\} \cap B)$$
$$\mathbb{E}[c(S(\pi_A)|B)] = \sum_{i\in A} q_{(1,\dots,i-1)} \cdot c(i|\{1,\dots,i-1\} \cup B)\,.$$

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

$$\sum_{i\in A} q_{(1,\dots,i-1)} \cdot [c(i|\{1,\dots,i-1\} \cup B) - c(i|\{1,\dots,i-1\})] \tag{2}$$
$$+\sum_{i\in B} q_{(1,\dots,i-1)} \cdot [c(i|\{1,\dots,i-1\} \cap B) - c(i|\{1,\dots,i-1\})] \tag{3}$$
$$\le 0.$$

Example 3. It is instructive at this point to see an example. Assume that $\pi = (1,2,3,4,5)$, $A = \{1,2,4\}$, $B = \{3,5\}$. Then, we have

$$\mathbb{E}[c(S(\pi))] = c(1) + q_{(1)}c(2|\{1\}) + q_{(1,2)}c(3|\{1,2\}) + q_{(1,2,3)}c(4|\{1,2,3\}) + q_{(1,2,3,4)}c(5|\{1,2,3,4\}),$$

whereas we have

$$\mathbb{E}[c(S(\pi_B))] = q_{(1,2)}c(3) + q_{(1,2,3,4)}c(5|\{3\}) \text{ and}$$
$$\mathbb{E}[c(S(\pi_A)|B)] = c(1|\{3,5\}) + q_{(1)}c(2|\{1,3,5\}) + q_{(1,2,3)}c(4|\{1,2,3,5\}).$$

Thus, Inequality (1) for this example becomes

$$[c(1|\{3,5\}) - c(1)] + q_{(1)}[c(2|\{1,3,5\} - c(2|\{1\})] + q_{(1,2)}[c(3) - c(3|\{1,2\}]$$
$$+q_{(1,2,3)}[c(4|\{1,2,3,5\}) - c(4|\{1,2,3\}] + q_{(1,2,3,4)}[c(5|\{3\}) - c(5|\{1,2,3,4\})]$$
$$\le 0.$$

We now split into 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 $\{1,\dots,k-1\} \cup B = \{1,\dots,k-1\}$, and thus, the summand in line (2) corresponding to $i = k$ cancels out and equals zero. Therefore, if we denote $\pi^{(k)} := (1,\dots,k-1)$, then Inequality (1) is equivalent to

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

which holds by the induction hypothesis.

We now handle the case $k \in B$. Note that if $A = \emptyset$ 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 $|A| \ge 1$.

The rest of the proof involves a systematic manipulation of the inequality, where we repeatedly derive an expression that upper bounds the previously obtained expression (starting from the left-hand side of Inequality (1)). Mostly, we shall make repeated use of the following observation, which we term the "cancellation lemma." In the online version, 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 5 (Cancellation Lemma). *For every cost function $c : 2^{[n]} \to \mathbb{R}_{\ge 0}$, subset $T \subseteq [n]$, and elements $h, \ell \in [n] \setminus T$, we have*

$$c(h|T \cup \{\ell\}) - c(\ell|T \cup \{h\}) = c(h|T) - c(\ell|T).$$

Proof. The lemma holds because

$$\begin{aligned} &c(h|T \cup \{\ell\}) - c(\ell|T \cup \{h\}) \\ &\quad = [c(\{h\} \cup T \cup \{\ell\}) - c(T \cup \{\ell\})] - [c(\{\ell\} \cup T \cup \{h\}) - c(T \cup \{h\})] \\ &\quad = c(T \cup \{h\}) - c(T \cup \{\ell\}) \\ &\quad = [c(T \cup \{h\}) - c(T)] - [c(T \cup \{\ell\}) - c(T)] \\ &\quad = c(h|T) - c(\ell|T). \quad \square \end{aligned}$$

Denote $A = \{a_1, \dots, a_w\}$, where $w = |A| \geq 1$ and where the order $a_1, \dots, a_w$ is consistent with the relative ordering of A in π (i.e., $a_1 < \cdots < a_w$ (recall that we denoted $\pi = (1, \dots, k)$)). In the context of the example above, we get $w = 3$ and $a_1 = 1$, $a_2 = 2$, $a_3 = 4$. Using this notation, we can rewrite line (2) as follows:

$$\sum_{i=1}^{w} q_{(1,\dots,a_i-1)} \cdot [c(a_i|\{a_1, \dots, a_{i-1}\} \dot\cup B) - c(a_i|\{1, \dots, a_i - 1\})]. \tag{4}$$

For each $i = 1, \dots, w$, 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

$$q_{(1,\dots,k-1)} \cdot [c(k|B \setminus \{k\}) - c(k|(B \setminus \{k\}) \dot\cup \{a_1, \dots, a_{w-1}, a_w\})].$$

Additionally, 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:

$$q_{(1,\dots,a_w-1)} \cdot [c(a_w|\{a_1, \dots, a_{w-1}\} \dot\cup (B \setminus \{k\}) \dot\cup \{k\}) - c(a_w|\{1, \dots, a_w - 1\})].$$

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 brackets in the k-summand by

$$D := [c(k|B \setminus \{k\}) - c(k|(B \setminus \{k\}) \dot\cup \{a_1, \dots, a_{w-1}, a_w\})]$$

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

$$\begin{aligned} q_{(1,\dots,k-1)} \cdot D &\leq q_{(1,\dots,a_w-1)} \cdot D \\ &= q_{(1,\dots,a_w-1)} \cdot [c(k|B \setminus \{k\}) - c(k|(B \setminus \{k\}) \dot\cup \{a_1, \dots, a_{w-1}, a_w\})] . \end{aligned}$$

Now, the "q" factors in the modified k-summand and the a_w-summand are the same, and we apply the cancellation lemma on the induced difference

$$q_{(1,\dots,a_w-1)} \cdot [c(a_w|\{a_1, \dots, a_{w-1}\} \dot\cup (B \setminus \{k\}) \dot\cup \{k\}) - c(k|(B \setminus \{k\}) \dot\cup \{a_1, \dots, a_w\})],$$

where in the cancellation lemma statement, we plugged in $h = a_w$, $\ell = k$ and $T = (B \setminus \{k\}) \dot\cup \{a_1, \dots, a_{w-1}\}$. We remove the k and a_w terms from the unions in the above difference accordingly. In turn, the a_w-summand becomes

$$q_{(1,\dots,a_w-1)} \cdot [c(a_w|\{a_1, \dots, a_{w-1}\} \dot\cup (B \setminus \{k\})) - c(a_w|\{1, \dots, a_w - 1\})],$$

and the k-summand becomes

$$q_{(1,\dots,a_w-1)} \cdot [c(k|B \setminus \{k\}) - c(k|(B \setminus \{k\}) \dot\cup \{a_1, \dots, a_{w-2}, a_{w-1}\})].$$

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

$$q_{(1,\dots,a_{w-1}-1)} \cdot [c(a_{w-1}|\{a_1, \dots, a_{w-2}\} \dot\cup (B \setminus \{k\}) \dot\cup \{k\}) - c(a_{w-1}|\{1, \dots, a_{w-1} - 1\})].$$

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

After the application of the cancellation lemma, we remove the $\{k\}$ term from the union in the a_{w-1}-summand, which now becomes

$$q_{(1,\dots,a_{w-1}-1)} \cdot [c(a_{w-1}|\{a_1, \dots, a_{w-2}\} \dot\cup (B \setminus \{k\})) - c(a_{w-1}|\{1, \dots, a_{w-1} - 1\})].$$

We also remove the a_{w-1} term from the union in the k-summand, which becomes

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

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

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

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 explicitly shown in the above process that

$$\begin{aligned}&\mathbb{E}[c(S(\pi_A)|B)] + \mathbb{E}[c(S(\pi_B))] - \mathbb{E}[c(S(\pi))]\\ &\le \mathbb{E}[c(S(\pi_A^{(k)})|B\setminus\{k\})] + \mathbb{E}[c(S(\pi_{B\setminus\{k\}}^{(k)}))] - \mathbb{E}[c(S(\pi^{(k)}))],\end{aligned}$$

and the bottom expression is upper bounded by zero by the induction hypothesis. This concludes the proof of Lemma 4. □

The remainder of the proof of Lemma 2 proceeds as follows. By Lemma 3, we have $\pi^N \setminus \pi^Y \neq \emptyset$ because otherwise, $\pi^N \subseteq \pi^Y$, which implies $p_{(\pi^N)} \le p_{(\pi^Y)}$. Lemma 3 also implies that $p_{(\pi^Y)} < 1$ (i.e., π^Y does not contain a deterministic box). To prove Lemma 2, we show that there exists a nonempty impulsive substrategy 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 $\pi^N = A\dot{\cup}B$ and that $A \neq \emptyset$. We can write π^N as a concatenation of contiguous substrategies made up of boxes from A or B as follows: $\pi^N = (B^{\text{pre}}, A^1, B^1, \ldots, A^k, B^k, A^{\text{suff}})$, where $A = (\dot{\cup}_{i=1}^k A^k)\dot{\cup}A^{\text{suff}}, B = B^{\text{pre}}\dot{\cup}(\dot{\cup}_{i=1}^k B^k)$. The only substrategies 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 $k \ge 1$ as otherwise, $A = \emptyset$, and we get a contradiction).

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

Claim 4. We have $\mathsf{u}_M(\pi^A|\pi^Y) < 0$.

Proof. Assume toward contradiction that $\mathsf{u}_M(\pi^A|\pi^Y) \ge 0$. Then, by Claim 3, we also have $\mathsf{u}_Y(\pi^A|\pi^Y) \ge 0$. 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 $|\pi^Y| + |\pi^N|$. Formally, consider the strategy obtained from π^* by replacing the strategy π^Y with (π^Y, π^A). Then, the utility obtained does not decrease because

$$\mathsf{u}_Y((\pi^Y, \pi^A)) = \mathsf{u}_Y(\pi^Y) + q_{(\pi^Y)}\mathsf{u}_Y(\pi^A|\pi^Y) \ge \mathsf{u}_Y(\pi^Y).$$

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

Observe that

$$\mathsf{u}_N[(\pi^Y, \pi^A)] \le \mathsf{u}_N(\pi^N) \le \mathsf{u}_N(\pi_A^N|B) + \mathsf{u}_N(\pi_B^N) \le \mathsf{u}_N(\pi_A^N|\pi^Y) + \mathsf{u}_N(\pi_B^N),$$

where the first inequality holds by the optimality of π^N for the scenario where $V_r = 0$, the second inequality holds by Lemma 4, and the third inequality holds by submodularity of the cost function c because $B \subseteq \pi^Y$. Now, because the strategy (π^Y, π^A) is a superset of π^N, then in particular, we have

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

where the second equality holds by Observation 4. Therefore, the chain of inequalities above implies

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

where the first and last equalities hold by Claim 3. Because $\mathsf{u}_M[(\pi^Y, \pi^A)] = \mathsf{u}_M(\pi^Y) + q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y)$, then the above inequality implies

$$\mathsf{u}_M(\pi^Y) - \mathsf{u}_M(\pi^N_B) \le \mathsf{u}_M(\pi^N_A|\pi^Y) - q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y). \tag{5}$$

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 is deferred to Appendix D.

