Pandora's Box Problem with Order Constraints

Shant Boodaghians
University of Illinois
Urbana, IL, USA
boodagh2@illinois.edu

Federico Fusco
Sapienza University, DIAG
Rome, Italy
fuscof@diag.uniroma1.it

Philip Lazos
Sapienza University, DIAG
Rome, Italy
lazos@diag.uniroma1.it

Stefano Leonardi
Sapienza University, DIAG
Rome, Italy
leonardi@diag.uniroma1.it

June 1, 2020

Abstract

The Pandora's Box Problem, originally formalized by Weitzman in 1979, models selection from a set of options each with stochastic parameters, when evaluation (i.e. sampling) is costly. This includes, for example, the problem of hiring a skilled worker, where only one hire can be made, but the evaluation of each candidate is an expensive procedure.

Weitzman showed that the Pandora's Box Problem admits a simple and elegant solution which considers the options in decreasing order of the value it which opening has exactly zero marginal revenue. We study for the first time this problem when the order in which the boxes are opened is constrained, which forces the threshold values to account for both the depth of search, as opening a box gives access to more boxes, and breadth, as there are many directions to explore in. Despite these difficulties, we show that greedy optimal strategies exist and can be efficiently computed for *tree-like* order constraints.

We also prove that finding approximately optimal adaptive search strategies is NP-hard when certain matroid constraints are used to further restrict the set of boxes which may be opened, or when the order constraints are given as reachability constraints on a DAG. We complement the above result by giving approximate adaptive search strategies based on a connection between optimal adaptive strategies and non-adaptive strategies with bounded adaptivity gap for a carefully relaxed version of the problem.

1 Introduction

Stochastic search is an important problem in many fields of application, and has seen theoretical study. Modelling evaluation as costly captures the commonly observed trade-off between exploration and exploitation, and is also a natural assumption in many settings. Furthermore, it is common to only select one of the many evaluated alternatives, such as when searching for a skilled employee, or choosing a house. These ideas were formalized as the *Pandora's Box Problem* by Weitzman, in 1979 [26]. In this setting, we are presented with various alternatives modeled by a set of boxes $B = \{b_1, \ldots, b_n\}$, where box b_i costs c_i to open, and has random payoff X_i, whose distribution is known. A strategy π is a rule which determines adaptively whether to terminate the search, and if not, which box to open next. The goal is to choose a strategy π which maximizes, in expectation, the following objective:

$$\max_{i \in S(\pi)} X_i - \sum_{i \in S(\pi)} c_i,$$

where $S(\pi)$ is the set of boxes opened by strategy π. Only *one* reward can be kept in the end, but we must pay for all opened boxes.

Despite the broad range of search strategies available, Weitzman showed that the solution to this problem boils down to a simple but subtle strategy. Each box is assigned a *reservation value* ζ_i satisfying the equation $\mathbb{E}\left[\max\{0, (X_i - \zeta_i)\} - c_i\right] = 0$. This is precisely the value that the player would need to have collected, in order to be indifferent between proceeding or not, and is used as a proxy for the value of the box. The optimal strategy is to greedily open the boxes in descending order of reservation value, and to stop when there is no box left or when the maximum reward seen in the past is greater than the reservation value of the next box. Note that the order of exploration is not adaptive, but the stopping time is.

In this paper, we enrich the classical Pandora's Box Problem by adding restrictions on the order in which boxes can be accessed. To motivate this, consider the task faced by a funding agency or a research and development department. There may be different high-level directions to explore, with differing endpoints. Is it better to follow a longer line of research with a higher probability of achieving moderately interesting results, or take a riskier but shorter path which could contain a miraculous discovery at the end? In both cases, the final goal as well as the way to reach it will greatly affect the costs and payoffs generated.

With these extra constraints, the search for the optimal strategy becomes much more complex. We seek to determine the order in which to assess the options and when to stop searching, given prior knowledge of the values, the structure of the constraints and the realizations of the observed rewards. Weitzman's classical greedy strategy [26] does not apply: opening an expensive box with little reward may be needed in order to access more valuable and inexpensive boxes.

As a motivating case for this line of work, we show that when the order constraints are tree-like — that is when the boxes are nodes of a tree, and can only be opened after their parents are — then optimal search procedures may be found. This also applies for forests. Furthermore, we investigate the relation between such order constraints and other models of the Pandora's Box Problem which have been studied in the past, such as in [22]. We formalize these results in the following section.

1.1 Overview of Results and Methods

In this paper we focus on the Pandora's Box Problem with constraints on the order of exploration. In what follows, *polynomial time* means polynomial in the number of boxes and the size of the support of the random variables. This implicitly suggests that the random variables are discrete and have support size polynomial in the number of boxes. This assumption is not restrictive: we extend the techniques of Guo et al. [10] to prove that a polynomial number of samples is enough to get an ϵ additive approximate of the solution for the Pandora's Box Problem with order constraints. Moreover, these bounds are tight for the Tree-Constrained Pandora's Box Problem. To approximate the performance of optimal policies within an additive $\epsilon > 0$, with probability $1 - \delta$, for n boxes with rewards and costs supported in [0,1], it suffices to learn each reward distribution from $\tilde{O}(n^3/\epsilon^3)$ samples. For the special case of tree constraints, it suffices to take $\tilde{O}(\frac{n}{\epsilon^2})$ samples. We have hidden $\text{poly}(\log(\frac{n}{\epsilon\delta}))$ terms. See Appendix A for details.

These ideas allow us to proceed with the analysis while assuming without loss of generality that the random variables have support $poly(n)$. We further note that this result suggests the methods in this paper are robust to imprecise knowledge of the distribution of the rewards on the boxes.

1.1.1 The Tree-Constrained Pandora's Box Problem.

We first consider the Tree case, where we present an optimal strategy with a nice structure analogous to Weitzman's.

Theorem 1. *When order constraints are given by a rooted tree over the boxes, there exists an optimal-in-expectation strategy of the following form: first, label each box with a "threshold", an order-aware analog of Weitzman's reservation value. Then:*

- *From the boxes that can be opened next, choose the one with the largest threshold.*

- *Terminate if the max observed value exceeds this threshold, otherwise open the box and repeat.*

Furthermore, the optimal thresholds can be computed in polynomial time.

Notice that such threshold strategies are simple and intuitive, and enforce the desirable property that the order of the exploration is fixed up to tie-breaking. Furthermore, that an optimal strategy has this form is surprising, as this need not hold for general constraints, as discussed below. The definition of these thresholds has to address the following additional complications:

(i) **Depth:** The value of a box is not only given by its reward and cost, but also from the possibilities its opening makes accessible. This effect propagates level after level: even the deepest of the leaves can influence the decision to open the root.

(ii) **Breadth:** A key feature is the order in which the boxes are opened. It is difficult to model and optimize the interplay between different explored branches of a tree, as distant directions of exploration must be compared at every time step.

To overcome these difficulties, the first step is to generalize the reservation values used by Weitzman to the setting where the boxes are constrained to be opened in a fixed order $[b_1, b_2, \ldots, b_n]$. These values must take into account the *future* as well as the present. Our solution consists in defining the threshold value of the generic box i according to a random *stopping time* $\tau^*(y, i)$, which indicates the last box that will be opened *playing optimally* given that the player has already found reward y and is in front of box i. We call this threshold value z_i. It relates to the original definition as follows:

$$\mathbb{E}\left[(X_i - \zeta_i)_+ - c_i\right] = 0 \quad \longleftrightarrow \quad \mathbb{E}\left[\left(\max_{j=i}^{\tau^*(z_i, i)} X_j - z_i\right)_+ - \sum_{j=i}^{\tau^*(z_i, i)} c_j\right] = 0,$$

where $(a)_+ := \max\{a, 0\}$. From these stopping times and threshold values we can infer that certain runs of boxes are essentially treated as one big box (which we refer to as a *macrobox*). If the algorithm decides to enter it, the exploration will either finish inside this macrobox or a decision will be made to enter the next one after having exhausted it. Moreover, these threshold values can be computed in polynomial time by a dynamic programming procedure.

For a single line, this reasoning may seem straightforward, but this property still holds when the constraints consist of a union of disjoint parallel lines. A naïve dynamic program would not be effective, as the state space is exponential in the number of lines. However, the threshold strategy which uses the reservation values computed for each line independently is still optimal: the algorithm will always enter the best available macrobox and either terminate search inside of it, or move on to another one, possibly from a different line. Surprisingly, the same approach works for *trees and forests*. Proceeding from leaves to the root it is possible to linearize the trees and use the definition of reservation value to induce a threshold strategy which is indeed optimal.

1.1.2 Impossibility and Hardness Results.

Unfortunately, these results do not extend to slightly more general constraint structures. We show below that it is NP-hard to approximate an optimal solution to the problem with some types of matroid constraints or more general order constraints. Remark that the standard notion of approximation — *i.e.* finding a solution whose performance is at least a multiplicative factor of the optimal solution in expectation — is not meaningful in this setting: a hard example (*e.g.* the hardness proof of Section 5) can be modified by adding a large-cost-no-payoff dummy box at the root of the tree. For the right cost, optimal strategies would have positive revenue, but approximately optimal strategies would have negative revenue.

We define below a modified notion of approximately optimal solution, which avoids these concerns. Note that similar metrics have been used before, in works such as [8, 12, 23].

Definition 1 (Approximately Optimal Solutions). *In this paper, we consider a notion which approximates only the reward term while paying similar costs. Formally, we seek a strategy $\hat{\pi}$ such that for any other rule π,*

$$\mathbb{E}\left[\max_{i \in S(\hat{\pi})} X_i - \sum_{i \in S(\hat{\pi})} c_i\right] \geqslant C \cdot \mathbb{E}\left[\max_{i \in S(\pi)} X_i\right] - \mathbb{E}\left[\sum_{i \in S(\pi)} c_i\right]$$

for some universal constant factor $C \in (0, 1]$, where $S(\pi)$ is the (random) set opened by strategy π.

Theorem 2. *Consider the Pandora's Box Problem with order constraints when either (i) a matroid constraint is added to the tree constraint or (ii) the tree constraint is generalized to a DAG. In either case, it is NP-hard to find a 0.9997-approximately optimal solution, and furthermore, the optimal solution need not have a fixed order of exploration.*

This approximation hardness is shown in the sense of Definition 1, but extends to the classical notion, since it is stronger. We finally remark that proving hardness of approximation for stochastic problems needs to address the effects of the randomness on the objective and the search trajectory, and as such is challenging. Moreover, the structure of the $\max_S X_i - \sum_S c_i$ objective makes optimal *random* solutions difficult to "hide", in a standard combinatorial sense.

1.1.3 Approximation Results.

The hardness result above, along with recent work on modified versions of the Pandora's Box Problem [2, 7, 16, 21, 22], motivates the study of approximation algorithms for the more general case of order constraints. Of the above citations, the ones closest to our setting are [21, 22], where the author reduces the Pandora's Box Problem in the presence of downwards-closed constraints to adaptive maximization of non-negative submodular functions. The key concept in their work is the *adaptivity gap*, i.e. the ratio between the best adaptive solution and the best non-adaptive one for this new problem. We similarly show the following.

