

Pandora’s Problem with Deadlines

Ben Berger¹, Tomer Ezra², Michal Feldman^{3,4}, Federico Fusco⁵

¹ Offchain Labs, Inc.

² Simons Laufer Mathematical Sciences Institute

³ Tel Aviv University

⁴ Microsoft ILDC

⁵ Department of Computer, Control
and Management Engineering “Antonio Ruberti”,
Sapienza University of Rome

Abstract

Pandora’s problem is a fundamental model that studies optimal search under costly inspection. In the classic version, there are n boxes, each associated with a known cost and a known distribution over values. A strategy inspects the boxes sequentially and obtains a utility that equals the difference between the maximum value of an inspected box and the total inspection cost. Weitzman (1979) presented a surprisingly simple strategy that obtains the optimal expected utility.

In this work we introduce a new variant of Pandora’s problem in which every box is also associated with a publicly known *deadline*, indicating the final round by which its value may be chosen. This model captures many real-life scenarios where alternatives admit deadlines, such as candidate interviews and college admissions. Our main result is an efficient threshold-based strategy that achieves a constant approximation relative to the performance of the optimal strategy for the deadlines setting.

Introduction

Pandora’s problem, introduced in a seminal paper by Weitzman (1979), models the search for a good alternative among different options, under costly evaluation. An instance of the problem consists of n boxes, each hiding some reward drawn from a known and independent distribution. Each box also comes with a known evaluation cost.

In the original variant, the decision maker opens the boxes sequentially in any desired order (which may be adaptive), until she decides to halt. The goal is to maximize the obtained utility, which is the maximum observed value across all opened boxes, minus the sum of their costs. This captures many real-life scenarios such as interviewing job candidates and choosing colleges, where the decision-maker has to balance between the need to explore many options in order to guarantee a good result, and to minimize the total exploration cost. Remarkably, Weitzman provided a simple and efficient strategy which obtains the optimal expected utility.

The original model described above assumes that all rewards are available throughout the entire procedure. However, this may not necessarily be the case in practice, as some alternatives might become unavailable at some point in time. For instance, job candidates might eventually defer to other

employment opportunities if the recruiter took too long to offer them the job.

Motivated by such applications, in this work we introduce a variant of Pandora’s problem where each box, in addition to its cost and value distribution, is also associated with a known *deadline*, after which it is no longer available. This constraint complicates the problem significantly, as it limits the exploration orders that can be chosen by the decision maker. Furthermore, it may introduce interference between boxes that are otherwise compatible. For instance, suppose there are two favorable candidates that an employer would have liked to interview one after the other, but both of them are available only in the very next time slot. Weitzman’s solution would not work in this case, as the employer must choose only one of them to interview.

Our results. While Weitzman’s algorithm is not applicable in our model, we devise a simple strategy that attains a constant approximation to the optimal utility attainable by any algorithm.

Result 1 (see Corollary 1). There is an efficient algorithm that achieves a 0.15-approximation to the optimal expected utility in Pandora’s Problem with Deadlines.

We further show how to generalize this result to the more restricted *time slots* setting — where in addition to the exploitation deadline, each box is also associated with the set of time slots in which it is allowed to be inspected.

Result 2 (see Theorem 1). There is an efficient algorithm that achieves a 0.15-approximation to the optimal expected utility in Pandora’s Problem with Time Slots.

Our results add to a large body of work that studies Pandora’s problem under different restrictions on exploration and exploitation (see Related Work).

Technical Challenges and Techniques. Solving Pandora’s problem entails finding a balance in the well known exploration-exploitation dilemma: inspecting new boxes is costly, but may offer the opportunity to collect better reward. The presence of deadlines and time slots imposes extra constraints on the order in which boxes can be opened, as well as on the “reward collection” process. For instance, continuing the exploration process may impede the decision maker to collect some very good observed reward whose deadline is close. To address the last challenge, we focus on strategies

that always take the reward that was seen last. This allows us to treat the deadlines only as a constraint to the order of exploration, and not on the exploitation. Then, we make the crucial observation that the family of sets of boxes for which there exists a feasible exploration order forms a transversal matroid. This allows us to exploit adaptivity gap results for submodular maximization (Bradac, Singla, and Zuzic 2019), and find a feasible subset of the boxes \tilde{S} that admits the following property: there exists a strategy that commits in advance to opening only boxes in \tilde{S} without sacrificing more than a constant fraction of the optimal utility. Furthermore, \tilde{S} admits a threshold such that for any ordering of it, selecting the first box that exceeds this threshold attains a constant approximation to the optimal utility. Finally, we compute this threshold (using techniques from Esfandiari et al. 2019) and the feasible ordering which exists by definition, and this defines our strategy.