Claim 5. 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 (nonempty) strategy $(A^1, \dots, A^j)$ to π^Y without losing utility. This would constitute a contradiction to the assumption that π^* maximizes $|\pi^Y| + |\pi^N|$.

Claim 6. We have

$$\sum_{j=1}^{k} q_{(B^{\mathsf{pre}}, B^1, \dots, B^{j-1})} p_{(B^j)} \mathsf{u}_M((A^1, \dots, A^j)|\pi^Y) \ge 0.$$

Proof. Consider Inequality (5). The expression on the left-hand side satisfies

$$0 \le \mathsf{u}_Y(\pi^Y) - \mathsf{u}_Y(\pi^N_B) = \mathsf{u}_M(\pi^Y) - \mathsf{u}_M(\pi^N_B), \tag{6}$$

where the equality holds by Claim 3 (recall that $B \subseteq \pi^Y$) and the inequality holds because π^Y is the optimal substrategy for the scenario where $V_r = v_r$.

The combination of Inequalities (5) and (6) implies that

$$\mathsf{u}_M(\pi^N_A|\pi^Y) - q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y) \ge \mathsf{u}_M(\pi^Y) - \mathsf{u}_M(\pi^N_B) \ge 0. \tag{7}$$

On the other hand, we have

$$\begin{aligned}
&\mathsf{u}_M(\pi^N_A|\pi^Y) - q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y)\\
&\le \mathsf{u}_M(\pi^N_A|\pi^Y) - q_{(\pi^B)}\mathsf{u}_M(\pi^A|\pi^Y)\\
&= \mathsf{u}_M[(B^{\mathsf{pre}}, A^1, B^1, \dots, A^k, B^k, A^{\mathsf{suff}})_A|\pi^Y]\\
&- q_{(B^{\mathsf{pre}}, B^1, \dots, B^k)}\mathsf{u}_M[(A^1, \dots, A^k, A^{\mathsf{suff}}|\pi^Y)]\\
&= \left[\sum_{i=1}^{k} q_{(B^{\mathsf{pre}}, A^1, B^1, \dots, A^{i-1}, B^{i-1})}\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\right.\\
&\qquad \left. + q_{(B^{\mathsf{pre}}, A^1, B^1, \dots, A^k, B^k)}\mathsf{u}_M(A^{\mathsf{suff}}|\pi^Y \cup A \setminus A^{\mathsf{suff}})\right]\\
&\qquad - q_{(B^{\mathsf{pre}}, B^1, \dots, B^k)}\left[\sum_{i=1}^{k} q_{(A^1, \dots, A^{i-1})}\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\right.\\
&\qquad \left. + q_{(A^1, \dots, A^k)}\mathsf{u}_M(A^{\mathsf{suff}}|\pi^Y \cup A \setminus A^{\mathsf{suff}})\right],
\end{aligned}$$

where the first inequality holds because $q_{(\pi^B)} \ge q_{(\pi^Y)}$ (because $\pi^B \subseteq \pi^Y$) and because by Claim 4, $\mathsf{u}_M(\pi^A|\pi^Y) < 0$. Now, for every i, we have

$$q_{(B^{\mathsf{pre}}, A^1, B^1, \dots, A^{i-1}, B^{i-1})} = q_{(A^1, \dots, A^{i-1}, B^{\mathsf{pre}}, B^1, \dots, B^{i-1})} = q_{(A^1, \dots, A^{i-1})} \cdot q_{(B^{\mathsf{pre}}, B^1, \dots, B^{i-1})}.$$

Thus, we can factor out $q_{(A^1,\ldots,A^{i-1})}$ in the chain above and cancel out the "A^{suff}" term, and by also plugging Inequality (7), we get

$$
\begin{aligned}
0 &\le \sum_{i=1}^{k} q_{(A^1,\ldots,A^{i-1})}\cdot(q_{(B^{\mathsf{pre}},B^1,\ldots,B^{i-1})} - q_{(B^{\mathsf{pre}},B^1,\ldots,B^k)})\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{i=1}^{k} q_{(A^1,\ldots,A^{i-1})}\cdot q_{(B^{\mathsf{pre}},B^1,\ldots,B^{i-1})}(1-q_{(B^i,\ldots,B^k)})\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{i=1}^{k} q_{(A^1,\ldots,A^{i-1})}\cdot q_{(B^{\mathsf{pre}},B^1,\ldots,B^{i-1})}p_{(B^i,\ldots,B^k)}\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{i=1}^{k} q_{(A^1,\ldots,A^{i-1})}\cdot q_{(B^{\mathsf{pre}},B^1,\ldots,B^{i-1})}\left(\sum_{j=i}^{k} q_{(B^i,\ldots,B^{j-1})}p_{(B^j)}\right)\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{j=1}^{k}\sum_{i=1}^{j} q_{(A^1,\ldots,A^{i-1})}\cdot q_{(B^{\mathsf{pre}},B^1,\ldots,B^{i-1})}q_{(B^i,\ldots,B^{j-1})}p_{(B^j)}\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{j=1}^{k} q_{(B^{\mathsf{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}\sum_{i=1}^{j} q_{(A^1,\ldots,A^{i-1})}fb\mathsf{u}_M(A^i|\pi^Y \cup \{A^j\}_{j\in[i-1]})\\
&= \sum_{j=1}^{k} q_{(B^{\mathsf{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}\mathsf{u}_M((A^1,\ldots,A^j)|\pi^Y).
\end{aligned}
$$

This concludes the proof. □

Note that all of the factors $q_{(B^{\mathsf{pre}},B^1,\ldots,B^{j-1})}p_{(B^j)}$ are strictly positive because 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 $\mathsf{u}_M((A^1,\ldots,A^j)|\pi^Y)$, for $j\in[k]$, is nonnegative. Choose some j that satisfies this (such a j exists because $k\ge 1$ by Claim 5). To conclude the proof, we would like to say that we can concatenate $(A^1,\ldots,A^j)$ to π^Y without decreasing the utility, thus obtaining the desired contradiction, analogously to what we did in the proof of Claim 4. The (small) problem is that there might be boxes $i\in(A^1,\ldots,A^j)$ for which $v_i < v_r$.

To get around this, we note that it cannot be the case that all boxes $i\in(A^1,\ldots,A^j)$ satisfy $v_i < v_r$ because the contribution of these boxes to $\mathsf{u}_M((A^1,\ldots,A^j)|\pi^Y)$ is strictly negative. Consider then the impulsive strategy with dummies $\pi^{\mathsf{rand}} = (A^1,\ldots,A^j)_{\{i|v_i\ge v_r\}\cap(A^1,\ldots,A^j)}$, which is obtained from $(A^1,\ldots,A^j)$ 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 $\mathsf{u}_M(\pi^{\mathsf{rand}}|\pi^Y)\ge 0$. Furthermore, by the observation above, this strategy has nonempty 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 $\mathsf{u}_M(\pi^{\mathsf{diff}}|\pi^Y)\ge 0$ and which satisfies $v_i\ge v_r$ for every $i\in\pi^{\mathsf{diff}}$. This, in turn, implies $\mathsf{u}_Y(\pi^{\mathsf{diff}}|\pi^Y)\ge 0$ by Claim 3. We now concatenate π^{diff} to π^Y without decreasing the utility, analogously to what we did in the proof of Claim 4, and we get a contradiction to the definition of π^*. This concludes the proof of Lemma 2.

5. Reduction to Bernoulli Instances

In this section, we show how Theorem 4 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 4.

5.1. 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 $\kappa_\epsilon := \min\{\kappa\ge 0 | \sum_{i=1}^{n}\mathbb{E}[(V_i-\kappa)^+]\le\epsilon\}$. Such a constant κ_ϵ is well defined for every $\epsilon>0$ because $\sum_{i=1}^{n}\mathbb{E}[(V_i-\kappa)^+]$ is a monotone continuous decreasing function in κ, and the limit as κ approaches infinity is zero.[7] 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 because κ_ϵ is finite, then the support of the new set of distributions is finite. Finally, the output of $\mathcal{T}^\epsilon$ is $\mathcal{T}^\epsilon(D_1, \ldots, D_n, c) = (D_1^\epsilon, \ldots, D_n^\epsilon, c)$.

Lemma 6. *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 $\mathsf{u}(\mathcal{I};\pi) \le \mathsf{u}(\mathcal{I}^\epsilon;\pi') + 2\epsilon$.*
2. *For every strategy π' on instance $\mathcal{I}^\epsilon$, there exists a strategy π on instance $\mathcal{I}$ such that $\mathsf{u}(\mathcal{I};\pi) \ge \mathsf{u}(\mathcal{I}^\epsilon;\pi')$. Furthermore, if π' is a fixed-order strategy, then there exists such a fixed-order strategy π.*

Proof. The two instances $\mathcal{I}$ and $\mathcal{I}^\epsilon$ are on the same boxes; the only difference is that the random variables V_i of $\mathcal{I}$ are discretized as $\bar{V}_i$ in $\mathcal{I}^\epsilon$. To prove the first part of the proposition, we show how to construct a strategy π' using π. When π prescribes to open box i, strategy π' does it and observes a realization $\bar{V}_i = \bar{v}_i$. Then, one can draw V_i, according to the distribution D_i conditioned on the event that $\epsilon \cdot \lfloor \frac{\min(V_i, \kappa_\epsilon)}{\epsilon} \rfloor = \bar{v}_i$ and keep playing as if π saw the realization of V_i. It is clear that if π' would have received the value of V_i (instead of $\bar{V}_i$), then π' would have had the same performance as π. As this is not always the case, we bound the difference of the two utilities partitioning the analysis in two cases. If the chosen V_i is at most κ_ϵ, then $V_i - \bar{V}_i \le \epsilon$. Otherwise, by the choice of κ_ϵ, if we do not count V_i in this event, we lose at most an additional ϵ. Formally, if we denote by i^* the box with the maximal simulated value observed by π', then

$$
\begin{aligned}
\mathsf{u}(\mathcal{I}^\epsilon;\pi') &= \mathsf{u}(\mathcal{I};\pi) - \mathbb{E}\left[V_{i^*} - \epsilon \cdot \left\lfloor \frac{\min(V_{i^*}, \kappa_\epsilon)}{\epsilon} \right\rfloor\right] \\
&= \mathsf{u}(\mathcal{I};\pi) - \mathbb{E}\left[\left(V_{i^*} - \epsilon \cdot \left\lfloor \frac{V_{i^*}}{\epsilon} \right\rfloor\right) \cdot \mathbb{I}_{\{V_{i^*} \le \kappa_\epsilon\}}\right] - \mathbb{E}\left[\left(V_{i^*} - \epsilon \cdot \left\lfloor \frac{\kappa_\epsilon}{\epsilon} \right\rfloor\right) \cdot \mathbb{I}_{\{V_{i^*} > \kappa_\epsilon\}}\right] \\
&\ge \mathsf{u}(\mathcal{I};\pi) - \mathbb{E}[\epsilon \cdot \mathbb{I}_{\{V_{i^*} \le \kappa_\epsilon\}}] - \mathbb{E}[(V_{i^*} - \kappa_\epsilon) \cdot \mathbb{I}_{\{V_{i^*} > \kappa_\epsilon\}}] - \mathbb{E}\left[\left(\kappa_\epsilon - \epsilon \cdot \left\lfloor \frac{\kappa_\epsilon}{\epsilon} \right\rfloor\right) \cdot \mathbb{I}_{\{V_{i^*} > \kappa_\epsilon\}}\right] \\
&\ge \mathsf{u}(\mathcal{I};\pi) - \mathbb{E}[\epsilon \cdot \mathbb{I}_{\{V_{i^*} \le \kappa_\epsilon\}}] - \left[\sum_i \mathbb{E}[(V_i - \kappa_\epsilon) \cdot \mathbb{I}_{\{V_i > \kappa_\epsilon\}}]\right] - \mathbb{E}[\epsilon \cdot \mathbb{I}_{\{V_{i^*} > \kappa_\epsilon\}}] \\
&\ge \mathsf{u}(\mathcal{I};\pi) - 2\epsilon,
\end{aligned}
$$

where the first two inequalities follow from the fact that for every x, it holds that $x - \epsilon \lfloor \frac{x}{\epsilon} \rfloor \le \epsilon$ and the last inequality is by definition of κ_ϵ.