Theorem 3. *Consider the Pandora's Box Problem with constraints modeled by some prefix closed family — a generalization of order constraints, defined in Section 2, then for every adaptive strategy π, there exists a non-adaptive strategy, i.e. a feasible set S, such that the following holds[1]:*

$$\mathbb{E}\left[\max_{i \in S} X_i - \sum_{i \in S} c_i\right] \geqslant \tfrac{1}{2}\mathbb{E}\left[\max_{i \in S(\pi)} X_i\right] - \mathbb{E}\left[\sum_{i \in S(\pi)} c_i\right] \qquad (1)$$

This result effectively reduces the problem of approximating adaptive strategies to the problem of selecting optimal non-adaptive sets. It should be noted, however, that in full generality of matroid constraints or precedence constraints, this could be intractable. We therefore follow an alternative approach. We show that there exists an *adaptive* strategy whose performance is better than that of every non-adaptive set, simultaneously.

Theorem 4. *For any tree-constraint, and any further downwards-closed constraint on the set of boxes that can be opened, there exists an adaptive strategy $\hat{\pi}$ such that for any fixed set S, $\hat{\pi}$ performs better in expectation than non-adaptively opening S. Note that $\hat{\pi}$ does not depend on S. Furthermore, when the downwards-closed constraints are given by generalized knapsack constraints, or any "sufficiently oblivious" matroid constraint (as defined in Section 4), the strategy $\hat{\pi}$ can be computed efficiently.*

1.2 Related Work

As discussed above, the starting point for this theory is the 1979 Weitzman's paper [26],which was at the time a generalization of preceding results in special cases, namely Kadane [13], Kadane et al. [14]. In the following years it was highlighted the similarity with the multi-armed bandits problem, which uses the notion of Gittins index as a similar cutoff value (e.g. Frostig and Weiss [9], Weber et al. [24]). Indeed the reservation value of the classical Pandora's Box Problem is a version of Gittins index, Dumitriu et al. [6]. Weiss [25], and in particular Keller and Oldale [15], deal with a similar problem, the branching bandit process. The branching process resembles the Tree-Constrained Pandora's Box Problem, though in that model, the process does not terminate, and the revenue is measured as an infinite-horizon discounted sum of payoffs, even if some finite horizon results are showed. The main difference with our works is that we focus on maximizing the largest reward minus the exploration costs, an objective function they cannot capture with their techniques, moreover their solution is defined with an implicit formula that becomes rapidly cumbersome as a function

---

[1] To further justify the benchmark of Definition 1, we show in Appendix B that approximation in the traditional sense with a non-adaptive set is impossible, as there exist examples where this classical adaptivity gap is arbitrarily large.

of the height of the tree. Interestingly enough they prove the optimality of a threshold strategy which is, in spirit, quite similar to ours.

Kleinberg et al. [17], borrowing from the language of finance, introduce the *covered call value* of a box, which is the minimum of the reservation value, and its true (random) value. They show that the expected performance of any search strategy is at most the expected covered call value of the last kept box. This inequality is tight for any method which immediately terminates search when the value of an opened box is greater than its reservation value. This novel point of view on Pandora's Box Problem started a new interest in the problem. Olszewski and Weber [18] investigate the existence of moving threshold strategies to address more general objective functions in the Pandora's setting, while Doval [5] and Beyhaghi and Kleinberg [2] analyze a setting in which a box can be chosen without paying its cost while retaining its expected reward. We highlight that these techniques do not extend to order-constrained settings.

Singla [21, 22] exploits the notion of *surrogate* box, an analog of the *covered call value*, to reduce the Pandora's Box Problem in the presence of downwards-closed constraints to adaptive maximization of non-negative submodular functions, and bound the *adaptivity gap* of this problem, *i.e.* the ratio between the best adaptive solution and the best oblivious one as in Gupta et al. [11] and later Bradac et al. [3]. Recently Kleinberg and Kleinberg [16] and Esfandiari et al. [7] studied a connection between the Pandora's Box Problem and another well known optimal stopping problem, termed the Prophet Inequalities.

2 Model and Preliminaries

In this section, we formally present our model, and give preliminaries. Recall, as the player, we adaptively open a constraint-satisfying set of boxes, paying for each one opened, and learn the (random) value of each box only after paying the cost of opening. The final payoff received is the largest value observed. This is formalized below.

The Pandora's Box Problem with order constraints. We are given a set of boxes $B = \{b_1, \ldots, b_n\}$, where b_i costs c_i to open, and has random payoff X_i, whose distribution is known. The $\{X_i\}_{i=1}^n$ are independent and need not be identically distributed.

A strategy π is a rule which determines, at any integer time $t \geqslant 0$, whether to terminate the search and, if not, which box to open next. The strategy may depend on the time t, the values observed in the past, the structure of the problem and some extra randomness. We use equivalently the terms strategy, rule and policy.

Let $S_t(\pi)$ denotes the (random) set of boxes that have been opened before time t (included) by strategy π, and let τ_π be the stopping time given by the same strategy. We use the shortcut $S(\pi)$ to denote $S_{\tau_\pi}(\pi)$, the final set of opened boxes following strategy π. Given constraint-set $\mathcal{F} \subseteq 2^B$, π is said to be $\mathcal{F}$-feasible if $S_t(\pi) \in \mathcal{F}$ with probability 1, for all t. Our goal is to choose a policy π^* which maximizes, in expectation, the following objective:

$$\mathbb{E}\left[\max_{i \in S(\pi^*)} X_i - \sum_{i \in S(\pi^*)} c_i\right]$$

Such strategies are called optimal.

Threshold Strategies. A rule π is said to be a threshold strategy if it pre-computes a collection of *threshold values*, and greedily opens the boxes following these values, stopping when the amount earned is greater than the threshold of all remaining legal moves. Formally, the strategy is defined by a *threshold function* $z : B \to \mathbb{R}$ and works as in Algorithm 1 below.

Observe that, given some consistent tie-breaking rule, the order of exploration is fixed, as the next box to consider only depends on the reservation values. We remark that any threshold function is defined *a priori*, *i.e.* it does not depend on the observed rewards, but only on the costs, the random distributions, and $\mathcal{F}$. The following facts follow by definition:

Claim 1. *Let π be a threshold strategy, where ties in the thresholds are solved arbitrarily but consistently. Then the following hold true:*

5

ALGORITHM 1: Threshold strategy

Data: Distributions of the random rewards, box costs and a threshold function $z : B \rightarrow \mathbb{R}$
$S_0 \leftarrow \varnothing$, $y \leftarrow 0$, $t \leftarrow 0$
while $y < \max\{z(b) | b \in B \backslash S_t \text{ and } \{b\} \cup S_t \in \mathcal{F}\}$ **do**
$\quad$ Let $\hat{b} \in \arg\max\{z(b) | b \in B \backslash S_t \text{ and } \{b\} \cup S_t \in \mathcal{F}\}$, tie-breaking arbitrarily
$\quad$ Open box $\hat{b}$, observe reward $\hat{X}$ and pay cost $\hat{c}$
$\quad$ $S_{t+1} \leftarrow S_t \cup \{\hat{b}\}$, $y \leftarrow \max\{y, \hat{X}\}$, $t \leftarrow t + 1$
end
---

1. ***Fixed order:*** *Following* π, $\mathbb{P}(b_i \text{ is opened before } b_j) > 0 \implies \mathbb{P}(b_j \text{ is opened before } b_i) = 0$, *for all* $i \neq j$.

2. ***Efficiency:*** *If the threshold function* z *is efficiently computable, then so is* π.

Order Constraints   In this paper, we focus on *order constraints*, where some boxes are required to be opened after others. Such constraints can be represented by a rooted tree T called precedence tree, whose nodes are the boxes. A box may only be opened if its parent has already been explored. Without loss of generality we can assume that T is connected, *i.e.* it is actually a tree, and that there is a unique node $r \in B$, which is the starting box — if not, it suffices to add a dummy root box with no cost and no reward. Formally, then, given the precedence tree T, the feasible sets $\mathcal{F}_T$ are:

$$\mathcal{F}_T := \{S \subseteq B \,|\, \forall u \in S \backslash \{r\}, \text{ PARENT}(u) \in S\}$$

Prefix closed Constraints   Order constraints may be seen as a special case of the more general *prefix-closed* constraints, which simply assert that for any legal sequence of moves, any truncation of this sequence is also legal. Formally, given a set of boxes B and a set $\mathscr{C}$ of possible orders of exploration, we say that $\mathscr{C}$ is prefix closed if for every $C \in \mathscr{C}$, every prefix of C is also in $\mathscr{C}$. Note that the order constraints defined above are a special case of this. Furthermore, intersecting any combination of order and downwards closed constraints results in some prefix closed family.

Distributional Assumptions   As discussed in the beginning of Section 1.1, it is not restrictive to assume that the random variables $\{X_i\}_{i=1,\dots,n}$ are discrete, and are supported on $s < \infty$ values. When we say an algorithm runs in polynomial time, we mean polynomial in s and n.

Preliminaries   Before proceeding with the details of the analysis, we would like to note that the Pandora's Box Problem with order constraints admits a naïve, albeit exponential-time solution: it suffices to solve a dynamic program whose states are all pairs (S, y) where $S \subseteq B$ is a set of boxes, and $y \in \mathbb{R}$ is the max value observed.

The literature on Markov Decision Processes (e.g. [20]) allows us to fix basic properties of optimal strategies: there exists an optimal strategy π^* which is a Markovian policy mapping states to actions, *i.e.* the optimal next box to open is deterministic function of the state (S, y).

Notation   In the following we use interchangeably $\max(a, b)$ or $a \vee b$ to denote the largest between two reals a and b. For the smallest we use $\min(a, b)$ or $a \wedge b$. As a further simplification $(a)_+ := \max\{a, 0\}$. We use the following simple equality repeatedly: $a \vee b - b = (a - b)_+$.

3 Optimal Search on Tree Constraints

In this section, we present and analyze an optimal-in-expectation search procedure for the Tree-Constrained Pandora's Box Problem. The classical result of Weitzman [26] shows that the unconstrained problem may be solved by a simple threshold strategy as discussed above. Formally,

Definition 2 (Reservation Value and Pandora's Rule). *Given a box b with cost $c > 0$ and nonnegative random reward X, we define the reservation value ζ of b as the smallest solution to $c = \mathbb{E}\left[(X - \zeta)_+\right]$. It can be shown that if X has finite mean, then the reservation value is well defined. The threshold strategy using the reservation values as thresholds is termed* Pandora's Rule.

The power of this strategy is that the reservation value depends only on the single box, allowing us to consider each box separately, leaving the problem dramatically more tractable. In Sections 3.1 and 3.2, we present a solution to the Pandora's Box Problem with order constraints given by one directed line, and a collection of disjoint, directed lines, respectively. In Section 3.3, we show that, in fact, solving the problem for generalized, rooted tree constraints follows immediately as a corollary, and as such, the heart of the technical contribution lies in Section 3.2.

3.1 Pandora's Box Problem on a single Line

We introduce here the simplest order constraint, given by an ordered path. This constraint may seem trivial at first, since the order of exploration is fixed, and it suffices to determine the stopping time, but it illustrates the main difficulty of order constraints: the intrinsic value of a box is not given *only* by its cost and its random reward, but also by the other boxes that are made available after its opening. Consider, for example, the line consisting of one box with cost but no reward, followed by another with reward but no cost. In this sense, a naïve threshold strategy in the sense of Weitzman's result does not immediately suffice.

As mentioned, solving the problem on the single line is not of great consequence. However, the concepts introduced in its solution are very informative for the following, so we present them separately for clarity of exposition.

Let $\mathcal{L} = [b_1, b_2, b_3, ..., b_n]$ be the ordered set of boxes, where for all i, b_i costs c_i to open, and gives random reward X_i, as usual. As an additional constraint we have that the boxes must be opened in order of their indices: box b_i can be opened only if it is the root or box b_{i-1} has already been opened. After having opened the first i boxes and collected some reward, it is a simple exercise of Dynamic Programming to determine whether to open the $(i + 1)$-th box and proceed optimally, or to terminate. The DP had $O(ns)$ states for n boxes with reward supported on s values.