Related Work. The two works closest to ours are Singla (2018) and Esfandiari et al. (2019). Singla (2018) considers a variant of Pandora’s problem that constrains the set of boxes that can be opened and selected, and offers constant factor approximation algorithms for many natural classes of downward closed constraints. Note that these constraints apply only to the set of boxes that can be opened, and not the order in which they can be opened. Esfandiari et al. (2019) study an online version of Pandora’s problem where the boxes have to be considered online, one after the other in some predetermined order. Once a box arrives, the decision maker chooses whether to open it (and possibly halt and take its realized reward), or discard it forever. Their model is strictly subsumed by ours, as it coincides with the special case of disjoint time slots but does not capture (natural) situations where the decision maker has some freedom in choosing the order of exploration (still respecting deadline constraints).

Weitzman’s seminal paper has given rise to many natural variants of Pandora’s Problem that have been investigated over the years: non-obligatory inspection (Doval 2018; Beyhaghi and Kleinberg 2019; Beyhaghi and Cai 2023; Fu, Li, and Liu 2023), precedence constraints (Boodaghians et al. 2023), commitment constraints (Fu, Li, and Xu 2018; Segev and Singla 2021), interdependent valuations (Chawla et al. 2020, 2023), and combinatorial costs (Berger et al. 2023). Finally, Pandora’s problem has also been studied from the learning perspective, both in the sample complexity framework (Guo et al. 2021), and in online learning (Gergatsouli and Tzamos 2022; Gatmiry et al. 2022).

Model and Preliminaries

In Pandora’s problem there are n boxes, each containing a hidden random variable V_i which is drawn from a distribution X_i . These distributions are independent, non-negative and publicly known. Each box is also associated with a known cost $c_i \in \mathbb{R}_{\geq 0}$. The *reservation value* of a box i , denoted r_i , is the number that satisfies the equation $\mathbb{E}[(V_i - r_i)^+] = c_i$, where the expectation is with respect to the variables V_i , which are drawn from X_i , and we use the notation $(\cdot)^+$ to denote $\max(\cdot, 0)$. The notion of reserva-

tion value is the crucial ingredient for Weitzman’s rule: the simple greedy strategy that is optimal for the original version of Pandora’s problem (Weitzman 1979). The rule works as follows: it opens boxes in decreasing order of reservation value, and stops when the reservation value of the next box is smaller than the largest reward experienced so far (or there is no other box left).

Pandora’s Problem with Deadlines. We introduce a variant of Pandora’s problem where each box i is also associated with a known *deadline* $d_i \in \mathbb{N}$, indicating the final time point by which it can be chosen. Formally, an instance of the problem is denoted by $\mathcal{I} = (X_i, c_i, d_i)_{i \in [n]}$ ^{*}, including all available information, namely, the box distributions, costs and deadlines.

A *strategy* π for \mathcal{I} opens the boxes in a sequential manner, revealing the hidden values inside. At each round t (starting from the first round $t = 1$), the strategy may choose to open any unopened box i , or otherwise the strategy may choose to halt. We say that a box is *expired* at a given round t if its deadline has passed before that round, *i.e.*, $d_i < t$. Otherwise, we say that the box is *active*.

The utility attained is the expected difference between the largest observed value of a box that is active at the last round of the process — *i.e.*, upon halting (or 0 if no such box exists), and the sum of the costs of all opened boxes. The strategy may be adaptive, *i.e.*, the decisions at each round may depend on the values observed along the procedure. Note that the classic Pandora’s problem introduced by Weitzman (*i.e.*, without deadlines) is a special case of our setting where $d_i = n$ for every box i , but that the greedy solution based on reservation value does not work, as it does not capture the central notion of deadlines. Given a strategy π for \mathcal{I} , we introduce the following notation:

- $S(\pi)$ denotes the (random) ordered tuple of boxes opened by π .
- $V(\pi)$ denotes the maximum value of an active box observed by π :

$$V(\pi) = \max_{\substack{i \in S(\pi), \\ |S(\pi)| \leq d_i}} V_i.$$

We also refer to this value as the reward obtained by π .

- $u(\pi) = \mathbb{E}[V(\pi)] - \mathbb{E}\left[\sum_{i \in S(\pi)} c_i\right]$ denotes the expected utility (*i.e.*, value minus cost) achieved by π .

Note that the quantities defined above depend on both the given instance and the strategy. Since only active boxes can be chosen, we may assume without loss of generality that strategies inspect only boxes that are active in the corresponding round.