The second part of the proposition follows by the simple observation that one can run the following strategy π. First, calculate κ_ϵ. Upon the arrival of $V_i = v_i$, calculate $\bar{v}_i = \epsilon \cdot \lfloor \frac{\min(v_i, \kappa_\epsilon)}{\epsilon} \rfloor$, then run strategy π' as if observed the value $\bar{v}_i$. As the actual value of v_i is always at least the value $\bar{v}_i$, it holds that $\mathsf{u}(\mathcal{I};\pi) \ge \mathsf{u}(\mathcal{I}^\epsilon;\pi')$. The "furthermore" part follows immediately by the structure of strategy π. □

5.2. 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 zero is in the union of the supports supp; then, we can rename the elements of the union of the supports in an increasing order $\mathsf{supp} = \{v_1, \ldots, v_m\}$, where $0 = v_1 < v_2 < \cdots < v_m$. 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 $\frac{\mathbb{P}(V_i = v_j)}{\mathbb{P}(V_i \le v_j)}$ and zero otherwise (where $\frac{0}{0}$ is interpreted as zero). Let $c' : 2^{[n]\times[m]} \to \mathbb{R}_{\ge 0}$ 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.

Lemma 7. *If c is submodular, then c' obtained by transformation $\mathcal{T}$ is also submodular.*

Proof. We prove that for any pair of sets S, T such that $S \subseteq T \subseteq [n] \times [m]$ and any pair $(i,j) \notin T$, it holds that $c'(\{(i,j)\} \cup S) - c'(S) \ge c'(\{(i,j)\} \cup T) - c'(T)$. Let $A_S := \{i' \mid \exists j' \in [m] \text{ such that } (i',j') \in S\}$, $A_T := \{i' \mid \exists j' \in [m] \text{ such that } (i',j') \in T\}$, and note that $A_S \subseteq A_T$. If $i \in A_S$, then it holds that $c'(\{(i,j)\} \cup S) = c'(S)$ and $c'(\{(i,j)\} \cup T) = c'(T)$; thus, $c'(\{(i,j)\} \cup S) - c'(S) = 0 = c'(\{(i,j)\} \cup T) - c'(T)$. Else, it holds that $c'(\{(i,j)\} \cup S) - c'(S) = c(A_S \cup \{i\}) - c(A_s) \ge c(A_T \cup \{i\}) - c(A_T) = c'(\{(i,j)\} \cup T) - c'(T)$, where the inequality follows by submodularity of c. □

In Appendix E, we show that $\mathcal{T}$ also maintains other important family of cost functions, namely matroid rank functions (Claim E.1), gross-substitutes functions (Claim E.2), coverage functions (Claim E.3), XOS (Claim E.4), and subadditive (Claim E.5) cost function but not budget additive (Claim E.6). For formal definitions of such classes, we refer to Appendix A.

We next show that the new instance $\mathcal{I}' = \mathcal{T}(\mathcal{I})$ is equivalent to $\mathcal{I}$ in the following sense.

Lemma 8. *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 $\mathsf{u}(\mathcal{I};\pi) \le \mathsf{u}(\mathcal{I}';\pi')$.*
2. *For every strategy π' on instance $\mathcal{I}'$, there exists a strategy π on instance $\mathcal{I}$ such that $\mathsf{u}(\mathcal{I};\pi) \ge \mathsf{u}(\mathcal{I}';\pi')$. Furthermore, if π' is impulsive, then there exists such a fixed-order strategy π.*

Proof. For the first part of the proposition, consider the following π' that simulates π; every time π opens a box i, the strategy π' opens the corresponding set of boxes $\{(i,j)\}_j$ that are created by the transformation (in arbitrary order). Let $\bar{u}_{i,j}$ be the realized value from box (i,j); then, π' behaves as if π observed the value $u_i = \max_j \bar{u}_{i,j}$. The distribution of u_i is exactly D_i because

$$\mathbb{P}(u_i = v_j) = \frac{\mathbb{P}(V_i = v_j)}{\mathbb{P}(V_i \le v_j)} \cdot \prod_{k>j}\left(1 - \frac{\mathbb{P}(V_i = v_k)}{\mathbb{P}(V_i \le v_k)}\right) = \mathbb{P}(V_i = v_j) \cdot \frac{\prod_{k>j}\mathbb{P}(V_i < v_k)}{\prod_{k\ge j}\mathbb{P}(V_i \le v_k)}$$
$$= \mathbb{P}(V_i = v_j) \cdot \frac{\prod_{k>j}\mathbb{P}(V_i < v_k)}{\mathbb{P}(V_i \le v_m) \cdot \prod_{k>j}\mathbb{P}(V_i < v_k)} = \mathbb{P}(V_i = v_j),$$

where the third equality is because $\mathbb{P}(V_i \le v_k) = \mathbb{P}(V_i < v_{k+1})$. Thus, the strategy guarantees the same expected utility (as both the distributions of the costs and the values are the same in π and π'). For the second part of the proposition, given a strategy π', consider the strategy π that simulates π' by the following process. Whenever π' tries to open $D_{i,j}$, if D_i was not already open, then open D_i (otherwise, do not open anything). If the value v_k was observed from D_i and $k = j$, then π behaves as if the value v_j was observed from $D_{i,j}$. If $k < j$, then π behaves as if the value of zero was observed from $D_{i,j}$. Otherwise, $(k > j)$; then, π draws a sample $s_{i,j}$ from $D_{i,j}$ and behaves as if this value was observed. The probability overall that π simulates that $D_{i,j}$ was nonzero is

$$\mathbb{P}(V_i = v_j) + \mathbb{P}(V_i > v_j) \cdot \frac{\mathbb{P}(V_i = v_j)}{\mathbb{P}(V_i \le v_j)} = \frac{\mathbb{P}(V_i = v_j)}{\mathbb{P}(V_i \le v_j)}.$$

The cost of π is always the same as that of π', but its value can only be larger (because π never pretends to see a larger value than what it actually observed). The "furthermore" part follows by observing that box i is opened when the first box of the form (i,j) is supposed to be opened by π'. Thus, box i is opened before box i' if and only if the first copy of i is opened before the first copy of i' in the instance $\mathcal{I}'$. □

In Section 4, 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 2. *For every instance $\mathcal{I} = (D_1, \dots, D_n, c)$, 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 an optimal strategy for $\mathcal{I}$ among the strategies with a fixed order. Assume toward 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 Lemma 6, there exists π_1 such that $\mathsf{u}(\mathcal{I};\pi^*) \le \mathsf{u}(\mathcal{I}^\epsilon;\pi_1) + 2\epsilon$. Let $\mathcal{I}' = \mathcal{T}(\mathcal{I}^\epsilon)$. Then, by Lemma 8, there exists π_2 such that $\mathsf{u}(\mathcal{I}^\epsilon;\pi_1) \le \mathsf{u}(\mathcal{I}';\pi_2)$. By Theorem 4, there exists an impulsive strategy π_3 such that $\mathsf{u}(\mathcal{I}';\pi_2) \le \mathsf{u}(\mathcal{I}';\pi_3)$. By Lemma 8, there exists a fixed order π_4 such that $\mathsf{u}(\mathcal{I}^\epsilon;\pi_4) \ge \mathsf{u}(\mathcal{I}';\pi_3)$, and by Lemma 6, there exists a fixed order π_5 such that $\mathsf{u}(\mathcal{I};\pi_5) \ge \mathsf{u}(\mathcal{I}^\epsilon;\pi_4)$. All together, we have

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

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

6. Computational Results

In this section, we consider Pandora's *decision problem*; given an instance $\mathcal{I} = (D_1, \dots, D_n, c)$, decide whether there exists a strategy π for $\mathcal{I}$ that achieves positive utility (i.e., $\mathsf{u}(\mathcal{I};\pi) > 0$). 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 $c(S)$.

We show that the decision problem cannot be solved with a polynomial number (in n) of queries, even by a randomized algorithm.

Theorem 3. *Let $\varepsilon > 0$, and let $\mathcal{A}$ be any (possibly randomized) algorithm for Pandora's decision problem that uses at most poly-(n) queries to the given submodular cost function. Then, there exists an instance $\mathcal{I}$ such that $\mathcal{A}$ outputs the correct answer on $\mathcal{I}$ with probability at most $0.5 + \varepsilon$.*[8]

The rest of the section is organized as follows. In Section 6.1, we construct a baseline cost function, denoted c_0, and a family of other cost functions such that c_0 cannot be distinguished from a random cost function from the family by using only polynomially many cost queries. Furthermore, the construction is designed such that c_0 is weakly greater (point wise) than any function in the family.

In Section 6.2, we construct distributions $D_1, \ldots, D_n$, which together with the cost functions from Section 6.1, induce a baseline instance, denoted $\mathcal{I}_0$, and a family of other instances to Pandora's problem. We prove that no strategy for $\mathcal{I}_0$ attains strictly positive utility, whereas for any instance in the other family, there exists a strategy that attains strictly positive utility.

Finally, in Section 6.3, we use these instances to prove Theorem 3.

6.1. A Family of Indistinguishable Cost Functions

To formalize our argument, we use the notion of distinguishability of submodular functions (Svitkina and Fleischer [28]). We say that an algorithm with oracle access to a set function distinguishes between two cost functions c_1 and c_2 if it produces different outputs when the oracle function is c_1 versus when the oracle function is c_2.

We construct a family of cost functions and a baseline cost function that with high probability, are indistinguishable using polynomially many cost queries, similarly to the construction of Svitkina and Fleischer [28].