However, the goal of this section is to illustrate structural properties which we will use in the general tree setting. It is clear that the set of possible strategies in this simple setting coincides with the set of all (random) stopping times τ, with respect to the filtration given by $X_1, X_2, \ldots, X_n$. For any $x \in \mathbb{R}_+$, τ and $i = 0, 1, \ldots, n, n + 1$ we denote (x, i) as the state in which box b_i is the next accessible box and x is the largest reward uncovered so far, and $\tau(x, i)$ as the (random) stopping time conditioned on being in that state, i.e. conditioned on the events $\tau \geqslant i - 1$ and $\max_{j=1}^{i-1} X_j = x$.

Before, and in what follows, we assume sum and the max operations over empty sets to have value 0. We can define the expected future reward following τ, starting in state (x, i), as:

$$\phi^\tau(x, i) := \mathbb{E}\left[\max\left\{x, \max_{j=i}^{\tau(x,i)} X_j\right\} - \sum_{j=i}^{\tau(x,i)} c_j\right],$$

where, by convention, we set $\phi^\tau(x, i)$ to be 0 when τ is ill-defined, *i.e.* the event $\tau(x, i)$ is conditioning on occurs with probability 0. In addition to that we define $\Phi(x, i) = \max_\tau \phi^\tau(x, i)$.

The objective is to choose a τ^* which maximizes $\phi^\tau(0, 0)$, corresponding to $\Phi(0, 0)$. By the remarks in the Preliminaries, we restrict ourselves to stopping times which are deterministic functions of the X_i's, i.e. the decision to stop depends deterministically only on the structure of the problem and the realizations of the rewards experienced so far.

Remark 1. *It is natural to ask whether a fully-deterministic, i.e. non adaptive, stopping time is a valid strategy. It turns out that such an approach may be arbitrarily worse than an adaptive stopping time. We provide an example in Appendix B, which resembles ones found in [21], and the extended version of [22].*

The notion of conditional stopping time defined above allows us to give the following definition, which is analogous to Definition 2:

Definition 3. *Let $\mathcal{L} = [b_1, b_2 \ldots, b_n]$ be a line of n boxes, then for every $i = 1, \ldots, n$ we can define the generalized reservation value of box b_i, denoted z_i, as the smallest solution to*

$$\mathbb{E}\left[\left(\max_{j=i}^{\tau^*(z_i, i)} X_j - z_i \right)_+ - \sum_{j=i}^{\tau^*(z_i, i)} c_j \right] = 0 \tag{2}$$

where $\tau^*(x, k)$ is an optimal random stopping time given that the largest reward sampled in the past has been x and the player has just opened box b_{k-1}, or nothing if $k = 1$.

Whereas the ζ value in Definition 2 was effectively the value collected in the past for which we were indifferent between opening a box or not, the z_i value in this definition is the past collected value for which we are indifference between proceeding (optimally) along the line of boxes or not. The following claim ensures that this is well-defined and its proof can be found in Appendix C.

Claim 2. *Definition 3 is well posed, in that the smallest solution of (2) exists and does not depend on the choice of τ^*. Also, if $z_i > 0$, some optimal stopping time $\tau^*(z_i, i)$ does not stop at $i - 1$. Finally, z_i is the value for which we are indifferent between stopping and proceeding optimally.*

We highlight here that, as z_i is defined, the optimal strategy $\tau^*(z_i, i)$ is ambiguous: we are indifferent between stopping ($\tau^* = i - 1$), and proceeding ($\tau^*(z_i, i) = \tau^*(z_i \vee X_i, i + 1)$). By convention we will always refer to the latter. A brief calculation confirms that this new definition contains as a special case the classical reservation value, viewing single boxes as length-1 lines. We will henceforth omit the term "generalized" in reference to reservation values, when clear from context. We call the threshold strategy associated to these reservation values Generalized Pandora's Rule.

Theorem 5. *The Generalized Pandora's Rule for the Line is optimal and can be computed in polynomial time and space.*

Proof. It follows by definition, and by Claim 2, that it is in our interest to proceed if the largest value seen is less than z_i, and to stop, if the largest value seen is greater. Thus, any deviation from this threshold strategy is sub-optimal. Furthermore, $\Phi(x, i)$ may be computed by a simple dynamic program solved in decreasing order of i. The reservation price of box i is the smallest point in the column for i where $\Phi(x, i) = x$. $\qquad\square$

We can define the procedure COMPUTETHRESHOLD$(b, \mathcal{L})$ as outputting the exact reservation price of b if it were added as a prefix to $\mathcal{L}$. This can be done in polynomial space and time by simply referring to the dynamic programming table for Φ and following Definition 3. This simple function plays a crucial role in the design of the final algorithm for the Tree constrained case.

We conclude this section by giving some properties of the reservation values:

Claim 3. *Given a line $\mathcal{L} = [b_1, \ldots, b_n]$ the following statements hold true for every $i = 1, \ldots, n$:*

1. *z_i can only increase if something is added at the end of the line $\mathcal{L}$;*

2. *For every i let $d(i)$ be the $\min\{t \geq i | z_{t+1} < z_i \text{ or } t = n\}$, then z_i depends only on the prefix $[b_i, \ldots, b_{d(i)}]$. If $d(i) = i$, then z_i depends only on b_i, i.e. $z_i = \zeta_i$.*

3. *If $z_i < z_j$ for all $j > i$, then the optimal stopping time $\tau^*(z_i, i)$ given by the Generalized Pandora's Rule does not depend on z_i. In particular $\tau^*(z_i, i) = \tau^*(y, i)$ for all $y \in [0, z_i]$.*

Proof. The first property derives from the fact that z_i is a fixed point of $\Phi(\cdot, i)$, and hence if something is added at the end of the line the expected revenue can only increase. For the other two claims, it is sufficient to observe that if an optimal play starts with some y, then $\tau^*(y, i) < j$ for all $j > i$ such that the reservation value $z_j < y$. $\qquad\square$

3.2 Pandora's Box Problem on a Union of Lines

We wish to generalize to a union of disjoint lines. Observe that this setting contains both the classical Pandora's Box Problem, and the single-line-constrained case as sub-problems. In essence, it captures both the breadth of exploration from the unconstrained setting, and the dependence on depth of the intrinsic value of a box of the single line case.

Formally, we have k paths $\mathcal{L}_1, \ldots \mathcal{L}_k$, where path $\mathcal{L}_i$ consists of boxes $b_1^i, \ldots, b_{n_i}^i$ which can only be opened in increasing order of subscript. Box b_j^i costs c_j^i to open, (generalized) reservation value z_j^i, and gives random reward X_j^i, following a known distribution. The reservation values are defined as in the previous section, taking only the line that the box belongs to as context, *i.e.* for any box b_j^i in $\mathcal{L}_i$, the value z_j^i is computed as if only $\mathcal{L}_i$ existed. Furthermore, by Claim 3, z_j^i depends only on $[b_j^i, \ldots, b_{d^i(j)}^i] \subseteq \mathcal{L}_i$, where $d^i(j)$ is the equivalent of $d(j)$ for line i. Since one must consider the interactions between the different lines, the dimension of the naïve dynamic program is exponential in the number of paths, so we need to be more clever. Surprisingly, we prove that the optimal strategy is still a threshold strategy, and the z_j^i's are exactly the correct thresholds. The rest of this section is a proof of this fact. The main effort in the proof is indeed to find a way to *decouple* the first box of a line from the rest of that line.

Theorem 6. *The Generalized Pandora's Rule is optimal for the Pandora's Box Problem on Union of Lines and can be implemented in polynomial time and space.*

Proof. As the reservation values are identical to the previous subsection, algorithmic results follow immediately. It remains to show optimality, which we will do by induction on the number of boxes yet to open.

If there is only one box remaining, then this is a special case of the unconstrained Pandora's Box Problem and we know that it is optimal to follow the Generalized Pandora's Rule since it coincides with the Pandora's Rule. For the induction step, without loss of generality, we may re-label the sequences such that the first box in every line is labelled 1, and $z_1^1 \geqslant z_1^2 \geqslant \cdots \geqslant z_1^k$. There are three actions to consider: (1) stopping, (2) opening b_1^1 first, and (3) opening b_1^i first for some i such that $z_1^1 > z_1^i$. We remark that the decision will depend on both the boxes to open and the largest value seen in the past, denoted y.

We begin by showing that stopping is optimal if and only if $y \geqslant z_1^1$. If $y < z_1^1$, we know that even the suboptimal strategy of opening box b_1^1 and playing only on $\mathcal{L}_1$ ignoring other lines is better than stopping. If $y \geqslant z_1^1$ and we open b_1^i for any i, then by induction, the optimal strategy is to go on exploring $\mathcal{L}_i$ without the possibility of changing line; as the z_1^j's are too small. This again contradicts the $k = 1$ case.

Suppose, then, that $y < z_1^1$, and that we decide to open b_1^i, for $i > 1$.

Case I: $z_1^1 \geqslant \cdots \geqslant z_1^i > y$. By induction, after opening b_1^i, the optimal strategy is to continue the exploration along $\mathcal{L}_i$ until the reservation value z_j^i becomes less than z_1^1 (or the reward exceeds the next reservation value), then to go along $\mathcal{L}_1$. Let $\mathcal{P}$ be the prefix of $\mathcal{L}_1$ with reservation values greater than z_1^1. By Claim 3, we note that z_1^1 depends only on $\mathcal{P}$. Let $\mathcal{Q}$ be the prefix of $\mathcal{L}_i$ including b_1^i, and extended to contain all reservation values greater than z_1^1, and let $\mathcal{Q}'$ be the prefix of $\mathcal{L}_i \backslash \mathcal{Q}$ containing all z values greater than z_1^1. Again, Claim 3 implies z_1^i depends only on $\mathcal{Q} \cup \mathcal{Q}'$. Define $z_\mathcal{Q}$ to be the reservation value of box b_1^i considering only prefix $\mathcal{Q}$. By Claim 3, $z_\mathcal{Q} \leqslant z_1^i$. For simplicity, denote also $z_\mathcal{P} = z_1^1$ and the optimal stopping times in the two prefixes $\mathcal{P}$ and $\mathcal{Q}$ as $\tau_\mathcal{P}$ and $\tau_\mathcal{Q}$. The heart of the proof is that we can treat the two prefixes $\mathcal{P}$ and $\mathcal{Q}$ as single *macro-boxes*, with random costs and random rewards. Let z_S be then the largest between $z_\mathcal{Q}$ and all the reservation values of the boxes accessible after having exhausted $\mathcal{P}$ and $\mathcal{Q}$.

Consider the following (suboptimal) executions of the algorithm:
– STRATEGY A begins by exploring $\mathcal{P}$, and if at the end of the exploration the largest reward is greater than z_S, then stops, otherwise explores $\mathcal{Q}$ as if no reward was found while exploring $\mathcal{P}$.
– STRATEGY B begins by exploring $\mathcal{Q}$, then plays optimally, that is it explores $\mathcal{P}$.