Pandora’s Problem with Time Slots. It is possible to further generalize the deadline model to scenarios where each box has a feasible set of rounds (not necessarily contiguous) at which it is available for exploration. This captures many real-life scenarios such as interviewing job candidates,

^{*}For any positive integer x , we denote with $[x]$ the set of the first x numbers: $[x] = \{1, \dots, x\}$

where each employee has a set of days when she is available for the interview. Formally, there are m days, and n boxes, where each box i , in addition to having a deadline d_i , a hidden value $V_i \sim X_i$, and an inspection cost c_i , has a set of feasible days $T_i \subseteq [m]$. An instance of Pandora’s Problem with Time Slots is then characterized as follows: $\mathcal{I} = (X_i, c_i, d_i, T_i)_{i \in [n]}$. Each day $t = 1, \dots, m$, the decision-maker can choose one box i such that $t \in T_i$ (i.e., the box is available for exploration on this day) or she can skip inspecting a box at that day. If the decision-maker chooses to open some box i then the cost c_i is incurred, and V_i is observed. For any strategy π it is possible to define, similarly to what is done for the deadline model, the ordered tuple of opened boxes $S(\pi)$, the maximum value $V(\pi)$, and the decision-maker’s utility u , which is the maximum observed value that did not expire minus the sum of the inspection costs. This new model is more general than the previous one, as deadlines are captured by setting $m = n$, and $T_i = [n]$ for every box i .

Efficiency. It is clear that both new versions of Pandora’s Problem can be solved exactly by exponential time algorithms via dynamic programming; therefore our goal is to find algorithms that are efficient and provide expected utility that at least approximates the one achievable by the optimal strategy. To simplify the analysis of our algorithms, we assume access to an *exp-max* oracle that, given in input a set of boxes S , outputs the expected maximum of the random variables $\min\{V_i, r_i\}$ for $i \in S$. When the support of the distributions is polynomially supported, this expected maximum can be computed directly, while in general it can be estimated via sampling (we refer to Guo et al. (2021) for further details on how sample-based discrete distributions can be used to approximate continuous ones in our setting).

Submodular Optimization and Matroid Constraints. An important ingredient for our analysis is submodular maximization subject to matroid constraints. We report here some of the basic definitions and relevant results; we refer the interested reader to Part IV of Schrijver et al. (2003) for further details.

A pair (E, \mathcal{M}) is called a matroid if E is a finite set and \mathcal{M} is a nonempty collection of subsets of S satisfying the following two properties:

- (Downward-closure) If $S \in \mathcal{M}$ and $T \subseteq S$, then $T \in \mathcal{M}$
- (Exchange property) If $S, T \in \mathcal{M}$ and $|T| < |S|$, then $T \cup \{x\} \in \mathcal{M}$, for some $x \in S \setminus T$.

Example 1 (Transversal Matroids). Starting from a bipartite graph $G = (L, R; E)$, it is possible to construct a matroid (L, \mathcal{M}) on the left nodes L as follows: $S \subseteq L$ belongs to \mathcal{M} if there exists a matching in G such that S constitutes the endpoints of the matching in L .

A set function $f : 2^E \rightarrow \mathbb{R}_{\geq 0}$ is submodular if, for any subsets $T \subseteq S \subseteq E$ and element $x \notin S$, the following inequality holds: $f(S \cup \{x\}) - f(S) \leq f(T \cup \{x\}) - f(T)$. A set function is monotone (with respect to inclusion), if for any subsets $T \subseteq S \subseteq E$ it holds that $f(T) \leq f(S)$. The state-of-the-art for monotone submodular maximization subject to a matroid constraint is an algorithm by Vondrák

(2008) which combines a continuous greedy approach with the pipage rounding technique; we report here its properties.

Lemma 1 (Vondrák 2008). For every monotone submodular function f defined on a base set E , and for every matroid (E, \mathcal{M}) and precision constant $\varepsilon \in (0, 1)$, there exists an efficient randomized algorithm that finds a set $\tilde{S} \in \mathcal{M}$ such that $\mathbb{E}[f(\tilde{S})] \geq (1 - 1/e - \varepsilon)f(S^*)$, where $S^* = \arg \max_{S \in \mathcal{M}} f(S)$.

The running time of a submodular optimization algorithm is typically measured in terms of the calls to two specific oracles: the value oracle (which returns the value $f(S)$ on given set S) and the independence oracle (which returns whether $S \subseteq E$ belongs to \mathcal{M}). The algorithm in Vondrák (2008) is efficient as it issues a number of oracle calls that is polynomial in $1/\varepsilon$ and in the cardinality of the base set E .

Set and sequence constraints. Pandora’s problem with deadlines or time slots imposes constraints on the boxes that may be chosen by an algorithm. In what follows, we introduce relaxations of the two models which are used in our analysis. These relaxations impose restrictions only on the time at which boxes can be inspected, but not on the time at which they can be chosen, as in the deadlines or time slots models. Hence, in the models presented below, the value obtained from a strategy is the maximum value of an inspected box, as in Weitzman’s original model.