Let X be a set of n boxes, and let $\alpha = \lceil \ln n \cdot \frac{\sqrt{n}}{5} \rceil$ and $\beta = \lceil \frac{\ln^2 n}{5} \rceil$. We introduce the baseline cost function $c_0 : 2^X \to \mathbb{R}_{\geq 0}$, which maps a set $S \subseteq X$ to $c_0(S) = \min\{|S|, \alpha\}$. Then, for any subset $R \subseteq X$ of boxes with $|R| = \alpha$, we define the cost function $c_R : 2^X \to \mathbb{R}_{\geq 0}$ on a set S as follows:

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

where R^C is the complementary set of R in X. 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.

Lemma 9. *Let A be any deterministic algorithm that has oracle access to a set function and that uses at most poly-(n) queries, where n is the size of the element set. Then, for any sufficiently large n, we have*

$$\mathbb{P}(\mathsf{A} \text{ distinguishes } c_{\mathcal{R}} \text{ from } c_0) \leq n^{-1},$$

where the probability is taken over the randomness of $\mathcal{R}$.

Proof. First, we show that for any deterministic set S, the event that $c_0(S) \neq c_R(S)$ is negligible, with respect to a random draw of $R \sim \mathcal{R}$. We observe that the probability of this event is maximized for sets S of cardinality α, so it suffices to restrict attention to such sets. To see this, we consider two cases. If $|S| \geq \alpha$, then $c_0(S) \neq c_R(S)$ if and only if $|S \cap R^C| < \alpha - \beta$, which is more likely to hold for small $|S|$. Conversely, if $|S| \leq \alpha$, then $c_0(S) \neq c_R(S)$ when $|S| > |S \cap R^C| + \beta$, which is equivalent to $|S \cap R| > \beta$; it is easy to see that the latter condition is more likely when $|S|$ is large. Moreover, for S of cardinality α, the function c_0 and the realized function c_R disagree if and only if $|S \cap R| > \beta$.

To simplify the calculations, consider a set R' that is obtained independently from R, sampling each element with probability α/n. We have

$$\begin{aligned}
\mathbb{P}(|S \cap R'| > \beta) &= \sum_{k=0}^{n} \mathbb{P}(|R'| = k)\mathbb{P}(|S \cap R'| > \beta \mid |R'| = k) \\
&\geq \mathbb{P}(|R'| = \alpha)\mathbb{P}(|S \cap R'| > \beta \mid |R'| = \alpha) \\
&\geq \frac{1}{n^2}\mathbb{P}(|S \cap R'| > \beta \mid |R'| = \alpha),
\end{aligned}$$

where the last inequality follows from the fact that $|R'|$ can attain $n + 1 < n^2$ different values and $|R'| = \alpha$ is the most likely of them. We can use this argument on R' to upper bound the probability of the event that $|S \cap R| > \beta$

in a simpler way:

$$\mathbb{P}(|S \cap R| > \beta) = \mathbb{P}(|S \cap R'| > \beta \| R'| = \alpha) \leq n^2 \cdot \mathbb{P}(|S \cap R'| > \beta).$$

We can now focus on upper bounding the term in the right-hand side by using the Chernoff bound. The expected cardinality of $S \cap R'$ (with respect to the random choice of R') is $\mu = \frac{\alpha \cdot |S|}{n} = \frac{\alpha^2}{n}$, whereas $\beta = 5\mu$. We thus have

$$\mathbb{P}(|S \cap R'| > \beta) < \left(\frac{e^\delta}{(1+\delta)^\delta}\right)^\mu = \left(\frac{e^4}{25}\right)^{\alpha^2/n} \leq 0.851^{\alpha^2/n}.$$

It follows that for all S such that $|S| = \alpha$, we have

$$\mathbb{P}(|S \cap R| > \beta) \leq n^2 \cdot 0.851^{\alpha^2/n}.$$

Consider now any deterministic algorithm A that performs at most polynomially many cost queries as in the statement of the lemma. Consider also the computation path it follows when the oracle function is the baseline cost function c_0. Both A and c_0 are deterministic; thus, this is a single computation path. Along this path, the algorithm performs at most a polynomial number of cost queries (say, at most n^t for some constant t), and each cost query distinguishes c_0 from c_R with probability at most $n^2 \cdot 0.851^{\alpha^2/n}$ (as we have shown in the first part of the proof); thus, by the union bound, we get

$$\mathbb{P}(\mathsf{A} \text{ distinguishes } c_0 \text{ from } c_{\mathcal{R}}) \leq n^{t+2} \cdot 0.851^{\alpha^2} \leq n^{t+2+\ln(0.851)\ln n} < \frac{1}{n},$$

where the last inequality holds for all $n > \frac{t+3}{-\ln(0.851)}$. The latter condition specifies what we mean by "sufficiently large n" in the statement of the theorem; for any fixed deterministic algorithm A that performs $O(n^t)$ queries, taking $n > \frac{t+3}{-\ln(0.851)}$ gives the desired claim for A. □

We remark that by the proof of Lemma 9, n^{-1} as a bound on the probability that 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 $n > \frac{a+2+b}{-\ln(0.851)}$.

6.2. A Corresponding Family of Instances

As we show next, the family of submodular cost functions introduced in the previous section 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.

Consider the i.i.d. weighted Bernoulli distributions $D_1, \ldots, D_n$, whose values are distributed as follows; the value of every box is $M = 5\beta > 0$ with probability $p = \frac{1}{\alpha}$ and zero otherwise.

For each $R \subseteq [n]$ with $|R| = \alpha$, we define the instance $\mathcal{I}_R = (D_1, \ldots, D_n, C_R)$ obtained from the above distributions and the cost function c_R (see 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 has nonzero marginal cost, whereas 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.

Lemma 10. *For any sufficiently large n, no strategy for $\mathcal{I}_0$ attains strictly positive utility, whereas for any R, there exists a strategy for $\mathcal{I}_R$ that attains strictly positive utility.*

Proof. We first establish the second part of the lemma. Consider any $\mathcal{I}_R$ and the (optimal) strategy π^R for $\mathcal{I}_R$, which knows the specific set R. The strategy π^R opens the boxes in R one after the other (in any order) and halts when the value M is realized for the first time and otherwise, when all boxes in R are exhausted.

The expected reward of π^R is computed as follows. The value M is achieved if at least one of the α Bernoulli boxes in R is realized, yielding an expected value of $M(1-(1-p)^\alpha)$. On the other hand, their total cost is at most β. We get

$$\mathsf{u}(\mathcal{I}_R; \pi^R) \geq M\left(1 - \left(1 - \frac{1}{\alpha}\right)^\alpha\right) - \beta \geq 5\beta\left(1 - \frac{1}{e}\right) - \beta > 0.$$

We next establish the first part of the lemma. First, there exists a deterministic strategy that is optimal for $\mathcal{I}_0$ (see Section 2); thus, we restrict attention to deterministic strategies. Second, because all boxes are symmetric, there exists an optimal strategy that is impulsive (see Section 4.1) (i.e., it commits to a subset S of the boxes and opens them sequentially in an arbitrary fixed order until M is realized (or until all boxes in S have been opened)). Note that $c_0(S)$ depends only on the cardinality of S, so all of the orderings are equivalent. Let π^S denote this strategy.

To conclude the proof, we show that for every set S, we have $\mathsf{u}(\mathcal{I}_0;\pi^S) \le 0$. We distinguish between four cases depending on the cardinality of S.

Case 1. $|S| \ge \alpha$. Such a policy opens all of the boxes in S (with expected reward $M(1-(1-p)^{|S|})$) but pays only for the first α of them. Using similar reasoning as above, we get

$$\begin{aligned}\mathsf{u}(\mathcal{I}_0;\pi^S) &= M(1-(1-p)^{|S|}) - p\cdot\sum_{i=1}^{\alpha} i(1-p)^{i-1} - \alpha(1-p)^{\alpha}\\ &\le M - \alpha\left(1-\frac{1}{\alpha}\right)^{\alpha} \le M - \frac{\alpha}{4} = 5\beta - \frac{\alpha}{4} < 0,\end{aligned}$$

where the last inequality follows from the fact that $\alpha > 20\beta$ (recall that $\alpha \in \Theta(\ln n\sqrt{n})$, whereas $\beta \in \Theta(\ln^2 n)$).

Case 2. $|S| \in \{21\beta,\dots,\alpha-1\}$. Note that for $|S| < \alpha$, the cost function c_0 is simply additive; thus,

$$\begin{aligned}\mathsf{u}(\mathcal{I}_0;\pi^S) &= M(1-(1-p)^{|S|}) - p\cdot\sum_{i=1}^{|S|} i\cdot(1-p)^{i-1} - |S|(1-p)^{|S|}\\ &\le M - |S|\left(1-\frac{1}{|S|}\right)^{|S|} \le M - \frac{|S|}{4} \le -\frac{\beta}{4} < 0.\end{aligned}$$

Case 3. $0 < |S| < 21\beta$. We have

$$\begin{aligned}\mathsf{u}(\mathcal{I}_0;\pi^S) &= M(1-(1-p)^{|S|}) - p\cdot\sum_{i=1}^{|S|} i\cdot(1-p)^{i-1} - |S|(1-p)^{|S|}\\ &\le M - (M+|S|)(1-p)^{|S|}\\ &\le M - (M+|S|)(1-p|S|)\\ &= |S|(pM + p|S| - 1)\\ &= p|S|\left(M + |S| - \frac{1}{p}\right)\\ &\le \frac{|S|}{\alpha}(26\beta - \alpha) < 0,\end{aligned}$$

where the third line uses the Bernoulli inequality and the last two transitions use the definitions of α and β, the fact that n is sufficiently large, and the condition of case 3 (i.e., $0 < |S| < 21\beta$).

Case 4. $|S| = 0$. This case corresponds to the strategy that does not do anything, whose utility is clearly zero. □

6.3. Proof of Theorem 3

In this section. we finally prove Theorem 3. Intuitively, because it is not possible to distinguish between c_R and the baseline c_0 with polynomially many queries, then it is also not possible to distinguish between $\mathcal{I}_R$ and $\mathcal{I}_0$.

Proof of Theorem 3. Let $\mathcal{A}$ be as in the theorem statement. $\mathcal{A}$, which is possibly randomized, is just a distribution over deterministic algorithms A, each of which queries the given cost functions at most *poly*-(n) times. For any $R \subseteq X$ of cardinality α and any deterministic algorithm A, we denote with $\mathcal{E}(\mathsf{A}, R)$ the event that A returns a different output when given $\mathcal{I}_R$ versus $\mathcal{I}_0$ as input.