Notice that B is the optimal execution of the algorithm, by induction, under the assumption that the first box opened is b_1^i. This allows for a more compact representation: for any initial value w, we can define

the following random variables:

$$X_{\mathcal{P}}(w) := \max_{\ell=1}^{\tau_{\mathcal{P}}(w)} X_\ell^{\mathcal{P}}, \quad c_{\mathcal{P}}(w) := \sum_{\ell=1}^{\tau_{\mathcal{P}}(w)} c_\ell^{\mathcal{P}}, \quad X_{\mathcal{Q}}(w) := \max_{\ell=1}^{\tau_{\mathcal{Q}}(w)} X_\ell^{\mathcal{Q}}, \quad c_{\mathcal{Q}}(w) := \sum_{\ell=1}^{\tau_{\mathcal{Q}}(w)} c_\ell^{\mathcal{Q}},$$

where the superscripts $\mathcal{P}$ and $\mathcal{Q}$ for the rewards and the costs simply specify to which prefix the boxes belong to. With this notation the expected revenue following strategy A is:

$$\mathbb{E}\left[-c_{\mathcal{P}}(y) + X_{\mathcal{P}}(y)\mathbb{I}_{X_{\mathcal{P}}(y) \geq z_S} + \mathbb{I}_{X_{\mathcal{P}}(y) < z_S}(-c_{\mathcal{Q}}(y) + \mathbb{I}_{X_{\mathcal{Q}}(y) \geq z_S}X_{\mathcal{Q}}(y))\right] +$$
$$+ \mathbb{E}\left[\mathbb{I}_{X_{\mathcal{P}}(y) < z_S}\mathbb{I}_{X_{\mathcal{Q}}(y) < z_S}\Phi_S(X_{\mathcal{Q}}(y) \vee X_{\mathcal{P}}(y) \vee y)\right],$$

where $\Phi_S(\cdot)$ here is the optimal expected revenue after exhausting both $\mathcal{P}$ and $\mathcal{Q}$.

For strategy B, it is convenient to note that the prefix $\mathcal{P}$ is opened only if in prefix $\mathcal{Q}$ no reward is greater than $z_{\mathcal{P}}$. This means that the stopping times $\tau_{\mathcal{P}}(y)$ and $\tau_{\mathcal{P}}(y \vee X_{\mathcal{Q}}(y))$ are the same in this case. Hence, $X_{\mathcal{P}}(y) = X_{\mathcal{P}}(X_{\mathcal{Q}} \vee y)$ and $c_{\mathcal{P}}(y) = c_{\mathcal{P}}(X_{\mathcal{Q}} \vee y)$. This is the main step: for both the strategies, if both the prefixes are activated, then the stopping time of the second one is independent to the realizations in the previous prefix. Now the reward due to B is

$$\mathbb{E}\left[-c_{\mathcal{Q}}(y) + X_{\mathcal{Q}}(y)\mathbb{I}_{X_{\mathcal{Q}}(y) \geq z_{\mathcal{P}}} + \mathbb{I}_{X_{\mathcal{Q}}(y) < z_{\mathcal{P}}}\mathbb{I}_{X_{\mathcal{P}}(y) \geq z_S}(X_{\mathcal{Q}}(y) \vee X_{\mathcal{P}}(y))\right] +$$
$$+ \mathbb{E}\left[-c_{\mathcal{P}}(y)\mathbb{I}_{X_{\mathcal{Q}}(y) < z_{\mathcal{P}}} + \mathbb{I}_{X_{\mathcal{P}}(y) < z_S}\mathbb{I}_{X_{\mathcal{Q}}(y) < z_S}\Phi_S(X_{\mathcal{Q}}(y) \vee X_{\mathcal{P}}(y) \vee y)\right].$$

If we observe the two expected revenues, we note that the last term is equal in both. Since we intend to compare the two revenues, this term may be ignored. Moreover, the dependence of all the stopping times is only on y, so we can omit it in the future. The remainder of this proof is similar in spirit to [26]. For simplicity of notation, we introduce the following shorthand:

$$\pi_{\mathcal{P}} := \mathbb{I}_{X_{\mathcal{P}} > z_{\mathcal{P}}}, \quad \lambda_{\mathcal{P}} := \mathbb{I}_{z_{\mathcal{P}} > X_{\mathcal{P}} > z_S}, \quad \pi_{\mathcal{Q}} := \mathbb{I}_{X_{\mathcal{Q}} > z_{\mathcal{P}}}, \quad \lambda_{\mathcal{Q}} := \mathbb{I}_{z_{\mathcal{P}} > X_{\mathcal{Q}} > z_S}.$$

A few observations: $\pi_{\mathcal{P}}\lambda_{\mathcal{P}} = \pi_{\mathcal{Q}}\lambda_{\mathcal{Q}} = 0$, *i.e.* the events are mutually exclusive, and $\pi_{\mathcal{P}} + \lambda_{\mathcal{P}}$ is the event that strategy A will stop after having explored $\mathcal{P}$ (by definition of strategy A). With this notation, and ignoring the common term, we have:

$$A = \mathbb{E}\left[-c_{\mathcal{P}} + (\lambda_{\mathcal{P}} + \pi_{\mathcal{P}})X_{\mathcal{P}} + (1 - \pi_{\mathcal{P}} - \lambda_{\mathcal{P}})(-c_{\mathcal{Q}}) + (1 - \pi_{\mathcal{P}} - \lambda_{\mathcal{P}})(\lambda_{\mathcal{Q}} + \pi_{\mathcal{Q}})X_{\mathcal{Q}}\right],$$
$$B = \mathbb{E}\left[-c_{\mathcal{Q}} + X_{\mathcal{Q}}\pi_{\mathcal{Q}} + \lambda_{\mathcal{Q}}(-c_{\mathcal{P}} + \pi_{\mathcal{P}}X_{\mathcal{P}} + \lambda_{\mathcal{P}}(X_{\mathcal{P}} \vee X_{\mathcal{Q}}) + X_{\mathcal{Q}}(1 - \pi_{\mathcal{P}} - \lambda_{\mathcal{P}}))\right] +$$
$$+ \mathbb{E}\left[(1 - \pi_{\mathcal{Q}} - \lambda_{\mathcal{Q}})(-c_{\mathcal{P}} + \pi_{\mathcal{P}}X_{\mathcal{P}} + \lambda_{\mathcal{P}}X_{\mathcal{P}})\right].$$

Computing the difference we get to:

$$A - B = \mathbb{E}\left[\pi_{\mathcal{Q}}(\pi_{\mathcal{P}}X_{\mathcal{P}} - c_{\mathcal{P}}) + (\lambda_{\mathcal{P}} + \pi_{\mathcal{P}})(c_{\mathcal{Q}} - \pi_{\mathcal{Q}}X_{\mathcal{Q}}) + (\lambda_{\mathcal{Q}} + \pi_{\mathcal{Q}})\lambda_{\mathcal{P}}X_{\mathcal{P}} - \lambda_{\mathcal{P}}\lambda_{\mathcal{Q}}(X_{\mathcal{P}} \vee X_{\mathcal{Q}})\right].$$

At this point we plug in the definition of reservation values using the independence of the two prefixes (they are independent because the strategies are designed in such a way that the stopping times on different prefixes are independent from the realization of the other one, given that they are both played):

$$\mathbb{E}[c_{\mathcal{P}}] = \mathbb{E}[(X_{\mathcal{P}} - z_{\mathcal{P}})_+] = \mathbb{E}[(X_{\mathcal{P}} - z_{\mathcal{P}})\pi_{\mathcal{P}}]$$
$$\mathbb{E}[c_{\mathcal{Q}}] \geq \mathbb{E}[(X_{\mathcal{Q}} - z_{\mathcal{Q}})\pi_{\mathcal{Q}} + (X_{\mathcal{Q}} - z_{\mathcal{Q}})\lambda_{\mathcal{Q}}] \qquad\qquad (\mathbb{I}_{X_{\mathcal{Q}} \geq z_{\mathcal{Q}}} \geq \pi_{\mathcal{Q}} + \lambda_{\mathcal{Q}}).$$

So we get:

$$A - B \geq \mathbb{E}\left[\pi_{\mathcal{Q}}\pi_{\mathcal{P}}(z_{\mathcal{P}} - z_{\mathcal{Q}}) - \lambda_{\mathcal{P}}\pi_{\mathcal{Q}}z_{\mathcal{Q}} + \lambda_{\mathcal{Q}}(\lambda_{\mathcal{P}} + \pi_{\mathcal{P}})(X_{\mathcal{Q}} - z_{\mathcal{Q}})\right]$$
$$+ \mathbb{E}\left[(\pi_{\mathcal{Q}} + \lambda_{\mathcal{Q}})\lambda_{\mathcal{P}}X_{\mathcal{P}} - \lambda_{\mathcal{P}}\lambda_{\mathcal{Q}}(X_{\mathcal{P}} \vee X_{\mathcal{Q}})\right]$$
$$\geq \mathbb{E}\left[\lambda_{\mathcal{P}}\pi_{\mathcal{Q}}(X_{\mathcal{P}} - z_{\mathcal{Q}}) + \lambda_{\mathcal{Q}}\pi_{\mathcal{P}}(X_{\mathcal{Q}} - z_{\mathcal{Q}}) + \lambda_{\mathcal{P}}\lambda_{\mathcal{Q}}(X_{\mathcal{P}} + X_{\mathcal{Q}} - z_{\mathcal{Q}} - X_{\mathcal{Q}} \vee X_{\mathcal{P}})\right],$$

where in the last inequality we used the fact that $z_\mathcal{P} \geqslant z_\mathcal{Q}$. Recalling the definitions of the λ's and π's, the only term that is not clearly positive is the rightmost term. But $X_\mathcal{P} + X_\mathcal{Q} - X_\mathcal{P} \vee X_\mathcal{Q} = X_\mathcal{P} \wedge X_\mathcal{Q}$, which is greater than $z_\mathcal{Q}$, unless $\lambda_\mathcal{P} \lambda_\mathcal{Q} = 0$. Thus, we conclude $A - B \geqslant 0$, as desired.

Case II: If $z_1^1 > y \geqslant z_1^i$, then consider the following modified instance: decrease the cost c_1^i of box b_1^i in such a way that the new z_1^i value now lies (strictly) between z_1^1 and y. Let Δ denote this change in cost. For the sake of mathematical analysis, we will allow negative cost. Denote as $\phi'(A)$, $\phi(B)$, and $\phi'(B)$ the expected performances of STRATEGY A on the modified instance, STRATEGY B on the original instance, and STRATEGY B on the modified instance, respectively. Define STRATEGY C as making the same decisions as A would in the modified instance — as a function of the observed values — but while playing in the original instance. Let $\phi(C)$ be its expected performance.

Since B always opens b_1^i, we have that $\phi'(B) - \phi(B) = \Delta$. However, A can choose to open b_1^i, depending on the observed random variables, so $\phi'(A) - \phi(C)$ is an expectation over 0 and Δ. Thus, $\phi'(A) - \phi(C) \leqslant \phi'(B) - \phi(B)$. However, we have shown in Case I that $\phi'(A) \geqslant \phi'(B)$, and so this implies that $\phi(C) \geqslant \phi(B)$. Since C is suboptimal, this implies that B is suboptimal, as desired. $\square$

3.3 Tree-Constrained Pandora's Box Problem

In the last subsection we have solved the problem on multiple parallel lines. The main difficulty lied in the interplay between the different lines in the optimal strategy. This was solved by proving that the lines can be divided into *macroboxes* which behave like single boxes. The work of the previous section is in fact enough to prove that the Generalized Pandora's Rule is optimal even for trees.

We begin by extending the definition of reservation values beyond lines. One might naturally try to extend stopping times to a more general exploration rule. Instead, we observe that if there is an optimal threshold strategy on a subtree, then it is equivalent to view it as a line constraint, following the threshold ordering. This highlights the power of the concept of macro-boxes: they are not only a feature of the analysis, but enable us to decompose a tree into a line.