We first make a distinction between a *sequence* constraint and a *set* constraint. A sequence constraint is a collection of *valid tuples* of boxes, namely, a collection of inspection orderings that a strategy may undertake. A set constraint is a collection of *valid sets* of boxes (that does not take into account the inner ordering of the boxes). We say that a sequence constraint \mathcal{F} is *prefix-closed* if for any tuple $(i_1, \dots, i_k) \in \mathcal{F}$ and any $j = 0, 1, \dots, k$ we have $(i_1, \dots, i_j) \in \mathcal{F}$. A set constraint \mathcal{F} is *downward-closed* if for any set $S \in \mathcal{F}$ and any subset $T \subseteq S$ we have $T \in \mathcal{F}$.

The main sequence constraint that we consider in this work is EXP-SEQ, where inspection times respect the boxes’ deadlines and the time slots constraints (if any). Formally, it is defined by

$$\text{EXP-SEQ} = \bigcup_{t=0}^m \{(i_1, \dots, i_t) \mid j \in T_{i_j} \cap [d_{i_j}] \forall j \in [t]\}.$$

Given a sequence (respectively, set) constraint \mathcal{F} , we say that a tuple (resp., set) of boxes S is *feasible* under \mathcal{F} if $S \in \mathcal{F}$. A strategy π is consistent with a constraint \mathcal{F} , if $S(\pi) \in \mathcal{F}$ with probability 1.[†]

We denote by EXP-SET the set constraint induced by EXP-SEQ. In other words, a set of boxes S is feasible under EXP-SET if there exists some ordering of S which is compatible with the given deadlines and possibly time slots, i.e., some ordering of S which is in EXP-SEQ. The following lemma shows that EXP-SET is a transversal matroid.

Lemma 2. EXP-SET is a transversal matroid constraint.

[†]In the case that \mathcal{F} is a set constraint, $S(\pi)$ –although technically a tuple– is interpreted as the corresponding set.

Proof. Consider the following bipartite graph, where the vertices on the left are the boxes, and the vertices on the right are the integers $1, \dots, m$. Each box i has edges to times $\{1, \dots, d_i\} \cap T_i$. A subset of boxes S is feasible under EXP-SET, (i.e., there exists an inspection order compatible with the deadlines) if and only if the sub-graph induced by S admits a matching of cardinality $|S|$. Therefore, by definition, EXP-SET is a transversal matroid. \square

Given a constraint \mathcal{F} , we denote by $\text{OPT}_{\mathcal{F}}$ the performance of an optimal strategy for Pandora’s problem with a feasibility constraint \mathcal{F} , i.e., this is the highest utility attainable by a strategy consistent with \mathcal{F} . Clearly, the TIME-SLOTS constraint induced by the Pandora’s problem with time slots is stronger than the corresponding EXP-SEQ constraint, which is in turn stronger than EXP-SET. Therefore, the following inequality follows:

$$\text{OPT}_{\text{EXP-SET}} \geq \text{OPT}_{\text{EXP-SEQ}} \geq \text{OPT}_{\text{TIME-SLOTS}} \quad (1)$$

The same holds for DEADLINE. The following example shows how Formula (1) may be strict even in the special case of deadlines.

Example 2. Consider an instance of the Pandora’s problem with deadlines composed by two boxes with deadlines $d_1 = 1$, $d_2 = 2$ and costs $c_1 = 2$, $c_2 = 1$. The first box has value $V_1 = 20$ with probability $1/5$ and value $V_1 = 10$ otherwise. The second box has value $V_2 = 20$ with probability $1/5$ and value 0 otherwise. If we look at the sequence constraint, we note that it is not possible to open the box 1 after the box 2, therefore $\text{EXP-SEQ} = \{\emptyset, (1), (2), (1, 2)\}$. On the contrary, it is possible to open the box 2 after box 1, thus the EXP-SET does not impose any actual constraint on the set of boxes that can be opened: $\text{EXP-SET} = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$. In other words, EXP-SET is equivalent to the original version of Pandora’s problem.

In the EXP-SEQ setting there is a small number of strategies to consider, as once box 2 is opened then box 1 cannot be opened afterwards due to its deadline. A straightforward calculation shows that the optimal strategy in the EXP-SEQ setting is to first open box 1, halt if $V_1 = 20$ and open box 2 otherwise. The expected cost of this strategy is $2 + 4/5 = 14/5$. The expected reward of this strategy is

$$20 \cdot \left(\frac{1}{5} + \frac{4}{5} \cdot \frac{1}{5}\right) + 10 \cdot \frac{4}{5} \cdot \frac{4}{5} = \frac{68}{5}$$

Thus, the expected utility is $\text{OPT}_{\text{EXP-SEQ}} = (68-14)/5 = 54/5$.