Recall that $\mathcal{R}$ is a set of cardinality α drawn uniformly at random. Denote with $\mathcal{R}and$ (Det, respectively) the set of all of the randomized (deterministic, respectively) efficient algorithms for Pandora's decision problem. Yao's principle gives the following:

$$\min_R \mathbb{P}(\mathcal{E}(\mathcal{A},R)) \le \min_R \max_{\mathcal{A}'\in\mathcal{R}and} \mathbb{P}(\mathcal{E}(\mathcal{A}',R)) \le \max_{\mathsf{A}\in\mathsf{Det}} \mathbb{P}(\mathcal{E}(\mathsf{A},\mathcal{R})). \tag{9}$$

Consider the rightmost term. Each deterministic algorithm A is an algorithm with oracle access to the underlying submodular cost function, which in the event $\mathcal{E}(\mathsf{A},\mathcal{R})$, gives different outputs on $\mathcal{I}_R$ and $\mathcal{I}_0$. In other words, in the event $\mathcal{E}(\mathsf{A},\mathcal{R})$, A distinguishes between $c_{\mathcal{R}}$ from c_0 given that the rest of the input is identical. From

Equation (9), we have then

$$\min_R \mathbb{P}(\mathcal{E}(\mathcal{A}, R)) \le \max_{\mathsf{A}\in \mathrm{Det}} \mathbb{P}(\mathcal{E}(\mathsf{A}, \mathcal{R})) \le \frac{1}{n} \le \varepsilon, \tag{10}$$

where the second inequality follows from Lemma 9 for any n sufficiently large.

What we have shown so far is that there exists a set R such that $\mathcal{A}$ 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 by Lemma 10. Now, let $\mathcal{G}_0$ ($\mathcal{G}_R$, respectively) be the event that $\mathcal{A}$ is correct on input $\mathcal{I}_0$ ($\mathcal{I}_R$, respectively). 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:

$$\begin{aligned}\mathbb{P}(\mathcal{G}_R) &= \mathbb{P}(\mathcal{G}_R \cap \mathcal{E}(\mathcal{A}, R)) + \mathbb{P}(\mathcal{G}_R \setminus \mathcal{E}(\mathcal{A}, R))\\ &\le \mathbb{P}(\mathcal{E}(\mathcal{A}, R)) + \mathbb{P}(\mathcal{G}_0^C) \le \varepsilon + (0.5-\varepsilon) = 0.5.\end{aligned}$$

To see why the previous formula holds, we study separately the two summands. The event $\mathcal{G}_R \cap \mathcal{E}(\mathcal{A}, R)$ is clearly contained in $\mathcal{E}(\mathcal{A}, R)$, and we know that its probability is smaller than ε by Equation (10). The event $\mathcal{G}_R \setminus \mathcal{E}(\mathcal{A}, R)$, on the other hand, is disjoint from $\mathcal{G}_0$; in fact, we know that if $\mathcal{A}$ 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 $\mathbb{P}(\mathcal{G}_0) \ge 0.5+\varepsilon$; thus, its complementary event has at most a probability $0.5-\varepsilon$ of being realized. □

The previous result directly implies that there does not exist an algorithm that uses polynomially many queries, which given a problem instance, outputs a strategy that gives any positive approximation to the expected utility of the optimal strategy. Assume by contradiction that such an algorithm exists; then, it would be easy to construct an algorithm for the decision problem that violates the previous theorem.

7. Conclusion and Future Directions

In this paper, we initiate the study of Pandora's problem with a combinatorial cost function and study to what extent the simplicity of Weitzman's solution extends beyond additive cost functions. We show that the structural simplicity carries over to submodular cost functions but not to XOS cost functions. Namely, Pandora's problem with submodular cost functions admits an optimal strategy that is fixed order. From a computational perspective, we prove that no polynomial-time approximation algorithm for Pandora's problem with submodular cost function exists, even for the subclass of matroid rank functions.

We remark that our hardness of approximation result extends to other variants of Pandora's problem. For example, consider the nonobligatory inspection setting, where the decision maker is allowed to select a box without opening it (and not paying its cost). One cannot get the equivalent of Theorem 3 in this setting because a strictly positive utility is always attainable (by selecting the box with the maximum expectation without opening it). However, one can show (by using the same family of instances described in Section 6.1) that achieving any constant approximation with a polynomial number of cost queries is impossible. This is true because for the baseline instance $\mathcal{I}_0$, the best strategy is to select one box without opening it; this achieves a utility of $\frac{M}{\alpha}$, which is asymptotically smaller than the utility attained by the best strategy from any other instance $\mathcal{I}_R$, which is $\Omega(\beta)$.

Another well-studied variant of Pandora's problem is the minimization variant, where the goal of the decision maker is to minimize the chosen observed value plus the total inspection cost. The minimization problem is equivalent to the maximization problem in terms of hardness of finding an optimal strategy because the minimization version is reducible to the maximization version by considering the value distributions with flipped signs and vice versa. However, when considering the hardness of finding approximately optimal strategies, the minimization version is intuitively easier because both parts of the objective function have the same sign.

Unfortunately, achieving a constant approximation is impossible as in the nonobligatory inspection model. This can be shown by adapting our example and replacing the distribution of the boxes to be α with probability $1-\frac{\beta}{\alpha}$ and zero otherwise. In the bad instance $\mathcal{I}_0$, no strategy pays a cost of less than $\Omega(\alpha)$, whereas in every good instance $\mathcal{I}_R$, the strategy that opens the cheap set of boxes R pays $O(\beta)$.

Our work suggests intriguing directions for future research. In particular, many of the variants of Pandora that have been studied under the classic model of additive cost functions can be studied under combinatorial cost functions. Obvious examples are settings beyond single choice and also, studying whether the structural simplicity shown in this paper extends to the nonobligatory inspection setting.

Furthermore, some computational problems remain open. For example, it is not clear whether there exists a polytime algorithm that finds an optimal (or an approximately optimal) strategy for various classes of cost functions, such as budget-additive and supermodular cost function.

Appendix A. Classes of Combinatorial Functions

In this appendix, we recall the definitions and some properties of matroid rank functions, gross-substitute functions, and coverage functions.

A.1. Matroids

A family of subsets $\mathcal{M} \subseteq 2^X$ of a base set X is called a matroid if the following two properties hold.

- *Downward closure*. If $A \in \mathcal{M}$ and $B \subseteq A$, then $B \in \mathcal{M}$.
- *Augmentation property*. If $A, B \in \mathcal{M}$ and $|A| > |B|$, then there exists $x \in A \setminus B$ such that $B \cup \{x\} \in \mathcal{M}$.

Subsets in $\mathcal{M}$ are called independent sets.

A.2. Matroid Rank Functions

Given any matroid $\mathcal{M}$, it is possible to define its associated rank function $r_{\mathcal{M}}: 2^X \to \mathbb{N}$ as the cardinality of the largest independent set:

$$r_{\mathcal{M}}(A) = \max_{B \subseteq A}\{|B| \,|\, B \in \mathcal{M}\}.$$

Recall the definitions of the two functions c_0 and c_R in Section 6. A base set V of n elements is given as well as two integers $\alpha > \beta$ smaller than n and a subset $R \subseteq V$, with $|R| = \alpha$; we have

$$c_0(S) = \min\{|S|, \alpha\}, \quad c_R(S) = \min\{|S|, \alpha, |S \cap R^C| + \beta\}.$$

Cost function c_0 is the rank function of the matroid $\mathcal{M}_0 = \{S \subseteq X \,|\, |S| \le \alpha\}$ (thus, it is also submodular). For what concerns c_R, instead of explicitly exhibiting the relative matroid, we show that it respects two properties that ensure that there exists a matroid $\mathcal{M}_R$ on X, of which c_R is indeed the rank function.

Theorem A.1 (Schrijver et al. [25, Theorem 39.8]). *Let V be a set, and let $r: 2^X \to \mathbb{N}$. Then, r is the rank function of a matroid if and only if for all $T, U \subseteq V$,*

i. $r(T) \le r(U) \le |U|$ *if* $T \subseteq U$ *and*
ii. $r(U \cap T) + r(U \cup T) \le r(U) + r(T)$.

It is immediate to verify that c_R respects condition (i), whereas condition (ii) is equivalent to submodularity, which holds for c_R.

A.3. Gross Substitutes

A function $f: 2^X \to \mathbb{R}_{\ge 0}$ is *gross substitutes* if for any two vectors $\boldsymbol{p}, \boldsymbol{p}' \in \mathbb{R}^n_{\ge 0}$ (with $|X| = n$) such that $\boldsymbol{p}' \ge \boldsymbol{p}$ (component wise) and any $S \subseteq X$ such that $S \in \arg\max_{T \subseteq X} f(T) - \sum_{i \in T} p_i$, there is a $S' \subseteq X$ such that $S' \in \arg\max_{T \subseteq X} f(T) - \sum_{i \in T} p'_i$ and $\{i \in S \,|\, p'_i = p_i\} \subseteq S'$. It is known that the class of gross-substitute functions strictly contains that of matroid rank functions (Balkanski and Leme [2]).

An alternative definition of gross-substitute functions is given by Reijnierse et al. [24] and is as follows; a valuation function f has the gross-substitutes property on set X if and only if it is submodular, and for all sets $S \subseteq X$ and all distinct $i, j, k \in X \setminus S$, the following multiset does not have a unique max:

$$\{f(\{i,j\}|S) + f(k|S), f(i|S) + f(j,k|S), f(j|S) + f(\{i,k\}|S)\}. \tag{A.1}$$

A.4. Coverage

A function $f: 2^X \to \mathbb{R}_{\ge 0}$ is a *coverage* if there exist a set of elements E, a mapping of the elements to weights $w: E \to \mathbb{R}_{\ge 0}$, and a mapping $g: E \to 2^X$ such that $f(S) = \sum_{j \in E} w(j) \cdot \mathbb{I}_{\{\exists i \in S \cap g(j)\}}$.

Appendix B. Pandora's Problem with Order Constraints

In this appendix, we show how some relevant features of Pandora's problem with order constraints (Boodaghians et al. [8]) are captured by suitable instances of Pandora's problem with combinatorial cost functions. When the underlying precedence graph is a tree, we show in Propositions B.1 and B.2 that it is possible to construct a gross-substitutes cost function such that the two problems (order constraint and combinatorial cost) are equivalent. This construction exhibits a nontrivial class of gross-substitutes cost functions for which an optimal strategy with a fixed order can be computed efficiently. Then, in Claim B.1, we borrow an instance from Boodaghians et al. [8] for which every optimal strategy is not fixed order, and we translate it into an instance of Pandora's problem with a subadditive cost function that has the same property. This constitutes an alternative proof that adaptivity may be necessary under subadditive costs (other than the proof of Theorem 1).

B.1. Pandora's Problem with Order Constraints

An instance of Pandora's problem with order constraint is defined by n boxes and a precedence graph G. Each box b_i contains a random reward with distribution D_i and has (additive) cost c_i. The precedence graph is a directed acyclic graph whose vertices are the boxes; to simplify the presentation, we assume the existence of an auxiliary box b_0 (that has no cost and contains zero reward) that is the unique root of G. The decision maker selects (possibly in an adaptive way) which boxes to open sequentially, with the constraint that in order to open a box b_i, at least one of the boxes corresponding to its parent nodes needs to have been already opened. The goal, like in the standard Pandora's problem, is to maximize the utility (i.e., maximum value minus sum of costs of opened boxes).

B.2. Tree Constraints

Given any rooted directed tree and subset of nodes S, we define $cl(S)$, the closure of S, as the minimal connected set of nodes containing S and the root. Clearly, the closure is monotone under inclusion (i.e., if $S \subseteq T$, then $cl(S) \subseteq cl(T)$). Alternatively, the closure of S can be characterized as the union of all of the paths from the root to the nodes in S. We can use the closure operator to define a cost function c_G as follows: $c_G(S) = \sum_{b_i \in cl(S)} c_i$.