This requires the introduction of the function $\text{MERGE}(\mathcal{L}_1, \mathcal{L}_2)$ which takes as input two lines along with the information on the reservation values of their boxes, and outputs the line obtained by their merging according to decreasing reservation values, maintaining the relative orders of boxes in the same line. With this in mind, the algorithm is presented formally in Algorithm 2.

Having defined the Generalized Reservation prices for the tree case, we show the following:

Theorem 7. *The Generalized Pandora's Rule is optimal for the Tree-Constrained Pandora's Box Problem and can be computed in polynomial time and space.*

Proof. We will again do this by induction on the number of un-opened boxes. If there is only one box, everything follows as before. If we are in a state where there are multiple available subtrees to continue along, let $T_1, \ldots, T_k$ be these subtrees, and $b_1, \ldots, b_k$ be their respective roots. By induction, after opening any root box b_i, there is an ordering $\prec_i$ induced on the remaining boxes in all subtrees, for the optimal strategy to explore. Observe that restricted to any subtree T_j, the nodes of T_j are ordered the same in all $\prec_i$ orderings, including $\prec_j$. Thus, we may define $\mathcal{L}_j$ as the line that represents this common ordering of the nodes of T_j; it is clear by induction that the multi-tree problem on $T_1, \ldots, T_k$ is no more profitable than the multi-lines problem on $\mathcal{L}_1, \ldots, \mathcal{L}_k$. Furthermore, the reservation value for b_j at the head of T_j is the same as the reservation value for entering $\mathcal{L}_j$, again by induction. Thus, this theorem is a corollary of the previous. $\square$

Remark 2. *It is clear from the previous proof that the Generalized Pandora's Rule can be easily applied to Forest-Constrained instances. As we have already mentioned, it is sufficient to add a dummy root with no cost and no reward pointing to the roots of the actual trees in the forest to recover an equivalent Tree-Constrained Pandora's Box Problem.*

ALGORITHM 2: Pandora's Rule for Tree

Data: Distributions of the random rewards, box costs and tree constraint T.

Initialize queue *leaves*;

for $i \leftarrow 1$ **to** n **do**
    if *box b_i is a leaf* **then**
        initialize line $\mathcal{L}_i = [b_i]$ and enqueue b_i in *leaves*
    end
end

while *leaves is not the empty queue* **do**
    dequeue box b_ℓ from *leaves*;
    Initialize an empty line $\mathcal{L}$
    for b_j *in* CHILDREN(b_ℓ) **do**
        $\mathcal{L} \leftarrow$ MERGE$(\mathcal{L}, \mathcal{L}_j)$;   /* Taking reservation values into account */
    end
    $z_\ell \leftarrow$ COMPUTETHRESHOLD$(b_\ell, \mathcal{L})$ and $\mathcal{L}_\ell \leftarrow [b_j] + \mathcal{L}$
    if $\ell \neq 1$ *and z has been computed for all the children of* PARENT(b_ℓ) **then**
        enqueue PARENT(b_ℓ) in *leaves*
    end
end

return the generalized reservation values z_j, and the linearized tree $\mathcal{L}_1$

---

4   Adaptivity Gaps and Approximation Beyond Tree Constraints

In the previous section, we sought to design exactly optimal policies, and required exactly comparing the performance of alternative strategies. As we will see in Section 5, we can not hope to do so for more general constraints, as the problem becomes NP-hard to approximate. For this reason, we seek instead to find approximately optimal solutions. We present in this section approximation algorithms for some cases of the Pandora's Box Problem with order constraints. Following recent literature on stochastic probing [3, 11, 21, 22], we will go through an *adaptivity gap* route, arguing that for any adaptive strategy, there exists a non-adaptive strategy — *i.e.* pre-computing a fixed set and opening it obliviously — which approximates its performance. Therefore, the optimal non-adaptive strategy is a good approximation of the optimal adaptive strategy.

However, as our setting is very broad, and captures many of the complexities of stochastic submodular optimization, it is not likely that an optimal non-adaptive set will be easy to find. Instead, following an approach similar to [1], we give a single adaptive strategy which performs, in expectation, better than *any* fixed set, and yields therefore a good approximation for the optimal adaptive strategy. We begin with the adaptivity gaps. Recall the statement of Theorem 3:

Theorem 3. *Consider the Pandora's Box Problem with constraints modeled by some prefix closed family — a generalization of order constraints, defined in Section 2, then for every adaptive strategy π, there exists a non-adaptive strategy, i.e. a feasible set S, such that the following holds:*

$$\mathbb{E}\left[\max_{i \in S} X_i - \sum_{i \in S} c_i\right] \geqslant \tfrac{1}{2}\mathbb{E}\left[\max_{i \in S(\pi)} X_i\right] - \mathbb{E}\left[\sum_{i \in S(\pi)} c_i\right] \tag{1}$$

The proof of this theorem closely follows [3], but it is short, and so we include it here for completeness. As Bradac et al. [3] gives the proof in a more general setting, this proof should in theory extend to other objective functions beyond the max-of-all-entries objective that we have been using.

Proof. The proof is a relatively simple, but clever, idea introduced in the adaptivity gap upper-bound of [3]. The idea is to show that if we choose a set at random, according to the distribution induced by $S(\pi^*)$ from the randomness on the rewards, then this set will perform well in expectation over both the random set,

and the random rewards. It follows that there must exist some set which performs at least as well as this in expectation.

Formally, we wish to show that if we randomly sample the value of each box twice, choose optimal boxes adaptively for one of the samples, but measure revenue using the other samples, we lose only a factor 2 in the expectation of the $\max_{i \in S} X_i$ term. As for the $\sum_{i \in S} c_i$ term, we are opening the same set, so they cost the same. Note that this considers only feasible sets S, by definition.

To this end, let $X_1, \ldots, X_n$ be the random payoff values of the boxes, and let $Z_1, \ldots, Z_n$ be respectively identically distributed copies of the X_i's, sampled independently. Fix an optimal adaptive strategy π, and let $\pi(S, y) \in [n]$ denote the choice of the next box to open after having opened S, and observing largest value y. Let $\mathbb{S}(\pi, X | S, y)$ be the (random) final set that π opens when it chooses to terminate, if it starts with set S and total y. We denote

$$\mu_Z(S, y, y') := \mathbb{E}\left[(-y' + \max\{Z_i : i \in \mathbb{S}(\pi, X | S, y), \ i \notin S\})_+\right]$$

the expected future gain when playing according to the X_i values starting in state (S, y), but measuring revenue with the Z_i's from state (S, y'). Note that $\mu_X(\varnothing, 0, 0)$ is the expected revenue of playing according to the adaptive strategy, and $\mu_Z(\varnothing, 0, 0)$ is the expected revenue of randomly picking a set according to the Z_i's.

We wish to show $\mu_X(S, y, y') \leqslant 2\mu_Z(S, y, y')$, by induction on the set S, as it ranges over all subsets, in decreasing order of cardinality. Note that if (S, y) is such that the policy π will choose to terminate, then both values are $y - y'$. Otherwise, fix S, y, and y', and let $p := \pi(S, y)$. We have

$$\begin{aligned}
\mu_X(S, y, y') &= \mathbb{E}\left[(X_p - y')_+ + \mu(S + p, \ y \vee X_p, \ y' \vee X_p)\right] \\
&\leqslant \mathbb{E}\left[((X_p \vee Z_p) - y')_+ + \mu(S + p, \ y \vee X_p, \ y' \vee (X_p \vee Z_p))\right] \\
&\leqslant \mathbb{E}\left[(X_p - y')_+ + (Z_p - y')_+ + \mu(S + p, \ y \vee X_p, \ y' \vee (X_p \vee Z_p))\right] \\
&= \mathbb{E}\left[2(Z_p - y')_+ + \mu(S + p, \ y \vee X_p, \ y' \vee (X_p \vee Z_p))\right]
\end{aligned}$$

Where the first inequality asserts that earning more up front an only help, and the last equality holds by linearity of expectation and the identical distributions of X and Z. Furthermore,

$$\begin{aligned}
\mu_Z(S, y, y') &= \mathbb{E}\left[(Z_p - y')_+ + \mu_Z(S + p, \ y \vee X_p, \ y' \vee Z_p)\right] \\
&\geqslant \mathbb{E}\left[(Z_p - y')_+ + \mu_Z(S + p, \ y \vee X_p, \ y' \vee (X_p \vee Z_p))\right]
\end{aligned}$$

Since $\mu(S, y, y')$ is non-increasing in y'. By linearity of expectation, and by induction on S, we get $\mu_Z(\varnothing, 0, 0) \geqslant \frac{1}{2}\mu_X(\varnothing, 0, 0)$, as desired. $\qquad\square$

With this result in hand, it remains to show that we can develop an *adaptive* strategy which performs at least as well as every non-adaptive strategy, in expectation. We will take advantage of the fact that we are working in a tree-constraint, and label the boxes with a pre-order of the nodes of the tree. We will denote the index of box b as i_b. Recall that, by the properties of a pre-order, we have that for all b, if b has k descendants in the tree, then the descendants of b are exactly those boxes indexed by $i_b + 1, i_b + 2, \ldots, i_b + k$. This allows us to keep track of which boxes can legally be opened if we choose to not open b, since we may simply jump ahead in the pre-order.

We wish to use this fact to design a simple dynamic program computing the best adaptive strategy among all which only consider boxes following the pre-order. The pre-order allows us to use an index in the order to store the tree-constraint information on S, but it remains to efficiently encode information regarding the matroid constraint. To this end, we define here a characterization of all constraints with "oblivious feasibility oracles":

Definition 4. *A constraint on the feasible sets S of boxes to open is said to have an* oblivious feasibility oracle *if it is characterized by a set function $D(S)$ with the following properties:*

1. *$\{D(S) : S \subseteq [n]\}$ is supported on polynomial in n values,*

2. *For any S and $u \notin S$, $D(S+u)$ is efficiently computable knowing only $D(S)$ and u, and*

3. *For any S, it can be efficiently determined whether S is feasible knowing only $D(S)$.*

To illustrate this notion, we take as an example a *generalized knapsack constraint*, where every box b is assigned a vector $\boldsymbol{w}_b \in \mathbb{Z}_+^d$, and we have a capacity vector $\boldsymbol{m} \in \mathbb{Z}_+^d$. Here d is a constant. A set S is feasible if $\sum_{b \in S} \boldsymbol{w}_b \leq \boldsymbol{m}$, taken componentwise. The function $D(S)$ is simply $\sum_{b \in S} \boldsymbol{w}_b$, and we require that the entries of $\boldsymbol{m}$ be polynomial in n.

Note that this generalized knapsack constraint includes, as a special case, knapsack constraints, cardinality constraints, and even partition matroids with $O(1)$ partitions.

We will define the function $\Psi(i, y, D)$ recursively below, which denotes the expected revenue if we start at position i in the sequence, having already collected y, with feasibility oracle value D. As a base case, $\Psi = 0$ when $D(S) = D$ implies S is not feasible, and $\Psi = y$ when $i = n + 1$. Otherwise, let $\text{NEXT}(i)$ denote the first position after i in the pre-order on the tree such that $\text{NEXT}(i)$ is not a descendent of i. Then

$$\Psi(i, y, D(S)) := \max \begin{cases} y \\ \Psi(\text{NEXT}(i), y, D(S)) \\ -c_i + \mathbb{E}\left[(X_i - y)_+\right] + \mathbb{E}\left[\Psi(i+1, \ y \vee X_i, \ D(S+i))\right] \end{cases}$$

Since the X_i's take only polynomially many values, then this function can be computed in polynomial time, by definition of D. We can also simultaneously compute the associated adaptive policy π as in Algorithm 3 below. Let Y be the set of all possible values attained by all the X_i's, and $\mathcal{D}$ be all possible values attained by $D(S)$. The θ function returned by the algorithm determines the policy: if we are in state (S, y), and the max index of an element in S is i, then $\pi(S, y) := \theta(i, y)$.