On the other hand, if box 2 is opened first then there is clearly no point in opening box 1 in the second round. Therefore an optimal strategy for the DEADLINE setting will open only one box, and between opening just box 1 or box 2 the former is strictly better. We thus have

$$\text{OPT}_{\text{DEADLINE}} = 20 \cdot \frac{1}{5} + 10 \cdot \frac{4}{5} - 2 = 10 < \text{OPT}_{\text{EXP-SEQ}}.$$

As for the EXP-SET setting, since the deadlines do not impose any constraints, the optimal strategy for this setting is to follow Weitzman’s rule which for this instance would be to first open box 2 and then open box 1 in case V_2 is observed to be 0. Similarly to the calculations above, the expected utility for this strategy is

$$\text{OPT}_{\text{EXP-SET}} = \frac{55}{5} > \text{OPT}_{\text{EXP-SEQ}}$$

Our Algorithm

In this section we present our main result: an efficient strategy for the Pandora’s Problem with Time Slots which yields an expected utility that is a constant factor approximation to the optimal one. As a corollary, we obtain the same approximation guarantee for the simpler problem of deadlines.

We achieve our approximation result using a *threshold strategy*. A threshold strategy for a generic instance $\mathcal{I} = (X_i, c_i, d_i, T_i)_{i \in [n]}$ of the Pandora’s Problem with Time Slots is defined by a threshold τ , and a feasible inspection ordering $\sigma = (i_1, \dots, i_k)$ for some $k \in [n]$ (by convention, we introduce dummy boxes to model the choice of “skipping turn” without opening any box). Such a strategy opens boxes according to σ and stops when $V_{i_j} \geq \tau$, or after all boxes in σ have been opened. Note, σ is feasible if any time step $j \leq k$ falls before deadline d_{i_j} and belongs to the admissible time slots T_{i_j} . Given an instance of the problem, we compute threshold τ and ordering σ in three steps:

Finding the right set of boxes. Start computing for each box i its reservation value r_i , and define the random variable $Y_i = \min\{V_i, r_i\}$. The function f that maps a set of boxes $S \subseteq [n]$ to the expected value of the maximum Y_i in S is then well defined:

$$f(S) = \mathbb{E} \left[\max_{i \in S} Y_i \right]. \quad (2)$$

It is an easy exercise to see that f is a monotone submodular set function over the set of boxes $[n]$. In the previous Section we showed that the set constraint EXP-SET is a transversal matroid over $[n]$, so it is possible to find a subset \hat{S} of the boxes that is feasible and has good value (Lemma 1). It is possible to find \hat{S} efficiently as the value oracle for the submodular function f is provided by the exp-max oracle for the underlying random variables, while the independence oracle can be obtained via a maximum matching routine on the bipartite graph that defines the transversal matroid.

Finding the right order. We know that \hat{S} is a feasible solution according to EXP-SET, therefore it is possible to find an ordering $\hat{\sigma}$ of the boxes that respects the EXP-SEQ family of constraints. The ordering $\hat{\sigma}$ can be computed efficiently by reducing the problem to finding a perfect matching in a bipartite graph (one side will have the boxes, and the other side the time slots).

Finding the right threshold. Now that we have the (non-adaptive) sequence $\hat{\sigma}$ according to which the boxes in \hat{S} will be opened, we need to find a stopping rule (in our case governed by a certain threshold $\hat{\tau}$) that balances the exploration-exploitation trade off, while taking into account the deadlines of the various boxes. This can be done by setting:

$$\hat{\tau} = \frac{1}{2} \mathbb{E} \left[\max_{i \in \hat{S}} \{Y_i\} \right] = \frac{1}{2} f(\hat{S}).$$

This choice draws inspiration from the literature on prophet inequalities, and has already been proposed in the Pandora’s setting in Esfandiari et al. (2019). Note, $\hat{\tau}$ can be computed via a single call to the exp-max oracle.

Theorem 1. For any constant $\varepsilon \in (0, 1)$, there exists a threshold strategy for the Pandora’s Problem with Time Slots that can be computed efficiently and provides a $(\frac{e-1}{4e} - \varepsilon)$ -approximation to the optimal expected utility.

Proof. We have already argued that it is possible to compute a threshold $\hat{\tau}$ and order $\hat{\sigma}$ in time that is polynomial in $1/\varepsilon$ and n ; so we only need to prove the approximation guarantee. Recall, we denote with $\text{OPT}_{\text{TIME-SLOTS}}$ the expected utility guaranteed by the optimal strategy, while $\text{OPT}_{\text{EXP-SET}}$ is the expected utility achievable when the deadlines and time slots constraints are removed from the exploitation and only limit the feasible subset of boxes that can be explored. For convenience, we denote with ALG the expected utility of the threshold strategy with $\hat{\sigma}$ and $\hat{\tau}$.