Proposition B.1. *If G is a tree, then c_G is gross substitutes.*

Proof. Let $S \subseteq T$ be any two subsets of the nodes. Moreover, let b be any element not in T; to prove submodularity, we need to show that

$$c_G(S \cup \{b\}) - c_G(S) \geq c_G(T \cup \{b\}) - c_G(T).$$

If $b \in cl(T)$, then the inequality trivially holds, so we need to argue only about elements outside $cl(T)$ (and therefore, also outside $cl(S)$).

Because b is not in $cl(S)$, there exists a unique path $s_0, s_1, \ldots, s_k$ such that $s_0 \in S$, $s_i \notin S$ for all $i \in [k]$ and $s_k = b$, which connects b to S. This means that $c_G(S \cup \{b\}) - c_G(S) = \sum_{i=1}^{k} c_i$. Now, because $b \notin cl(T)$, there is also a path connecting T to b. Because there is only one path from the root to b, we have that this second path is a suffix of $s_0, s_1, \ldots, s_k$, starting at a certain $s_\ell \in T$. This means that $c_G(T \cup \{b\}) - c_G(T) = \sum_{i=\ell}^{k} c_i$ for some $0 < \ell \leq k-1$. This concludes the proof of the submodularity.

To argue about gross substitutes, we use the equivalent definition of gross substitute that we introduced at the end of Appendix A. We define P_i, P_j, and P_k as the paths that connect S (and hence, its closure) to i, j, and k, respectively. Let $P_{i,j,k} = P_i \cap P_j \cap P_k$, and let $\hat{S} = S \cup P_{i,j,k}$. It holds that

$$\begin{cases} c_G(i,j|S) + c_G(k|S) = c_G(i,j|\hat{S}) + c_G(k|\hat{S}) + 2 \cdot c_G(P_{i,j,k}|S) \\ c_G(i|S) + c_G(j,k|S) = c_G(i|\hat{S}) + c_G(j,k|\hat{S}) + 2 \cdot c_G(P_{i,j,k}|S) \\ c_G(j|S) + c_G(i,k|S) = c_G(j|\hat{S}) + c_G(i,k|\hat{S}) + 2 \cdot c_G(P_{i,j,k}|S). \end{cases}$$

Therefore, the multiset in Equation (A.1) has a unique max if and only if the following has a nonunique maximum:

$$\{c_G(i,j|\hat{S}) + c_G(k|\hat{S}), c_G(i|\hat{S}) + c_G(j,k|\hat{S}), c_G(j|\hat{S}) + c_G(i,k|\hat{S})\}.$$

If $P_i \cap P_j = P_i \cap P_k = P_j \cap P_k = P_{i,j,k}$, then the three terms share the same value:

$$\begin{aligned} c_G(i,j|S) + c_G(k|S) &= c_G(i|S) + c_G(j,k|S) \\ &= c_G(j|S) + c_G(i,k|S) = c_G(P_i|S) + c_G(P_j|S) + c_G(P_k|S). \end{aligned}$$

Else, without loss of generality, $(P_i \cap P_j) \neq P_{i,j,k}$, and $P_k \cap (P_i \cup P_j) = P_{i,j,k}$; so, we get that

$$\begin{aligned} c_G(P_i \cup P_j|S) + c_G(P_k|S) &< c_G(i|S) + c_G(j,k|S) \\ &= c_G(j|S) + c_G(i,k|S) = c_G(P_i|S) + c_G(P_j|S) + c_G(P_k|S). \quad \square \end{aligned}$$

Fix any instance $\mathcal{I}^G = (D_1, \ldots, D_n, c_1, \ldots, c_n, G)$ of Pandora's problem with precedence constraint on G and any instance $\mathcal{I} = (D_1, \ldots, D_n, c_G)$ of Pandora's problem with combinatorial cost function c_G. The following proposition establishes an equivalence relation between the two problems.

Proposition B.2. *For any optimal strategy π_G^* for $\mathcal{I}^G$, the strategy π^* for $\mathcal{I}$ that opens the same boxes in $\mathcal{I}$ is optimal.*

Moreover, for any optimal strategy π^* for $\mathcal{I}$, consider the strategy π_G^* for $\mathcal{I}^G$ that whenever π^* opens box i, opens the boxes of $cl(i)$ that were not already opened, then strategy π_G^* is optimal.

Note that the transformations defined in Proposition B.2 preserve the property of using a fixed order. By Boodaghians et al. [8], we know that there always exists a fixed-order strategy for $\mathcal{I}^G$ that is optimal and efficiently computable; this implies that there is a fixed-order strategy for $\mathcal{I}$ (same boxes and cost c_G) that is optimal for it.

Proof. Fix any instance $\mathcal{I}^G$ of Pandora's problem with order constraint on a tree and the corresponding instance $\mathcal{I}$ of Pandora's problem with the same boxes but without order constraint and gross substitutes cost function c_G. Given any strategy π^G for $\mathcal{I}^G$, it is immediate to design a strategy π for $\mathcal{I}$ that attains at least the same expected utility; π opens exactly the same boxes as π^G (and thus, earns exactly the same reward), and given the structure of the cost function, it pays the same cost.

Conversely, fix any strategy π for $\mathcal{I}$; we can associate a strategy π^G for $\mathcal{I}^G$ that yields at least the same utility. It opens the same boxes as π but before exploring a costly box i (if it is not already opened), always exhausts the boxes within the unique path to box i according to G. All the boxes along this path have a zero marginal cost. Clearly, the latter strategy yields larger utility as it gives larger rewards (as it opens more boxes) but pays at most the same cost (by submodularity). □

B.3. General Order Constraints

If the precedence graph G is not a tree, we cannot prove an equivalence relation as above. However, we can translate the (nontree) example given in Boodaghians et al. [8, theorem 6 of the extended version] into an instance of our model with a subadditive cost function, which admits no fixed-order strategy that is optimal.

Claim B.1. There exists an instance of Pandora's problem with subadditive cost c that admits no fixed-order strategy that is optimal.

Proof. Consider the following instance on four boxes. The rewards are weighted Bernoulli defined as follows:

$$V_1 = \begin{cases} 100 & \text{w.p. } 1/3 \\ 2.5 & \text{w.p. } 1/3 \\ 0 & \text{w.p. } 1/3 \end{cases} \qquad V_2 = 2 \qquad V_3 = \begin{cases} 3 & \text{w.p. } 1/2 \\ 0 & \text{w.p. } 1/2 \end{cases} \qquad V_4 = \begin{cases} 6 & \text{w.p. } 1/2 \\ 0 & \text{w.p. } 1/2. \end{cases}$$

The cost function is defined as follows. Box 1 has a marginal cost of zero given any other subset of boxes. For the other boxes, the cost is given by $c(\{2,3,4\}) = c(\{2,3\}) = 2.1$, $c(\{3\}) = c(\{3,4\}) = 1.1$, $c(\{2\}) = c(\{4\}) = c(\{2,4\}) = 1$. One can easily verify that this cost function is subadditive (but not submodular). This function is an adaptation of Boodaghians et al. [8, theorem 9] to our setting.

One can verify that the unique optimal strategy is the following. Open box 1. If $V_1 = 100$, halt. Else, if $V_1 = 2.5$, open box 4. If $V_4 = 6$, halt. Else ($V_4 = 0$), open box 3, and halt. Finally, if $V_1 = 0$, open box 4. If $V_4 = 6$, halt. Else, open box 2, and halt. Moreover, one can also verify that no fixed-order strategy obtains the same expected utility or more. □

Appendix C. Missing Proof from Section 2

Observation 1. For every permutation σ, there exists an optimal strategy with fixed order σ that is a fixed-order strategy with thresholds.

Proof. Let $\sigma : [n] \to [n]$ be a permutation, and for the sake of readability, we assume without loss of generality that σ is the identity permutation. Consider any strategy π that is restricted to inspecting the boxes in $[n]$ in the order $1, \dots, n$. Let $i \in [n]$, and let us assume that the strategy π is in the stage where it has already inspected the boxes $[i-1]$, and it now has to decide whether to inspect box i or halt. Let $x \geq 0$ be the maximum value among those observed in the previous rounds.

Given these parameters, we define $f_i(x)$ to be the maximum extra utility attainable by any strategy from this point onward. If π decides to halt, then this extra utility is zero. If π decides to open box i, then the contribution of that box to the utility is $(V_i - x)^+ - c(i|[i-1])$, and π moves to round $i+1$, from which the maximum extra utility attainable is $f_{i+1}(\max(V_i, x))$. We thus get the following backward-recursive relation:

$$f_i(x) = \max[0, \mathbb{E}[(V_i - x)^+ - c(i|[i-1]) + f_{i+1}(\max(V_i, x))]].$$

The base function f_{n+1} is defined to be the constant function 0, and the maximum utility attainable in total by any strategy with this order is $f_1(0)$.

Now, note that $f_i(x)$ is a monotone decreasing function of x. Furthermore, f_i is 1-Lipshtitz and thus, continuous. In fact, fix any two values $x < y$, and let $g_i(y)$ be the extra utility (with respect to y) attainable playing the same strategy underlying $f_i(x)$ (i.e., playing optimally pretending to have as largest reward found so far the value x and not y). It is easy to see that $f_i(x) - g_i(y)$ is at most $y - x$; thus, we have the following:

$$f_i(x) - f_i(y) \leq f_i(x) - g_i(y) \leq y - x,$$

where the first inequality follows from the suboptimality of the strategy followed in g_i given reward y.

Consider the set $C_i = \{x \in \mathbb{R}_{\geq 0} | f_i(x) = 0\}$. For every I, we define

$$t_i = \begin{cases} \infty & C_i = \emptyset \\ \min C_i & \text{otherwise.} \end{cases}$$

Because f_i is continuous and monotonic, then t_i is well defined. Note also that $f_i(x) = 0$ if and only if $x \geq t_i$. Therefore, the strategy π that for every round I, halts if and only if the maximum value observed thus far, x, is at least t_i is an optimal strategy. This concludes the proof because the strategy that we described is a fixed-order strategy with thresholds t_i. □

Appendix D. Missing Proofs from Section 4

In this appendix, we provide missing proofs from Section 4.