ALGORITHM 3: Approximately Optimal Adaptive Strategy

Data: Pre-ordering $b_1, \ldots, b_n$, Oblivious feasibility oracle D, box costs, and random payoffs.

for $i \leftarrow n + 1$ **to** 1 **do**
    for $y \in Y$, $D \in \mathcal{D}$ **do**
        if D *is infeasible* **then** $\Psi(i, y, D) \leftarrow 0$;
        else if $i = n + 1$ **then** $\theta(n, y) \leftarrow \text{TERMINATE}$ **and** $\Psi(n + 1, y, D) \leftarrow y$;
        else
            $\text{OPEN} \leftarrow -c_i + \mathbb{E}\left[(X_i - y)_+\right] + \mathbb{E}\left[\Psi(i+1, \ y \vee X_i, \ D(S+i))\right]$
            $\text{SKIP} \leftarrow \Psi(\text{NEXT}(i), y, D)$
            $\Psi(i, y, D) \leftarrow \max\{y, \text{OPEN}, \text{SKIP}\}$
            if $\max\{y, \text{OPEN}, \text{SKIP}\} = y$ **then** $\theta(i - 1, y) \leftarrow \text{TERMINATE}$;
            if $\max\{y, \text{OPEN}, \text{SKIP}\} = \text{OPEN}$ **then** $\theta(i - 1, y) \leftarrow i$;
            if $\max\{y, \text{OPEN}, \text{SKIP}\} = \text{SKIP}$ **then** $\theta(i - 1, y) \leftarrow \theta(\text{NEXT}(i), y)$;
        end
    end
end

Claim 4. *The strategy returned by Algorithm 3 is at least as good as any non-adaptive strategy.*

Proof. This can be seen by induction on the i variable of the dynamic program. Let $S_{-j} := S \cap \{n - j, \ldots, n\}$. We wish to show that $\Psi(n - j, y, D(S \backslash S_{-j})) \geq \mathbb{E}\left[\max_{i \in S_{-j}} X_i - c(S_{-j})\right]$ for all y, by induction on j. Note that for $j = 0$, both values are equal to the revenue of set S. For $j > 0$, regardless of whether $n - j + 1 \in S$, Ψ takes the max over including it and not including it, and by induction, the following Φ term performs better than $S_{-(j-1)}$ in expectation. $\square$

Combining Claim 4 with Theorem 3 gives us the following result:

Theorem 8. *For the Tree-Constrained Pandora's Box Problem augmented with oblivious-feasibility-oracle matroid constraints, we can efficiently compute a policy $\hat{\pi}$ such that for any π,*

$$\mathbb{E}\left[\max_{i\in S(\hat{\pi})} X_i - \sum_{i\in S(\hat{\pi})} c_i\right] \geqslant \tfrac{1}{2}\mathbb{E}\left[\max_{i\in S(\pi)} X_i\right] - \mathbb{E}\left[\sum_{i\in S(\pi)} c_i\right]$$

5 Impossibility and Hardness Results

In this Section, we show the impossibility results outlined in Section 2. We first show that, when the precedence graph is not a tree, then there may not exist an optimal strategy which has a threshold structure. We then show the approximation hardness of solving the Pandora's Box Problem with both general order constraints and $\mathcal{F} = \mathcal{F}_T \cap \mathcal{I}_\mathcal{M}$ where T is a tree and $\mathcal{M}$ is a matroid.

The Pandora's Box Problem with order constraints where constraints are given by a DAG G , requires that a box only be opened once one at least one of its in-neighbours in G is open.

5.1 Suboptimality of Threshold Strategies

Theorem 9. *The Pandora's Box Problem with order constraints need not admit an optimal threshold strategy, when the constraint graph is not a tree. Moreover the same holds for constraints $\mathcal{F}_T \cap \mathcal{I}_\mathcal{M}$, where T is a tree and $\mathcal{M}$ a matroid.*

Proof. Consider graph (a) in Figure 1 with the following parameters:

$$X_A = \begin{cases} 2.5 & \text{w.p. } \tfrac{1}{2} \\ 0 & \text{w.p. } \tfrac{1}{2} \end{cases}, \quad X_B = 2, \quad X_C = \begin{cases} 3 & \text{w.p. } \tfrac{1}{2} \\ 0 & \text{w.p. } \tfrac{1}{2} \end{cases}, \quad X_D = \begin{cases} 6 & \text{w.p. } \tfrac{1}{2} \\ 0 & \text{w.p. } \tfrac{1}{2} \end{cases},$$

$$c_A = 0, \qquad\qquad c_B = 1, \quad c_C = 1 - \tfrac{\varepsilon}{2}, \qquad c_D = 0.$$

For $\varepsilon \in [\tfrac{5}{4}, 2]$ it can be shown that it is optimal to start the exploration of the graph from A, then, depending on the realization of X_A it is optimal to open B (and then D) or to open C (then D and then possibly B). If we now consider an instance of the Pandora's Box Problem with order constraints on $\mathcal{F}_T \cap \mathcal{I}_\mathcal{M}$ where $\mathcal{I}_\mathcal{M}$ is all subsets of cardinality 4, and T is given by (b) in Figure 1, with boxes A, B and C and two copies E and F of D, then we inherit the results from (a). $\qquad\square$
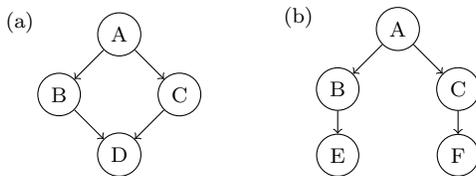

Figure 1: The order of optimal adaptive exploration is not fixed

5.2 Hardness of Approximation

As previously mentioned, we wish to show that it is NP-hard to approximate an optimal strategy for general order constraints and $\mathcal{F} = \mathcal{F}_T \cap \mathcal{I}_\mathcal{M}$ where T is a tree and $\mathcal{I}_\mathcal{M}$ are the independent sets of a matroid $\mathcal{M}$. Formally, we prove it is NP-hard to design a policy with approximately optimal rewards, for some constant. Approximation is taken in the sense of the previous section.

Theorem 10. *It is NP-hard to approximate within 0.9997 the optimal strategy to the Pandora's Box Problem with DAG constraints. It is sufficient for the DAG to have depth 2 and fan-in 3.*

We will be reducing from the problem of finding a minimum vertex cover on cubic graphs, which is known to be hard to approximate. We argue here that, since the constraint graph G has depth 2 and fan-in 3, this also implies hardness for tree-and-matroid constraints, $\mathcal{F}_T \cap \mathcal{I}$.

Corollary 1. *It is NP-hard to find the optimal strategy to the Pandora's Box Problem with order constraints with constraint $\mathcal{F} = \mathcal{F}_G \cap \mathcal{I}_\mathcal{M}$ where $\mathcal{M}$ can be any matroid on B even if G is restricted to be a tree. It suffices for $\mathcal{M}$ to be a partition matroid.*

Proof. The constraint graph G from Theorem 10 consists of sources and sinks, such that every sink had exactly 3 sources as its parents. Replace each sink with 3 identical copies (including costs and rewards), assign one to each parent, and restrict that at most one copy of each is opened. This is exactly a partition matroid constraint, and the resulting graph is a forest of depth-2 trees, and is equivalent in terms of exploration costs and rewards to the constraint graph G. $\square$

Finally, the proof of the Theorem is given below.

Proof of Theorem 10. It is known that it is NP-hard to approximate the minimum vertex cover of a cubic graph within a factor of $\approx 1.0012 =: 1 + \epsilon_0$ [4, 19]. Let $G = (V, E)$ be a hard-to-approximate instance, and let $n := |V|$, and $m := |E|$. Let α be such that the optimal vertex cover has size αm. Observe, since G is cubic, that $m := \frac{3}{2}n$, and $\alpha \geqslant \frac{1}{3}$. Furthermore, any greedy independent set must have at least $\frac{1}{3+1}n$ nodes, which implies that its complement is a vertex cover of size at most $\frac{3}{4}n = \frac{1}{2}m$. Thus, $\alpha \in [\frac{1}{3}, \frac{1}{2}]$.

We construct, now, the constraint graph D. The nodes of D will be labelled by $V \cup E$, where the V nodes will be the sources of the DAG, each having cost 1 and reward 0, and the E nodes will be the sinks of the DAG, each having cost 0 and reward βm with probability $\frac{c}{m}$ and 0 otherwise, for constants $\beta, c > 0$ which we will choose later. There is an edge connecting any vertex-box v to each edge-box e such that e in incident to v. Since G is cubic, this implies that D has depth 2 and fan-in 3, as required in the theorem statement.

Any optimal strategy must take the following form: (1) Fix an ordering on the boxes labelled by V, (2) Pay to open the next vertex-box in the order, then reveal the $\leqslant 3$ unopened edge-boxes which it reveals. (3) Repeat until the reward has been collected. Suppose that the i-th vertex-box we pay for allows us to open $0 \leqslant n_i \leqslant 3$ new edge-boxes, and $N_i := \sum_{j=1}^{i-1} n_j$. Then the expected max reward will be $\beta m \cdot (1 - (1 - \frac{c}{m})^m)$, and the expected cost will be

$$\mathbb{E}\left[\# V \text{ boxes opened}\right] = \sum_{i=1}^{n} \mathbb{P}\left(\text{opening} \geqslant i \text{ boxes}\right) = \sum_{i=1}^{n} (1 - \tfrac{c}{m})^{N_i}$$

Observe, without loss of generality, $n_{i-1} \geqslant n_i$ for all i, as swapping the $(i-1)$-st and i-th boxes will only increase N_i and leave N_{i+1} and onwards unchanged, reducing the expected cost. Thus, in any fixed order, after this swapping, there must exist numbers k_3, k_2, and k_1, such that

$$n_1 = n_2 = \cdots = n_{k_3} = 3, \quad n_{k_3+1} = n_{k_3+2} = \cdots = n_{k_3+k_2} = 2, \quad n_{k_3+k_2+1} = 0 \cdots = n_{k_3+k_2+k_1} = 1$$

Note that $3k_3 + 2k_2 + k_1 = m$, and that the vertex cover has size $k_3 + k_2 + k_1$. Setting $r = (1 - \frac{c}{m})$, we have that the expected cost becomes

$$\sum_{i=1}^{k_3} r^{3i} + \sum_{i=1}^{k_2} r^{3k_3+2i} + \sum_{i=1}^{k_1} r^{3k_3+2k_2+i} = \tfrac{r^3}{1-r^3}\left(1 - r^{3k_3}\right) + \tfrac{r^2}{1-r^2}\left(r^{3k_3} - r^{3k_3+2k_2}\right) + \tfrac{r}{1-r}\left(r^{3k_3+2k_2} - r^m\right)$$

In the remainder of the proof, we will bound the values of k_3, k_2, and k_1, for optimal and sub-optimal vertex covers, and show that the difference in expected cost is at least a constant factor of the expected reward. Since it is NP-hard to approximate the vertex cover, this will imply that is it NP-hard to approximate the optimal strategy for the Pandora's Box Problem on D.