We know, by Equation (1), that the utility of the optimal strategy for the time slots problem is dominated by that achievable under the weaker EXP-SET constraints. So, as a first step in the analysis, we recall a result by Kleinberg and Kleinberg (2018) which relates $\text{OPT}_{\text{EXP-SET}}$ with a specific instance of the stochastic probing problem (Asadpour and Nazerzadeh 2016). In this latter problem, a decision-maker faces the same set of boxes with the random values Y_1, \dots, Y_n that can be opened subject to the prefix-closed constraint EXP-SET, and her utility is the maximum observed value. At each round the decision maker can (adaptively) choose one box to open, as long as the feasibility constraint is satisfied. Let π^* be an optimal strategy for the above problem, and denote by $S(\pi^*)$ the random set of variables Y_i which were probed by the algorithm, the following Lemma holds (we report the proof for completeness).

Lemma 3 (Kleinberg and Kleinberg 2018). The following inequality holds:

$$\mathbb{E} \left[\max_{i \in S(\pi^*)} \{Y_i\} \right] \geq \text{OPT}_{\text{EXP-SET}}.$$

Proof of Lemma 3. Let π be an optimal strategy for the Pandora’s problem under inspection constraints EXP-SET, i.e., π achieves utility $\text{OPT}_{\text{EXP-SET}}$. For every $i \in [n]$, let I_i be the indicator function for box i being inspected by π , and let A_i be the indicator function for box i being chosen by π . Then $\text{OPT}_{\text{EXP-SET}}$, the expected utility achieved by π equals:

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n A_i V_i - I_i c_i \right] &= \mathbb{E} \left[\sum_{i=1}^n A_i V_i - I_i \mathbb{E} [(V_i - r_i)^+] \right] \\ &= \sum_{i=1}^n \mathbb{E} [A_i V_i] - \mathbb{E} [I_i (V_i - r_i)^+] \\ &\hspace{10em} (I_i \text{ and } V_i \text{ are independent}) \\ &= \sum_{i=1}^n \mathbb{E} [A_i Y_i + A_i (V_i - r_i)^+ - I_i (V_i - r_i)^+] \\ &\leq \sum_{i=1}^n \mathbb{E} [A_i Y_i + I_i (V_i - r_i)^+ - I_i (V_i - r_i)^+] \\ &\hspace{10em} (I_i = 0 \text{ implies } A_i = 0) \\ &= \sum_{i=1}^n \mathbb{E} [A_i Y_i] \leq \mathbb{E} \left[\max_{i \in S(\pi)} \{Y_i\} \right], \end{aligned}$$

where the last inequality follows because only one box in $S(\pi)$ is chosen. To conclude the proof we note that by definition of π^* (i.e., the adaptive strategy that maximizes Y_i) we have $\mathbb{E} [\max_{i \in S(\pi)} \{Y_i\}] \leq \mathbb{E} [\max_{i \in S(\pi^*)} \{Y_i\}]$, concluding the proof. \square

Our focus has now shifted towards upper bounding this new quantity $\mathbb{E} [\max_{i \in S(\pi^*)} \{Y_i\}]$. This optimization task is extremely challenging, as one needs to compare against all *adaptive* strategies, which may change the set of boxes to open according to the realized Y_i . To overcome this issue, we lean on an elegant result by Bradac, Singla, and Zucic (2019) that holds for general monotone submodular function (we are using the function f defined in Equation (2)), which bounds the the gap in performance between adaptive and non-adaptive strategies (the so-called *adaptivity gap*). As the proof of this result is involved and technical, we only report the statement.

Lemma 4 (Bradac, Singla, and Zucic 2019). The following inequality holds:

$$\max_{S \in \text{EXP-SET}} \mathbb{E} \left[\max_{i \in S} \{Y_i\} \right] \geq \frac{1}{2} \cdot \mathbb{E} \left[\max_{i \in S(\pi^*)} \{Y_i\} \right].$$

We have proved in the previous Section that the family of constraints EXP-SET is a transversal matroid, so it is possible to apply the approximation guarantees of Lemma 1 for the submodular function f as in Equation (2) and the matroid constraint EXP-SET, to argue that the set \hat{S} output by continuous greedy enjoys the following property:

$$\mathbb{E} \left[\max_{i \in \hat{S}} \{Y_i\} \right] \geq \left(1 - \frac{1}{e} - \varepsilon\right) \max_{S \in \text{EXP-SET}} \mathbb{E} \left[\max_{i \in S} \{Y_i\} \right]. \tag{3}$$

Let’s recap what we achieved so far: starting from Lemma 3 and then with Lemma 4 and finally Equation (3), we proved that the expected utility of the optimal solution is “close” to the value of the submodular function f on \hat{S} . In the next step of the analysis we prove how the threshold strategy characterized by $\hat{\sigma}$ and $\hat{\tau}$ manages to achieve an expected utility that is comparable with $f(\hat{S})$. This passage is analogous to what is done in the Prophet Pandora setting in Beyhaghi and Kleinberg (2019). Once the set of boxes \hat{S} is chosen, and the compatible ordering $\hat{\sigma}$ is fixed, the crucial observation is that $f(\hat{S}) = \mathbb{E} [\max_{i \in \hat{S}} Y_i]$ is the *performance of the prophet* in the prophet inequality problem[‡], therefore the following bound holds (we add our version of the proof for completeness):

Lemma 5 (Esfandiari et al. 2019). The following inequality holds:

$$\text{ALG} \geq \frac{1}{2} \cdot \mathbb{E} \left[\max_{i \in \hat{S}} \{Y_i\} \right].$$

[‡]In the prophet inequality problem, a decision-maker observes random variables given in some order, and needs to select one of them in an immediate and irrevocable manner. The goal is to maximize the chosen value.