Claim 3. Let $\pi \subseteq [n]$ be an impulsive strategy. Then, for any set of boxes $T \subseteq [n]$ such that $T \cap \pi = \emptyset$, we have

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

Proof. The first bullet is clearly implied by Observation 3 and the fact that for every $i \in [n]$, we have

$$v_i - v_r \le (v_i - v_r)^+ \le v_i$$

The second bullet holds because

$$\begin{aligned}\mathsf{u}_N(\pi|T) &= \mathbb{E}\left[\max_{i\in S(\pi)}(V_i)\right] - \mathbb{E}[c(S(\pi)|\{r\} \cup T)] \\ &= p_{(\pi)} \cdot \mathbb{E}\left[\max_{i\in S(\pi)}(V_i)|\exists i \in \pi \text{ s.t. } V_i = v_i\right] - \mathbb{E}[c(S(\pi)|\{r\} \cup T)] + p_{(\pi)}v_r - p_{(\pi)}v_r \\ &= p_{(\pi)} \cdot \mathbb{E}\left[\max_{i\in S(\pi)}(V_i - v_r)|\exists i \in \pi \text{ s.t. } V_i = v_i\right] - \mathbb{E}[c(S(\pi)|\{r\} \cup T)] + p_{(\pi)}v_r \\ &= \mathsf{u}_M(\pi|T) + p_{(\pi)}v_r \qquad .\end{aligned}$$

Finally, the third bullet is immediately implied by the fact that $(v_i - v_r)^+ = v_i - v_r$ for every $i \in \pi^Y$ and by the fact that

$$\begin{aligned}\mathsf{u}_Y(\pi|T) &= \mathbb{E}\left[\max_{i\in S(\pi)}(V_i - v_r)^+\right] - \mathbb{E}[c(S(\pi)|\{r\} \cup T)] \\ &= p_{(\pi)} \cdot \mathbb{E}\left[\max_{i\in S(\pi)}(V_i - v_r)^+|\exists i \in \pi \text{ s.t. } V_i = v_i\right] - \mathbb{E}[c(S(\pi)|\{r\} \cup T)] \quad \square\end{aligned}$$

Claim 5. There exist boxes $a \in A, b \in B$ such that π^N inspects b only after inspecting a (i.e., $k \ge 1$).

Proof. Assume toward contradiction that $k = 0$. In this case, we have $\pi^N = (B^{\mathsf{pre}}, A^{\mathsf{suff}}) = (\pi^B, \pi^A)$. We split into cases. (a) π^B is a permutation of π^Y, and (b) π^B is strictly included in π^Y.

Case a. Note that we have $\mathsf{u}_N(\pi^N) = \mathsf{u}_N(\pi^B) + q_{(\pi^B)}\mathsf{u}_N(\pi^A|\pi^B)$. Furthermore, we have $p_{(\pi^Y)} = p_{(\pi^B)}$ because π^B and π^Y are permutations of each other. Thus, because π^N is optimal for the scenario that $V_r = 0$, then in particular, we must have $\mathsf{u}_N(\pi^B) \ge \mathsf{u}_N(\pi^Y)$ as otherwise, we could replace π^N by the strategy (π^Y, π^A) and strictly improve utility; note that the fact that π^B and π^Y are permutations of one another also implies $q_{(\pi^B)}\mathsf{u}_N(\pi^A|\pi^B) = q_{(\pi^Y)}\mathsf{u}_N(\pi^A|\pi^Y)$.

Similarly to the proof of Lemma 3, from here we get

$$\mathsf{u}_N(\pi^Y) = \mathsf{u}_Y(\pi^Y) + p_{(\pi^Y)}v_r \ge \mathsf{u}_Y(\pi^B) + p_{(\pi^B)}v_r = \mathsf{u}_N(\pi^B) \ge \mathsf{u}_N(\pi^Y),$$

where the first and second equalities hold by Claim 3 and the first inequality holds by the optimality of π^Y for the scenario that $V_r = v_r$. Thus, all expressions in the above chain are equal, and in particular, we have $\mathsf{u}_N(\pi^Y) = \mathsf{u}_N(\pi^B)$. This implies that the strategy π' obtained from π^* by replacing π^B with π^Y is also optimal. Now, consider the impulsive strategy $\pi'' = (\pi^Y, r, \pi^A)$. Because $v_i \ge v_r$ 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 than that of π', again for any realization of the boxes. Thus, the impulsive strategy π'' is optimal as well, a contradiction.

Case b. Note that in this case, because π^Y (and therefore, also π^B) does not have a deterministic box, then in particular, we have $q_{(\pi^B)} - q_{(\pi^Y)} > 0$. Now, the fact that $\pi^N = (\pi^B, \pi^A)$ implies

$$\mathsf{u}_M(\pi^N_A|\pi^Y) = q_{(\pi^B)}\mathsf{u}_M(\pi^A|\pi^Y) \tag{D.1}$$

because in the impulsive strategy with dummy boxes $(\pi^B, \pi^A)_A$, the strategy π^A is executed if and only if all of the dummy boxes corresponding to B are not realized, which happens with probability $q_{(\pi^B)}$. From here, we get

$$\begin{aligned}0 &\le \mathsf{u}_Y(\pi^Y) - \mathsf{u}_Y(\pi^B) \\ &= \mathsf{u}_M(\pi^Y) - \mathsf{u}_M(\pi^B) \\ &= \mathsf{u}_M(\pi^Y) - \mathsf{u}_M(\pi^N_B) \\ &\le \mathsf{u}_M(\pi^N_A|\pi^Y) - q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y) \\ &= q_{(\pi^B)}\mathsf{u}_M(\pi^A|\pi^Y) - q_{(\pi^Y)}\mathsf{u}_M(\pi^A|\pi^Y) \\ &= (q_{(\pi^B)} - q_{(\pi^Y)})\mathsf{u}_M(\pi^A|\pi^Y),\end{aligned}$$

where the first inequality holds because π^Y is optimal for the scenario where $V_r = v_r$, the first equality holds by Claim 3, the second equality holds because the impulsive strategy with dummy boxes $(\pi^B, \pi^A)_B$ exactly coincides with the strategy π^B, the second inequality holds by Inequality (5), and the third equality holds by Equation (D.1).

Because $q_{(\pi^B)} - q_{(\pi^Y)} > 0$, then this implies $\mathsf{u}_M(\pi^A|\pi^Y) \ge 0$, a contradiction to Claim 4. $\square$

Appendix E. Missing Proofs from Section 5

In this appendix, we show that the reduction $\mathcal{T}$ preserves several additional properties, similarly to what is done in Lemma 7.

Claim E.1. If c is a matroid rank function, then c' obtained by transformation $\mathcal{T}$ is also MRF.

Proof. Cost function c is MRF; thus, there exists a matroid $\mathcal{M} \subseteq 2^{[n]}$ such that for every $S \subseteq [n]$, $c(S) = r_{\mathcal{M}}(S)$. For a set $S \subseteq [n] \times [m]$, let $A_S = \{i | \exists j \text{ such that } (i,j) \in S\}$, and let $B_S^i = |\{j | (i,j) \in S\}|$. Consider the family of subsets $\mathcal{M}' \subseteq 2^{[n]\times[m]}$, where

$$\mathcal{M}' = \{S | A_S \in \mathcal{M}, B_S^i \leq 1 \ \forall i \in [n]\}.$$

We claim that $\mathcal{M}'$ is a matroid. It is easy to verify that $\mathcal{M}'$ is downward closed, so we only need to verify the augmentation property (see Appendix A). Let $S, T \in \mathcal{M}'$ such that $|S| > |T|$. $\mathcal{M}$ is a matroid; thus, because $A_S, A_T \in \mathcal{M}$, and $|A_S| > |A_T|$, there exists $i \in A_S \setminus A_T$ such that $A_T \cup \{i\} \in \mathcal{M}$. This implies that there exists j such that $(i,j) \in S$. Therefore, because (i) $A_{T\cup\{(i,j)\}} = A_T \cup \{i\}$, (ii) $B^{i'}_{T\cup\{(i,j)\}} = B^{i'}_T$ for $i' \neq i$, and (iii) $B^i_{T\cup\{(i,j)\}} = 1$, it holds that $T \cup \{(i,j)\} \in \mathcal{M}'$. The function c' is the rank function of the matroid $\mathcal{M}'$. □

Claim E.2. If c is gross substitutes, then c' obtained by transformation $\mathcal{T}$ is also gross substitutes.

Proof. Given $T \subseteq [n] \times [m]$, let $A_T = \{i | \exists j, (i,j) \in T\}$.

To prove the claim, we first prove the following two observations. Given a vector of prices $\boldsymbol{p}$ for $[n] \times [m]$, let $\boldsymbol{q}$ be the vector of prices for $[n]$ such that $q_i = \min_j p_{i,j}$. As a convention, we denote with $p(S)$ the sum of all of the prices of the elements in S. Then,

$$S \in \arg\max_{T\subseteq[n]\times[m]} c'(T) - p(T) \Rightarrow A_S \in \arg\max_{T\subseteq[n]} c(T) - q(T) \tag{E.1}$$

$$q(A_S) = p(S) \wedge A_S \in \arg\max_{T\subseteq[n]} c(T) - q(T) \Rightarrow S \in \arg\max_{T\subseteq[n]\times[m]} c'(T) - p(T). \tag{E.2}$$

We first observe that (i) for every $S \subseteq [n] \times [m]$, it holds that $c'(S) = c(A_S)$ and $q(A_S) \leq p(S)$. (ii) For each set $A \subseteq [n]$, then for the set $T = \{(i, \text{argmin}_{j\in[m]} p_{i,j}) | i \in A\}$, it holds that $p(T) = q(A)$.

Equation (E.1) follows because given a set $S \in \text{argmax}_{T\subseteq[n]\times[m]} c'(T) - p(T)$, assume toward contradiction that there exists $A \subseteq [n]$ such that $c(A) - q(A) > c(A_S) - q(A_S)$; then, if we let $T = \{(i, \arg\min_{j\in[m]} p_{i,j}) | i \in A\}$, it holds that $p(T) = q(A)$, and then, $c'(T) - p(T) = c(A) - q(A) > c(A_S) - q(A_S) \geq c'(S) - p(S)$, which is a contradiction.

Equation (E.2) follows because given $S \subseteq [n] \times [m]$ that satisfies the conditions of the equation, for every $S' \subseteq [n] \times [m]$, it holds that

$$c'(S') - p(S') \leq c(A_{S'}) - q(A_{S'}) \leq c(A_S) - q(A_S) = c'(S) - p(S).$$

To conclude the proof, we need to argue that for any vectors of prices $\boldsymbol{p}, \boldsymbol{p}'$ over $[n] \times [m]$ such that $\boldsymbol{p}' \geq \boldsymbol{p}$ component wise and for every set $S \in \text{argmax}_T c'(T) - p(T)$, there exists $S' \in \text{argmax}_T c'(T) - p'(T)$ that contains $R = \{(i,j) \in S | p_{i,j} = p'_{i,j}\}$. It holds that $p(S) = q(A_S)$ (otherwise, $T = \{(i, \arg\min_{j\in[m]} p_{i,j}) | i \in A_s\}$ has the same value with lower price). Let $q'_i = \min_j p'_{i,j}$. Because for every $i \in A_R$, it holds that $q_i = q'_i$, and by gross substitutes of c, it holds that there exists $A \in \arg\max_{T\subseteq[n]} c(T) - q'(T)$ that contains A_R. It holds that $q'(A_R) = p'(R) = p(R)$ because otherwise, there exists R' that $A_{R'} = A_R$, and $p(R') < p(R)$, contradicting the optimality of S. Thus, for $S' = R \cup \{(i, \text{argmin}_{j\in[m]} p_{i,j}) | i \in A \setminus A_R\}$, it holds that $q'(A_{S'}) = p(S')$ and $A_{S'} \in \text{argmax} c(T) - q'(T)$; then, by Equation (E.2), we found S' that is in the demand and contains R, which concludes the proof. □

Claim E.3. If c is coverage, then c' obtained by transformation $\mathcal{T}$ is also coverage.