Let S^* be an optimal vertex cover of size αm, and let S' be any vertex cover of size $\geqslant (1 + \epsilon_0)\alpha m$. Let k_3^*, k_2^*, and k_1^*, be as above for the set S^*, and k_3', k_2', and k_1' be similarly for S'. We wish to lower-bound

the cost of opening S^*, and upper-bound the cost of opening S', by bounding the possible values of the k^*'s and k''s obtained by sub- and super-optimal orderings, respectively. Note that we can trade off $k_3 + k_1$ for $2k_2$ to increase the expected cost, and vice versa. Since $\alpha \leqslant \frac{1}{2}$, then for S^*, it will suffice to assume $k_1^* = 0$, and increase k_3^* as α approaches $\frac{1}{3}$. For S', it will suffice to assume $k_2' = 0$. With the constraints on the k^*'s and the k''s, this gives

$$\begin{cases} k_3^* = (1 - 2\alpha)m \\ k_2^* = (3\alpha - 1)m \end{cases} \qquad \begin{cases} k_3' = \frac{1}{2}(1 - (1 + \epsilon_0)\alpha)m \\ k_1' = \frac{1}{2}(3(1 + \epsilon_0)\alpha - 1)m \end{cases}$$

Furthermore, in the expected cost expression above, we get

$$\mathbb{E}\left[\text{cost}(S^*)\right] \leqslant \frac{r^3}{1-r^3}\left(1 - r^{3k_3^*}\right) + \frac{r^2}{1-r^2}\left(r^{3k_3^*} - r^m\right) = \frac{r^3}{1-r^3} - \frac{r^2}{1-r^2}(r^m) + r^{3k_3^*}\left[\frac{r^2}{1-r^2} - \frac{r^3}{1-r^3}\right] \tag{3}$$

$$\mathbb{E}\left[\text{cost}(S')\right] \geqslant \frac{r^3}{1-r^3}\left(1 - r^{3k_3'}\right) + \frac{r}{1-r}\left(r^{3k_3'} - r^m\right) = \frac{r^3}{1-r^3} - \frac{r}{1-r}(r^m) + r^{3k_3'}\left[\frac{r}{1-r} - \frac{r^3}{1-r^3}\right] \tag{4}$$

Combining (3) and (4), the difference $\Delta := \mathbb{E}\left[\text{cost}(S') - \text{cost}(S^*)\right]$ is at least

$$\Delta \geqslant -r^m\left[\frac{r}{1-r} - \frac{r^2}{1-r^2}\right] - \frac{r^3}{1-r^3}\left(r^{3k_3'} - r^{3k_3^*}\right) + \frac{r}{1-r}r^{3k_3'} - \frac{r^2}{1-r^2}r^{3k_3^*}$$

$$= \left(r^{3k_3'} - r^{3k_3^*}\right)\left[\frac{r^2}{1-r^2} - \frac{r^3}{1-r^3}\right] + \left[\frac{r}{1-r^2}\right]\left(r^{3k_3'} - r^m\right)$$

Recalling the values of k_3^* and k_3' above, expanding $r = 1 - \frac{c}{m}$, and first taking MacLaurin series around "$\frac{c}{m}$"$= 0$ for the terms in square brackets, then Taylor series for the terms in round brackets, we have

$$\Delta \geqslant \left(\frac{m}{6c} + O(1)\right)\left((1 - \frac{c}{m})^{3k_3'} - (1 - \frac{c}{m})^{3k_3*}\right) + \left(\frac{m}{2c} + O(1)\right)\left((1 - \frac{c}{m})^{3k_3'} - (1 - \frac{c}{m})^m\right)$$

$$= \frac{m}{6c}\left(e^{-3c(1-\alpha-\epsilon_0\alpha)/2} - e^{-3c(1-2\alpha)}\right) + \frac{m}{2c}\left(e^{-3c(1-\alpha-\epsilon_0\alpha)/2} - e^{-c}\right) + O(1)$$

$$= \frac{m}{6c}\left(4e^{-3c(1-\alpha-\epsilon_0\alpha)/2} - 3e^{-c} - e^{-3c(1-2\alpha)}\right) + O(1) \tag{5}$$

Setting $c = (2\epsilon_0)/(3\alpha)$, and for convenience, denoting $A := 1/\alpha$, we have

$$\Delta \geqslant \frac{m}{4A\epsilon_0}\left(4e^{-\epsilon_0(A-1-\epsilon_0)} - 3e^{-c} - e^{-3c(1-2\alpha)}\right) + O(1) \tag{6}$$

Recalling that $A \in [2,3]$. For $\epsilon_0 = 0.0012$ as in [4, 19], it can be shown that the function is non-increasing in A on its domain, and plugging $A = 3$, we numerically have[2] $\Delta \geqslant 0.000399 \cdot m$.

It remains to determine the ratio of the difference in expected cost to the expected reward. Recall that we have set the reward to be βm with probability $\frac{c}{m}$, and 0 otherwise. Since it costs 1 to open a box, and we wish to ensure that even when there is a single edge-box remaining, it is in our interest to open the box, we must set $\beta = \frac{1}{c}$. Recall, then, that the expected reward will be

$$\beta m \cdot \left(1 - (1 - \frac{c}{m})^m\right) = m \cdot \frac{1}{c}(1 - e^{-c} + O(\frac{1}{m})) = m + O(1)$$

Thus, an approximation for the Pandora's Box Problem which additively approximates the cost within a 1.00039 factor of the revenue implies an approximation algorithm for vertex cover on cubic graphs within a factor of < 1.0012, which is not possible unless $P = NP$. This concludes the proof with a multiplicative constant of $1 - 0.00039 < 0.9997$ in the sense of Definition 1. $\qquad\square$

---

[2] More generally, taking the second derivative in ϵ_0 suffices to show that the right hand side is strictly convex in ϵ_0, since $A \in [2,3]$, and its derivative is 0 when $\epsilon_0 = 0$. This ensures the constant is a positive function of ϵ_0.

6   Conclusion and Further Directions

We have shown that solving the Pandora's Box Problem with order constraints admits an efficiently computable optimal solution, for tree-like order constraints. We further showed that unless P=NP, there is no PTAS when the constraints are slightly generalized, and complement this result by showing an approximation algorithm for oblivious matroid constraints on top of tree-like precedence constraints. This latter result was shown by upper-bounding the adaptivity gap, and giving methods for beating optimal non-adaptive strategies. It is clear then that extending this problem to more general constraints, or even more general objective functions, can be done by giving more general approximation algorithms for the non-adaptive problem.

References

[1] Aris Anagnostopoulos, Ilan Reuven Cohen, Stefano Leonardi, and Jakub Lacki. Stochastic graph exploration. In *ICALP*, volume 132 of *LIPIcs*, pages 136:1–136:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[2] Hedyeh Beyhaghi and Robert Kleinberg. Pandora's problem with nonobligatory inspection. In *EC*, pages 131–132. ACM, 2019.

[3] Domagoj Bradac, Sahil Singla, and Goran Zuzic. (near) optimal adaptivity gaps for stochastic multi-value probing. In *APPROX-RANDOM*, volume 145 of *LIPIcs*, pages 49:1–49:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[4] Andrea E. F. Clementi and Luca Trevisan. Improved non-approximability results for minimum vertex cover with density constraints. *Theor. Comput. Sci.*, 225(1-2):113–128, 1999.

[5] Laura Doval. Whether or not to open pandora's box. *J. Econ. Theory*, 175:127–158, 2018.

[6] Ioana Dumitriu, Prasad Tetali, and Peter Winkler. On playing golf with two balls. *SIAM J. Discret. Math.*, 16(4):604–615, 2003.

[7] Hossein Esfandiari, Mohammad Taghi Hajiaghayi, Brendan Lucier, and Michael Mitzenmacher. Online pandora's boxes and bandits. In *AAAI*, pages 1885–1892. AAAI Press, 2019.

[8] Moran Feldman. Guess free maximization of submodular and linear sums. In *WADS*, volume 11646 of *Lecture Notes in Computer Science*, pages 380–394. Springer, 2019.

[9] Esther Frostig and Gideon Weiss. Four proofs of gittins' multiarmed bandit theorem. *Annals OR*, 241 (1-2):127–165, 2016.

[10] Chenghao Guo, Zhiyi Huang, Zhihao Gavin Tang, and Xinzhi Zhang. Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora's problem. *CoRR*, abs/1911.11936, 2019.

[11] Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and XOS functions. In *SODA*, pages 1688–1702. SIAM, 2017.

[12] Chris Harshaw, Moran Feldman, Justin Ward, and Amin Karbasi. Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 2634–2643. PMLR, 2019.

[13] Joseph B Kadane. Quiz show problems. *Journal of Mathematical Analysis and Applications*, 27(3): 609–623, 1969.

[14] Joseph B Kadane, Herbert A Simon, et al. Optimal strategies for a class of constrained sequential problems. *The Annals of Statistics*, 5(2):237–255, 1977.

[15] Godfrey Keller and Alison Oldale. Branching bandits: a sequential search process with correlated pay-offs. *J. Econ. Theory*, 113(2):302–315, 2003.

[16] Jon M. Kleinberg and Robert Kleinberg. Delegated search approximates efficient search. In *EC*, pages 287–302. ACM, 2018.

[17] Robert D. Kleinberg, Bo Waggoner, and E. Glen Weyl. Descending price optimally coordinates search. In *EC*, pages 23–24. ACM, 2016.

[18] Wojciech Olszewski and Richard Weber. A more general pandora rule? *J. Econ. Theory*, 160:429–437, 2015.

[19] Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991.

[20] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1994.

[21] Sahil Singla. *Combinatorial Optimization Under Uncertainty: Probing and Stopping-Time Algorithms*. PhD thesis, PhD thesis, Carnegie Mellon University, 2018.

[22] Sahil Singla. The price of information in combinatorial optimization. In *SODA*, pages 2523–2532. SIAM, 2018.

[23] Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and super-modular optimization with bounded curvature. *Math. Oper. Res.*, 42(4):1197–1218, 2017.

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

[25] Gideon Weiss. Branching bandit processes. *Probab. in the Eng. and Info. Sciences*, 2(3):269–278, 1988.

[26] Martin L Weitzman. Optimal search for the best alternative. *Econometrica*, pages 641–654, 1979.

A Learning Pandora

In this section we apply the techniques in [10] to prove that a polynomial number of samples from the random variables $\{X_i\}_{i=1}^n$ is enough to solve with a good approximation the Pandora's Box Problem with order constraints generalizing to any prefix-closed constraint. Furthermore we prove that, for the Tree-Constrained Pandora's Box Problem, a linear number of samples is enough and it is tight.

The learning procedure to achieve this goal is indeed quite simple: For a fixed $\epsilon > 0$, we take an $O(\epsilon)$-grid of the interval $[0, 1]$, and for sufficiently many samples, learn the empirical distribution on these grid points. We then compute the optimal search policy using this empirical distribution. That this is an ϵ-approximation is a straightforward application of standard techniques, and a proof is given in Appendix C2 of [10]. Formally,

- For each box b, let X_b^ϵ be the random variable obtained by rounding down the reward X_b to the nearest multiple of $c \cdot \epsilon$;

- Given N *i.i.d.* samples of X_b^ϵ, let $\hat{X}_b$ be the random variable distributed according to the empirical distribution;

- Output the strategy $\hat{\pi}$ which is optimal with respect to the $\hat{X}_b$'s.

We will show that for constants $\epsilon, \delta > 0$, and a sufficiently large N depending on ϵ and δ, $\hat{\pi}$ will be an additive ϵ-approximation of the optimal policy with probability $1 - \delta$.