Proof of Lemma 5. Consider the prophet inequality problem with random variables $\{Y_i\}_{i \in \tilde{S}}$ that arrive in the given order $\hat{\sigma}$. It is well known (Kleinberg and Weinberg 2019) that the threshold algorithm with the above threshold $\hat{\tau} = 1/2 \mathbb{E} [\max_{i \in \tilde{S}} \{Y_i\}]$ achieves expected utility (with respect to the Y_i) which is at least $\hat{\tau}$. For each box i , the probability that ALG halts upon consideration of box i is exactly $\Pr(Y_i \geq \hat{\tau})$, the same as the corresponding probability with respect to the threshold algorithm for the prophet inequality problem.

Therefore, to conclude the proof of the Lemma, it is enough to show that the utility obtained by ALG from each box i conditioned on that box being considered, denoted $u(i)$, is at least as much as the conditional utility that the aforementioned prophet inequality threshold algorithm obtains from box i . Note that the latter equals $\mathbb{E} [Y_i \cdot \mathbb{I}_{\{Y_i \geq \tau\}}] = \mathbb{E} [\min\{V_i, r_i\} \cdot \mathbb{I}_{\{Y_i \geq \tau\}}]$. Therefore, to finish the proof we establish the following inequality:

$$u(i) \geq \mathbb{E} [\min\{V_i, r_i\} \cdot \mathbb{I}_{\{Y_i \geq \tau\}}] \quad (4)$$

Assume that ALG considers box i in its execution. We first consider the case $r_i < \tau$. In this case, ALG does not open box i , implying that the utility achieved from that box is $u(i) = 0$. But in this case we also have $\mathbb{I}_{\{Y_i \geq \tau\}} = 0$, implying in particular that $\mathbb{E} [\min\{V_i, r_i\} \cdot \mathbb{I}_{\{Y_i \geq \tau\}}] = 0$, and we are done.

We now consider the case $r_i \geq \tau$. Recall that in this case, upon considering box i , ALG inspects the box (incurring its cost), and stops and obtains the value V_i if $V_i \geq \tau$. Thus, in this case the utility $u(i)$ is at least

$$\begin{aligned} & \mathbb{E} [-c_i + V_i \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [-\mathbb{E} [\max\{V_i - r_i, 0\}] + V_i \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [\mathbb{E} [\min\{r_i - V_i, 0\}] + V_i \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [\min\{r_i - V_i, 0\} + V_i \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [\min\{r_i - V_i, 0\} \cdot \mathbb{I}_{\{V_i \geq \tau\}} + V_i \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [(\min\{r_i - V_i, 0\} + V_i) \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [\min\{r_i, V_i\} \cdot \mathbb{I}_{\{V_i \geq \tau\}}] \\ &= \mathbb{E} [\min\{V_i, r_i\} \cdot \mathbb{I}_{\{\min\{V_i, r_i\} \geq \tau\}}] \end{aligned}$$

The fourth and last equalities in the chain above hold since we are in the case $r_i \geq \tau$. \square

We have all the ingredients to conclude the proof of the Theorem, as chaining the results in Lemmas 3 and 4, Inequality (3) and Lemma 5 yields the desired bound (together with rescaling the precision ε used in pipage rounding). \square

Since the Pandora’s Problem with Deadlines is a special case of the Time Slots model, we have immediately the following Corollary. Note, the only difference in the algorithm is that the ordering $\hat{\sigma}$ can be easily computed starting from the set of boxes \tilde{S} by simply sorting them in increasing order of deadlines.

Corollary 1. For any constant $\varepsilon \in (0, 1)$, there exists a threshold strategy for the Pandora’s Problem with Deadlines that can be computed efficiently and provides a $(\frac{e-1}{4e} - \varepsilon)$ -approximation to the optimal expected utility.

Discussion and Further Directions

In this paper we have introduced the Pandora’s Problem with Deadlines and Time Slots, where each box is associated with a deadline and, possibly, with a subset of the time slots where the exploration is allowed. Carefully combining techniques from submodular maximization, stochastic probing and prophet inequalities we have constructed an efficient algorithm that is guaranteed to achieve at least a 0.15 fraction of the optimal utility. A natural open question is whether this approximation factor is tight for efficient algorithms or not. Since the three “approximation steps” in our approach are known in general to be tight, namely the factor $(1 - 1/e)$ for submodular maximization, the adaptivity gap of $1/2$ and the factor $1/2$ for prophet inequalities, in order to improve on the 0.15 term a new –possibly more direct– approach is needed.