Proof. Let E, w, and g be elements and the mappings that are defined in the definition of coverage with respect to c. We construct $g' : E \to 2^{[n]\times[m]}$ such that E, w, and g' define c' as a coverage. In particular, we have

$$g'(e) = g(e) \times [m], \ \forall e \in E. \quad \square$$

Claim E.4. If c is XOS, then c' obtained by transformation $\mathcal{T}$ is also XOS.

Proof. Cost function c is XOS on $[n]$; this means that there exist additive functions $a^1, \ldots, a^\ell$ over $[n]$ such that for every $S \subseteq [n]$, $c(S) = \max_{t\in[\ell]} a^t(S)$. We want to show that c' is XOS over $[n] \times [m]$. To this end, we construct the following family of additive functions over $[n] \times [m]$; for every $t \in [\ell]$ and $r \in [m]^n$, define the additive function $a^{t,r}(i,j) = a^t(i) \cdot \mathbb{I}_{\{r_i = j\}}$. We conclude the proof by observing that

$$c'(S) = \max_{t\in[\ell], r\in[m]^n} a^{t,r}(S), \ \forall S \subseteq [n] \times [m]. \quad \square$$

Claim E.5. If c is subadditive, then c' obtained by transformation $\mathcal{T}$ is also subadditive.

Proof. We need to prove that for any pair of sets $S, T \subseteq [n] \times [m]$, it holds that $c'(S) + c'(T) \geq c'(S \cup T)$. Let $A_S = \{i | \exists j \in [m] \text{ such that } (i,j) \in S\}$, $A_T = \{i | \exists j \in [m] \text{ such that } (i,j) \in T\}$, and $A_{S\cup T} = \{i | \exists j \in [m] \text{ such that } (i,j) \in S \cup T\}$. Then, it holds that

$$c'(S) + c'(T) = c(A_S) + c(A_T) \geq c(A_{S\cup T}) = c'(S \cup T),$$

where the inequality is by subadditivity of c and because $A_S \cup A_T = A_{S\cup T}$. □

Claim E.6. Transformation $\mathcal{T}$ does not maintain budget additive.

Proof. Consider the case where $n = 3, m = 2$, and $c(S) = \min(|S|, 2)$ (a symmetric function with three values that equal one and a budget of two, and we create two copies of every original box). c' must define a value of one for all boxes and a budget of two, but then, $c'(\{(1,1),(1,2)\}) = 2$, which should be one by the definition of the transformation. □

Endnotes

[1] The exploration order of a strategy π is adaptive (to the realizations of the boxes) if there exist boxes i, j in which there are realizations (of all boxes) that π opens box i before box j and that π opens box j before box i.

[2] Notably, for the larger class of subadditive cost functions, we find that an example demonstrating the necessity of adaptive order can be induced by an (seemingly unrelated) example that has been given in a completely different model of Pandora's box under constrained exploration order (Boodaghians et al. [8]) (see Claim B.1 in Appendix B).

[3] Usually, a strategy will be described as an algorithm that calculates the next decision of the strategy given the current history rather than explicitly describing the function π.

[4] By equivalent, we mean that there is a natural one-to-one correspondence between strategies of the instance and the subinstance that lead to the same (marginal) utility.

[5] If $S(\pi) = \emptyset$, then $\max_{i \in S(\pi)} V_i = 0$.

[6] A distribution is weighted Bernoulli if its support is $\{0, w\}$ for some value $w \geq 0$.

[7] Note that the generic random variable $(V_i - \kappa)^+$ is nonnegative and smaller than V_i, which has finite expectation. Moreover, $(V_i - \kappa)^+$ converges point wise to zero. Hence, $\mathbb{E}[(V_i - \kappa)^+]$ converges to zero by the dominated convergence theorem for random variables.

[8] We remark that even a demand oracle to the cost function, in the sense of Blumrosen and Nisan [7], 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 Appendix A for definitions of gross substitutes and matroid rank functions.

References

[1] Alaei S, Makhdoumi A, Malekian A (2021) Revenue maximization under unknown private values with non-obligatory inspection. Biró P, Chawla S, Echenique F, eds. *EC '21 Proc. 2021 ACM Conf. Econom. Comput.* (ACM, New York), 27–28.

[2] Balkanski E, Leme RP (2020) On the construction of substitutes. *Math. Oper. Res.* 45(1):272–291.

[3] Bechtel C, Dughmi S, Patel N (2022) Delegated Pandora's box. Pennock DM, Segal I, Seuken S, eds. *EC '22 Proc. 2022 ACM Conf. Econom. Comput.* (ACM, New York), 666–693.

[4] Berger B, Ezra T, Feldman M, Fusco F (2024) Pandora's problem with deadlines. Wooldridge MJ, Dy JG, Natarajan S, eds. *Thirty-Eighth AAAI Conf. Artificial Intelligence, AAAI 2024* (AAAI Press, Palo Alto, CA), 20337–20343.

[5] Beyhaghi H, Cai L (2023) Pandora's problem with nonobligatory inspection: Optimal structure and a PTAS. Beyhaghi H, Cai L, eds. *Proc. 55th Annual ACM Sympos. Theory Comput., STOC 2023* (ACM, New York), 803–816.

[6] Beyhaghi H, Kleinberg R (2019) Pandora's problem with nonobligatory inspection. Karlin AR, Immorlica N, Johari R, eds. *Proc. 2019 ACM Conf. Econom. Comput., EC 2019* (ACM, New York), 131–132.

[7] Blumrosen L, Nisan N (2007) *Algorithmic Game Theory* (Cambridge University Press, Cambridge, UK).

[8] Boodaghians S, Fusco F, Lazos P, Leonardi S (2023) Pandora's box problem with order constraints. *Math. Oper. Res.* 48(1):498–519.

[9] Chawla S, Gergatsouli E, McMahan J, Tzamos C (2023) Approximating Pandora's box with correlations. Megow N, Smith AD, eds. *Approximation, Randomization, Combinatorial Optimization. Algorithms Techniques, {APPROX/RANDOM} 2023*, LIPIcs, vol. 275 (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Wadern, Germany), 26:1–26:24.

[10] Chawla S, Gergatsouli E, Teng Y, Tzamos C, Zhang R (2020) Pandora's box with correlations: Learning and approximation. Irani S, ed. *61st IEEE Annual Sympos. Foundations Computer Sci., FOCS 2020* (IEEE, Piscataway, NJ), 1214–1225.

[11] Doval L (2018) Whether to open Pandora's box. *J. Econom. Theory* 175:127–158.

[12] Dumitriu I, Tetali P, Winkler P (2003) On playing golf with two balls. *SIAM J. Discrete Math.* 16(4):604–615.

[13] Esfandiari H, Hajiaghayi MT, Lucier B, Mitzenmacher M (2019) Online Pandora's boxes and bandits. *Thirty-Third AAAI Conf. Artificial Intelligence, AAAI 2019* (AAAI Press, Palo Alto, CA), 1885–1892.

[14] Fu H, Li J, Liu D (2023) Pandora box problem with nonobligatory inspection: Hardness and approximation scheme. Saha B, Servedio RA, eds. *Proc. 55th Annual ACM Sympos. Theory Comput., STOC 2023* (ACM, New York), 789–802.

[15] Fu H, Li J, Xu P (2018) A PTAS for a class of stochastic dynamic programs. Chatzigiannakis I, Kaklamanis C, Marx D, Sannella D, eds. *45th Internat. Colloquium Automata, Languages, Programming, ICALP 2018*, LIPIcs, vol. 107 (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Wadern, Germany), 56:1–56:14.

[16] Gatmiry K, Kesselheim T, Singla S, Wang Y (2024) Bandit algorithms for prophet inequality and Pandora's box. Woodruff DP, ed. *Proc. 2024 ACM-SIAM Sympos. Discrete Algorithms, SODA 2024* (SIAM, Philadelphia), 462–500.

[17] Gergatsouli E, Tzamos C (2022) Online learning for min sum set cover and Pandora's box. Chaudhuri K, Jegelka S, Song L, Szepesvári C, Niu G, Sabato S, eds. *Internat. Conf. Machine Learn., ICML 2022*, Proceedings of Machine Learning Research, vol. 162 (PMLR, New York), 7382–7403.

[18] Gergatsouli E, Tzamos C (2023) Weitzman's rule for Pandora's box with correlations. Oh A, Naumann T, Globerson A, Saenko K, Hardt M, Levine S, eds. *Advances in Neural Information Processing Systems*, vol. 36 (Curran Associates Inc., Red Hook, NY), 12644–12664.

[19] Guo C, Huang Z, Tang ZG, Zhang X (2021) Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and Pandora's problem. Belkin M, Kpotufe S, eds. *Conf. Learn. Theory, COLT 2021*, Proceedings of Machine Learning Research, vol. 134 (PMLR, New York), 2248–2288.

[20] Kleinberg JM, Kleinberg R (2018) Delegated search approximates efficient search. Tardos E, Elkind E, Vohra R, eds. *Proc. 2018 ACM Conf. Econom. Comput.* (ACM, New York), 287–302.

[21] Kleinberg RD, Waggoner B, Weyl EG (2016) Descending price optimally coordinates search. Conitzer V, Bergemann D, eds. *Proc. 2016 ACM Conf. Econom. Comput., EC'16* (ACM, New York), 23–24.

[22] Lehmann B, Lehmann D, Nisan N (2006) Combinatorial auctions with decreasing marginal utilities. *Games Econom. Behav.* 55(2):270–296.

[23] Olszewski W, Weber R (2015) A more general Pandora rule? *J. Econom. Theory* 160:429–437.

[24] Reijnierse H, van Gellekom A, Potters JAM (2002) Verifying gross substitutability. *Econom. Theory* 20(4):767–776.

[25] Schrijver A (2003) *Combinatorial Optimization: Polyhedra and Efficiency*, vol. 24 (Springer, Berlin).

[26] Segev D, Singla S (2021) Efficient approximation schemes for stochastic probing and prophet problems. Biró P, Chawla S, eds. *EC '21 Proc. 22nd ACM Conf. Econom. Comput.* (ACM, New York), 793–794.

[27] Singla S (2018) The price of information in combinatorial optimization. Czumaj A, ed. *Proc. Twenty-Ninth Annual ACM-SIAM Sympos. Discrete Algorithms, SODA 2018* (SIAM, Philadelphia), 2523–2532.

[28] Svitkina Z, Fleischer L (2011) Submodular approximation: Sampling-based algorithms and lower bounds. *SIAM J. Comput.* 40(6):1715–1737.

[29] Weber R (1992) On the Gittins index for multiarmed bandits. *Ann. Appl. Probab.* 2(4):1024–1033.

[30] Weitzman ML (1979) Optimal search for the best alternative. *Econometrica* 47(3):641–654.