General Constraints. We recall that we have assumed rewards and costs are bounded in $[0, 1]$. It is an immediate corollary of Theorem 1 and 7 in [10] that, for the Pandora's Box Problem with *any* constraints, it suffices to have

$$N \geqslant C_1 \cdot \frac{n^3}{\epsilon^3} \log(\frac{n}{\epsilon\delta})$$

for some universal constant $C_1 > 0$.

Tree Constraints. We extend the techniques of [10] to show that linearly many samples are sufficient for tree constraints. The n^3 term in the previous bound comes from the fact that, when rewards and costs are bounded in $[0, 1]$, the total performance of a strategy must lie in $[-n, 1]$, requiring the ϵ value to be normalized by n. They use more specialized concentration bounds to get around this issue, which we extend to our setting.

The goal is to show that, for an optimal algorithm, the performance over time forms a *submartingale*. Equivalently, one should only open a box if, in expectation, the revenue is increasing. This is not true at face value, as it is often necessary to open bad boxes to allow us to move onto better boxes. We will use the notion of *macro-boxes* which were used in the proof of Theorem 6, where we showed the generalized Pandora's Rule was optimal for the Tree-Constrained Pandora's Box Problem.

Definition 5 (Macro-Boxes). *Let $\mathcal{L}_T = (b_1, b_2, \ldots, b_n)$ denote the optimal order of exploration given by the Generalized Pandora's Rule on T, and assume z_i is the generalized reservation value for box b_i in this order. Construct a sequence of indices as follows: $j_1 = 1$, and for all $i \geqslant 1$, j_{i+1} is the first index $j > j_i$ such that $z_j \leqslant z_{j_i}$. Then we say that the i-th macro box is given by the collection of boxes $\{b_{j_i}, b_{j_i+1}, \ldots, b_{j_{i+1}-1}\}$.*

Claim 5. *Let $j_1, \ldots$ be as in Definition 5, Let $S_i(\pi^*)$ denote the (random) set obtained by following the optimal policy π^* only until index j_i, and define the random variable*

$$M_i := \max_{\ell \in S_i(\pi^*)} X_\ell - \sum_{\ell \in S_i(\pi^*)} c_\ell.$$

Then the M_i's form a submartingale.

Proof. We must show that for all $i \geqslant 1$,

$$\mathbb{E}\left[M_{i+1} | M_i, \ldots, M_1\right] \geqslant M_i.$$

By Claim 3, and by the definition of the j_i's, we have that $z_{j_i}, \ldots, z_{j_{i+1}-1}$ remain unchanged if we truncate the sequence $\mathcal{L}_T$ to end at $b_{j_{i+1}-1}$. But in this case, M_{i+1} is simply the performance of π^* on the whole set. By definition of reservation values, we then have that $\mathbb{E}[M_{i+1}|M_i] \geqslant M_i$. $\qquad\square$

Another necessary condition for the result in [10] is the *strong monotonicity* of the problem.

Definition 6 (First-Order Stochastic Dominance). *Random vector $\boldsymbol{X}'$ stochastically dominates $\boldsymbol{X}$ if, for every component i, and every $x \in \mathbb{R}$, we have $\mathbb{P}(X_i' \geqslant x) \geqslant \mathbb{P}(X_i \geqslant x)$.*

Definition 7 (Strong Monotonicity). *A problem is* strong monotone *if for any random variable $\boldsymbol{X}$, and any random variable $\boldsymbol{X}'$ which dominates $\boldsymbol{X}$, Letting π^* be the optimal policy for the distribution on $\boldsymbol{X}$, we have that the performance of π^* over $\boldsymbol{X}'$ is at most the performance over $\boldsymbol{X}$.*

The following is a direct corollary of Appendix C3 in [10], when viewed over the macro-boxes, as we have a fixed order of exploration, and reservation prices.

Claim 6. *The Tree-Constrained Pandora's Box Problem is strongly monotone.*

These two previous claims imply that Lemma 25 in [10] apply. Thus, for the Tree-Constrained Pandora's Box Problem, it suffices to take
$$N \geqslant C_2 \cdot \tfrac{n}{\epsilon^2} \log^2(\tfrac{1}{\epsilon}) \log(\tfrac{n}{\epsilon}) \log(\tfrac{n}{\epsilon\delta})$$
for some universal constant C_2.

Lower Bounds. This latter result is tight up to poly $\log(\tfrac{n}{\epsilon\delta})$ terms: [10] show that it takes at least $\Omega(\tfrac{n}{\epsilon^2})$ samples to get the desired degree of accuracy.

B The Adaptivity Gap of the Pandora's Box Problem is unbounded

In this section we present a counterexample showing that the adaptivity gap for the Pandora's Box Problem is unbounded. This is inspired by an example found in [21].

Let $p > 0$, and consider n identical boxes with cost $c = 1 - p/2$, and reward $\frac{1}{p^2}$ with probability p^2, 0 otherwise. Since $c < 1$, the adaptive optimal strategy is to open boxes until you get the reward, which guarantees reward $\frac{1}{p^2}$ and costs $\frac{c}{p^2}$ in expectation, for a total expected revenue of $\frac{1}{p^2}(1-c) = \frac{1}{2p}$

We now consider the non-adaptive strategy which opens k boxes. It earns $\frac{1}{p^2}$ with probability $1-(1-p^2)^k$ and 0 otherwise, and pays ck. Note that $(1-p^2)^k$ is convex in k, and so $\frac{1}{p^2}(1-(1-p^2)^k)$ is concave, and thus has non-increasing derivatives. At $k = \frac{1}{p}$, we have

$$\frac{\mathrm{d}}{\mathrm{d}k}\left[-ck + \tfrac{1}{p^2}(1-(1-p^2)^k)\right] = -c - (1-p^2)^k \cdot \frac{\ln(1-p^2)}{p^2}$$

$$\leqslant -c + \frac{-(-p^2)+p^4}{p^2} \cdot (1-p^2)^k \qquad\qquad \text{if } p^2 \leqslant \tfrac{1}{2} \qquad (7)$$

$$\leqslant -c + (1+p/8)(1-p^2)^k \qquad\qquad \text{if } p \leqslant \tfrac{1}{2}$$

$$\leqslant -c + (1+p/8)e^{-p}$$

$$\leqslant -(1-p/2) + (1+p/8)(1-0.632p) \qquad\qquad (8)$$

$$\leqslant (\tfrac{1}{2} - 0.632 + 0.125)p = -0.007p < 0$$

Where (7) holds because $\ln(1+x) \geqslant x - x^2$ for $x \in [-\tfrac{1}{2}, 0]$: at $x = 0$, $\ln(1+x) = x - x^2 = 0$, and $\frac{\mathrm{d}}{\mathrm{d}x}\ln(1+x) = \frac{1}{1+x} \leqslant \frac{\mathrm{d}}{\mathrm{d}x}x - x^2 = 1 - 2x$ over the domain, by convexity of the former. (8) holds since

$e^{-p} \leqslant 1 - (1 - 1/e)p$ for $p \in [0, 1]$, again by convexity. Hence, the derivative is negative, and we conclude the optimum is attained on $1 \leqslant k < 1/p$. However, $(1 - p^2)^k \geqslant 1 - kp^2$, so

$$-ck + \tfrac{1}{p^2}(1 - (1 - p^2)^k) \leqslant -ck + \tfrac{1}{p^2}kp^2 = (1 - c)k$$

Since $1 - c = p/2$, and $k \leqslant 1/p$, this upper-bounds the total revenue by $\frac{1}{2}$.

Recalling that the adaptive strategy earned on average $1/2p$, then the ratio of the two revenues is $\frac{1}{p}$. Taking $p \to 0$, this suggests that the adaptivity gap can be arbitrarily large. $\qquad\square$

We remark that this counterexample works also for the single line constrained case as the boxes are all equal and the order is irrelevant.

C Proof of Claim 2

Recall, we have denoted the expected future reward following τ, given that the player has just opened box b_{i-1} and the largest reward that has been sampled in the past has been x, as:

$$\phi^\tau(x, i) := \mathbb{E}\left[\max\left\{x, \ \max_{j=i}^{\tau(x,i)} X_j\right\} - \sum_{j=i}^{\tau(x,i)} c_j\right],$$

and denoted $\Phi(x, i) = \max_\tau \phi^\tau(x, i) = \phi^{\tau^*(x,i)}(x, i)$, for all $0 \leqslant i \leqslant n$. Observe that Φ is a non decreasing function in its first argument.

For the first part of the Claim we are going to prove that $\forall i = 1, \ldots, n$ the functions

$$H_i(x) = \mathbb{E}\left[\left(\max_{j=i}^{\tau^*(x,i)} X_j - x\right)_+ - \sum_{j=i}^{\tau^*(x,i)} c_j\right] = \Phi(x, i) - x, \tag{9}$$

admit at least one zero each in $[0, \infty)$ and that the each solution set admits a minimum. In order to do so it is sufficient to show that the H_i are continuous and monotone non increasing: it is then straightforward to conclude, since $H_i(0) = \Phi(0, i) \geqslant 0$ and $\lim_{x \to \infty} H_i(x) = 0$.

For any two numbers $b \geqslant a \geqslant 0$ and i we have that

$$H_i(b) - H_i(a) = \Phi(b, i) - \Phi(a, i) + a - b \leqslant \mathbb{E}\left[\max_{j=i}^{\tau^*(b,i)} X_j \vee b - \max_{j=i}^{\tau^*(b,i)} X_j \vee a\right] - b + a \leqslant 0, \tag{10}$$

where in the first inequality we used the fact that the stopping time $\tau^*(b, i)$ is sub-optimal for $\phi^\tau(a, i)$. In particular (10) means that $\Phi(\cdot, i)$ is 1-Lipschitz and hence continuous. So we can claim the continuity of all the $H_i(\cdot)$ as compositions of continuous functions.

We also argue that from (9) it is clear that the H_i (and hence the solutions z_i) do not depend on the particular stopping time used in the definition (which may not be unique), but only on the optimal values $\Phi(x, i)$, which are unique.

Let's now focus on the second part, we want to prove that for every i there exists a $\tau^*(z_i, i)$ which is different from $i - 1$, as long as $z_i > 0$. We recall that we are only considering stopping times that depends deterministically on the realizations, hence the event $\tau(x, i)$ opens or not box b_i has either probability 1 or 0. Let's fix any i for which $z_i > 0$ and let $\overline{\tau}_i$ be the best strategy between all those that open box b_i, we want to prove by contradiction that

$$\delta = \Phi(z_i, i) - \phi^{\overline{\tau}_i}(z_i, i) = 0$$

So let's assume $\delta > 0$. First we note that for every $\varepsilon > 0$ such that $z_i - \varepsilon > 0$ we must have $\tau^*(z_i - \varepsilon, i) \geqslant i$, because otherwise z_i would not be minimal in the solution set of $H_i(\cdot) = 0$.

Moreover it holds that $\phi^{\tau^*(z_i - \varepsilon, i)}(z_i, i) \leqslant \phi^{\overline{\tau}_i}(z_i, i)$ by definition of $\overline{\tau}_i$. Using the Lipschitz property of $\Phi(\cdot, i)$ and $\phi^{\tau^*(z_i - \varepsilon, i)}(\cdot, i)$, we hence get the following contradiction:

$$0 < \delta \leqslant \Phi(z_i, i) - \phi^{\tau^*(z_i - \varepsilon, i)}(z_i, i) \pm \Phi(z_i - \varepsilon, i) \leqslant 2\varepsilon \to 0 \text{ for } \varepsilon \to 0. \qquad \square$$