Acknowledgments

Ben Berger’s work was done while being a Ph.D. student at Tel-Aviv University. Feldman and Berger received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), by an Amazon Research Award, and by the NSF-BSF (grant number 2020788), and by a grant from the Tel Aviv University Center for AI and Data Science (TAD). Federico Fusco is partially supported by ERC Advanced Grant 788893 AM-DROMA “Algorithmic and Mechanism Design Research in Online Markets”, PNRR MUR project PE0000013- FAIR”, and PNRR MUR project IR0000013-SoBigData.it. Part of the work of Federico Fusco was done while visiting Michal Feldman at Tel Aviv University. This work was also partially supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778, while Tomer was in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI) in Berkeley, California, during the Fall 2023 semester.

References

- Asadpour, A.; and Nazerzadeh, H. 2016. Maximizing stochastic monotone submodular functions. *Management Science*, 62(8): 2374–2391.
- Berger, B.; Ezra, T.; Feldman, M.; and Fusco, F. 2023. Pandora’s Problem with Combinatorial Cost. In *EC*, 273–292. ACM.
- Beyhaghi, H.; and Cai, L. 2023. Pandora’s Problem with Nonobligatory Inspection: Optimal Structure and a PTAS. In *STOC*, 803–816. ACM.
- Beyhaghi, H.; and Kleinberg, R. 2019. Pandora’s Problem with Nonobligatory Inspection. In *EC*, 131–132. ACM.
- Boodaghians, S.; Fusco, F.; Lazos, P.; and Leonardi, S. 2023. Pandora’s Box Problem with Order Constraints. *Mathematics of Operations Research*, 48(1): 498–519.
- Bradac, D.; Singla, S.; and Zuzic, G. 2019. (Near) Optimal Adaptivity Gaps for Stochastic Multi-Value Probing.

- In *APPROX-RANDOM*, volume 145 of *LIPICs*, 49:1–49:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Chawla, S.; Gergatsouli, E.; McMahan, J.; and Tzamos, C. 2023. Approximating Pandora’s Box with Correlations. In *APPROX/RANDOM*, volume 275 of *LIPICs*, 26:1–26:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Chawla, S.; Gergatsouli, E.; Teng, Y.; Tzamos, C.; and Zhang, R. 2020. Pandora’s Box with Correlations: Learning and Approximation. In *FOCS*, 1214–1225. IEEE.
- Doval, L. 2018. Whether or not to open Pandora’s box. *Journal of Economic Theory*, 175: 127–158.
- Esfandiari, H.; Hajiaghayi, M. T.; Lucier, B.; and Mitzenmacher, M. 2019. Online Pandora’s Boxes and Bandits. In *AAAI*, 1885–1892. AAAI Press.
- Fu, H.; Li, J.; and Liu, D. 2023. Pandora Box Problem with Nonobligatory Inspection: Hardness and Approximation Scheme. In *STOC*, 789–802. ACM.
- Fu, H.; Li, J.; and Xu, P. 2018. A PTAS for a Class of Stochastic Dynamic Programs. In *ICALP*, volume 107 of *LIPICs*, 56:1–56:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Gatmiry, K.; Kesselheim, T.; Singla, S.; and Wang, Y. 2022. Bandit Algorithms for Prophet Inequality and Pandora’s Box. *CoRR*, abs/2211.08586.
- Gergatsouli, E.; and Tzamos, C. 2022. Online Learning for Min Sum Set Cover and Pandora’s Box. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 7382–7403. PMLR.
- Guo, C.; Huang, Z.; Tang, Z. G.; and Zhang, X. 2021. Generalizing Complex Hypotheses on Product Distributions: Auctions, Prophet Inequalities, and Pandora’s Problem. In *COLT*, volume 134 of *Proceedings of Machine Learning Research*, 2248–2288. PMLR.
- Kleinberg, J. M.; and Kleinberg, R. 2018. Delegated Search Approximates Efficient Search. In *EC*, 287–302. ACM.
- Kleinberg, R.; and Weinberg, S. M. 2019. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113: 97–115.
- Schrijver, A.; et al. 2003. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer.
- Segev, D.; and Singla, S. 2021. Efficient Approximation Schemes for Stochastic Probing and Prophet Problems. In *EC*, 793–794. ACM.
- Singla, S. 2018. The Price of Information in Combinatorial Optimization. In *SODA*, 2523–2532. SIAM.
- Vondrák, J. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, 67–74. ACM.
- Weitzman, M. L. 1979. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641–654.